HANDBOOK OF THE ECONOMICS OF EDUCATION



Edited by

ERIC A. HANUSHEK

STEPHEN MACHIN

LUDGER WOESSMANN



Amsterdam • Boston • Heidelberg • London New York • Oxford • Paris • San Diego San Francisco • Singapore • Sydney • Tokyo North-Holland is an imprint of Elsevier



INTRODUCTION TO THE SERIES

The aim of the Handbooks in Economics series is to produce Handbooks for various branches of economics, each of which is a definitive source, reference, and teaching supplement for use by professional researchers and advanced graduate students. Each Handbook provides self-contained surveys of the current state of a branch of economics in the form of chapters prepared by leading specialists on various aspects of this branch of economics. These surveys summarize not only received results but also newer developments, from recent journal articles and discussion papers. Some original material is also included, but the main goal is to provide comprehensive and accessible surveys.

The Handbooks are intended to provide not only useful reference volumes for professional collections but also possible supplementary readings for advanced courses for graduate students in economics.

KENNETH J. ARROW and MICHAEL D. INTRILIGATOR

North-Holland is an imprint of Elsevier Radarweg 29, PO Box 211, 1000 AE Amsterdam, The Netherlands The Boulevard, Langford Lane, Kidlington, Oxford OX5 1GB, UK 225 Wyman Street, Waltham, MA 02451, USA

© 2011 Elsevier B.V. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or any information storage and retrieval system, without permission in writing from the publisher. Details on how to seek permission, further information about the Publisher's permissions policies and our arrangements with organizations such as the Copyright Clearance Center and the Copyright Licensing Agency, can be found at our website: *www.elsevier.com/permissions*.

This book and the individual contributions contained in it are protected under copyright by the Publisher (other than as may be noted herein).

Notices

Knowledge and best practice in this field are constantly changing. As new research and experience broaden our understanding, changes in research methods, professional practices, or medical treatment may become necessary.

Practitioners and researchers must always rely on their own experience and knowledge in evaluating and using any information, methods, compounds, or experiments described herein. In using such information or methods they should be mindful of their own safety and the safety of others, including parties for whom they have a professional responsibility.

To the fullest extent of the law, neither the Publisher nor the authors, contributors, or editors, assume any liability for any injury and/or damage to persons or property as a matter of products liability, negligence or otherwise, or from any use or operation of any methods, products, instructions, or ideas contained in the material herein.

British Library Cataloguing in Publication Data

A catalogue record for this book is available from the British Library

Library of Congress Cataloging-in-Publication Data

Library of Congress Control Number: 2007310346

ISBN: 978-0-444-53444-6

For information on all North-Holland publications visit our website at *books.elsevier.com*

Typeset by: diacriTech, India

Printed and bound in the United States of America 11 12 13 14 15 10 9 8 7 6 5 4 3 2 1



CONTRIBUTORS

Mathilde Almlund

Department of Economics, University of Chicago, 1126 East 59th Street, Chicago, Illinois 60637, USA.

Gary S. Becker

Department of Economics, University of Chicago, 1126 East 59th Street, Chicago, Illinois 60637, USA.

Eric Bettinger

School of Education, Stanford University, CERAS 522, 520 Galvez Mall Stanford, CA 94305, USA.

John Bound

Department of Economics, University of Michigan, Ann Arbor, MI 48109-1220, USA.

Angela L. Duckworth

Positive Psychology Center, University of Pennsylvania, 3701 Market St, Philadelphia, PA, USA.

Christian Dustmann

Economics Department, University College London, Gower Street, London WC1E6BT, UK.

Oded Galor

Department of Economics, Brown University, Providence, RI 02912, USA.

Albrecht Glitz

Department of Economics and Business, Universitat Pompeu Fabra, Ramon Trias Fargas, 25-27, 08005 Barcelona, Spain.

Gerhard Glomm

Wylie Hall, Department of Economics, Indiana University, Wylie Hall, Bloomington, IN 47405, USA.

Eric Hanushek

Hoover Institution, Stanford University, Stanford, CA 94305-6010, USA.

James J. Heckman

Department of Economics, University of Chicago, 1126 East 59th Street, Chicago, Illinois 60637, USA.

Tim Kautz

Department of Economics, University of Chicago, 1126 East 59th Street, Chicago, Illinois 60637, USA.

Edwin Leuven

CREST (ENSAE), 3, avenue Pierre Larousse, 92245 MALAKOFF Cedex, France.

Lance Lochner

Department of Economics, Social Science Centre, University of Western Ontario, London, Ontario N6A 5C2, Canada.

Stephen Machin

Department of Economics, University College London, Gower Street, London WC1E 6BT, UK.

Derek Neal

Department of Economics and The Committee on Education, University of Chicago, 1126 East 59th Street, Chicago, IL 60637, USA.

Hessel Oosterbeek

Amsterdam School of Economics, University of Amsterdam, Roetersstraat 11, Room 3.60, Building E, 1018 WB Amsterdam, The Netherlands.

B. Ravikumar

Research Department, Federal Reserve Bank of St. Louis, St. Louis, MO 63166, USA.

Ioana C. Schiopu

Department of Economics, ESADE Business School, Av. de la Torre Blanca, 59 E-08172 Sant Cugat, Barcelona, Spain.

Sarah E. Turner

Department of Economics, University of Virginia, 2015 Ivy Road, Charlottesville, VA 22904-4182, USA.

Ludger Woessmann

University of Munich, Ifo Institute for Economic Research, Poschingerstrasse 5, 81679 Munich, Germany.

REFLECTIONS ON THE ECONOMICS OF EDUCATION

Gary S. Becker

University of Chicago and the Hoover Institution

Modern research on the economics of education began in the 1950s with research by T. W. Schultz, Jacob Mincer, Sherwin Rosen, myself, and some others, although there are earlier precedents, including analyses by Adam Smith, Alfred Marshall, and Milton Friedman. This new literature treats education as an investment that has both costs and returns. The returns analyzed are principally the increase in earnings due to greater amounts of schooling. The costs included tuition, fees, and other direct expenses from schooling, and the earnings foregone by being in school rather than at work.

The research in the 1950s already considered the economics of education to be part of the general field of investments in human capital, where individuals and societies acquired knowledge, skills, and information by spending money and time on schooling, job training, health, and other investments. The concept of "human capital" was even more controversial than that of the economics of education. It was considered demeaning to individuals because it was claimed to involve treating people as machines or houses rather than as real individuals with emotions and feelings.

I pondered for a while this controversy over the term human capital in deciding the name of the book that was to be published in 1964. I felt that human capital is the best description of what the book analyzes, yet I was mindful that the term would put off some potential readers. I eventually decided to bite the bullet and did indeed name the book *Human Capital*. However, I hedged my risk by also including a long subtitle *A Theoretical and Empirical Analysis, with Special Reference to Education* (Becker 1993).

Eventually, the term human capital became fully accepted not only by academics but also by the media and politicians. I remember being both pleased and surprised when the cover page of a *Business Week* issue in 1988 featured the term "human capital" in bold letters as it described the focus of its principle story for that week. My book soon began to be called simply *Human Capital*, with no reference at all to the subtitle.

My first inkling on the politics of the attitudes by educators to research on the economics of education came when I was invited rather shortly after publication of *Human Capital* to speak about my work before a meeting of the superintendents of large school districts in the United States. During the Q&A session after presentation of my research results on the economic effects of education, one superintendent after another got up to denounce this work as being too materialistic and ignore crucial cultural and humanistic sides of education. Then, Benjamin Willis, a long time and controversial leader of the Chicago school system, got up and chastised his fellow superintendents. He said that the work indicting that education increased earnings would give the superintendents another weapon in justifying their request for more funding from school boards and legislatures that controlled spending on schools. What is better, he argued, than appealing to the practical interests of parents and others? After his speech, one after another of the superintendents began to see the great virtue of work on the earnings' effects of education!

During the past 50 years, studies on the economics of education and other human capital have boomed at a rate I never in my wildest dreams anticipated when I published *Human Capital* in 1964. Economists, sociologists, and others using concepts from human capital theory have considered education from preschool levels to graduate degree levels and have calculated the returns to different skills learned in school. Researchers have also analyzed the education of both boys and girls, and men and women, and they have considered the effects of different incentive structures within schools and the degree of competition among schools for students, teachers, and funds. Important research has also treated expenditures on health as investments in a particular and crucial form of human capital. They have analyzed investments by firms and by workers on the job, the amount and effects of education provided online, and the variety of courses offered in adult education programs. Any reasonably complete bibliography on human capital and the economics of education would cover many thousands of articles and books.

I will concentrate my remaining comments on an important development that occurred almost entirely after the early generation of research on the economics of education; namely, the link between family economics and education economics. My *Human Capital* book, and the other initial human capital literature, mentioned the family as influencing investments in education and other human capital, but paid little systematic attention to the links between these investments and the family.

This began to change with research that related the labor force participation of married women to the earnings of their husbands, fertility levels, and other family variables (see Mincer (1962) and the earlier book by Long (1958)). Before long, research on labor force participation began to analyze why the average woman earns considerably less than the average man. In addition to market discrimination against women, human capital–based analysis argued that women earned less because they dropped out of college to get married—in 1970, women in the United States received only a little more than 40% of the 4-year college degrees. Furthermore, women had less incentive to invest on the job and in other skills that increased earnings because their labor force participation rates and hours worked were so much below those of men.

It had long been known that children from poorer and less educated families were much less likely to complete high school and college than were children from better-off families. Further research indicted not only that poorer families could less afford to finance their children's education but also that less educated families provided much smaller early childhood home investments that involved reading to children and other ways to increase the learning and noncognitive skills of children. This meant that children from poorer families were by age six already far behind than children from more educated families in their reading and other skills, and also in their noncognitive development, such as ability to pay attention to what teachers were saying.

Other research showed that families with only one parent present, usually the mother, made fewer investments of time and money in their children, and that divorced fathers (not divorced mothers) invest less in their children before as well as after their divorce. In addition, the degree of intergenerational mobility in occupations and earnings between parents and children has been shown to be greatly affected by the degree of transmission of education from parents to children that families with larger number of children generally invest less in the education and other human capital of each child, and that college-educated parents, especially college-educated mothers, have many fewer children than mothers who drop out of high school. Moreover, educated men and women are more likely to stay married, and that the strong negative effect of a women's education on her likelihood of marrying that prevailed in the past has now largely disappeared in the United States and other rich countries.

Higher education has boomed throughout the world during the past three decades in many poorer and in all rich nations. An important part of the explanation for this development is that new technologies, such as computers and the Internet, increased the demand for persons with college education since college graduates more easily utilize and adapt to these technologies. Other important developments explaining the greater incentives to get a higher education are the shift to high-skilled services, such as the education and health sectors, and away from manufacturing and increased globalization that helped spread the demand for these new technologies throughout the world.

Higher education of women has especially been booming in recent years so that in almost all rich countries, and in many not so rich countries, women are now much more likely to graduate from college than are men. In the United States, for example, the fraction of all 4-year college degrees obtained by women increased from only 43% in 1970 to almost 60% in 2010. Women in the United States now receive about 60% of Masters' degrees and slightly more than half of all PhD degrees. The rapid growth in higher education of women is related to the worldwide decrease in fertility during the past few decades, great improvements in contraceptives, and delays in age at marriage. Also important is that the median female is a better student than the median male, partly because noncognitive skills are typically both higher and less variable for women than for men (see Becker, Hubbard, and Murphy (2010)).

Although early research on the links between family and education mainly dealt with the United States and other developed countries, the past couple of decades have paid much more attention to this connection also in poor and developing countries. The contribution of even young children to family resources explains why many families in poor countries take their children out of elementary schools, and especially out of secondary schools so that the children can work and thereby contribute to family income. The interaction between the quantity and quality of children also implies that families in poorer countries invest less in each of their children because they have many more children than do families in richer countries.

Research on human capital in general, and the economics of education in particular, has remained an exciting field for over half a century in good part because the issues discussed are so vital to wellbeing and so important in public policy discussions. However, also important in maintaining this vitality is that the field has been mainly neither theoretical nor empirical. Instead, human capital analysis from the beginning involved a close dialog between theory and empirical work. The theory suggested what to analyze empirically, although empirical findings caused modifications and extensions of the theoretical approach. This dynamic interaction between theory and data analysis helped evolve economic research on human capital and on education in new and interesting directions.

The increased importance of knowledge, skills, and information in the modern economy means that human capital and education will be even more important in the future than it has been during the past 50 years. As long as studies of education and other human capital continue to have important feedbacks between the theory and empirical findings, I anticipate that human capital analysis will stay relevant and important and will generate great excitement. The articles in these handbook volumes convey some of the excitement that continues in research on the economics of education.

REFERENCES

- Becker, G.S., 1993. Human Capital: A Theoretical and Empirical Analysis, with Special Reference to Education, third ed. University of Chicago Press, Chicago Print.
- Becker, G.S., Hubbard, W.H., Murphy, K.M., 2010. Explaining The Worldwide Boom in Higher Education of Women. J. Human Capital 4 (3), pp. 203–241, Print.
- Long, C.D., 1958. The Labor Force Under Changing Income and Employment. Princeton University Press, Princeton, Print.
- Mincer, J., 1962. Labor Force Participation of Married Women: A Study of Labor Supply. Aspects of Labor Economics. Universities-National Bureau Committee for Economic Research, No. 14. Princeton University Press (for NBER), Princeton, Print.
- Nussbaum, B., Bernstein, A., Ehrlich, E., Garland, S.B., Pennar, K., 1988. Needed: Human Capital. Business Week (19 September 1988), pp. 100–103, Print.

INTRODUCTION TO HANDBOOK

The economics of education has flourished as a research field since the publication of Volumes I and II of this Handbook. There has been a big upsurge in new research by economists on education and education policy. Although this study has predominantly had an empirical leaning, the chapters of this volume make clear that theory also plays a continuing role in the development of the field.

This economic research continues to be highly practical, with an explicit aim to understand better how education is acquired, how it affects economic and social outcomes of interest, and how it can inform public policy. The chapters in this new volume (Volume IV), and in companion Volume III that has been released recently, focus upon this new research and its grounding with past economics of education research.

There are several reasons for the explosion of research in this area. First is the very significant improvement in data availability and quality. There are now more and better data to address core questions in the economics of education. This includes international data on test scores, high-quality administrative data, and rich register data in the Scandinavian countries. A second, related, reason for the upsurge of work, including publishing in the top journals of the economics profession, is the expansion of research separately accumulating in many countries around the world. The coverage of the new volumes—in terms of choice of topics, authorship, and editorship—is meant to reflect this globalization of research in the economics of education. A third aspect is the use of new methodological approaches that overlap with significant developments that have been made in other areas of economics. Finally, there has been a heightened policy relevance related to economics of education research. Many governments demand more evidence-based policy, and this has been particularly true in terms of education.

All of these have resulted in the economics of education being a thriving and burgeoning specialty within economics. The topics covered by the chapters of the current volume step into some of the most obvious gaps that have become evident with the newly emerging research. Although there are antecedents to the work in each of the chapters, recent research has taken the ideas into new and productive areas. This is mirrored in the introductory chapter by Gary Becker, one of the founding fathers of the economics of education as we know it today, who received the Nobel Prize in Economics in 1992 for his groundbreaking research on investments in human capital, among other things. Becker reflects on the ground covered by the economics of education over the past 50 years, with its exciting topics, its astonishing breadth of thousands of works, and its increasing relevance in times when knowledge and skills form the centerpiece of modern economies. He places a particular focus on recent developments in the research on the links between family and education, which also emerges in many chapters of the new couple of volumes of this Handbook. Becker also stresses that the economics of education has throughout been characterized by a dynamic interaction between theory and empirical work, a topic clearly reflected in the selection of chapters in this volume.

The first five specific chapters of this volume cover topics on the role of education for different outcomes. These chapters respectively focus on personality traits, nonproduction benefits, possible mismatch in the labor market, migration, and comparative development. The next two chapters deal with two important institutional features of school systems, namely performance pay and vouchers. The final two chapters address outcomes in higher education and the political economy of education funding, respectively.

We are grateful to all authors contributing to the new volume, as we appreciate (from own experience) how much of a task it is to produce a Handbook chapter that both covers the existing literature and provides ideas that lead into the future. Their expertise, enthusiasm, and hard work are highly appreciated. We also gratefully acknowledge the professional support in the *Handbooks in Economics* series, especially by the general editors Kenneth Arrow and Michael Intriligator and by Scott Bentley, Kathleen Paoni, Heather Tighe, Stacey Walker, and others at Elsevier. We also thank CESifo, which provided financial support and facilities to hold the inaugural meeting of the CESifo research network's Economics of Education area in Munich in September 2009 where initial drafts of the chapters of Volumes III and IV were presented and discussed.

Education is widely recognized as an important determinant of a wide range of economic and social outcomes. Through use of rich data and study of issues of high contemporary policy relevance, the economics of education is one of the primary areas of research attraction across the economics profession, appealing to new PhD students and experienced researchers alike. The significant bodies of research studied in this volume suggest that it is highly likely that this volume will not be the last in this series, as study of education acquisition and its economic and social impact will undoubtedly remain a fertile research ground for the foreseeable future.

> Eric A. Hanushek Stephen Machin Ludger Woessmann February 2011



Personality Psychology and Economics¹

Mathilde Almlund^{*}, Angela Lee Duckworth^{**}, James Heckman^{*,***}, and Tim Kautz^{*}

*Department of Economics, University of Chicago, Chicago, IL 60637

**University of Pennsylvania, Philadelphia, Pennsylvania, PA 19104

***University College Dublin, American Bar Foundation

Contents

1.	Introduction	4
2.	Personality and Personality Traits: Definitions and a Brief History of Personality Psychology	7
	2.1. A Brief History of Personality Psychology	10
	2.1.1 The Pioneers of Personality Psychology	10
	2.1.2 The Person-Situation Debate, Its Lingering Influence in Economics, and the Subsequent	
	Flourishing of Personality Psychology	12
3.	Conceptualizing Personality and Personality Traits within Economic Models	14
	3.1. An Approach Based on Comparative Advantage	15
	3.2. Allowing for Multitasking	17
	3.3. Identifying Personality Traits	18
	3.4. Extensions of the Roy Model	20
	3.5. Adding Preferences and Goals	21
	3.6. Adding Learning and Uncertainty	22
	3.7. Definition of Personality within an Economic Model	23
	3.8. Life Cycle Dynamics	27
	3.9. Relationship of the Model in This Section to Existing Models in Personality Psychology	28

This research was supported by grants from NIH R01-HD054702, R01-HD065072, and K01-AG033182; the University of Chicago; the Institute for New Economic Thinking (INET); A New Science of Virtues: A Project of the University of Chicago; the American Bar Foundation; a conference series from the Spencer Foundation; the JB & MK Pritzker Family Foundation; the Buffett Early Childhood Fund; the Geary Institute, University College Dublin, Ireland; and an anonymous foundation. The opinions expressed in this report are those of the authors and do not necessarily reflect the views of any of the funders. Amanda Agan and Pietro Biroli are major contributors to this essay through their surveys of the effect of personality on crime (presented in Web Appendix A7.2) and health (presented in Web Appendix A7.1), respectively. We are grateful to Pia Pinger for her analyses of the German Socio-Economic Panel (GSOEP) survey data. We have benefited from comments received from Amanda Agan, Dan Benjamin, Pietro Biroli, Dan Black, Daniel Cervone, Deborah Cobb-Clark, Flavio Cunha, Kathleen Danna, Thomas Dohmen, Steven Durlauf, Joel Han, Moshe Hoffman, John Eric Humphries, Miriam Gensowski, Bob Krueger, Jongwook Lee, Xiliang Lin, Dan McAdams, Terrance Oey, Lawrence Pervin, Pia Pinger, Armin Rick, Brent Roberts, Molly Schnell, Bas ter Weel, and Willem van Vliet. We also benefited from a workshop at the University of Illinois, Department of Psychology, on an early draft of this paper and presentations of portions of this paper at the Spencer/INET workshop at the University of Chicago, December 10-11, 2010, and at the IZA workshop on Cognitive and Non-Cognitive Skills, January 25-27, 2011, Bonn, Germany. Additional material that supplements the text is presented in a Web Appendix (http://jenni.uchicago.edu/personality_economics/). Parts of this paper build on an earlier study by Borghans, Duckworth, Heckman, and ter Weel (2008).

1

4.	Measuring Personality	30
	4.1. Linear Factor Models	30
	4.2. Discriminant and Convergent Validity	31
	4.3. Predictive Validity	33
	4.4. Faking	35
	4.5. The Causal Status of Latent Variables	36
5.	Implementing the Measurement Systems	36
	5.1. Cognition	37
	5.1.1 Fluid versus Crystallized Intelligence	39
	5.1.2 Predictive Validity of Tests of Cognition	40
	5.2. Personality Traits	43
	5.3. Operationalizing the Concepts	47
	5.3.1 IQ Tests	47
	5.4. Personality Constructs	48
	5.5. Alternatives to the Big Five	50
	5.5.1 Self-Esteem and Locus of Control Are Related to Big Five Emotional Stability	52
	5.5.2 Relating the Big Five to Measures of Psychopathology	53
	5.6. IQ and Achievement Test Scores Reflect Incentives and Capture Both	5.4
	Cognitive and Personality Traits	54
~	5.7. The Evidence on the Situational Specificity Hypothesis	63
0.	Personality and Preference Parameters	05
	6.1. Evidence on Preference Parameters and Corresponding Personality Measures	CO 73
	6.2. Mapping Frederences into Fersonality 6.3. Do Mossured Professores Parameters Predict Pool World Rehavior?	73
	64 Integrating Traits into Economic Models	74
	641 Traits as Constraints	75
	6.4.2 Traits as Preferences	70
	65 Summary of Section 6 and Some Concluding Thoughts	84
7.	The Predictive Power of Personality Traits	84
	7.1. An Overview of the Main Findings	89
	7.2. Educational Attainment and Achievement	90
	7.2.1 Educational Attainment	90
	7.2.2 Course Grades	99
	7.2.3 Standardized Achievement Test Scores	101
	7.2.4 Where Course Grades and Standardized Achievement Test Scores Diverge	103
	7.3. Labor Market Outcomes	106
	7.4. Personality and Health	112
	7.5. Crime	116
8.	Stability and Change in Personality Traits and Preferences	117
	8.1. Broad Evidence on Changes in Traits over the Life Cycle	118
	8.2. Evidence on Ontogenic and Sociogenic Change	122
	8.3. External Changes to Biology	126
	8.3.1 Brain Lesion Studies	126
	8.3.2 Chemical and Laboratory Interventions	127

	8.4. The Evidence on the Causal Effects of Parental Investment, Education, and Interventions	128	
	8.4.1 Evidence of Change in Traits from Other Studies of Parental Investment	131	
	8.4.2 The Effects of Schooling on Cognitive and Personality Traits	132	
	8.4.3 Evidence from Interventions	134	
	8.4.4 Evidence from Psychotherapy	139	
	8.5. Stability of Economic Preference Parameters	139	
	8.6. Summary of Section 8	150	
9.	Summary and Conclusions	150	
Re	References		

Abstract

This chapter explores the power of personality traits both as predictors and as causes of academic and economic success, health, and criminal activity. Measured personality is interpreted as a construct derived from an economic model of preferences, constraints, and information. Evidence is reviewed about the "situational specificity" of personality traits and preferences. An extreme version of the situationist view claims that there are no stable personality traits or preference parameters that persons carry across different situations. Those who hold this view claim that personality psychology has little relevance for economics.

The biological and evolutionary origins of personality traits are explored. Personality measurement systems and relationships among the measures used by psychologists are examined. The predictive power of personality measures is compared with the predictive power of measures of cognition captured by IQ and achievement tests. For many outcomes, personality measures are just as predictive as cognitive measures, even after controlling for family background and cognition. Moreover, standard measures of cognition are heavily influenced by personality traits and incentives.

Measured personality traits are positively correlated over the life cycle. However, they are not fixed and can be altered by experience and investment. Intervention studies, along with studies in biology and neuroscience, establish a causal basis for the observed effect of personality traits on economic and social outcomes. Personality traits are more malleable over the life cycle compared with cognition, which becomes highly rank stable around age 10. Interventions that change personality are promising avenues for addressing poverty and disadvantage.

Keywords

Personality Behavioral Economics Cognitive Traits Wages Economic Success Human Development Person-situation Debate

1. INTRODUCTION

The power of cognitive ability in predicting social and economic success is well documented.² Economists, psychologists, and sociologists now actively examine determinants of social and economic success beyond those captured by cognitive ability.³ However, a substantial imbalance remains in the scholarly and policy literatures in the emphasis placed on cognitive ability compared to other traits. This chapter aims to correct this imbalance. It considers how personality psychology informs economics and how economics can inform personality psychology.

A recent analysis of the Perry Preschool Program shows that traits other than those measured by IQ and achievement tests causally determine life outcomes.⁴ This experimental intervention enriched the early social and emotional environments of disadvantaged children of ages 3 and 4 with subnormal IQs. It primarily focused on fostering the ability of participants to plan tasks, execute their plans, and review their work in social groups.⁵ In addition, it taught reading and math skills, although this was not its main focus. Both treatment and control group members were followed into their 40s.⁶

Figure 1.1 shows that, by age 10, the mean IQs of the treatment group and the control group were the same. Many critics of early childhood programs seize on this and related evidence to dismiss the value of early intervention studies.⁷ Yet on a variety of measures of socioeconomic achievement, the treatment group was far more successful than the control group.⁸ The annual rate of return to the Perry Program was in the range 6–10% for boys and girls separately.⁹ These rates of return are statistically significant and above the returns to the US stock market over the postwar period.¹⁰ The intervention changed something other than IQ, which produced strong treatment effects. Heckman, Malofeeva, Pinto, and Savelyev (first draft 2008, revised 2011) show that the personality

² See, e.g., the studies cited in Becker (1964) and the discussion of ability bias in Griliches (1977).

³ See Bowles, Gintis, and Osborne (2001a) and Borghans, Duckworth, Heckman, and ter Weel (2008) for reviews of the literature in economics. Marxist economists and sociologists (e.g., Bowles and Gintis (1976) and Mueser (1979), respectively) pioneered the analysis of the impact of personality on earnings. Mueller and Plug (2006) estimate empirical relationships between personality traits and earnings, schooling and occupational attainment. Hartog (1980, 2001) relates the Big Five personality factors to earnings. Van Praag (1985) draws on the psychology literature to analyze economic preferences. Van Praag and Van Weeren (1988) and Borghans, Duckworth, Heckman, and ter Weel (2008) link economics with psychology.

⁴ We draw on the research of Heckman, Malofeeva, Pinto, and Savelyev (first draft 2008, revised 2011). See Weikart, Epstein, Schweinhart, and Bond (1978); Sylva (1997); Schweinhart et al. (2005); and Heckman, Moon, Pinto, Savelyev, and Yavitz (2010a) for descriptions of the Perry program.

⁵ Sylva (1997) shows that the Perry Program has important features that are shared with programs designed to foster selfcontrol in children, e.g., *Tools of the Mind* (Bodrova and Leong, 2001).

⁶ Plans are underway to follow the Perry sample through age 50.

⁷ See the Westinghouse study of Head Start (Project Head Start, 1969).

⁸ See Heckman, Malofeeva, Pinto, and Savelyev (first draft 2008, revised 2011) and Heckman, Moon, Pinto, Savelyev, and Yavitz (2010a).

⁹ See Heckman, Moon, Pinto, Savelyev, and Yavitz (2010b).

¹⁰ See DeLong and Magin (2009) for estimates of the return on equity.



Figure 1.1 Perry Preschool Program: IQ, by Age and Treatment Group.

Notes: IQ measured on the Stanford–Binet Intelligence Scale (Terman and Merrill, 1960). The test was administered at program entry and at each of the ages indicated. Source: Cunha, Heckman, Lochner, and Masterov (2006) and Heckman and Masterov (2007) based on data provided by the High Scope Foundation.

traits of the participants were beneficially improved in a lasting way.¹¹ This chapter is about those traits.

Personality psychologists mainly focus on empirical associations between their measures of personality traits and a variety of life outcomes. Yet for policy purposes, it is important to know mechanisms of *causation* to explore the viability of alternative policies.¹² We use economic theory to formalize the insights of personality psychology and to craft models that are useful for exploring the causal mechanisms that are needed for policy analysis.

We interpret personality as a *strategy function* for responding to life situations. Personality traits, along with other influences, produce measured personality as the output of personality strategy functions. We discuss how psychologists use measurements of the performance of persons on tasks or in taking actions to identify personality traits and cognitive traits. We discuss fundamental identification problems that arise in applying their procedures to infer traits.

Many economists, especially behavioral economists, are not convinced about the predictive validity, stability, or causal status of economic preference parameters or personality traits. They believe, instead, that the constraints and incentives in situations

¹¹ We discuss this evidence in Section 8. The traits changed were related to self-control and social behavior. Participants of both genders had better "externalizing behavior," while for girls there was also improvement in Openness to Experience. See Heckman, Malofeeva, Pinto, and Savelyev (first draft 2008, revised 2011). Duncan and Magnuson (2010) offer a different interpretation of the traits changed by the Perry experiment. But both analyses agree that it was not a boost in IQ that improved the life outcomes of Perry treatment group members.

¹² See Heckman (2008a).

almost entirely determine behavior.¹³ This once popular, extreme situationist view is no longer generally accepted in psychology. Most psychologists now accept the notion of a stable personality as defined in this chapter.¹⁴ Measured personality exhibits both stability and variation across situations.¹⁵

Although personality traits are not merely situation-driven ephemera, they are also not set in stone. We present evidence that both cognitive and personality traits evolve over the life cycle, but at different rates at different stages. Recently developed economic models of parental and environmental investment in children help to explain the evolution of these traits.

This chapter addresses the following specific questions, which we pose here and answer in the concluding section:

- **1.** How can we fit psychological constructs of personality into an economic framework? Can conventional models of preferences in economics characterize the main theories in personality psychology?
- 2. What are the main measurement systems used in psychology for representing personality and personality traits, and how are they validated? How are different systems related to each other? What is the relationship between standard measures of personality and measures of psychopathology and child temperament?
- **3.** What is the relationship between economic preference parameters and psychological measurements?
- **4.** How stable across situations and over the life cycle are preference parameters and personality traits?
- 5. What is the evidence on the predictive power of cognitive and personality traits?
- 6. What is the evidence on the causal power of personality on behavioral outcomes?
- **7.** Can personality be altered across the life cycle? Are interventions that change personality traits likely fruitful avenues for policy?
- **8.** Do the findings from psychology suggest that conventional economic theory should be enriched?

This chapter is organized as follows. Section 2 presents a definition of personality that captures central ideas in the literature on personality psychology. It also presents a brief history of personality psychology and the person-situation debate that paralyzed the field for 20 years and that still influences behavioral economics. Section 3 defines measured personality as a response function using an economic model of preferences, expectations, and constraints. Our model distinguishes measured personality from personality traits. We interpret personality as a response function mapping variables

¹³ See Thaler (2008) for an example of this point of view.

¹⁴ See, e.g., Mischel and Shoda (1995, 2008).

¹⁵ McAdams (2006, p. XVIII); Funder (2009); Mischel (2009); Roberts (2007, 2009); and Revelle, Wilt, and Condon (2011) discuss the stability question.

that characterize traits and situations to manifest (measured) personality. Our definition formalizes various definitions of personality used in the literature on personality psychology and facilitates the analysis of personality using the tools of economics. We sketch a dynamic model of trait formation.

Section 4 discusses alternative criteria that psychologists use to define traits. It examines the strengths and limitations of each approach. We link our abstract definition to linear factor models that are commonly used to identify personality and cognitive traits.

Section 5 presents the main systems used to measure personality and cognition and discusses the relationship among the systems. We illustrate a nonidentification result developed in Section 3 by showing how scores on IQ tests are greatly affected by incentives and context. We present additional evidence showing that the scores on achievement tests depend on cognitive and personality measurements, with a substantial predictive role for personality measures. Measures of "IQ" commonly used in economics and social science conflate measures of cognition and personality.

Section 6 discusses economic preferences and examines the evidence relating economic preference parameters to psychological parameters. Section 7 surveys the evidence on the predictive validity of personality measures for education, crime, health, and labor market outcomes. The material presented in the main text summarizes a large and growing empirical literature. A Web Appendix presents additional detail on the literature relating cognition and personality in each of these areas of economic and social life.¹⁶

Section 8 presents evidence on the *causal* impact of personality on outcomes and evidence on the stability and malleability of personality traits and preferences. We extend the theoretical framework for trait formation introduced in Section 3 and discuss a corresponding measurement system. We discuss the evidence from intervention studies. Section 9 concludes with provisional answers to the eight questions.

2. PERSONALITY AND PERSONALITY TRAITS: DEFINITIONS AND A BRIEF HISTORY OF PERSONALITY PSYCHOLOGY

Personality psychology attempts to describe the whole person.¹⁷ It considers both universal traits and individual differences. It examines the ways in which people are unique. As a sign of its breadth, personality psychology considers cognitive functioning as one aspect of personality.

¹⁶ The Web Appendix can be found online at http://jenni.uchicago.edu/personality_economics/. Amanda Agan and Pietro Biroli are authors of some of these surveys.

¹⁷ Cervone and Pervin (2009) provide a clear introduction to personality psychology.

In considering the content of personality psychology, it is helpful to distinguish *personality traits, personality as a response function*, and *measured personality*. Personality is a response function that maps personality traits to measured (manifest) personality.

One leading personality psychologist defines *personality traits* in the following way:

"Personality traits are the relatively enduring patterns of thoughts, feelings, and behaviors that reflect the tendency to respond in certain ways under certain circumstances."

(Roberts, 2009, p. 140)

This definition, or closely related versions, is used throughout the personality psychology literature.¹⁸ We formalize these notions in Section 3.

Roberts' definition of personality traits refers to the stability of certain patterns of behavior—actions or responses to situations that people take, including patterns of thoughts or feelings. Perceptions, expectations of future events, and preferences may shape behavior, feelings, and thoughts. In this way, cognitive activities help to determine measured personality.

There are many different models of personality.¹⁹ A prototypical model that captures many features of a wide class of models in personality psychology is one due to Roberts (2006). He presents the schematic displayed in Fig. 1.2 to relate personality traits to behavior.²⁰ He distinguishes mental abilities from personality traits (the items in the boxes will be discussed in later sections of this chapter). These, along with preferences (motives, interests, and values) and narratives (the stories people tell themselves in organizing their lives and making meanings of them), shape one's identity and reputation, including the views of the person by others and the person's perception of how others perceive him or her. Identity and reputation in turn shape the roles of individuals in the economy and the society and the larger culture to which they belong. *Personality* is the *system* of relation-ships that map traits and other determinants of behavior into measured actions.

In Roberts' vision of personality, feedback processes operate among all components of Fig. 1.2. Thus, his broad conception of personality includes the possibility that identity shapes traits and abilities, perhaps through a mechanism such as epigenetics, in which environment affects gene expression.²¹ Measured personality results from interactions

¹⁸ However, some personality psychologists use this or a very similar definition to define *personality* and not personality traits. Thus, Cervone and Pervin (2009) define personality as "...psychological qualities that contribute to an individual's enduring and distinctive patterns of thinking, feeling and behaving" (p. 8). Another definition in a graduate text on personality by McAdams emphasizes context more strongly: "Personality is a patterning of dispositional traits, characteristic adaptations, and integrative life stories set in culture and shaped by human nature." (McAdams, 2006). In this chapter, we define personality as a property of a system of equations, and measured personality is the output of those equations.

¹⁹ See the models in John, Robins, and Pervin (2008).

²⁰ Graphical models like Fig 1.2 are the rule in personality psychology. Explicit formal models are rare. Section 3 presents a formal model.

²¹ See, e.g., Rutter (2006a).

9



Figure 1.2 Roberts' Model of Personality as the Output of a System. *Source: Roberts (2006).*

among components of the system. Personality traits are one determinant of personality and need to be carefully distinguished from the full expression of personality, which is generated by traits interacting with other factors. Personality is seen as a system of behaviors, thoughts, and feelings that emerge from the interacting components.

In Section 3, we formalize aspects of Roberts' framework for personality within an economic model of production, choice, and information. Figure 1.2 presages our discussion of a basic identification problem discussed in Section 3. Measurements and behaviors that arise from responses to incentives and interactions with culture are used to infer personality traits and abilities. Personality traits and cognitive abilities, along with the other "units of analysis" in Fig. 1.2, produce the observed behaviors that are used to infer the generating traits. To infer traits from behaviors requires "parsing out" or standardizing for all of the other contributing factors that produce the observed behavior—a challenging task. The inability to parse and localize behaviors that depend on a single trait or ability leads to a fundamental identification problem. Behavior depends on incentives created by situations. Accurately measuring personality traits requires standardizing for the situation.

2.1. A Brief History of Personality Psychology²²

Interest in how individual behavior differs in common situations is as old as human history. The importance of personality traits for determining educational outcomes was recognized by the creators of the first IQ tests. Alfred Binet, architect of the first modern intelligence test that became the Stanford–Binet IQ test, noted that performance in school

"... admits of other things than intelligence; to succeed in his studies, one must have qualities which depend on attention, will, and character; for example a certain docility, a regularity of habits, and especially continuity of effort. A child, even if intelligent, will learn little in class if he never listens, if he spends his time in playing tricks, in giggling, is playing truant."

(Binet and Simon, 1916, p. 254)

At about the same time that Binet was writing, Charles Spearman, best known for his work on "g"—a unitary factor that is claimed to capture the structure of intelligence—along with his student, Edward Webb, undertook studies of "character" because of "the urgency of its practical application to all the business of life" (Webb, 1915, p. 1). Spearman and Webb concluded that many positive aspects of character shared a relation to what modern personality psychologists term "Conscientiousness."²³ This general factor, which Spearman and Webb chose to call "persistence of motives," meaning "consistency of action resulting from deliberate volition, or will," was distinct from a general intelligence factor (Webb, 1915, p. 60).²⁴

Arthur Jensen, an intellectual heir of Spearman, who is widely regarded as a proponent of g as an explanatory factor of success and failure in many domains of life, writes

"What are the chief personality traits which, interacting with g relate to individual differences in achievement and vocational success? The most universal personality trait is conscientiousness, that is, being responsible, dependable, caring, organized and persistent."

(Jensen, 1998, p. 575)

2.1.1 The Pioneers of Personality Psychology

Over the past century, interest in personality among psychologists has fluctuated dramatically. During the first half of the twentieth century, many of the most prominent psychologists (e.g., Gordon Allport, Raymond Cattell, Hans Eysenck, Charles Spearman, Lewis Terman) were vigorously engaged in the study of individual differences in behaviors and traits. Psychologists studied personality traits along with intelligence, interests, and motivation and measured differences and similarities within and across individuals.

²² See Revelle, Wilt, and Condon (2011) for an informative history of personality psychology.

²³ Here and elsewhere through this essay, we capitalize personality traits.

²⁴ Many other psychologists who developed and promoted IQ tests expressed similar sentiments. See the Web Appendix Section A2.1.

A systematic approach to the study of personality was conceived by early psychologists who believed that the most important dimensions on which human beings differed would be captured in natural language. These personality pioneers extracted words from the English dictionary that characterized individual differences between people (e.g., irritable, proud), after eliminating synonyms and words not associated with traits. They designed and administered studies of trait inventories to large samples of individuals and applied the same factor analytic methods developed by Galton, Spearman, Binet, Pearson, Cattell, and Thorndike to these assessments in order to isolate g to identify the structure of cognitive abilities.

The fruits of several decades of research in this tradition beginning in the 1970s have produced a widely (but not universally) shared taxonomy of traits, known as the Big Five, that is arrived at through factor analysis of observer and self-reports of behaviors.²⁵ The Big Five posits a hierarchical organization for personality traits, with five factors at the highest level and progressively more narrowly defined traits (or facets) at lower levels.

Table 1.1 presents the Big Five traits. They are Openness to Experience (also called Intellect or Culture), Conscientiousness, Extraversion, Agreeableness, and Neuroticism (also called Emotional Stability).²⁶ The Big Five factors represent personality traits at the broadest level of abstraction. They summarize a large number of distinct, more specific, personality facets.

The Big Five traits are defined without reference to any context (i.e., situation). This practice leads to an identification problem that we discuss in Section 3. The behaviors

Trait	Definition of Trait*
I. Openness to Experience (Intellect)	The tendency to be open to new aesthetic, cultural, or intellectual experiences.
II. Conscientiousness	The tendency to be organized, responsible, and hardworking.
III. Extraversion	An orientation of one's interests and energies toward the outer world of people and things rather than the inner world of subjective experience; characterized by positive affect and sociability.
IV. Agreeableness	The tendency to act in a cooperative, unselfish manner.
V. Neuroticism (Emotional Stability)	Neuroticism is a chronic level of emotional instability and proneness to psychological distress. Emotional stability is predictability and consistency in emotional reactions, with absence of rapid mood changes.

Table 1.1 The Big Five Traits

*From the American Psychological Association Dictionary (2007).

²⁵ Goldberg (1993), Barenbaum and Winter (2008), John and Srivastava (1999), and Krueger and Johnson (2008) discuss the Big Five.

²⁶ The acronym OCEAN is sometimes used to summarize these traits.

used to measure the traits are also determined by factors other than the Big Five traits. John (1990), Goldberg (1993), and Costa and McCrae (1992a) present evidence that most of the variables used to assess personality traits in academic research in the field of personality psychology can be mapped into one or more of the dimensions of the Big Five. They argue that the Big Five are the longitude and latitude of personality traits, by which all more narrowly defined traits may be categorized (see also Costa and McCrae, 1992a). We discuss the Big Five further in Section 5, where we also consider alternative measurement systems.

2.1.2 The Person-Situation Debate, Its Lingering Influence in Economics, and the Subsequent Flourishing of Personality Psychology

In 1968, Walter Mischel published a monograph entitled *Personality and Assessment*, challenging the most important theoretical assumptions and empirical findings of personality psychology. An acrimonious "person-situation" debate ensued, which pitted those who favored situational factors as explaining behavior against those who considered personality traits as more consequential. During this time, considered by many to be a dark age in the history of personality psychology, the general zeitgeist favored experimental social psychological approaches that focused on the importance of the situation compared to the individual traits featured in personality psychology.

Mischel noted that correlations between behavioral task measures of personality and questionnaire measures seldom, if ever, exceeded 0.3.^{27,28} The implication of such within-individual behavioral heterogeneity suggested to Mischel that "the behaviors which are often construed as stable personality trait indicators are highly specific and depend on the details of the evoking situations and the response mode employed to measure them" (p. 37). Mischel wrote

"... with the possible exception of intelligence, highly generalized behavioral consistencies have not been demonstrated, and the concept of personality traits as broad dispositions is thus untenable."

(Mischel, 1968, p. 146)

Mischel went on to write that global (i.e., domain-general) traits (e.g., "impulsive," "confident") measured in one situation did not predict future behavior and outcomes in other situations. His view was that global traits, in attempting to summarize behavioral dispositions without regard to situational contingencies, were "excessively crude, gross units

²⁷ There is great irony that none of the correlations of cognitive measures with outcomes that are reported in Table A1 in the Web Appendix are as high as 0.3, but no one questions the power of cognition in predicting outcomes in social life. Few studies in social psychology show correlations as high as 0.2 (Richard, Bond, and Stokes-Zoota, 2003).

²⁸ Psychologists often work with standardized variables (variables normalized by standard deviations). They report correlations between standardized variables as "effect sizes."

to encompass adequately the extraordinary complexity and subtlety of the discriminations that people constantly make" (p. 301).

Mischel (2004) now suggests that behaviors can be consistent across time, but that the locus of consistency is to be found in highly contextualized if-situation/then-behavior contingencies (e.g., "*If* I feel threatened, *then* I am aggressive"). Variance across situations was, in Mischel's view, improperly treated by most personality psychologists as "error."²⁹ Indeed, in his view, the systematic variation of behavior across situations points to underlying motivations, beliefs, schemas, strategies, and other factors that collectively and interactively lead to coherence in any individual's measured personality. His revised view of personality is broadly consistent with Roberts' Fig. 1.2.

In Section 3, we formalize the "if-then" relationship using an economic model. We show that the person-situation debate boils down to an empirical question about the relative importance of person, situation, and their interaction in explaining behaviors. Although Mischel may have intended otherwise, proponents of the situationist view have used his monograph as ammunition in the battle against accepting evidence from personality psychology into economics. Like most heated debates in social science, this one occurred in the absence of much data. In Section 5, we discuss the body of evidence that has emerged over the past four decades on the existence of stable personality traits.

The debate over the relative importance of person and situation in the 1960s and 1970s reflected deeper currents in psychology and social science more generally, that still run strong. Behaviorism, associated with B. F. Skinner, was influential. It posited that experience explains all aspects of behavior. There was the widely held notion that situation and experience were all powerful—that people were born as blank slates.³⁰ This captured the interventionist spirit of the times. Interindividual heterogeneity in traits was ignored. Ross and Nisbett (1991) summarize the position of many social psychologists:

"Manipulations of the immediate social situation can overwhelm in importance the type of individual differences in personal traits or dispositions that people normally think of as being determinative of social behavior." (p. 14)

Many behavioral economists hold a similar view, and they often appeal to Mischel as a guiding influence. For example, in a recent roundtable discussion, Richard Thaler noted that

"The great contribution to psychology by Walter Mischel [...] is to show that there is no such thing as a stable personality trait."

(Thaler, 2008)

²⁹ That is, unobserved heterogeneity.

³⁰ Pinker (2002).

Many studies in behavioral economics attempt to establish inconsistency in behavior across situations, in violation of standard assumptions of stable preferences used in main-stream economics. For instance, several studies find very low correlations in risk-taking behavior across situations.³¹

Personality psychology survived the behaviorist assault and is a prospering field. A rich body of correlational evidence, which we summarize in Section 7, shows that for many outcomes, measured personality traits are predictive and, sometimes more predictive than standard measures of cognition, that traits are stable across situations, and situations also matter.

Mounting evidence that behavior has a biological basis suggests that personality is an important determinant of behavior. The evidence from behavioral genetics shows that measured personality traits are as heritable as cognitive traits. Studies in neuroscience show that alterations in brain structure and function through accidents, disease, and by experiments affect measured personality. They reinforce the evidence from heritability studies. This evidence and other evidence shows that measured personality is not situation-specific ephemera. We review this evidence in Section 8.

3. CONCEPTUALIZING PERSONALITY AND PERSONALITY TRAITS WITHIN ECONOMIC MODELS

Personality psychologists rarely use formal models to define or measure their constructs. In order to introduce their knowledge to economists, we formalize their frameworks. This makes the concepts of personality psychology more precise and provides a basis for measurement and policy analysis.

We introduce a series of progressively more comprehensive models to integrate concepts from personality psychology into economics.³² Roberts' framework (Fig. 1.2) captures the main features of the influential models used in personality psychology. We use it as a point of departure. Psychology adds new and often more nuanced descriptions of human behavior to the standard descriptions used in economics.

In the nineteenth century, economics and psychology were closely aligned. Economists then spoke of the "hedonic calculus" used by people weighing choices.³³ One of the advances made in neoclassical economics in the first half of the twentieth century was to

³¹ See, e.g., Slovic (1962); Kogan and Wallach (1967); Slovic (1972); Blais and Weber (2006); Johnson, Wilke, and Weber (2004); and Weber, Blais, and Betz (2002).

³² Borghans, Duckworth, Heckman, and ter Weel (2008) develop a variety of economic models for integrating personality psychology into economic models. We build on their analysis. We review these frameworks in Section A3 of the Web Appendix.

³³ See, e.g., Schumpeter (1954).

focus on choices and the objective (easily measured) factors (such as prices and incomes) that determine choices. Revealed preference became a central tool of economics and was implemented using the marginal rate of substitution between choices—a key parameter that emerged from the neoclassical revolution.³⁴ This parameter did not require measurable utility or knowledge of the mental states of the agents making choices. Mental states and measureable utility, once the province of economists, were eliminated by Occam's Razor.

Measurable utility was used in utilitarian economics but fell out of favor (see Samuelson, 1956, and Foster and Sen, 1997). Preferences that fulfilled criteria for rationality were consistent with utility functions that were determined up to monotonic transformations. Measurable utility returned in a specific fashion with analyses of decision making under uncertainty (see Savage, 1954).

Most economists view mental states as unnecessary baggage except insofar as they affect choices. Thus, the traits, abilities, and narratives used by Roberts in Fig. 1.2 are of interest to most economists only if they affect choices through preferences, constraints, and effects on information processing capabilities. Motives and values are captured in part by economic preference parameters. Until recently, "happiness" and "aggregate utility," as well as other subjective mental states that do not affect behavior (choices), were considered uninteresting to most economists.³⁵

Preferences, constraints, and expectations provide the most direct way to introduce psychological variables into economic models. We begin our analysis with a bare-bones approach that focuses on constraints. For example, cognitive and personality traits affect earnings capacity because they enhance productivity (see, e.g., Bowles, Gintis, and Osborne, 2001a), and at least up to a point, more of a trait can generate more resources that enlarge choice sets and hence directly affect behavior.

3.1. An Approach Based on Comparative Advantage

The Roy model (1951) of comparative advantage provides a useful starting point. Heckman, Stixrud, and Urzua (2006) use the Roy model to introduce psychological variables into the study of social and economic outcomes.³⁶ Personality traits are treated as endowments, and choices are determined by personality traits and other factors as they affect productivity in skills.

³⁴ See Hicks (1946).

³⁵ However, see the revival of utility measurement in the happiness literature (see Layard, 2005). Perceptions on which one does not act, included in the domain of psychology, have recently entered economic studies through the happiness literature.

³⁶ See Roy (1951), Heckman and Sedlacek (1985), and Heckman and Honoré (1990).

Agents can perform one of *J* tasks with productivity $P_j, j \in \{1, ..., J\}$. The productivity in task *j* depends on the traits of agents represented by θ and the "effort" they expend on the task, e_j :

$$P_j = \phi_j(\theta, e_j), \ j \in \mathcal{J} = \{1, \dots, J\}, \ e_j \in \mathcal{E}, \ \theta \in \Theta.$$

$$(1.1)$$

The traits are the endowments of agents that govern behavior. Examples of traits include height, personality characteristics, problem-solving ability, and strength. θ is a public good as it is available in the same amount for all tasks. Productivity also depends on effort e_i . Effort is assumed to be divisible and fixed in supply.

In much applied research, effort and traits are often assumed to be measured so that over the relevant range, assuming differentiability with respect to e_i and θ ,

$$\frac{\partial \phi_j}{\partial e_j} \ge 0$$
 and $\frac{\partial \phi_j}{\partial \theta} \ge 0$,

but neither condition is strictly required. Excess effort (overexertion; too much attention to detail) may be counterproductive so that function ϕ_j need not be monotonic in e_j , contrary to what is assumed here. Indeed, as discussed in Section 5, certain psychopathologies are associated with extreme levels of traits that are quite productive at normal levels. Different traits may have different productivities in different tasks, leading to comparative advantage in different tasks for people with different endowments.³⁷

Efforts may complement traits $\left(\frac{\partial^2 \phi_j}{\partial e_j \partial \theta} > 0\right)$ or may substitute for them $\left(\frac{\partial^2 \phi_j}{\partial e_j \partial \theta} < 0\right)$. A variety of intermediate cases might exist where some effort-trait relationships are complementary and others are substitution relationships. Some people may solve complex math problems with no effort, whereas others may have to allocate considerable time and effort to achieve the same result. Effort can be a vector (time, mental energy, attention), and it is assumed to be a divisible private good with the feature that the more that is applied to task *j*, the less is available for all other tasks at any point in time. $\sum_{j=1}^{J} e_j = \overline{e}$ where \overline{e} is the endowment of total effort. Baumeister, Bratslavsky, Muraven,

and Tice (1998) interpret self-control as a component of e that is fixed over given time periods. A person who exerts more self-control in one task may be less self-controlled in another task.

Let R_j be the reward per unit productivity in task *j*. In the first case we analyze, agents can productively engage in only one of the *J* tasks at any time. This restriction can be interpreted as a case in which effort can only be applied to a single task. A reward-maximizing

³⁷ Cattan (2011) shows that sociability has negative returns in some sectors but positive returns in other sectors.

agent with trait θ and endowment \overline{e} faces the problem of picking the maximal task to perform, \hat{j} where

$$\hat{j} = \underset{j \in \{1, \dots, J\}}{\operatorname{argmax}} \{ R_j \phi_j(\theta, \overline{e}) \}.$$
(1.2)

In this case, θ and \overline{e} play the same role. People with different effort and capability endowments will generally choose different tasks.^{38,39} Heckman, Stixrud, and Urzua (2006) show how persons with different endowments of personality and intelligence sort into different occupations and levels of schooling. People low in certain traits may have better endowments of effort and may compensate for their shortfall in ability by exerting effort. For certain tasks (e.g., creating new branches of mathematics), there may be threshold levels of θ such that for $\theta < \overline{\theta}_j$, $\phi_j(\theta, e_j) = 0$ for all $e_j < \overline{e}$. (The person needs a given level of trait θ , no matter how hard they try.) The higher R_j , the more likely will the person choose to perform task *j*. The particular choice of which *j* to perform depends on the productivity of traits in different tasks.

3.2. Allowing for Multitasking

More generally, at a point in time, people may perform multiple tasks.⁴⁰ A less discrete version of the Roy model builds on the same foundations, allows people to perform multiple tasks at any time, and postulates that $\phi_j(\theta, e_j)$ is concave and increasing in e_i .⁴¹ The agent chooses effort levels e_i across the J tasks to maximize total rewards:

$$\max_{\{e_j\}_{j=1}^{J}} \sum_{j=1}^{J} R_j \phi_j(\theta, e_j)$$
(1.3)

subject to
$$\sum_{j=1}^{J} e_j = \overline{e}.^{42}$$

⁴⁰ This, of course, depends on the time unit. Agents may be able to do only one task at one time if the time unit is defined finely enough.

³⁸ A straightforward extension works with utilities and not rewards so we define utility $U(P_1,...,P_J)$ and the agents picks the *j* that maximizes utility, with the other arguments zeroed out. Formally, define $d_{P_j} = 1$ if a person chooses to perform task *j*. Array the d_{P_j} into a vector d_P . Array the P_j into a vector *P*. Realized utility is thus $U(d_P \odot P)$ where \odot is a Hadamard (component-wise) product, i.e., a product of two vectors of the same length where the operation is such that the result is the product of the first element of one vector with the first element of the second vector and so forth for each component.

³⁹ See Heckman, Stixrud, and Urzua (2006); Cattan (2011); and the evidence in Section 7.

⁴¹ Failure of concavity can take us back to case I.

⁴² The first-order conditions for this problem are standard: $R_j \frac{\partial \phi_j}{\partial e_j} \ge \lambda$ and $e_j \ge 0$, j = 1,...,J, where λ is the vector of multipliers associated with the effort constraint. Some people may allocate no effort to some tasks. P_j may be zero if $e_j = 0$, but this is not strictly required. Again, it is straightforward to generalize this reward function to a general utility function $U(P_1,...,P_J)$.

As the reward for activity j, R_j , increases, everything else constant, the effort devoted to j will increase.^{43,44} This model is consistent with effort that compensates for shortfalls in endowments, as well as effort that reinforces initial endowments. The choice of effort depends on the pattern of complementarity and substitutability. Different situations may be associated with different rewards for the same task. Such variation can produce differences in performance across tasks of the sort featured in the person-situation debate discussed in Section 2. One needs to standardize for the incentives to exert effort across tasks and differences in the endowments of effort in order to use measurements of performance on tasks to identify traits, θ .

3.3. Identifying Personality Traits

Before considering more general models, it is useful to discuss basic identification problems that arise in simple settings and in more general models. At the current level of generality, all traits can potentially affect productivity in all tasks. However, some tasks may require only a single trait or a subset of all of the traits. Following a traditional dichotomy in psychology that is explicit in Roberts' Fig. 1.2, divide θ into "mental," μ and "personality," π traits: θ_{μ} and θ_{π} , each of which may in turn be a vector.⁴⁵

Psychological measurement systems sometimes use productivity measured in different tasks to identify θ_{μ} and θ_{π} .⁴⁶ This is the way Carroll (1993) defines mental ability where the task is performed on "mental" tests. To use performance on a task (or on multiple measures of the task) to identify a trait requires that performance on certain tasks (performance on a test, performance in an interpersonal situation, etc.) depends exclusively on one component of θ , say $\theta_{1,j}$. In that case,

$$P_j = \boldsymbol{\phi}_j(\boldsymbol{\theta}_{1,j}, \boldsymbol{e}_j).$$

Even if we can measure productivity P_j in task j, and only one component of θ affects P_j , to identify the level of a trait, one must control for the level of effort applied to j in order to use P_j to infer the level of $\theta_{1,j}$. That is, one must standardize for the effort at a benchmark level, say e^* , to use P_j to identify a measure of the trait that is uniform across different situations that elicit different levels of effort.⁴⁷

The activity of picking a task (or a collection of tasks) to measure a particular trait ($\theta_{1,j}$ in our example) is called *operationalization* in psychology. *Construct validity* refers to

⁴³ $\frac{\partial^2 \phi_j}{\partial \partial \partial e'_j} > 0$ is a force toward devoting more effort to task *j*. If effort is complementary with traits in all tasks as traits expand, more effort will be expended in those tasks that are relatively more complementary in effort.

⁴⁴ In case I, agents will pick *j*.

⁴⁵ Effort endowment might also be divided in the same fashion ($\bar{e}_{\mu}, \bar{e}_{\pi}$) but we do not explicitly develop this possibility.

⁴⁶ They also use observer reports and tests, which can be interpreted as observation on performance of tasks.

⁴⁷ A weaker notion is to achieve relative ranks of a trait. One can define the rank of a trait holding fixed the ranks of all other influences.

whether or not a purported measure of the trait constructed in the stage of operationalization correlates with measures deemed to represent the trait. Considerable judgment is required to operationalize a trait and independently validate it. There is clear danger of circularity. Economists should carefully scrutinize how the measures they borrow from psychology are operationalized and validated in that literature. We should not necessarily assume that the measures created in that field have been rigorously established. We discuss how major constructs are validated in Section 5.

Assuming that construct validity has been established, if effort is involved in the performance of a task used to uniquely define a trait, the measurement of performance must be standardized in order to use measured productivity, P_j , to identify the trait. Otherwise, the endowment of effort and all of the factors that contribute to the exertion of effort, including the reward to the task, R_j , will contaminate the estimate of the trait. Failure to adjust for effort produces the kind of variability across situations with different rewards that was much discussed in the person-situation debate. We present examples of such contamination of measurement by the operation of incentives on effort in Section 5.

Operationalization and construct validation clearly require heroic assumptions. Even if one adjusts for effort in a task, and thus adjusts for situational specificity, productivity in a task may depend on multiple traits. Thus, two components of θ (say $\theta_{1,\mu}$, $\theta_{1,\pi}$) may determine productivity in task *j*. Without further information, one cannot infer which of the two traits produces the productivity in *j*. But in general, even having two (or more) measures of productivity that depend on $(\theta_{1,\mu}, \theta_{1,\pi})$ is not enough to identify the separate components.

Consider the following case of two productivity measurements on tasks *j* and *k*:

$$P_{j} = \phi_{j}(\theta_{1,\mu}, \theta_{1,\pi}, e_{j})$$
$$P_{k} = \phi_{k}(\theta_{1,\mu}, \theta_{1,\pi}, e_{k}), \quad j \neq k$$

One might have such measurements if data are available on the productivity of the same person performing two different tasks. Standardize measurements at a common level of effort, $e_j = e_k = e^{*.48}$ If the functional forms of the $\phi_j(\cdot)$ and $\phi_k(\cdot)$ are known, and the system of equations satisfies a local rank condition, then one can solve for the pair $(\theta_{1,\mu}, \theta_{1,\pi})$ at $e^{*.49}$

$$\begin{bmatrix} \frac{\partial \phi_j}{\partial \theta} & \frac{\partial \phi_k}{\partial \theta} \end{bmatrix}_{e_j = e_k}^{\prime}$$

be nonvanishing in open neighborhoods around any point of solutions for θ (see, e.g., Buck, 2003).

⁴⁸ Note that if the support of e_i and e_k is disjoint, no e^* exists, and hence, no such standardization is possible.

⁴⁹ Let $\theta = (\theta_{1,\mu}, \theta_{1,\pi})$. Assume that the functional forms of $\phi_j(\cdot)$ and $\phi_k(\cdot)$ are known. Formally, a solution from P_j and P_k for $\theta_{1,\mu}$ and $\theta_{1,\pi}$ requires that the Jacobian of the system of equations for P_j and P_k ,

The rank condition might not be satisfied, and the functional forms ϕ_j and ϕ_k might not be known. The productivity functions need not be monotone in $\theta_{1,\mu}$ or $\theta_{1,\pi}$. Interacting systems might produce multiple equilibria so that the same values of θ produce different values of (P_j, P_k) .⁵⁰ Interacting systems might also have no solution.

Note that even if these problems do not arise, only the pair $(\theta_{1,\mu}, \theta_{1,\pi})$ is identified. One cannot (without further information) determine which component of the pair is $\theta_{1,\mu}$ or $\theta_{1,\pi}$. In Section 5, we present an example in which scores on achievement tests depend on both IQ and personality traits. In the absence of dedicated constructs (constructs that are generated by only one component of θ), an intrinsic identification problem arises in using measures of productivity in tasks to infer traits.⁵¹ A dedicated measurement for at least one component is an essential requirement for identification. Other components can be defined relative to that measurement.^{52,53}

3.4. Extensions of the Roy Model

Many empirical economists use the simple gross income maximizing framework of the Roy model to study the effects of personality on outcomes. The model is amended in many papers by including a cost $C_j(\theta, e_j)$ for obtaining rewards so that instead of criterion (1.2), the agent picks \hat{j} that maximizes the net reward

$$\hat{j} = \arg\max_{j \in \{1, \dots, J\}} \{R_j \phi_j(\theta, \overline{e}) - C_j(\theta, \overline{e})\}.$$

In the analogous extension for criterion (1.3), the agent maximizes

$$\sum_{j=1}^{J} R_j \phi_j(\theta, e_j) - C_j(\theta, e_j)$$

⁵⁰ Thus, there is a correspondence between (P_i, P_k) and θ , but no unique functional relationship.

- ⁵¹ There are various ways around this identification problem. For example, one might be able to choose configurations of data with low (or zero) values of one component. At high levels of effort, induced by a change in the reward, the effect of one component on productivity might vanish, etc.
- ⁵² This problem arises in linear factor models, but it is a more general problem. See, e.g., Anderson and Rubin (1956) for a definitive treatment of linear factor models. The scales in any factor model are arbitrary and are always defined with respect to a normalization (i.e., a dedicated measurement that defines the factor). The more general nonlinear model considered in the text faces the same problem.
- ⁵³ In general, without knowledge of the functional forms of the $\phi_j(\cdot)$, j = 1,...,J, the problem of solving for two measurements P_j , P_k to infer $\theta_{1,\mu}$ and $\theta_{1,\pi}$ at a common level of $e_j = e_k$ is intractable. Many alternative solutions are possible. The traditional factor analysis literature reviewed in Section 5 below assumes linearity of the $\phi_j(\cdot)$, j = 1,...,J. But even in that literature, attention focuses primarily on identifying the *distribution* of $(\theta_{1,\mu}, \theta_{1,\pi})$, not individual values $(\theta_{1,\mu}, \theta_{1,\pi})$ when P_j , j = 1,...,J is measured with error, although methods for solving for individual values of $(\theta_{1,\mu}, \theta_{1,\pi})$ and correcting for measurement error of the resulting estimates are available in the literature and are widely applied (see, e.g., Heckman, Malofeeva, Pinto, and Savelyev, first draft 2008, revised 2011; Savelyev, 2010; Heckman and Williams, 2011). Cunha, Heckman, and Schennach (2010) establish conditions under which it is possible to nonparametrically identify the functional form of $\phi_j(\cdot)$, j = 1,...,J and the distributions of $(\theta_{1,\mu}, \theta_{1,\pi})$ in the presence of measurement error on P_i , j = 1,...,J.

with respect to the choice of e_j . This extension creates a further identification problem—whether the trait identified arises from its role in costs, productivity, or both. The identification problem deepens when we allow the costs to be psychic costs as in Heckman and Sedlacek (1985); Cunha, Heckman, and Navarro (2005); or Heckman, Stixrud, and Urzua (2006); and attempt to separate productivity traits from preference traits.⁵⁴

The framework of the Roy model is widely used in recent analyses of the role of personality and cognition.⁵⁵ It has precedents in the work of Mandelbrot (1962), Heckman and Sedlacek (1985), and Heckman and Honoré (1990). In most applications, the $\phi_j(\theta, e_j)$ and $C_j(\theta, e_j)$ (or their logarithms) are assumed to be linear or log linear in θ and e_j . For example:

$$P_{j} = \alpha_{\theta}^{\prime} \theta + \alpha_{e}^{\prime} e_{j}$$
$$C_{j} = \beta_{\theta}^{\prime} \theta + \beta_{e}^{\prime} e_{j}.$$

The analyst models both the choice of the task and the output from the chosen task. A third (mixed) case can arise in which some clusters of tasks are mutually exclusive, so the agent can perform only one task within each cluster of tasks, but the agent can simultaneously engage in tasks across clusters.

3.5. Adding Preferences and Goals

Preferences and goals (see Fig. 1.2) may also shape effort.⁵⁶ This takes us to a fourth and more general case. There may be direct utility benefits or costs associated with exerting effort in each task. Array the effort across tasks in vector $e = (e_1, \ldots, e_J)$. Agents might also attach direct value to the productivity in tasks arrayed in vector $P = (P_1, \ldots, P_J)$ with reward R_j .

Output can produce income $\sum_{j=1}^{J} R_j P_j$, which can be spent on final consumption goods X with associated prices W. A utility function can be specified over X, P, and e with preference parameter vector $\psi \in \Psi$.⁵⁷ Thus, we write

$$U(X, P, e; \psi), \tag{1.4}$$

⁵⁴ Heckman and Navarro (2007) and Abbring and Heckman (2007) present conditions for identification of productivity and costs when there are direct measures of gross productivity, at least when there are measurements on P_j for individuals who select *j*.

⁵⁵ See, e.g., Heckman, Stixrud, and Urzua (2006); Heckman, Humphries, Urzua, and Veramendi (2011); Báron and Cobb-Clark (2010); and Cattan (2011).

⁵⁶ In some versions of the preceding models with costs, preferences can be embodied in psychic costs.

⁵⁷ Robson (1996, 2001) and Robson and Samuelson (2007, 2009) discuss the evolutionary origin of preference parameters.

where the agent maximizes (1.4) subject to the constraints

$$Y + R'P = W'X, \tag{1.5}$$

where *Y* is a flow of unearned income available to the agent in addition to his earnings from his productive activities, and

$$\sum_{j=1}^{J} e_j = \overline{e}.$$
(1.6)

Preference specification (1.4) captures the notions that agents have preferences over goods, agents may value the output of tasks in their own right, and agents may value the effort devoted to tasks.⁵⁸

The parameter ψ determines the trade-offs in preferences among X, P, and e. In one interpretation, subjective measures of well-being (Layard, 2005) attempt to directly measure (1.4).⁵⁹ Parameters that affect subjective well-being but not choices can be identified from the measures of well-being, but not from choices.⁶⁰

3.6. Adding Learning and Uncertainty

All of the preceding models can be extended to account for learning and uncertainty. Let \mathcal{I} be the information possessed by the agent, and let E denote mathematical expectation. An agent can be interpreted as making decisions based on

$$E[U(X, P, e; \psi); \mathcal{I}], \tag{1.7}$$

where ψ may be in the agent's information set (i.e., the agent knows his preferences).

Different theories specify different amounts of information available to agents. They might be uncertain about their preferences, ψ , traits, θ , the prices they face, W, the rewards to productivity, R, the outcomes of purchase decisions, X, and their endowments of effort, \overline{e} . The theory can be suitably modified to account for this uncertainty.

The use of the expectations operator begs the question of how agents construct the information set and how subjective expectations are formed. Psychological traits θ may affect information perception and processing. Several recent studies that apply personality traits to the economics of search suggest that agents with a higher perception of the

⁵⁸ Goods might also be direct arguments of the productivity functions, but, for simplicity, we do not analyze that case.

⁵⁹ However, the happiness literature is not strictly wedded to the notion that happiness is the same as our *U*, which is used only to characterize choice behavior.

⁶⁰ The model can readily be extended to cover more general cases. There is no need to impose the linear reward structure (R'P). The resources raised from productive tasks can be a nonlinear in P. Another simple extension of the model is the case in which there is no financial gain from engaging in tasks, but the agent receives a direct utility benefit from doing so. In this case, constraint (1.5) is redefined as Y = W'X, but P remains as an argument of the utility function. One might also introduce goods as inputs into the ϕ_i functions.

control they have over their lives have greater confidence in the arrival of job offers.⁶¹ Overconfidence may be a trait that causes persons to inflate their perceived productivity.⁶² A production function for information may depend on components of the trait vector, $\theta_{\mathcal{I}}$, and the effort devoted to acquire information, $e_{\mathcal{I}}$. Intelligent people may acquire information more readily than dull people. People more open to experience likely acquire more knowledge. Aggressive people may reduce their social interactions and impair their ability to learn from others. We discuss the evidence on how psychological traits affect information updating in Section 6.

One might object to the rationality and self-perception implicit in this formulation. As in Freud (1909, reprinted 1990), decision making might be made by a subconscious mind lacking self-perception. Decision making may be unconscious and agents may not recognize their desired goals. Nonetheless, constraints limit their revealed choice behavior. Borghans, Duckworth, Heckman, and ter Weel (2009) develop a model in which agents have random preferences and make choices at random within their feasible set. Variations in constraints drive the measured behavior of group averages but do not predict the behavior of any individual.

3.7. Definition of Personality within an Economic Model

Personality *traits* are the components of \bar{e} , θ , and ψ that affect behavior. One might define measured personality as the performance (the P_j) and effort (the e_j) that arise from solutions to any of the optimization problems previously discussed. Thus, the derived productivity and effort functions would constitute the systems generating measured personality as a response to constraints, information, and preferences, that is, as a system of functions that solve out for the P_j and e_j that agents choose in terms of their choice parameters.⁶³

This approach to defining personality would not capture the full range of behaviors or *actions* considered by personality psychologists as constituting manifestations of personality. The actions considered by psychologists include a variety of activities that economists normally do not study, for example, cajoling, beguiling, bewitching, charming, etc. Thus, in selling a house, various actions might be taken, for example, smiling, persuading people by reason, threatening, scowling, showing affection, etc. Actions also include emotions, feelings, and thoughts and are not restricted to be activities that promote physical productivity. Colloquially, "there are many ways to skin a cat," and the choice of which way to do so in any task defines the action taken.

⁶¹ McGee (2010) and Caliendo, Cobb-Clark, and Uhlendorff (2010).

⁶² See, e.g., Akerlof and Dickens (1982); Caplin and Leahy (2001); Köszegi (2006); and Möbius, Niederle, Niehaus, and Rosenblat (2010).

⁶³ As previously noted in a simpler setting, no solutions may exist or multiple solutions may exist (so, there is a system of correspondences) between traits and personality outcomes.

To capture these more general notions, we introduce the concept of "actions" that are broader than what is captured by *e*. Actions are *styles* of behavior that affect *how* tasks are accomplished. They include aspects of behavior that go beyond effort as we have defined it.

Any task can be accomplished by taking various actions. We denote the *i*th possible action to perform task *j* by $a_{i,j}$, $i \in \{1, ..., K_j\}$. Array the actions in a vector $a_j = (a_{1,j}, ..., a_{K_j,j}) \in A$. The actions may be the same or different across the tasks. Thus, one can smile in executing all tasks or one may smile in only some. The productivity of the agent in task *j* depends on the actions taken in that task:

$$P_j = \tau_j(a_{1,j}, a_{2,j}, \dots, a_{K_j,j}).$$
(1.8)

The actions themselves depend on traits θ and "effort" $e_{i,j}$:

$$a_{i,j} = \boldsymbol{\nu}_{i,j}(\boldsymbol{\theta}, \boldsymbol{e}_{i,j}), \tag{1.9}$$

where

$$\sum_{i=1}^{K_j} e_{i,j} = e_j \text{ and } \sum_{j=1}^J e_j = \overline{e}.$$

Less effort may be required to perform a given action if a person has endowment θ that favors performance of the action. For example, a naturally gregarious person may find it easier to engage in social interactions than others. Stated this way, actions generalize the notion of effort to a broader class of behavior. Analytically, they play the same role as effort, and some actions may be components of effort. There may be utility costs or benefits of effort exerted. A special case arises when there are increasing returns to effort in each action. In that case, the agent will simply apply all of his effort e_j in task j to the action that gives him the highest productivity, and the other possible actions are not taken.

Agents may have utility over actions beyond the utility derived from consuming the outputs of tasks. For example, an agent may prefer accomplishing a task by working hard rather than by cheating. Different beliefs, thoughts, and feelings may have different effects on outcomes. Introducing actions in this fashion allows for the possibility that some actions are valued in their own right and do not directly contribute to productivity in any of the J tasks. Let M be the index set for the set of possible actions, including actions that do not directly contribute to productivity. In this more general formulation

$$a_{i,m} = \mathbf{v}_{i,m}(\theta, e_{i,m}), m \in M,$$

where $\mathcal{A} \subseteq \mathcal{M}$.

We define utility over actions. Let a denote the choice of actions, some of which may not be associated with any particular task. Using the same information as used to characterize (1.7), the agent solves

$$\max E[U(a, X, P, e; \psi); \mathcal{I}]$$

with respect to X and e given the stated constraints. Actions may also directly affect \mathcal{I} , so the production of information can depend on θ , e, and a. The choice of which actions to take depends on goals and values (captured by ψ) and on the available information. Part of learning may consist of agents learning about the set of actions that are available to them, $\mathcal{A}(\mathcal{I})$.

One can extend the framework to introduce the effects of the situation in the person-situation debate, by considering specific situations represented by $h \in \mathcal{H}$. These situations are assumed to affect productivity by affecting the set of possible actions and hence the action taken. Thus, for a person with traits θ and effort vector e_j with action $a_{i,j}$, using the specification (1.9), the action function can be expanded to be dependent on situation h:

$$a_{i,j,h} = \nu_{i,j}(\theta, e_{i,j}, h),$$
 (1.10)

and productivity on a task can be specified as a function of the action taken to perform the task in situation h:

$$P_{j,h} = \tau_j(a_{1,j,h}, \dots, a_{K_j,j,h})$$
(1.11)

or by a more general specification where situation h, along with traits, has a direct effect on productivity in addition to their effects on actions taken:

$$P_{j,h} = \tau_j(\theta, a_{1,j,h}, \dots, a_{K_i,j,h}, h).^{64}$$
(1.12)

Situations could include physical aspects of the environment in which the agent is located or the network (and other social situations) in which the agent is embodied. The situation can include social factors such as peer effects.⁶⁵ Persons taking an achievement test sometimes perform worse if they are told that their scores will influence social perceptions of their group as is found in the stereotype threat literature.⁶⁶

The situation represents a key notion in the "person-situation" debate discussed in Section 2. Equations (1.10)–(1.12) capture the "if-then" notion of Mischel and Shoda (1995). Under specification (1.12), agents with the same actions, efforts, and traits may have different productivities. Failure to control for situation h, just like failure to control for effort, will contaminate identification of traits using measures of actions or productivities. Situations may be forced on the agents or may be chosen.⁶⁷

⁶⁴ A more general formulation would treat $h \in \mathcal{H}$ as mutually exclusive descriptions of situations and not claim to represent all situations by a base set of characteristics and would index all of the $\nu_{i,i}$ functions by h.

 $^{^{65}}$ Included in situation *h* might be the act of being observed by third parties and other possible sources of social interactions.

⁶⁶ Steele and Aronson (1998) and Sackett, Hardison, and Cullen (2004).

⁶⁷ At the cost of further notation, we could make the set of possible situations task specific.
Let $T \in \mathcal{T}$ be the vector of traits $(\theta, \psi, \overline{e})$. At any point in time, traits are endowments. In the general case, the solution to the constrained maximization problem involves choosing goods X, situation h, actions $a_{i,j}$, and efforts $e_j, j \in \{1, ..., J\}$ subject to the constraints. h is fixed if agents cannot choose the situation. For simplicity, we analyze this case. Relaxing this assumption is straightforward but is notationally more cumbersome.

In the case of fixed *h*, the solution to the maximization problem produces a set of response functions.⁶⁸ Preference parameters, ψ , characterize the trade-offs and goals that help shape manifest behavior. The agent's response functions (assumed to exist) are

$$X = X(R, W, T, h, Y, \mathcal{I})$$
(1.13)

$$e = e(R, W, T, h, Y, \mathcal{I})$$

$$(1.14)$$

$$a = a(R, W, T, h, Y, \mathcal{I}).$$
 (1.15)

Productivity P across tasks is derived from the actions, efforts, and traits of the agents.⁶⁹

The behaviors that constitute personality are defined as a *pattern of actions* in response to the constraints, endowments, and incentives facing agents given their goals and preferences. This interpretation incorporates the notion that personality is a system of functions. People may have different personalities depending on their trait endowments, constraints, and situations. Their actions—not the traits—constitute the data used to identify the traits.

Introducing actions widens the set of data from which one might infer the components of T. Personality psychologists often use actions (e.g., "dispositions") to infer traits. The same identification issues previously discussed continue to arise but now apply to a broader set of measurements.

As noted in the introduction to Section 2, many personality psychologists define personality as "enduring patterns of thoughts, feelings, and behaviors" that reflect tendencies of persons to respond in certain ways under certain circumstances. Our notion of action a is broad enough to encompass the wide array of behaviors considered by the personality psychologists. We previously defined personality traits T as generators of behavior.

One way to capture the notion of enduring actions is to average the *a* functions (1.15) for a person with a given trait vector T = t over situations and efforts. Thus, for a given task *j* and trait vector *t*, the average action for information set \mathcal{I} can be defined as

$$\overline{a}_{T,i,j,\mathcal{I}} = \int_{S_{T,\mathcal{I}}(h, e_{i,j})} \nu_{i,j}(\theta, e_{i,j}, h) g\left(h, e_{i,j} \mid T = (\theta, \psi, \overline{e}), \mathcal{I}\right) dh de_{i,j},$$

⁶⁸ The same warnings as previously issued apply. No solutions may exist or they may be multiple valued.

⁶⁹ In the case of *h* chosen, we get a system of derived demands for *X*, *h*, *a_{i,j}*, *e_j*.

where $S_{T,\mathcal{I}}(h, e_{i,j})$ is the support of $(h, e_{i,j})$ given T and \mathcal{I} , and $g(h, e_{i,j} | T = (\theta, \psi, \overline{e}), \mathcal{I})$ is the density of $(h, e_{i,j})$ given $T = (\theta, \psi, \overline{e})$ and information set \mathcal{I} . $\overline{a}_{T,i,j,\mathcal{I}}$ is the "enduring action" *i* of agents across situations in task *j* with information \mathcal{I} , that is, the average personality. Note that if $\nu_{i,j}$ is separable in *T* the marginal effect of personality trait vector θ is the same in all situations. One can define the "enduring traits" in a variety of ways, for example by averaging over tasks, *j*, situations, *h*, or both. Only under separability will one obtain the same marginal effect of θ . Epstein (1979) and a subsequent literature present evidence against nonseparability but in favor of an "enduring trait" that is common across situations.

3.8. Life Cycle Dynamics

The analysis in the preceding subsection was for a particular point in time (e.g., a period). Traits are not set in stone. In a dynamic setting, one can think of traits, T, information, \mathcal{I} , situations, h, and actions, a, as state variables that evolve through aging, experience, and investment. As a result of experience (including social interactions), situations, biology (ontogeny), and investment, traits may change over the life cycle. We briefly discuss the dynamics of trait and state formation, leaving a more complete discussion to Section 8.

To capture the evidence from a large and growing literature, we consider the dynamic evolution of traits.⁷⁰ Let T^{ν} be traits at age ν , $\nu \in \{0, ..., \nu\} = \nu$. Traits may change through family and self-investment (Cunha and Heckman, 2007, 2009), schooling, biology, or experience. Information \mathcal{I}^{ν} may be updated through various channels of learning. All task outputs, actions, and goods inputs may be time dated.

Investment in period v is an action or set of actions that an individual (or a person or group acting for the individual) may take in period v. Investments have dynamic effects. The technology of skill formation (Cunha and Heckman, 2007, 2009) captures the notion that traits may evolve in response to the inputs of a vector of investments (IN^v) and through aspects of the situation in which the agent is found, h^v where h^v is the vector of attributes of the situation:

$$T^{\nu+1} = \eta^{\nu} (\underbrace{T^{\nu}}_{\text{self-productivity}}, \underbrace{IN^{\nu}}_{\text{investment}}, h^{\nu}), \quad \nu \in \mathcal{V},$$
(1.16)

where the first set of arguments arises from self- and cross-productivity (skill begets skill; traits beget other traits, and traits cross-foster each other; see Cunha and Heckman, 2007, 2009). The second set of arguments arises from investment. Investment is a broad concept and includes parental nurturance, schooling, learning by doing, and learning by imitation, etc. The third set of arguments arises from the situation in which the person is placed.⁷¹

⁷⁰ We survey the evidence on the life cycle dynamics of traits in Section 8, focusing primarily on the traits θ that affect measured productivity.

⁷¹ The actions taken by agents might also enter as arguments to this technology.

Note that if elements of T^{ν} are augmented over the life cycle through investment and practice, the actions and efforts required to achieve a given task can change. Thus, if θ^{ν} is enhanced over time, the amount of effort required to perform a task may be reduced. In this way, we can model habit formation and capture the notion of *arete*, effortless performance of actions, discussed in Aristotle (1956).⁷²

As emphasized by Mischel and Shoda (1995) and Roberts and Jackson (2008), situations may change over time as a function of past actions, past situations, investment, information, and the like. We present this possibility by the following equation of motion:

$$h^{\nu+1} = \chi^{\nu}(h^{\nu}, IN^{\nu}, a^{\nu}). \tag{1.17}$$

Past actions may serve to determine the set of present situations. Those situations in turn may influence current actions.

Information \mathcal{I}^{ν} may change over the life cycle through experimentation and exogenous learning:

$$\mathcal{I}^{\nu+1} = \rho^{\nu}(\mathcal{I}^{\nu}, a^{\nu}, T^{\nu}, IN^{\nu}, h^{\nu}).$$
(1.18)

This learning mechanism incorporates the beliefs of agents about the available data. Thus, people may learn about their environments and themselves in part as a consequence of their own actions and in part as a consequence of the exogenous arrival of information. Equations of motion (1.16)–(1.18) are very general. We consider special cases of them used in the empirical literature in Section 8.

A rich and evolving literature investigates dynamic preferences when agents do not possess full knowledge of their future environments (see, e.g., Hansen, 2005; Hansen and Sargent, 2008; Rust, 2008; Epstein and Zin, 1989; Epstein and Schneider, 2003; Skiadas, 1998). That literature is too large to summarize in this chapter. Preferences need not be separable over time, and there may be time inconsistency of choices associated with hyperbolic discounting.⁷³ We discuss commonly used dynamic preference specifications in Section 6.

3.9. Relationship of the Model in This Section to Existing Models in Personality Psychology

Personality psychologists generally do not present formal models. The formalization, to our knowledge, in this section is the first mathematically precise definition of personality traits and measured personality. The models we have sketched in this section capture central features of leading models in personality psychology.

⁷² See Lear (2004). A habit can be defined as an effortless performance of a task, that is, an action that requires no effort. High levels of traits might allow people to perform actions effortlessly.

⁷³ See Kirby and Herrnstein (1995) and Gul and Pesendorfer (2004).

By its authors' own admission, the McCrae and Costa (2008) Five Factor Theory is not a fully articulated model. Their model emphasizes the role of traits (T) and, in particular, the Big Five factors, in producing outcomes and agent actions, and is sketchy about other details. Agents are assumed to learn about their own traits, but precise learning mechanisms are not discussed. Expression of traits is affected by the external environment and through social interactions in a not fully specified fashion. The concept of an evolving information set \mathcal{I}^{v} plays a central role in Five Factor Theory. People learn about their traits through actions and experience, but the exact mechanisms are not precisely formulated. Equation (1.18) captures these notions. Situations may also evolve as a function of actions and experience, but no role is assigned to investment in Five Factor Theory.

Thus, a restricted version of (1.17) formalizes aspects of the Five Factor Theory. The theory features "characteristic adaptations," which correspond to the actions and efforts of our model that also affect the productivity in tasks. The role of preferences is left unspecified. However, McCrae and Costa explicitly feature *rationality* (McCrae and Costa, 2008, p. 161) and reject the characterization of flawed human decision making that dominates social psychology and the field of behavioral economics that was spawned from social psychology. They explicitly reject a purely situationist explanation of the origin of actions, but they allow for situations to affect actions. Traits evolve through biological processes (ontogeny), but investment or experience do not affect the evolution of traits. Thus, the arguments of equation (1.16) are suppressed, but traits may still exogenously evolve as a function of age and the biology of the individual. Even though traits evolve as part of an exogeneous maturation process, persons may learn about themselves (their traits) by taking actions and by being acted on by the external environment.

"Social cognitive" theories are rivals to trait theories based on the Big Five.⁷⁴ Albert Bandura, Daniel Cervone, and Walter Mischel are central figures in this literature. Roberts' diagram (Fig. 1.2) captures key aspects of this theory, and Roberts himself can be viewed as a member of both camps. This line of thinking stresses the role of cognition in shaping personality and the role of social context in shaping actions and self-knowledge. Authors writing in this school of thought reject the "cognitive–noncognitive" distinction that is often used in economics. They view manifest personality as an outcome of cognitive processes. A major role is assigned to agency—individual goals and motives that produce actions. Their goals and motives are captured by our ψ . The arrival of information is captured by \mathcal{I} . Although the literature in personality psychology often contrasts these two schools of thought, they are not distinct to us. Only in one extreme version of the social-cognitive theory are traits entirely absent. In that version, agent behavior is entirely

⁷⁴ See Cervone and Pervin (2009).

shaped by situations. More generally, Mischel and Shoda (2008) focus on the role of situation in shaping actions, efforts, and productivities and allow for traits to influence actions. The "sociogenomic" model of Roberts and Jackson (2008) also considers the dynamics of personality formation.

Thus, both schools of thought accept specification (1.9) or its extension (1.10), and both would be comfortable with response systems (1.13)–(1.15). The relative importance of the factors emphasized by the two schools of thought can only be settled by empirical research. The social-cognitive theorists tolerate deviations from rationality in their theories, whereas trait theorists typically do not.

Both schools of thought entertain the possibility of learning about oneself. A major difference between the two groups comes in the role of investment in producing traits. The social-cognitive theorists feature investment and social interactions as direct determinants of traits that are assumed to evolve as a function of the experiences of agents. The trait theorists do not consider this possibility. Instead they emphasize self-learning about traits that evolve by fixed biological principles unrelated to the experiences of individuals.⁷⁵

4. MEASURING PERSONALITY

Unlike other personal traits, like height or weight, personality traits cannot be directly measured. Observed productivities, efforts, and actions are used to infer traits. This leads directly to the analysis of latent variables and to factor models that underlie much of the analysis of trait psychology. This is an area where psychology and the econometrics of measurement error, and latent variables more generally, fruitfully interact. Factor models underlie the concepts of validity of measurements that are used in psychology.

4.1. Linear Factor Models

Linear factor models are widely used in personality psychology and in psychometric models for mental test scores. We review the use of these models in psychology. Versions are already in widespread use in economics.⁷⁶ To capture essential points, we abstract from many of the issues discussed in Section 3. We consider measurements arising from productivity in tasks. We thus focus solely on the outputs of tasks, abstracting from actions, efforts, and situations. With suitable extensions of the notation used here, we can extend the factor model to the more general models discussed in Section 3.

⁷⁵ Cervone (2004) contrasts the two schools of thought.

⁷⁶ See, e.g., Heckman, Stixrud, and Urzua (2006); Heckman, Humphries, Urzua, and Veramendi (2011); Piatek and Pinger (2010); Cattan (2011); and Cunha and Heckman (2008).

We assume additive separability of the arguments of Eq. (1.1). The stripped-down model writes task performance of person n on task j, $P_{n,j}$, based on traits arrayed in a vector T_n in the following manner:

$$P_{n,j} = \mu_j + \lambda'_j T_n + \Delta_{n,j}, \ n = 1, \dots, N, \ j = 1, \dots, J,$$
(1.19)

where μ_j is the mean productivity in the *j*th task, λ_j is a vector of factor loadings, and $\Delta_{n, j}$ is other determinants of measured performance, including measurement errors. The number of components in T_n , L, has to be small relative to J(L < J) for the factor model to have any explanatory power. Otherwise for each task, one can create a unique factor, and the model becomes tautological. A purely cognitive task would be associated with zero values of the components of vector λ_j on elements of T_n that are associated with personality traits. Factor model (1.19) captures the notions that (a) latent traits, T_n , generate a variety of outcomes, (b) task outputs are imperfect measures of the traits, T_n , because $\Delta_{n,j}$ also determines task output, and (c) tasks other than tests or observer reports may also proxy the underlying traits, that is, latent traits generate both test scores and behaviors. A correlation of outcomes across tasks can arise because tasks depend on the same vector of traits.⁷⁷ Outcomes across tasks may be correlated even if the components of T_n are not.⁷⁸

4.2. Discriminant and Convergent Validity

In this simplified framework, most personality psychologists focus on observer- and self-reports as measures of $P_{n,j}$. The measurements are designed to capture a particular trait. As discussed in Section 3, the choice of which collection of tasks is used to measure a capability ("operationalization and construct validity") is an inherently subjective activity. Many psychologists take a pragmatic, empirical point of view. Traits are what the measurements used capture.⁷⁹ The danger with this empiricist definition is that it offers no guide to the choice of measurements, which are usually settled by conventions or intuitions.

The concept of "discriminant validity" of a collection of tasks (e.g., a set of test scores or a set of observer reports or measurements of productivities) is commonly used to test for construct validity. This approach exploits the notion that a particular battery of measurements captures a component of T_n , for example, $T_{n,l}$, l = 1,...,L, and not other components. Many measurements may be taken on $T_{n,l}$, and having multiple measurements helps to control for measurement error.

⁷⁷ The strength of the correlation depends in part on the magnitudes of λ_i and λ_k across the two tasks, j and k.

⁷⁸ Cunha, Heckman, and Schennach (2010) present a nonparametric identification analysis for a general nonseparable model allowing for measurement error in measures of performance. In the notation of Eq. (1.19), they nonparametrically identify the distribution of T_n and the distribution of $\Delta_{n,j}$, j = 1,...,J. The latter is identified without assuming full independence among the measurement errors.

⁷⁹ Borsboom, Mellenbergh, and Van Heerden (2003) compare the approach taken in Section 3 of defining traits a priori within a model with the operationalist approach (Bridgman, 1959) of defining a trait by whatever measurements are available on it. Operationalism begs the questions that arise in operationalization and construct validity.

All measurements are really just outcomes on a type of task, although the effort applied may vary greatly across tasks. The literature in psychology usually assigns a special status to tests, self-reports, and observer reports of latent traits and also uses direct measures of productivity, such as supervisor ratings.⁸⁰ Behaviors, tests, observer reports, and self-reports all can be used to proxy the underlying traits. These include repeated measurements on the same types of assessment mechanisms, as well as measurements on different behaviors and assessments that are assumed to be generated by common traits.⁸¹

A standard approach to defining constructs in personality psychology is based on factor analysis. This approach takes a set of measurements that are designed to capture a construct and measures within-cluster and across-cluster correlations of the measurements to isolate latent factors $T_{n,l}$, l = 1,...,L, or their distributions. The measurements and clusters of tests are selected on intuitive grounds or *a priori* grounds and not on the basis of any predictive validity in terms of real-world outcomes (e.g., success in college, performance on the job, earnings). This process led to the taxonomy of traits that became the Big Five. Because of the somewhat arbitrary basis of these taxonomies, there is some controversy in psychology about competing construct systems, which we discuss in Section 5. In practice, as we document below, the requirement of independence of the latent factors across constructs (lack of correlation of tests across clusters) is not easily satisfied.⁸² This fuels controversy among psychologists advocating competing taxonomies.

To state these issues more formally, let $P_{n,l}^q$ be the *q*th measurement on trait *l* for person *n*. Using a linear factor representation, the *q*th measurement of factor *l* for person *n* can be represented as

$$P_{n,l}^{q} = \mu_{l}^{q} + \lambda_{l}^{q} T_{n,l} + \epsilon_{n,l}^{q},$$

$$q = 1, \dots, Q_{l}, \quad n = 1, \dots, N, \quad l = 1, \dots, L.$$
(1.20)

The factor $T_{n,l}$ is assumed to be statistically independent of the "measurement errors," $\epsilon_{n,l}^q$, $q = 1, ..., Q_l$. Different factors are sometimes assumed to be independent $(T_{n,l}$ independent of $T_{n,l'}$ for $l \neq l'$). The measurement errors (or "uniquenesses") are usually assumed to be mutually independent within and across constructs.⁸³

In fact, measurement $P_{n,l}^q$ may depend on other components of T_n so that the measurement captures a composite of latent traits. A more general case is

$$P_{n,l}^{q} = \mu_{l}^{q} + (\lambda^{q})' T_{n} + \epsilon_{n,l}^{q}, \quad q = 1, \dots, Q_{l},$$
(1.21)

⁸⁰ See Groth-Marnat (2009).

⁸¹ Different measurements may load onto different traits.

⁸² Indeed, as documented in Section 7, the factors associated with personality are also correlated with some measures of cognitive factors, but not all.

⁸³ The literature in economics relaxes the independence assumptions. See Cunha and Heckman (2008) and Cunha, Heckman, and Schennach (2010) and the literature they cite. They present conditions under which independence can be eliminated and identification of factors is still possible.

where λ^q is a vector with possibly as many as *L* nonzero components. The $\epsilon_{n,l}^q$ are assumed to be independent of T_n and mutually independent within and across constructs (*l* and *l'* are two constructs). The task has *discriminant validity* for trait *l* if the only nonzero component of λ^q is λ_l^q . The μ_l^q and λ_l^q can depend on measured characteristics of the agent, Q_n .⁸⁴ The task has *convergent validity* if measures within the construct are highly correlated.

More precisely, conventional psychometric validity of a collection of items or test scores for different constructs has three aspects. (1) Factor T_l for construct l is statistically independent of factor $T_{l'}$ for construct $l' \neq l$, discriminant validity.⁸⁵ (2) A factor T_l is assumed to account for the intercorrelations among the items or tests within a construct l. (3) Item-specific and random-error variance are low (intercorrelations among items are high within a cluster).⁸⁶ Criteria (2) and (3) define *convergent validity*.⁸⁷

Oblique factor analysis picks factors and factor loadings that allow the factors to be correlated across traits. Its criterion is to maximize the correlation of measurements on a trait and minimize the correlation of measurements across traits, but not imposing that cross-trait correlation be zero. See Harman (1976) and Gorsuch (1983) for a discussion of alternative criteria in oblique factor analysis.

4.3. Predictive Validity

An alternative criterion for validating measurement systems is based on the predictive power of the tests for real-world outcomes, that is, on behaviors measured outside of the exam room or observer system. The Hogan Personality Inventory,⁸⁸ the California Personality Inventory, and the Minnesota Multiphasic Personality Inventory were all developed with the specific purpose of predicting real-world outcomes. Decisions to retain or drop items during the development of these inventories were based, at least in part, on the ability of items to predict such outcomes. This approach has an appealing concreteness about it. Instead of relying on abstract, *a priori* notions about domains of personality and subjectively defined latent factors generated from test scores and self- and observer-personality assessments, it anchors measurements in tangible, real-world outcomes and constructs explicit tests with predictive power. Yet, this approach has its own problems.

First, all measurements of factor $T_{n,l}$ can claim incremental predictive validity as long as each measurement is subject to error ($\epsilon_{n,l}^q \neq 0$). Proxies for $T_{n,l}$ can appear to be

⁸⁴ Hansen, Heckman, and Mullen (2004) show how to allow Q_n to depend on T_l and still identify the model.

⁸⁵ This is sometimes weakened to a condition of zero correlation.

⁸⁶ Cronbach's alpha is a widely used measure of intercorrelation among test scores, that is, a measure of importance of the variance of the $\epsilon_{n,l}^{q}$ uniquenesses relative to the variance of the factors. See Hogan, Hogan, and Roberts (1996) for a precise definition. Sijtsma (2009) discusses the severe limitations of Cronbach's alpha.

⁸⁷ Nothing in these standard testing procedures guarantees that the measurements that satisfy convergent and discriminant validity identify a single trait. Multiple traits operating in the same fashion across many outcomes would produce outcomes and factors that satisfy the criteria. The multiple traits would be captured into a single factor. Only if different traits differentially affect different outcomes can one identify different traits.

⁸⁸ See http://www.hoganassessments.com/products_services/hpi.aspx and also Hogan and Roberts (2001).

separate determinants (or "causes") instead of surrogates for an underlying one-dimensional construct or factor. Thus, suppose that measurement system (1.20) is the correct specification and that a set of measurements display both convergent and discriminant validity. As long as there are measurement errors in the measures for construct l, there is no limit to the number of proxies for $T_{n,l}$ that will show up as statistically significant predictors of an outcome.⁸⁹ For this reason, it is necessary to correct for measurement error in using predictive validity to identify and measure traits.

A second problem with this approach to validation is reverse causality. This is especially problematic when interpreting correlations between personality measurements and outcomes. Outcomes may influence personality measures, as well as the other way around. For example, self-esteem might increase income, and income might increase self-esteem. Measuring personality traits prior to measuring predicted outcomes does not necessarily solve this problem. For example, the anticipation of a future pay raise may increase present self-esteem.

Psychologists sometimes address the problem of reverse causality by using early measures of traits determined well before the outcomes are measured to predict later outcomes.⁹⁰ This approach is problematic if the traits the analyst seeks to identify evolve over time and the contemporary values of traits drive behavior. This practice trades a reverse causality problem with a version of an errors in variables problem. Early measures of the traits may be poor proxies for the traits that drive measured current behavior. In our review of the literature in Section 7, we distinguish studies that attempt to control for reverse causality and those that do not.

Heckman, Stixrud, and Urzua (2006) demonstrate the importance of correcting for reverse causality arising from schooling affecting traits and traits affecting schooling in interpreting the effects of personality tests on a variety of socioeconomic outcomes. Application of econometric techniques for determining the causal effects of factors on outcomes makes a distinctive contribution to psychology.

Many psychologists focus on prediction, not causality.⁹¹ Establishing predictive validity will often be enough to achieve the goal of making personnel assignment and student placement decisions.⁹² However, for policy analysis, including analyses of new programs designed to augment the skills of the disadvantaged, causal models are required in order to generate policy counterfactuals.⁹³

⁸⁹ This is a standard result in the econometrics of measurement error. See, e.g., Aigner, Hsiao, Kapteyn, and Wansbeek (1984).

⁹⁰ This approach is based on the *post hoc ergo propter hoc* fallacy.

⁹¹ There is a long tradition in psychology of conducting predictive analysis based on factor analysis (see, e.g., the essays in Cudeck and MacCullum, 2007), but, to our knowledge, there is no systematic treatment of the problem of reverse causality in that field.

⁹² See, e.g., Hogan and Roberts (2001); and Hogan and Hogan (2007).

⁹³ See Heckman (2008a).

The papers of Heckman, Stixrud, and Urzua (2006) and Cunha and Heckman (2008), develop frameworks for circumventing the problems that arise in using predictive validity to define and measure personality constructs. These frameworks recognize the problem of measurement error in the proxies for constructs. Constructs are created on the basis of how well latent factors predict outcomes. They develop frameworks for testing *discriminant validity*. They allow the factors across different clusters of constructs to be correlated and show how to test for the presence of correlations across the factors.

They use an extension of factor analysis to represent proxies for low-dimensional factors. They test for the number of latent factors required to fit the data and rationalize the proxies.⁹⁴ Generalizing the analysis of Hansen, Heckman, and Mullen (2004), Heckman, Stixrud, and Urzua (2006) allow for lifetime experiences and investments to determine, in part, the coefficients of the factor model and to affect the factor itself. Cunha, Heckman, and Schennach (2010) and Cunha and Heckman (2008) allow for the latent factor to determine investment and experience. They correct estimates of latent factors on outcomes for the effects of spurious feedback and separate proxies from factors. The factors are estimated to change over the life cycle as a consequence of experience and investment. We review these studies in Sections 7 and 8.

4.4. Faking

"Faking" may corrupt measurements designed to proxy latent factors. There are at least two types of false responses: those arising from impression management and those arising from self-deception (Paulhus, 1984). For example, individuals who know that their responses on a personality questionnaire will be used to make hiring decisions may deliberately exaggerate their strengths and downplay their weaknesses.⁹⁵ Subconscious motives to see themselves as virtuous may produce the same faking behavior, even when responses are anonymous. It is possible to fake Conscientiousness on a self-report questionnaire, whereas it is impossible to fake superior reasoning ability on an IQ test. To a lesser degree, a similar bias may also operate in cognitive tests. Persons who know that their test scores will affect personnel or admissions decisions may try harder. The literature on "stereotype threat" shows that the framing of an achievement test can affect the performance of the test taker.⁹⁶ Some evidence suggests that faking has a surprisingly minimal effect on predicting job performance.⁹⁷ Correcting for faking using

⁹⁴ Conti, Heckman, Lopes, and Piatek (2010) discuss alternative approaches to selecting the number of latent factors. See also Cragg and Donald (1997).

⁹⁵ See Viswesvaran and Ones (1999), Sternberg (2001), and Sternberg et al. (2000).

⁹⁶ See Steele and Aronson (1998) and Sackett, Hardison, and Cullen (2004).

⁹⁷ See Hough, Eaton, Dunnette, Kamp, and McCloy (1990); Hough and Ones (2002); and Ones and Viswesvaran (1998).

scales designed to measure deliberate lying does not seem to improve predictive validity.⁹⁸ Nevertheless, as noted in Section 3, when measuring cognitive and personality traits, one should standardize for incentives and environment.

The linear factor model does not capture a variety of interesting interactions among traits. Cunha, Heckman, and Schennach (2010) and the papers they cite develop a non-linear nonnormal factor analysis that allows for measurement errors to be correlated across measures and over time. We report estimates based on their nonlinear factor analyses in Section 8.

4.5. The Causal Status of Latent Variables

Some psychologists question the causal status of latent variables extracted from factor analyses of measurements across individuals.⁹⁹ Such factor analytic studies summarize interindividual variation but do not necessarily inform analysts about the effects of exogenously changing the factor in producing outcomes across individuals. In addition, variations of traits within persons may have very different effects than variations across persons.

The distinction between the effects of changing traits within and across persons is traditional in econometrics.¹⁰⁰ Econometric models that capture this distinction could be fruitfully applied to psychology, as can hierarchical linear models.¹⁰¹

These methods do not address the deeper problem that most of the estimates of "the effects" of psychological traits on outcomes (either from "within" or "across" studies) have no causal status. Structural equation methods have been used to estimate causal relationships using cross-person variation. They rely on the usual toolkit of simultaneous equations exclusion restrictions to secure identification.¹⁰² Standard experimental and econometric techniques for inferring causality from within-person changes have only recently been applied to estimate causal effects of personality.¹⁰³ We review this literature in Section 8.

5. IMPLEMENTING THE MEASUREMENT SYSTEMS

How do psychologists measure individual differences? In this section, we analyze the major measurement systems for cognition and personality. We examine the relative performance of cognition and personality in predicting a variety of outcomes. For cognition, there is a fairly well-established set of terminologies and conventions. Aptitude tests are designed to measure differences in the rates at which individuals learn (i.e., fluid intelligence).

⁹⁸ See Morgeson et al. (2007).

⁹⁹ See Borsboom, Mellenbergh, and Van Heerden (2003) and Cervone (2005).

¹⁰⁰ See, e.g., Mundlak (1978) and Hsiao (2003).

¹⁰¹ See Raudenbush and Bryk (2001).

¹⁰² See, e.g., Aigner, Hsiao, Kapteyn, and Wansbeek (1984) for a review of this classical literature.

¹⁰³ See, e.g., Cunha, Heckman, and Schennach (2010) and Heckman, Malofeeva, Pinto, and Savelyev (first draft 2008, revised 2011).

Achievement tests are designed to measure acquired knowledge (i.e., crystallized intelligence). For personality, a variety of alternative measurement systems are proposed, and this is a source of confusion. We attempt to compare and equate these systems of measurement. We link them to measures of childhood temperament and psychopathology, which are also used to describe individual differences. We note that the problems of operationalization and construct validity are present in analyzing any measures of traits.

5.1. Cognition

Intelligence (also called cognitive ability and general mental ability) is defined by psychologists to include the "ability to understand complex ideas, to adapt effectively to the environment, to learn from experience, to engage in various forms of reasoning, to overcome obstacles by taking thought" (Neisser et al., 1996, p. 77).¹⁰⁴ These are clearly distinct traits, and the literature distinguishes more finely among them. The term "IQ" is often used synonymously with intelligence but in fact refers specifically to scores on intelligence tests. Notwithstanding a century of active study and general agreement about the sorts of tasks on which more intelligent individuals perform better, the construct of intelligence "resists a consensual definition."¹⁰⁵

Scores on different tests of cognitive ability tend to be highly correlated, with half or more of the variance of diverse tests accounted for by a single general factor labeled "g" and more specific mental abilities loading on other factors.¹⁰⁶ g is widely interpreted as general mental ability.¹⁰⁷ An extreme version of g-theory that is no longer widely accepted is that g accounts for all the correlation among different tests of cognition.¹⁰⁸

Psychometricians have expanded this notion to create a hierarchy of "orders." The order of a factor indicates its generality in explaining a variety of tests of cognitive ability deemed to satisfy construct validity. Tests have different emphases (e.g., verbal ability, numeracy, coding speed, and other tasks). A first-order factor is predictive in all cognitive tasks, j = 1, ..., J in Eq. (1.19). In modern parlance, this general correlation is called "g" but it is no longer viewed as the sole predictor of cognitive test scores. A lower order factor is predictive of performance in only some tasks. Lower order factors can be correlated with the higher order factors and may be correlated with each other. They have independent predictive power from the higher order factors. Figure 1.3

¹⁰⁴ Psychologists have attempted to broaden the concept of intelligence beyond this list. Most notably, Gardner (2004) suggests that the notion of intelligence should also include creativity and the ability to solve practical, real-world problems. He includes in his theory of multiple intelligences, musical intelligence, kinesthetic intelligence, and interpersonal and intrapersonal intelligence, among others.

¹⁰⁵ See Wilhelm and Engle (2005).

¹⁰⁶ See Johnson, Bouchard, Krueger, McGue, and Gottesman (2004); Jensen (1998); Lubinski (2004); and Spearman (1904, 1927).

¹⁰⁷ See Gottfredson (2002).

¹⁰⁸ See, e.g., Carroll (1993).



Figure 1.3 A Hierarchical Scheme of General Intelligence and Its Components.

Source: Recreated from Ackerman and Heggestad (1997).

reports one possible partition of general intelligence due to Ackerman and Heggestad (1997), who summarize the work of Carroll (1993) on the multiple facets of intelligence.¹⁰⁹

5.1.1 Fluid versus Crystallized Intelligence

There is less agreement about the number and identity of lower order factors.¹¹⁰ Carroll (1993) proposed a general intelligence factor g and several more specific second-order factors, including, but not limited to, what Cattell (1971, 1987) dubbed crystallized and fluid intelligence. Crystallized intelligence, Cattell proposed, comprises acquired skills and knowledge and, thus, is partly dependent on educational opportunity and motivation. Fluid intelligence, by contrast, is a general "relation-perceiving ability" (p. 138). Cattell's student Rindermann (2007) elaborates

"Fluid intelligence is the ability to perceive complex relations, educe complex correlates, form concepts, develop aids, reason, abstract, and maintain span of immediate apprehension in solving novel problems in which advanced elements of the collective intelligence of the culture were not required for solution." (p. 462)

In contrast, crystallized intelligence is the same class of skills, "but in materials in which past appropriation of the collective intelligence of the culture would give one a distinct advantage in solving the problems involved" (p. 462).

Carroll (1993) and Horn and McArdle (2007) summarize the large body of evidence against the claim that a single factor *g* is sufficient to explain the correlation structure of achievement and intelligence tests.¹¹¹ Two pieces of evidence are worth highlighting. First, crystallized intelligence tends to increase monotonically for most of the life cycle, whereas fluid intelligence tends to peak in very early adulthood then to decline.¹¹² Second, the well-known Flynn effect, which documents the population-wide increase in performance on intelligence tests over the past half-century, is particularly dramatic for measures of fluid intelligence but much smaller for measures of crystallized intelligence.¹¹³ SAT scores have declined rather than increased over the same period, requiring a renorming in the 1990s.

The relative weighting of fluid versus crystallized intelligence varies among tests according to the degree to which prior experience is crucial to performance. These second-order

¹⁰⁹ Carroll's own organization of his evidence is somewhat different. See Carroll (1993, p. 626).

¹¹⁰ Carroll (1993) analyzed 477 data sets and estimated a structure with g as the highest order factor, eight second-order ability clusters, and over 70 more narrowly defined third-order abilities on a variety of different tests. Alternative hierarchical models, also with g as the highest-order factor, have been proposed (e.g., Lubinski, 2004, and Horn, 1970).

¹¹¹ Recent research by Ardila, Pineda, and Rosselli (2000) shows that more than one factor is required to summarize the predictive power of cognitive tests in economic data. This could be due to the existence of multiple intellective factors or because personality factors affect the measurement of cognitive factors as we discuss later on in this section.

¹¹² See McArdle, Hamagami, Meredith, and Bradway (2000).

¹¹³ See Dickens and Flynn (2001).

factors are not only correlated with the first-order factor g but also contribute additional explanatory power to predicting some clusters of test score outcomes. Achievement tests, such as the Armed Forces Qualifying Test used by economists and psychologists alike, are heavily weighted toward crystallized intelligence,¹¹⁴ whereas tests like the Raven Progressive Matrices (1962) are heavily weighted toward fluid intelligence.¹¹⁵ Several studies have shown that fluid intelligence is much more strongly related to g than are measures of crystallized intelligence.¹¹⁶ Moreover, lay intuitions of intelligence (i.e., what most people mean by "being smart") correspond more closely with the ability to learn than with possession of already acquired knowledge.¹¹⁷ Thus, it seems to us useful to reserve the term "intelligence tests" for tests that primarily measure fluid intelligence and the term "achievement tests" for tests that primarily measure crystallized intelligence. Some would argue that g has been usurped by fluid intelligence. A closer reading is that what is commonly meant by intelligence encompasses a number of distinct traits captured in the lower order factors of Fig. 1.3.

5.1.2 Predictive Validity of Tests of Cognition

How well do IQ and achievement tests predict success in life? This is a hard question to answer. Many different skills are required to achieve success in any task.¹¹⁸ Different tasks in life require different skills in different degrees.¹¹⁹ Table 1.2 shows the domains of validation and the estimated validities of a number of widely used tests of cognition. Note that the domains of validation differ greatly. For IQ tests, the validities are usually established by comparing test scores with other test scores or with grades in school and not success in life. Nevertheless, it is well established that standardized tests of ability and achievement predict objectively measured academic, occupational, and life outcomes.¹²⁰

The SAT college entrance exam is moderately successful in predicting grades in college, which the SAT was designed to do.¹²¹ However, high-school grades are better

¹²⁰ See Kuncel, Ones, and Sackett (2010).

¹¹⁴ See Roberts et al. (2000).

¹¹⁵ See Raven, Raven, and Court (1988). Conti and Pudney (2007) uses data on intelligence and achievement tests across nations to show that a single factor accounts for 94–95% of the variance across both kinds of tests. The high correlation between intelligence and achievement tests is in part due to the fact that both require cognitive ability and knowledge. Common developmental factors may affect both of these traits and that fluid intelligence promotes the acquisition of crystallized intelligence.

¹¹⁶ See Cattell (1971, 1987), Gustafsson (1988), and Kvist and Gustafsson (2008).

¹¹⁷ See Gottfredson (1998).

¹¹⁸ See Mandelbrot (1962).

¹¹⁹ See, e.g., Roy (1951); Mandelbrot (1962); Willis and Rosen (1979); Heckman and Sedlacek (1985); and Heckman, Stixrud, and Urzua (2006).

¹²¹ See Young and Kobrin (2001).

Table 1.2 Predictive Validities of Various Tests of Fluid and Crystallized Intelligence

Test	Domain over Which It Is Validated	Estimated Validities	Source	Notes
SAT	First-year college GPA	0.35–0.53	Kobrin, Patterson, Shaw et al. (2008)	
АСТ	Grades in early years of college	0.42	ACT, Incorporated (2007)	
Stanford–Binet	Correlations with other intelligence tests	0.77–0.87 with WISC-R	Rothlisberg (1987) and Greene, Sapp, and Chissom (1990)	
WISC (Wechsler Intelligence Scale for Children)	Correlations with academic achievement	WISC: 0.443–0.751 with WRAT tests, 0.482– 0.788 with first-grade grades, 0.462–0.794 with second-grade grades; WISC-R: 0.346–0.760 with WRAT tests, 0.358–0.537 with first- grade grades, 0.420– 0.721 with second-grade grades	Hartlage and Steele (1977)	WRAT, Wide Range Achievement Test; ranges are given because correlations vary by academic subject
WAIS (Wechsler Adult Intelligence Scale)	Correlations with other intelligence tests, achievement tests, and outcomes	0.67 (median) with verbal tests, 0.61 (median) with nonverbal tests, 0.69 with education attained, 0.38– 0.43 with college grades, 0.62 with high-school grades	Feingold (1982)	

Cognitive Achievement and IQ Tests

Continued

Table 1.2 Predictive Validities of Various Tests of Fluid and Crystallized Intelligence-continued

Test	Domain over Which It Is Validated	Estimated Validities	Source	Notes
Raven's Standard Progressive Matrices	Correlations with other intelligence tests	0.74–0.84 with WAIS-R	O'Leary, Rusch, and Guastello (1991)	
GATB (General Aptitude Test Battery)	Supervisor rating performance in training programs and in job performance	0.23–0.65	Hunter (1986)	Large range due to variety of jobs
ASVAB (Armed Services Vocational Aptitude Battery)	Performance in military training programs and military attrition rates	0.37–0.78 for training (mean = 0.56); –0.15 for attrition	Schmidt, Hunter, and Larson (1988) for performance in training programs; Sticht, Hooke, and Caylor (1982) for attrition rates	Large range in training correlations due to a variety of jobs
GED (General Educational Development)	Test difficulty is normed against graduating HS seniors. Test scores of high-school seniors and grades of high-school seniors	0.33–0.49 for HS Senior GPA	Technical Manual: 2002 Series GED Tests	
DAT (Differential Aptitude Tests)	Correlations with academic achievement	0.13–0.62 for college GPA	Omizo (1980)	Large range is due to varying validity of eight subtests of DAT
WIAT (Wechsler Individual Achievement Test)	Correlation with other achievement tests; teacher ratings of student achievement	0.80 with grade 4 CAT/2, 0.69 with grade 5 CAT/2, 0.83 with grade 6 CAT/2; 0.67 with teacher ratings	Michalko and Saklofske (1996)	CAT, California Achievement Test

Cognitive Achievement and IQ Tests

predictors of college performance.¹²² The rival American College Test (ACT) is validated in a similar fashion but uses broader measures of college performance, such as grades in higher years of college rather than freshman year grades.¹²³ The Graduate Record Exam is validated by performance in graduate school.¹²⁴ The Armed Forces Qualifying Test (AFQT) is validated by performance in the military. Performance is measured by success in military training schools and performance standardized tasks such as fixing a rifle or repairing a radio.¹²⁵ One can interpret *The Bell Curve* by Herrnstein and Murray and the flood of papers it stimulated as conducting validity studies of the AFQT using real-world outcomes of the sort studied in Tables A1 and A2 in the Web Appendix. The correlation of AFQT with wages is a moderate, r = 0.3. The General Aptitude Test Battery (GATB) predicts success at work as measured by supervisor ratings in over 12,000 occupations and participation in training programs.¹²⁶

5.2. Personality Traits

We have noted in Sections 2 and 3 that sharp contrasts between cognition and personality are not easy to make. Consider, for example, the so-called "quasi-cognitive" traits (Kyllonen, Walters, and Kaufman, 2005). These include creativity (Csikszentmihalyi, 1996), emotional intelligence (Mayer and Salovey, 1997), cognitive style (Stanovich, 1999; Perkins and Tishman, 2001), typical intellectual engagement (Ackerman and Heggestad, 1997), and practical intelligence (Sternberg et al., 2000). Furthermore, the Big Five factor of Openness to Experience has as facets curiosity ("ideas") and imagination ("fantasy") that are often associated with intellect and measured intelligence.¹²⁷ (See the entries under Openness in Table 1.3.) We note in Section 5.3 that personality can affect performance on tests of fluid intelligence. Personality traits also affect acquired skills and knowledge (i.e., crystallized intelligence).¹²⁸ A general pattern is higher correlation of personality tests with tests of crystallized knowledge (e.g., achievement tests).¹²⁹ For many personality traits and for measures of cognition that are based on fluid intelligence, the correlations are close to zero, as we note below.

Finally, consider the construct of *executive function*. "Cognitive control" and "executive function" are terms used interchangeably, primarily in the neuroscience literature. Both have been defined as the voluntary, effortful blocking of a habitual behavior in

¹²² See Bowen, Chingos, and McPherson (2009a) and Geiser and Santelices (2007). However, there is a potential problem with restriction on the range in many of these studies.

¹²³ See ACT, Incorporated (2007).

¹²⁴ See Kuncel and Hezlett (2007).

¹²⁵ See McHenry, Hough, Toquam, Hanson, and Ashworth (1990).

¹²⁶ Schmidt and Hunter (1983, 1998); Hartigan and Wigdor (1989); and McHenry, Hough, Toquam, Hanson, and Ashworth (1990).

¹²⁷ McCrae and Costa (1997a) and Noftle and Robins (2007).

¹²⁸ See Chamorro-Premuzic and Furnham (2005) for an extended discussion of this topic.

¹²⁹ See, e.g., Borghans, Golsteyn, Heckman, and Humphries (2011).

 Table 1.3 The Big Five Domains and Their Facets

Big Five Personality Factor	American Psychology Association Dictionary Description	Facets (and Correlated Trait Adjective)	Related Traits	Childhood Temperament Traits
Conscientiousness	"the tendency to be organized, responsible, and hardworking"	Competence (efficient) Order (organized) Dutifulness (not careless) Achievement striving (ambitious) Self-discipline (not lazy) Deliberation (not impulsive)	Grit Perseverance Delay of gratification Impulse control Achievement striving Ambition Work ethic	Attention/(lack of) distractibility Effortful control Impulse control/delay of gratification Persistence Activity*
Openness to Experience	"the tendency to be open to new aesthetic, cultural, or intellectual experiences"	Fantasy (imaginative) Aesthetic (artistic) Feelings (excitable) Actions (wide interests) Ideas (curious) Values (unconventional)	_	Sensory sensitivity Pleasure in low- intensity activities Curiosity
Extraversion	"an orientation of one's interests and energies toward the outer world of people and things rather than the inner world of subjective experience; characterized by positive affect and sociability"	Warmth (friendly) Gregariousness (sociable) Assertiveness (self-confident) Activity (energetic) Excitement seeking (adventurous) Positive emotions (enthusiastic)	_	Surgency Social dominance Social vitality Sensation seeking Shyness* Activity* Positive emotionality Sociability/affiliation
Agreeableness	"the tendency to act in a cooperative, unselfish manner"	Trust (forgiving) Straightforwardness (not demanding) Altruism (warm) Compliance (not stubborn) Modesty (not show-off) Tender-mindedness (sympathetic)	Empathy Perspective taking Cooperation Competitiveness	Irritability* Aggressiveness Willfulness

Continued

Big Five	American Psychology Association	Facets (and Correlated Trait	Related Traits	Childhood
Personality Factor	Dictionary Description	Adjective)		Temperament Traits
Neuroticism/ Emotional Stability	Emotional stability is "predictability and consistency in emotional reactions, with absence of rapid mood changes." Neuroticism is "a chronic level of emotional instability and proneness to psychological distress."	Anxiety (worrying) Hostility (irritable) Depression (not contented) Self-consciousness (shy) Impulsiveness (moody) Vulnerability to stress (not self-confident)	Internal vs. External Locus of control Core self-evaluation Self-esteem Self-efficacy Optimism Axis I psychopathologies (mental disorders) including depression and anxiety disorders	Fearfulness/behavioral inhibition Shyness* Irritability* Frustration (Lack of) soothability Sadness

Table 1.3 The Big Five Domains and Their Facets—continued

Notes: Facets specified by the NEO-PI-R personality inventory (Costa and McCrae, 1992b). Trait adjectives in parentheses from the Adjective Check List (Gough and Heilbrun, 1983).

*These temperament traits may be related to two Big Five factors.

Source: Table adapted from John and Srivastava (1999).

order to execute a less familiar behavior.¹³⁰ Some authors (e.g., Gray, 2004) also use the terms "cognitive control" and "self-control" interchangeably, though self-control is traditionally considered a personality trait rather than an aspect of cognition. Although tasks requiring executive function are related to questionnaire measures of self-control, the size of these associations is only about r = 0.11 - 0.14 (Duckworth and Schulze, 2009).

A region of the brain called the dorsolateral prefrontal cortex (PFC), in conjunction with the nearby region called the anterior cingulate cortex (ACC), are now understood as responsible for "executive control" over lower order processes.¹³¹ That is, executive control entails top-down, intentional control of behavior and is not necessary for the performance of simple, automatic tasks (Miller and Cohen, 2001). The PFC achieves structural and functional maturity later than other (e.g., sensorimotor) brain regions (Casey, Tottenham, Liston, and Durston, 2005). Specific executive functions attributed to the PFC include abstract reasoning, planning, decision making, working memory (the ability to keep the facts of a problem at hand), attention, conflict monitoring, task switching, and inhibition of prepotent (i.e., dominant, habitual) impulses. Although many functions have been attributed to the PFC, Miller (2000) notes that "there is little agreement on the cardinal prefrontal functions" (p. 449). Nevertheless, there is some consensus that one can distinguish between working memory on the one hand and response inhibition and task switching on the other (Garon, Bryson, and Smith, 2008, and Miyake, Friedman, Emerson, Witzki, and Howerter, 2000). This distinction is important because working memory is highly related to performance on measures of fluid intelligence.¹³² Being able to access all of the data about a problem is helpful in solving it. Thus, working memory is a common component of the constructs of both executive function and general intelligence.¹³³

Although the construct of executive function demonstrates the inadequacy of terms such as "cognitive" and "noncognitive," many personality traits nevertheless are conceptually and empirically easily distinguished from general cognitive ability. Most personality traits are in fact very weakly correlated with IQ (Webb, 1915; McCrae and Costa, 1994; Stankov, 2005; Ackerman and Heggestad, 1997). Thus, regardless of the terms used to describe individual differences that determine life outcomes, one thing is clear: human ability entails more than intelligence. Personality traits, however defined, do matter, and they have independent predictive power from standard measures of intelligence.

¹³⁰ Matsumoto and Tanaka (2004).

¹³¹ Notably, the volume of dorsolateral prefrontal cortex (PFC) is correlated with Big Five Conscientiousness (DeYoung et al., 2010).

¹³² Carpenter, Just, and Shell (1990) and Heitz, Unsworth, and Engle (2005).

¹³³ Friedman, Miyake, Corley et al. (2006).

5.3. Operationalizing the Concepts

Intelligence tests are routinely used in a variety of settings including business, education, civil service, and the military.¹³⁴ Psychometricians attempt to use test scores to measure a factor (a component of T in the notation of Section 4). The working hypothesis in the intelligence testing business is that specific tests measure only a single component of T and that tests with different "content domains" measure different components. We first discuss the origins of the measurement systems for intelligence, and we then discuss their validity.¹³⁵

5.3.1 IQ Tests

Modern intelligence tests have been used for just over a century, beginning with the decision of a French minister of public instruction to identify retarded pupils in need of specialized education programs. In response, Alfred Binet created the first IQ test.¹³⁶ Other pioneers in intelligence testing include Cattell (1890) and Galton (1883), both of whom developed tests of basic cognitive functions (e.g., discriminating between objects of different weights). These early tests were eventually rejected in favor of tests that attempt to tap higher mental processes. Terman (1916) adapted Binet's IQ test for use with American populations. Known as the Stanford–Binet IQ test, Terman's adaptation was, like the original French test, used primarily to predict academic performance. Stanford–Binet test scores were presented as ratios of mental age to chronological age multiplied by 100. IQ scores centered at 100 as the average are now conventional for most intelligence tests.

Wechsler (1939) noted two major limitations of the Stanford–Binet test. First, it was overly reliant on verbal skills and, therefore, dependent on formal education and cultural exposure. Second, the ratio of mental to chronological age was an inappropriate metric for adults (Boake, 2002). Wechsler created a new intelligence test battery divided into verbal subtests (e.g., similarities) and performance subtests (e.g., block design, matrix reasoning). He also replaced the ratio IQ score with deviation scores that have the same normal distribution at each age. This test, the Wechsler Adult Intelligence Scale (WAIS)—and, later, the Wechsler Intelligence Scale for Children (WISC)—produces two different IQ subscores, verbal IQ and performance IQ, which sum to a full-scale IQ score. The WAIS and the WISC have for the past several decades been by far the most commonly used IQ tests.

¹³⁴ Siegler (1992) provides a detailed overview of the different types of applications of psychological testing.

¹³⁵ See Roberts, Markham, Matthews, and Zeidner (2005) for a more complete history of intelligence testing.

¹³⁶ In 1904, La Société Libre pour l'Etude Psychologique de l'Enfant appointed a commission to create a mechanism for identifying these pupils in need of alternative education led by Binet. See Herrnstein and Murray (1994) for an overview of Binet's life and work.

Similar to Wechsler's Matrix Reasoning subtest, the Raven Progressive Matrices test is a so-called culture-free IQ test because it does not depend heavily on verbal skills or other knowledge explicitly taught during formal education. Each matrix test item presents a pattern of abstract figures.¹³⁷ The test taker must choose the missing part.¹³⁸ If subjects have not had exposure to such visual puzzles, the Raven test is an almost pure measure of fluid intelligence. However, the assumption that subjects are unfamiliar with such puzzles is not typically tested. It is likely that children from more-educated families or from more-developed countries have more exposure to such abstract puzzles (Blair, 2006). Our view is that to varying degrees, IQ and achievement tests reflect fluid intelligence, crystallized intelligence, and personality factors, such as motivation, to succeed on the test. We offer evidence on the effect of motivation on test scores below in Section 5.6.

5.4. Personality Constructs

Dominant theories of personality assume a hierarchical structure analogous to that found for intelligence. However, despite early efforts to identify a *g* for personality (e.g., Webb, 1915), even the most parsimonious personality models incorporate more than one factor. The most widely accepted taxonomy of personality traits is the Big Five.¹³⁹ The Big Five factors are obtained from conventional factor analysis using a version of Eq. (1.19) in which the "tests" are measures of different domains of personality based on observer reports or self-reports.

The Five-Factor model has its origins in Allport and Odbert's (1936) lexical hypothesis, which posits that the most important individual differences are encoded in language. Allport and Odbert combed English dictionaries and found 17,953 personality-describing words, which were later reduced to 4,504 personality-describing adjectives. Subsequently, several different psychologists working independently and on different samples concluded that personality traits can be organized into five superordinate factors.

Table 1.3 presents the Big Five factors that were discussed in Section 2. It summarizes the 30 lower level facets (six facets for each of five factors) identified in the Revised NEO Personality Inventory (NEO-PI-R, Costa and McCrae, 1992b). The acronym is shorthand for Neuroticism, Extroversion, Openness to Experience—Personality Inventory—Revised. Of course, these lower level facets (e.g., "impulsive") can be further subdivided into even more narrow traits ("impulsive about junk food," "impulsive about smoking"). The more narrowly defined a trait, the more specific are the contexts in which the trait is predictive. In parentheses in the third column of Table 1.3, we present

¹³⁷ See John and Srivastava (1999) for a discussion of the Raven test.

 $^{^{138}\,}$ See Fig. A1 in Section A5 of the Web Appendix.

¹³⁹ See Goldsmith et al. (1987) for an historical overview of the development of the Big Five.

other traits in each family. In the fifth column, we relate the Big Five to children's temperament traits studied by developmental psychologists.

Temperament is the term used by developmental psychologists to describe the behavioral tendencies of infants and children.¹⁴⁰ Because individual differences in temperament emerge so early in life, these traits have traditionally been assumed to be biological (as opposed to environmental) in origin.¹⁴¹ However, findings in behavioral genetics suggest that, like adult personality, temperament is only partly heritable, and as discussed in Section 8, both adult- and child-measured traits are affected by the environment.

Temperament is studied primarily by child and developmental psychologists, whereas personality is studied by adult personality psychologists. However, the past decade has seen some convergence of these two research traditions, and there is evidence that temperamental differences observed during the preschool years anticipate adult personality and interpersonal functioning decades later (e.g., Caspi, 2000; Newman, Caspi, Moffitt, and Silva, 1997; Shiner and Caspi, 2003). The traits displayed in column 5 of Table 1.3 have been associated both theoretically and empirically with adult personality traits.

Historically, many temperament researchers examined specific lower order traits rather than broader, higher level factors that characterize studies of adult intelligence and personality.¹⁴² Shiner (1998) suggests that "there is therefore a great need to bring order to this vast array of studies of single lower-level traits" (p. 320). Recently, taxonomies of temperament have been proposed that group lower order traits into higher order dimensions; several of these taxonomies resemble the Big Five (e.g., John, Caspi, Robins, and Moffitt, 1994; Putnam, Ellis, and Rothbart, 2001; Rothbart, Ahadi, and Evans, 2000; Shiner and Caspi, 2003). However, compared to adults, there seem to be fewer ways that young children can differ from one another. Child psychologists often refer to the "elaboration" or "differentiation" of childhood temperament into the full flower of complex, adult personality. The lack of direct correspondence between measures of temperament and measures of adult personality presents a challenge to researchers

¹⁴⁰ See Caspi and Shiner (2006) and Zentner and Bates (2008) for a discussion of varying perspectives on temperament, including a summary of points where major theorists converge.

¹⁴¹ Indeed, some psychologists use the term "temperament" to indicate all aspects of personality that are biological in origin. They study temperament in both children and adults.

¹⁴² Measuring temperament presents unique methodological challenges. Self-report measures, by far the most widely used measure for adult personality, are not appropriate for young children for obvious reasons. One strategy is to ask parents and teachers to rate the child's overt behavior (e.g., California Child Q-sort), but informants can only guess what a child might be thinking and feeling. Infants present a special challenge because their behavioral repertoire is so limited. One strategy is to place infants in a standard situation and code reactions under a standardized scenario (e.g., the Strange Situation, which is used to distinguish infants who are securely attached to their caregiver versus insecurely attached). Young children can be interviewed using puppets or stories. For obvious reasons, all measures of temperament are more difficult and more expensive to collect than adult self-report measures. This may explain the absence of large-sample studies of child temperament.

interested in documenting changes in personality over the full life cycle. Developing the required measures is an active area of research.

5.5. Alternatives to the Big Five

The Five-Factor model is not without its critics. Alternative systems have been proposed. For example, Eysenck (1991) offers a model with just three factors (i.e., Neuroticism, Extraversion, and Psychoticism). Cloninger (1987) and Tellegen (1985) offer different three-factor models. Figure 1.4 shows the commonalities across some competing taxonomies and also areas of divergence. Solutions with more factors can increase the prediction of outcomes including job performance, income, and change in psychiatric status (Mershon and Gorsuch, 1988). On the other hand, more parsimonious models in which the five factors are reduced to two "metatraits" have also been suggested (Digman, 1997). In addition to these controversies, the facet-level organization of any given Big Five factor is subject to debate and controversy.

Recent research suggests that the rush to accept the Big Five may be premature.¹⁴³ The first studies of the Big Five were based primarily on English-speaking samples. Although the Big Five structure appears to replicate across many cultures (McCrae and Costa, 1997b), studies of more diverse cultures show that taxonomies known as the Big Six (Ashton et al., 2004) or the Multi-Language Seven (ML7; Saucier, 2003), may better represent the personality domain. Although they add one or two dimensions to the Big Five and shift the meaning of the Big Five slightly, they are, however, not very different from the Big Five.¹⁴⁴

One of the most stinging criticisms of the Five-Factor model is that it is atheoretical (Block, 1995). It is derived from factor analysis of a variety of measures without any firm biological underpinnings. Although research is under way on determining the neural substrates of the Big Five (see Canli, 2006, and DeYoung et al., 2010), the finding that descriptions of behavior as measured by tests, self-reports, and reports of observers cluster reliably into five groups has not so far been satisfactorily explained by any scientifically grounded theory.

Some psychologists suggest that the categories are too crude to be useful. Estimates based on the Big Five factors obscure relationships between specific facets of the Big Five and outcomes.¹⁴⁵ Given that each Big Five factor is a composite of distinct facets, the predictive validities are diluted when analyses consider only factor-level aggregate scores. For instance, Paunonen and Ashton (2001) compared Big Five Conscientiousness and Openness to Experience with two related facets, need for achievement and

¹⁴³ This discussion draws heavily on Roberts (2006).

¹⁴⁴ See Roberts (2006) for a description of these shifts.

¹⁴⁵ Hough (1992), Hough and Oswald (2000), and Paunonen and Ashton (2001).



Figure 1.4 Competing Taxonomies of Personality.

Source: Reproduced from Bouchard and Loehlin (2001), with kind permission from Springer Science and Business Media.

need for understanding. In each comparison, the lower level facet predicted course grades among undergraduates better than its higher level factor measure.

The Five-Factor model is largely silent about motivation. In the notation of Section 3, ψ parameterizes preferences and goals. The omission of motivation (i.e., what people value or desire) from measures of Big Five traits is not complete, however. The NEO-PI-R, for example, includes as a facet "achievement striving." Individual

differences in motivation were more prominent in older (now rarely used) measures of personality. The starting point for Jackson's Personality Research Form (PRF; Jackson, 1974), for example, was Murray's (1938) theory of basic human drives. Included in the PRF are scales for (need for) play, order, autonomy, achievement, affiliation, social recognition, and safety. The Schwartz Values Survey (Schwartz, 1992) is another self-report measure of motivation, which yields scores on 10 different motivations including power, achievement, benevolence, and conformity. Some motivation theorists believe that one's deepest desires are unconscious and, therefore, may dispute the practice of measuring motivation using self-report questionnaires (see McClelland, Koestner, and Weinberger, 1989). For a brief review of this debate and an overview of how motivation and personality trait measures differ, see Roberts, Harms, Smith, Wood, and Webb (2006).

A practical problem facing the analyst who wishes to measure personality is the multiplicity of personality questionnaires. The proliferation of personality measures reflects, in part, the more heterogeneous nature of personality in comparison to cognitive ability, although, as we have seen, various types of cognitive ability have been distinguished in the literature.¹⁴⁶ The panoply of measures and constructs also points to the relatively recent and incomplete convergence of personality psychologists on the Big Five model, as well as the lack of consensus among researchers about identifying and organizing lower order facets of the Big Five factors (see DeYoung, Quilty, and Peterson, 2007, and Hofstee, de Raad, and Goldberg, 1992). For example, some theorists argue that impulsivity is a facet of Neuroticism (Costa and McCrae, 1992b), others claim that it is a facet of Conscientiousness (Roberts, Chernyshenko, Stark, and Goldberg, 2005), and still others suggest that it is a blend of Conscientiousness, Extraversion, and perhaps Neuroticism (Revelle, 1997). Figure 1.4 shows in italics facets of the Big Five whose classifications are under question. One reason for the proliferation of measures is the variety of alternative methodologies for verifying tests discussed in Section 4, which are not guaranteed to produce the same taxonomies.

5.5.1 Self-Esteem and Locus of Control Are Related to Big Five Emotional Stability The traits of self-esteem and locus of control deserve special attention since they are collected in large-sample longitudinal studies used by economists.¹⁴⁷ They are not part of the traditional Big Five typology. However, they can be related to it.

Self-esteem refers to an individual's subjective estimation of his or her own worth. An example item from the widely used Rosenberg Self-Esteem Scale (Rosenberg, 1989) asks respondents to indicate their agreement with the statement, "I feel that

¹⁴⁶ See, e.g., Carroll (1993).

¹⁴⁷ See, e.g., NLSY79-based studies Heckman, Stixrud, and Urzua (2006) and some of the models in Heckman, Humphries, Urzua, and Veramendi (2011). The German Socioeconomic Panel (GSOEP) also collects these measures.

I am a person of worth, at least on an equal plane with others." Locus of control refers to one's belief about whether the determinants of one's life events are largely internal or external. Those with an internal (as opposed to external) locus of control believe that life events are typically caused by their own actions. An example item from the widely used Rotter Locus of Control Scale (Rotter, 1966) requires respondents to choose between "Many of the unhappy things in people's lives are partly due to bad luck" and "People's misfortunes result from the mistakes they make."

For the most part, researchers who study self-esteem and locus of control have carried out their work in isolation of each other and without reference to the Big Five taxonomy. However, Judge and colleagues (Judge, Bono, Erez, and Locke, 2005; Judge, Erez, Bono, and Thoresen, 2002; Judge and Hurst, 2007) have proposed that locus of control, self-esteem, and Big Five Emotional Stability are indicators of a common construct, termed *core self-evaluations*. They point out that measures of these three traits, as well as generalized self-efficacy (the belief that one can act effectively to bring about desired results), demonstrate high convergent validity, and discriminant validity. There is a negligible gain in predictive power of adding components of the construct beyond the predictive power of the construct itself. Positive core self-evaluations indicate a generally positive and proactive view of oneself and one's relationship to the world. Accordingly, we have, in Table 1.3, associated aspects of core self-evaluations with the Big Five factors of Neuroticism and Emotional Stability.

5.5.2 Relating the Big Five to Measures of Psychopathology

Extreme manifestations of personality traits may be a form of mental illness. Thus, a very conscientious person may be viewed as an obsessive-compulsive person. It is of interest to consider how psychopathology may be characterized using the Big Five.

Psychopathology is defined by the APA dictionary as "patterns of behavior or thought processes that are abnormal or maladaptive." Used interchangeably with the terms mental illness and mental disorder, psychopathology is primarily studied by psychiatrists and clinical psychologists. Historically, the study of psychopathology was carried out in near complete isolation from the study of "normal" variation in personality. Very recently, however, several attempts have been made to integrate taxonomies of psychopathology and normal personality into a single framework. In particular, a compelling argument can be made for conceptualizing and measuring mental disorders as extreme variants of personality traits (see Krueger and Eaton, 2010, and ensuing commentary). This approach is quite revolutionary in the study of psychopathology in at least two ways. First, it takes a dimensional as opposed to categorical characterization of mental disorders. By a dimensional approach, psychologists mean that traits lie on an underlying continuum and are not discrete valued. Second, the recent research relies on structural validity (e.g., evidence of convergent and discriminant validity) rather than historical path dependency (e.g., diagnoses that persist because they are familiar to clinicians who learned about them during their training). The Diagnostic and Statistical Manual (DSM) of the American Psychiatric Association distinguishes between Axis I disorders, which are acute disorders requiring clinical attention (e.g., depression, schizophrenia), and Axis II disorders,¹⁴⁸ which are 10 personality disorders that are more chronic and, generally, less impairing of overall functioning. Research has documented that Big Five Neuroticism is a nonspecific correlate of various Axis I disorders and that various other reliable associations can be documented (e.g., the positive emotionality facet of Extraversion is associated with bipolar disorder). The direction of causality is difficult to ascertain in typical crosssectional studies (Bagby et al., 1997; Cloninger, Svrakic, Bayon, and Przybeck, 1999; Gunderson, Triebwasser, Phillips, and Sullivan, 1999). Twin studies demonstrate that the shared variance in mental disorders and personality traits is predominantly genetic, that is, common genetic antecedents give rise to certain mental disorders and personality traits.

More research has examined relations between Axis II disorders and normal personality variation. Several authors have proposed a Big Four taxonomy (the Big Five minus Openness to Experience). Watson, Clark, and Chmielewski (2008) proposed that a fifth factor called Oddity is needed to model traits related to eccentricity. Others have argued that the Big Five structure itself, without modification, can account for Axis II personality disorders (Widiger and Costa, 2002, and Widiger, Trull, Clarkin, Sanderson, and Costa, 2002). For instance, Widiger, Trull, Clarkin, Sanderson, and Costa (2002) suggest that all Axis II personality disorders can be "translated as maladaptively extreme variants of the 30 facets" of the NEO-PI-R. More recently, Samuel and Widiger (2008) completed a meta-analytic review of the relationships between facets of the NEO-PI-R and Axis II personality disorders, which we reproduce in Table 1.4. Notably, personality disorders relate to multiple facets spanning more than one Big Five factor.

5.6. IQ and Achievement Test Scores Reflect Incentives and Capture Both Cognitive and Personality Traits

We now elaborate on the discussion of Section 3 on the difficulty of isolating a pure measure of intelligence. Performance on intelligence and achievement tests depends, in part, on personality traits of the test taker, as well as their motivation to perform.¹⁴⁹ A smart child unable to sit still during an exam or uninterested in exerting much effort can produce spuriously low scores on an IQ test.

¹⁴⁸ Axis II also includes mental retardation.

¹⁴⁹ It is likely that performance on personality tests also depends on cognitive ability, but that is less well documented. For example, it is likely that more intelligent people can ascertain the rewards to performance on a personality inventory test. Motivation is sometimes, but not usually, counted as a personality trait.

	FFM Facet	Paranoid	Schizoid	Schizotypal	Antisocial	Borderline	Histrionic	Narcissistic	Avoidant	Dependent	Obsessive
Ν	Anxiousness	0.27	0.13	0.27	0.00	0.38	0.00	0.02	0.41	0.39	0.16
	Angry hostility	0.41	0.19	0.29	0.27	0.48	0.08	0.23	0.29	0.18	0.10
	Depressiveness	0.35	0.28	0.39	0.12	0.50	-0.06	0.03	0.53	0.41	0.09
	Self-consciousness	0.29	0.23	0.32	0.02	0.35	-0.11	-0.03	0.56	0.42	0.13
	Impulsiveness	0.15	0.00	0.17	0.27	0.34	0.17	0.14	0.14	0.17	-0.07
	Vulnerability	0.22	0.14	0.25	0.04	0.39	0.01	-0.01	0.40	0.43	0.03
Е	Warmth	-0.28	-0.42	-0.28	-0.13	-0.20	0.26	-0.07	-0.35	-0.03	-0.07
	Gregariousness	-0.20	-0.48	-0.25	0.02	-0.12	0.35	0.04	-0.42	-0.03	-0.16
	Assertiveness	-0.08	-0.22	-0.13	0.06	-0.09	0.27	0.19	-0.39	-0.21	-0.01
	Activity	-0.08	-0.25	-0.13	0.02	-0.10	0.25	0.09	-0.29	-0.12	0.03
	Excitement seeking	-0.01	-0.21	-0.04	0.25	0.06	0.27	0.16	-0.23	-0.06	-0.12
	Positive emotions	-0.27	-0.38	-0.26	-0.09	-0.26	0.23	-0.02	-0.39	-0.15	-0.09
Ο	Fantasy	0.00	-0.05	0.14	0.10	0.13	0.16	0.11	0.00	0.05	-0.09
	Aesthetics	-0.05	-0.06	0.07	0.00	0.05	0.10	0.04	-0.03	0.01	0.01
	Feelings	-0.02	-0.17	0.03	-0.02	0.09	0.18	0.05	-0.04	0.05	0.01
	Actions	-0.10	-0.13	-0.06	0.10	-0.03	0.12	0.04	-0.20	-0.13	-0.12
	Ideas	-0.03	0.00	0.09	0.04	-0.01	0.04	0.07	-0.05	-0.12	0.03
	Values	-0.05	-0.05	0.01	0.08	0.05	0.04	-0.01	-0.05	-0.04	-0.09
А	Trust	-0.45	-0.28	-0.31	-0.22	-0.29	0.05	-0.20	-0.29	-0.07	-0.08
	Straightforwardness	-0.24	-0.09	-0.16	-0.37	-0.21	-0.10	-0.31	-0.06	0.00	0.04
	Altruism	-0.21	-0.19	-0.15	-0.24	-0.18	0.02	-0.20	-0.12	0.03	0.04
	Compliance	-0.27	-0.08	-0.13	-0.32	-0.27	-0.12	-0.26	-0.02	0.10	0.01
	Modesty	-0.06	0.08	0.05	-0.17	0.03	-0.16	-0.37	0.20	0.16	0.02
	Tender-mindedness	-0.18	-0.11	-0.05	-0.19	-0.09	0.02	-0.17	-0.02	0.09	0.00
С	Competence	-0.13	-0.13	-0.18	-0.21	-0.29	-0.01	0.01	-0.23	-0.25	0.19
	Order	0.00	-0.02	-0.06	-0.18	-0.10	-0.05	-0.03	-0.03	-0.06	0.25
	Dutifulness	-0.10	-0.08	-0.10	-0.29	-0.22	-0.08	-0.10	-0.09	-0.08	0.25
	Achievement	-0.07	-0.13	-0.13	-0.19	-0.19	0.04	0.02	-0.19	-0.16	0.25
	striving										
	Self-discipline	-0.14	-0.12	-0.18	-0.25	-0.29	-0.04	-0.09	-0.22	-0.23	0.21
	Deliberation	-0.09	-0.02	-0.10	-0.38	-0.27	-0.16	-0.13	-0.01	-0.06	0.24

Table 1.4 Mean Correlations of Psychopathological Measures with the Big Five Traits

Notes: All values larger than r = 0.04 are significant at p < 0.05; correlations larger than 0.20 are marked in **boldface** type.

Source: Reproduced from Samuel and Widiger (2008).

It is sometimes claimed that IQ tests measure maximal performance, that is, that IQ scores reflect the application of the maximal capacity of the person to the test.¹⁵⁰ The analysis of Section 3 suggests that IQ scores should be standardized for effort. A series of studies conducted over the past 40 years support this concern.

These studies show that among individuals with low IQ scores, performance on IQ tests could be increased up to a full standard deviation by offering incentives such as money or candy, particularly on group-administered tests and particularly with individuals at the low end of the IQ spectrum.¹⁵¹ Engaging in complex thinking is effortful, not automatic (Schmeichel, Vohs, and Baumeister, 2003), and therefore, motivation to exert effort affects performance. Zigler and Butterfield (1968) found that early intervention (nursery school, for example) for low-SES kids may have a beneficial effect on motivation, not on cognitive ability per se. In their study, the benefits of intervention (in comparison to a no-treatment control group) on IQ were not apparent under testing conditions in which motivation to perform well was maximal. Raver and Zigler (1997) present further evidence on this point. Table 1.5 summarizes evidence that extrinsic incentives can substantially improve performance on tests of cognitive ability, especially among low-IQ individuals.¹⁵²

Segal (2008) shows that introducing performance-based cash incentives in a lowstakes administration of the coding speed test of the Armed Services Vocational Battery (ASVAB) increases performance substantially among roughly one-third of participants. Men with lower levels of Conscientiousness are particularly affected by incentives. Segal's work and a large body of related work emphasize heterogeneity in the motivations that affect human performance. Borghans, Meijers, and ter Weel (2008) show that adults spend substantially more time answering IQ questions when rewards are higher, but subjects high in Emotional Stability and Conscientiousness are less affected by these incentives. They already operate at a high level even without these incentives. Similarly, Pailing and Segalowitz (2004) find that an event-related potential (ERP) indexing the emotional response to making an error increases in amplitude when incentives are offered for superior test performance.¹⁵³ This effect is smaller for individuals high in Conscientiousness and Emotional Stability. Thus, IQ scores do not accurately reflect maximal intellectual performance for individuals who are low in Conscientiousness and Emotional Stability. Performance on IQ tests encodes, in part, how effective persons may be in application of their intelligence, that is, how people are likely to perform in a real-world setting. However, it is far from obvious that motivation on an exam and motivation in a real-world situation are the same.

¹⁵⁰ A leading psychometrician, Carroll (1993), does not accept the notion that IQ captures maximal effort.

¹⁵¹ The incentives for invoking effort vary across studies.

¹⁵² The studies do not include direct measures of personality traits.

¹⁵³ An ERP is an electrophysiological response of characteristic form and timing to a particular category of stimuli.

Table	1.5	Incentives	and	Performance	on	Intelligence	Tests
-------	-----	------------	-----	-------------	----	--------------	-------

Study	Sample and Study Design	Experimental Group	Effect Size of Incentive (in Standard Deviations)	Summary
Edlund (1972)	Between subjects study. 11 matched pairs of low-SES children; children were about one standard deviation below average in IQ at baseline	M&M candies given for each right answer	Experimental group scored <u>12 points</u> higher than control group during a second testing on an alternative form of the Stanford–Binet (about 0.8 standard deviations)	" a carefully chosen consequence, candy, given contingent on each occurrence of correct responses to an IQ test, can result in a significantly higher IQ score." (p. 319)
Ayllon and Kelly (1972) Sample 1	Within subjects study. 12 mentally retarded children (avg IQ, 46.8)	Tokens given in experimental condition for right answers exchangeable for prizes	6.25 points out of a possible 51 points on Metropolitan Readiness Test. $t = 4.03$	" test scores often reflect poor academic skills, but they may also reflect lack of motivation to do well in
Ayllon and Kelly (1972) Sample 2	Within subjects study. 34 urban fourth graders (avg IQ = 92.8)	Tokens given in experimental condition for right answers exchangeable for prizes	t = 5.9	the criterion test These results, obtained from both a population typically limited in skills and ability,
Ayllon and Kelly (1972) Sample 3	Within subjects study of 12 matched pairs of mentally retarded children	Six weeks of token reinforcement for good academic performance	Experimental group scored 3.67 points higher out of possible 51 points on a posttest given under standard conditions higher than at baseline; control group dropped 2.75 points. On a second posttest with incentives, expand control groups increased 7.17 and 6.25 points	as well as from a group of normal children (Experiment II), demonstrate that the use of reinforcement procedures applied to a behavior that is tacitly regarded as "at its peak" can significantly alter the level of performance of that behavior." (p. 483)

Continued

Study	Sample and Study Design	Experimental Group	Effect Size of Incentive (in Standard Deviations)	Summary
Clingman and Fowler (1976)	Within subjects study of 72 first and second graders assigned randomly to contingent reward, noncontingent reward, or no reward conditions	M&Ms given for right answers in contingent condition; M&Ms given regardless of correctness in noncontingent condition	Only among low-IQ (<100) subjects was there an effect of the incentive. Contingent reward group scored about 0.33 standard deviations higher on the Peabody Picture Vocabulary test than did no reward group	" contingent candy increased the IQ scores of only the "low-IQ" children. This result suggests that the high and medium-IQ groups were already functioning at a higher motivational level than children in the low- IQ group." (p. 22)
Zigler and Butterfield (1968)	Within and between subjects study of 52 low- SES children who did or did not attend nursery school were tested at the beginning and end of the year on Stanford–Binet Intelligence Test under either optimized or standard conditions to be consistent with previous sample descriptions.	Motivation was optimized without giving test-relevant information. Gentle encouragement, easier items after items were missed, and so on.	At baseline (in the fall), there was a full standard deviation difference (10.6 points and SD was about 9.5 in this sample) between scores of children in the optimized vs. standard conditions. The nursery group improved their scores, but only in the standard condition.	" performance on an intelligence test is best conceptualized as reflecting three distinct factors: (a) formal cognitive processes; (b) informational achievements that reflect the content rather than the formal properties of cognition, and (c) motivational factors that involve a wide range of personality variables. (p. 2) " the significant difference in improvement in standard IQ performance found between the nursery and non-nursery groups was attributable solely to motivational factors" (p. 10)

Table 1.5 Incentives and Performance on Intelligence Tests—continued

Study	Sample and Study Design	Experimental Group	Effect Size of Incentive (in Standard Deviations)	Summary
Breuning and Zella (1978)	Within and between subjects study of 485 <i>special</i> <i>education</i> high-school students all took IQ tests, then were randomly assigned to control or incentive groups to retake tests. Subjects were below average in IQ.	Incentives such as record albums, radios (<\$25) given for improvement in test performance	Scores increased by about 17 points. Results were consistent across the Otis- Lennon, WISC-R, and Lorge-Thorndike tests.	"In summary, the promise of individualized incentives contingent on an increase in IQ test performance (as compared with pretest performance) resulted in an approximate 17-point increase in IQ test scores. These increases were equally spread across subtests The incentive condition effects were much less pronounced for students having pretest IQs between 98 and 120 and did not occur for students having pretest IQs between 121 and 140." (p. 225)
Holt and Hobbs (1979)	Between and within subjects study of 80 delinquent boys randomly assigned to three experimental groups and one control group. Each exp group received a standard and modified administration of the WISC-verbal section.	Exp 1: Token reinforcement for correct responses; Exp 2: Tokens forfeited for incorrect responses (punishment), Exp 3: feedback on correct/incorrect responses.	1.06 standard deviation difference between the token reinforcement and control groups (inferred from $t = 3.31$ for 39 degrees of freedom).	"Knowledge of results does not appear to be a sufficient incentive to significantly improve test performance among below-average IQ subjects Immediate rewards or response cost may be more effective with below-average IQ subjects while other conditions may be more effective with average or above-average subjects." (p. 83)

Table 1.5 Incentives and Performance on Intelligence Tests—continued

Study	Sample and Study Design	Experimental Group	Effect Size of Incentive (in Standard Deviations)	Summary
Larson, Saccuzzo, and Brown (1994)	Between subjects study of 109 San Diego State University psychology students.	Up to \$20 for improvement over baseline performance on cognitive speed tests.	"While both groups improved with practice, the incentive group improved slightly more." (p.34) F(1,93) = 2.76, p < 0.05	Two reasons why incentive did not produce dramatic increase: few or no unmotivated subjects among college volunteers; information processing tasks are too simple for "trying harder" to matter.
Duckworth (2007)	Within subjects study of 61 urban low-achieving high- school students tested with a group-administered Otis- Lennon IQ test during their freshman year, then again two years later with a one-on-one (WASI) test.	Standard directions for encouraging effort were followed for the WASI brief test. Performance was expected to be higher because of the one-on-one environment.	Performance on the WASI as juniors was about 16 points higher than on the group-administered test as freshmen. Notably, on the WASI, this population looks almost "average" in IQ, whereas by Otis–Lennon standards they are low IQ. $t(60) = 10.67$, $p < 0.001$.	The increase in IQ scores could be attributed to any combination of the following: an increase in g due to schooling at an intensive charter school; an increase in knowledge or crystallized intelligence; an increase in motivation due to the change in IQ test format; and/or an increase in motivation due to experience at high- performing school.

Table 1.5 Incentives and Performance on Intelligence Tests—continued

Like low motivation, test anxiety can significantly impair performance (Hembree, 1988). That is, subjects do worse when they worry excessively about how they are performing and when their autonomic nervous system overreacts by increasing perspiration, heart rate, and so on. Because individuals who are higher in Big Five Neuroticism are more likely to experience test anxiety, there is another reason, beyond incentives, why Emotional Stability can impact IQ scores (Moutafi, Furnham, and Tsaousis, 2006).

Many IQ tests require factual knowledge acquired through schooling and life experience, which are, in part, determined by the motivation, curiosity, and persistence of the test taker. Thus, personality traits can also affect IQ scores indirectly through the knowledge acquired by individuals who are higher in Big Five Openness to Experience and Big Five Conscientiousness. Cunha and Heckman (2008) show a correlation between cognitive and personality factors of the order of r = 0.3. Hansen, Heckman, and Mullen (2004) and Heckman, Stixrud, and Urzua (2006) show how schooling and other acquired traits substantially *causally* affect measured cognitive and personality test scores. We discuss this research in Section 8. Cattell's investment theory (1971) anticipates recent findings that knowledge and specific complex skills depend not only on the fluid intelligence but also on the cumulative investment of effort and exposure to learning opportunities.

How, then, should one interpret a low IQ score? Collectively, the evidence surveyed here suggests that IQ test performance reflects not only pure intelligence but also personality traits (including anxiety), intrinsic motivation, and reactions to extrinsic incentives to perform well as indicated in our discussion of Section 3. It also reflects the knowledge acquired up to the date of the test, which reflects personality and motivational traits that affect the acquisition of knowledge. The relative impurity of IQ tests likely varies from test to test and individual to individual. Little effort to date has been made to standardize the context and incentives of tests. To capture pure intelligence, it is necessary to adjust for incentives, motivations, and context in which the measurements are taken, using the framework discussed in Section 3.

Just as personality traits and incentives can affect IQ scores, they can also affect standardized achievement tests that are commonly used as proxies for pure intelligence. Figures 1.5 and 1.6, below, show how scores on two achievement tests, the Armed Forces Qualifying Test (AFQT) and the Differential Aptitudes Test (DAT), are decomposed into IQ and personality measures.¹⁵⁴ We adjust by Rotter and Rosenberg in Fig. 1.5 and by the Big Five in Fig. 1.6.¹⁵⁵ A substantial portion of the variances in both AFQT scores and

¹⁵⁴ AFQT and DAT scores are highly correlated (r = 0.76). See Borghans, Golsteyn, Heckman, and Humphries (2011); Kilburn, Hanser, and Klerman (1998); Sticht (1995); and Wang (1993).

 $^{^{155}\,}$ The Big Five are not available in the NLSY79 data that have AFQT scores.


Figure 1.5 AFQT Decomposed by IQ, Rosenberg, and Rotter: (a) Not Controlling for Background Characteristics. (b) Controlling for Background Characteristics.

Notes: The data come from the NLSY79. Rotter was administered 1979. The ASVAB and Rosenberg were administered in 1980. To account for varying levels of schooling at the time of the test, scores have been adjusted for schooling at the time of the test conditional on final schooling using the method developed in Hansen, Heckman, and Mullen (2004). AFQT is constructed from the Arithmetic Reasoning, Word Knowledge, Mathematical Knowledge, and Paragraph Comprehension ASVAB subtests. IQ and GPA are from high-school transcript data. IQ is pooled across several IQ tests using IQ percentiles. GPA is the individual's core-subject GPA from ninth grade. Sample excludes the military oversample. Background variables include race and sex dummies, mother's highest grade completed, father's highest grade completed, southern residence at age 14, urban residence at age 14, living in a broken home at age 14, receiving newspapers in the household at age 14, receiving magazines in the household at age 14, and the household having a library card at age 14. Top 50% and bottom 50% are based on AFQT scores from the cross-sectional sample of the NLSY79.

Source: Borghans, Golsteyn, Heckman, and Humphries (2011).



Figure 1.6 DAT Scores and GPA Decomposed by IQ and Personality.

Notes: Data is from Stella Maris, a high school in the Netherlands. Students were administered part of a Raven's IQ test and personality questions based on the Big Five. DAT and GPA are from high-school records. Source: Borghans, Golsteyn, Heckman, and Humphries (2011).

DAT scores are explained by personality factors.¹⁵⁶ The variance explained is less than the variance independently explained by IQ scores, but it is still substantial. Furthermore, the facets are incrementally valid in that they explain the variance above and beyond the variance that IQ explains when all three are included in a regression. These findings caution the interpretation that these commonly used tests proxy mental ability. They likely proxy aspects of personality as well. Ironically, the measure of intelligence used by Herrnstein and Murray in *The Bell Curve* (1994) to predict a variety of social and economic outcomes is substantially affected by personality measures. We discuss evidence about personality and standardized achievement tests further in Section 7.

5.7. The Evidence on the Situational Specificity Hypothesis

Since the publication of Mischel's (1968) book, psychologists have addressed the situational specificity hypothesis, that is, the hypothesis that situations help explain variations across people in actions, effort, and behavior.¹⁵⁷ Boiled down to its essence, this hypothesis says little more than that situations affect actions and efforts in a nonlinear fashion, that is, that in Eqs (1.13)–(1.15), situational variables enter in a nonlinear fashion. This interaction effect gives rise to the Mischel and Shoda (1995) "if-then" relationship.

¹⁵⁶ The lower explained variance in the sample with DAT is likely a consequence of restriction on range. The DAT data come from a single school, whereas the AFQT data come from a national sample.

¹⁵⁷ For an early symposium in psychology on the person-situation debate, see Endler and Magnusson (1976).

An important paper by Epstein (1979) defines stability of personality generated by traits across situations using measurements that average across tasks. He notes that in the presence of nonlinearities, agents with the same traits will take different actions in different situations. In four different studies, he presents compelling empirical evidence that, averaging over tasks and situations at a point in time, persons act in a predictable fashion with a high level of reliability (R^2 of 0.6–0.8) of average behavior ("measured personality") across situations. He uses a variety of measures based on objective behaviors, self-ratings, and ratings by others. He also establishes consistency (high levels of correlation) across the different types of measures. In any given situations, stable patterns emerge. Fleeson (2001) and Moskowitz (1982) present additional evidence on this question. Fleeson and Noftle (2008) summarize a substantial body of evidence on the stability of behaviors across tasks and situations and the evidence of consistency of different measurements of personality (e.g., self-reports, observer reports).

In one of the most ambitious recent studies of this question, Borkenau, Mauer, Riemann, and Angleitner (2004) establish a correlation of 0.43 of personality traits measured by the Big Five (self-rated and observer-rated) across 15 very different tasks. The range of correlations is from 0.29 to 0.51.¹⁵⁸ Wood and Roberts (2006) present further evidence on the persistence of traits across a variety of situations. Roberts (2009) provides a valuable overview of the latest research. Funder (2008) provides another useful overview of the debate and the evidence on the existence of a stable personality trait that at a point in time predicts behavior in a variety of different situations. Mischel's (1968) claim that there is no stable personality trait across situations does not hold up against a large array of data.

A recent summary of the evidence on the person-situation debate is provided in a series of papers in the *Journal of Research in Personality* (January, 2009, Vol. 43) that offers a retrospective on the controversy. Virtually, all papers in that special issue acknowledge the existence of stable personality traits whose manifestations are affected by situations and incentives. The editors summarize the main message of the collected papers with the following words:

"All personality psychologists should be unified when it comes to asserting that personality differences are worthy of scientific study, that individual differences are more than just error variance and that not all behavior is simply a function of the situation."

(Lucas and Donnellan, 2009, p. 147)

¹⁵⁸ Achenbach, McConaughy, and Howell (1987) summarize correlations between children's problem behavior ratings by parents and teachers. Their meta-analysis produces an estimate of r = 0.28 and suggests consistency and variation in behavior and assessment across home and school situations. Whether this arises from parental bias or from situational specificity is not clear.

6. PERSONALITY AND PREFERENCE PARAMETERS

Measures of personality predict a wide range of life outcomes that economists study. However, with our current knowledge, it is difficult to relate them to economic preference parameters except, of course, when the traits are the parameters. Since personality psychologists define traits as "relatively" stable, person-specific determinants of behavior, preferences are the natural counterpart of these traits in economics. Preferences are unaffected by changes in constraints. Although personality might relate to preferences, the exact link is unclear. Do preferences generate measured personality? Does personality generate preferences? Or, are both generated from other, deeper, motivation parameters that are as yet unknown? The model in Section 3 links preferences to measured personality. This section reviews the empirical evidence linking preferences and personality and discusses the conceptual differences between the two.

Overall, the links between measures of personality and preferences are largely unexplored. However, some evidence suggests that social preferences can be linked to the Big Five. The links between traditional preferences, such as risk aversion and time discounting, and personality, remain largely unknown. Personality measures might allow economists to broaden the dimensions of preferences and could potentially resolve some apparent inconsistencies in observed choices that arise from commonly used preference specifications in economics.

6.1. Evidence on Preference Parameters and Corresponding Personality Measures

The aspects of preferences that receive the most attention in the economics literature time discounting, risk aversion, leisure preference, and social preferences¹⁵⁹—appear to have analogues in the literature in psychology. Table 1.6 presents the definitions of commonly used preferences, tasks, and survey questions that have been used to measure them, and an overview of how they relate to measures of personality. The table includes measures, as well as latent factors (see Section 4).

Since the 1960s, psychologists have used experiments to elicit time preference and risk preference (see, e.g., Mischel et al., 2010, and Slovic, 1962). A recent example is the Balloon Analogue Risk Task (BART) (Lejuez et al., 2002), a computer game in which participants make repeated choices between keeping a certain smaller monetary reward and taking a chance on an incrementally larger reward. In addition to the experimental measures, it is tempting to try to map preferences to more vaguely defined traits, but the precise mapping has not yet been made. Still, some speculation is useful. Time preference

¹⁵⁹ For a definition of these concepts and a discussion of measurement of preferences, see Table 1.6 and Sections A6.1 and A6.2 in the appendix.

Preference	Survey Questions and Experiments Used to Elicit Preference	Overview of Relationship to Personality Measures
<i>Time preference</i> : Preference over consumption in	<i>Delay discounting</i> : A participant is given a series of choices for whether he would prefer to receive smaller payments sooner versus larger payments later. The amounts and times vary	<i>Conceptual relationships</i> : Conscientiousness, Self- Control, Affective Mindfulness, Consideration of Future Consequences, Elaboration of
different time periods	across choices. The choices can be over hypothetical payoffs or real-stakes payoffs (see, e.g., Dohmen et al., 2011).	Consequences, Time Preference. Empirical relationships: Conscientiousness, Self-
	marshmallow. The experimenter leaves the room and tells the participant that he will receive a second marshmallow if he resists consuming the marshmallow until the	Consequences, Consideration of Future Consequences (Daly, Delaney, and Harmon, 2009)
	experimenter returns. The length of time that the participant waits is a measure of short-term discounting (see, e.g., Mischel et al., 2010).	Extraversion, Time Preference (Dohmen, Falk, Huffman, and Sunde, 2010). Agreeableness, Inhibitive Side of
	<i>Example survey question</i> : "How patient are you on a scale from 1 to 10?" (see GSOEP, 2008).	Conscientiousness (Anderson, Burks, DeYoung, and Rustichini, 2011).

Table 1.6 Measures of Standard Preference Parameters and Analogous Measures in the Psychology Literature

Continued

Preference	Survey Questions and Experiments Used to Elicit Preference	Overview of Relationship to Personality Measures
Risk aversion: Preference over different states of the world	Lottery choice task: A participant is given a series of choices between a safe amount of money and a lottery. The lottery remains the same across choices, whereas the safe amount varies. The lowest safe amount for which the participant prefers the lottery is a measure of risk aversion. The choices can be over hypothetical payoffs or real-stakes payoffs (see, e.g., Dohmen et al., 2011). <i>Devil's Task (Slovic's risk task)</i> : A participant sequentially chooses between ten "switches" or urns associated with hidden payoffs. The participant is told that nine of the switches are associated with a reward and one of them results in a loss of all previous winnings. Once a participant chooses a switch, he cannot flip the same switch again. The participant can elect to stop picking switches at any time. The number of switches chosen is a measure of risk aversion (see, e.g., Slovic, 1966). <i>Balloon Analogue Risk Task (BART)</i> : The participant is given a computerized task in which he is presented with a series of "balloons" that can be inflated by "pumping" the balloon. The participant receives potential earnings each time he pumps a balloon. At any point, the participant can stop pumping, realize the potential earnings, and move to the next balloon. After a threshold number of pumps each balloon "explodes," and the participant receives nothing. The threshold varies across balloons, and participants are not told the distribution of thresholds (see, e.g., Lejuez, Aklin, Zvolensky, and Pedulla, 2003). <i>Example survey question</i> : "How willing are you to take risks, in general?" (see, e.g., Dohmen et al., 2011).	Conceptual relationships: Impulsive Sensation Seeking, Balloon Analogue Risk Task. Empirical relationships: Sensation Seeking (Zuckerman, 1994; Eckel and Grossman, 2002). Openness (Dohmen, Falk, Huffman, and Sunde, 2010). Neuroticism, Ambition, Agreeableness (Borghans, Golsteyn, Heckman, and Meijers, 2009). Balloon Analogue Risk Task (Lejuez, Aklin, Zvolensky, and Pedulla, 2003). Neuroticism, Inhibitive Side of Conscientiousness (Anderson, Burks, DeYoung, and Rustichini, 2011).

Table 1.6 Measures of Standard Preference Parameters and Analogous Measures in the Psychology Literature—continued

Preference	Survey Questions and Experiments Used to Elicit Preference	Overview of Relationship to Personality Measures
<i>Leisure:</i> Preference over consumption and leisure	<i>Payments for working</i> : The participant is given a choice to work at different wages. Their reservation wage is their preference for leisure. The choices can be over hypothetical payoffs or real-stakes payoffs (see, e.g., Borghans, Meijers, and ter Weel, 2008).	Conceptual relationships: Achievement Striving, Endurance, Industriousness. Empirical relationships: Inconsistent with psychological measures of leisure preferences (Borghans, Meijers, and ter Weel, 2008).
<i>Altruism:</i> Unconditional kindness <i>Inequity aversion:</i> Value of equality in payoffs	<i>Dictator game</i> : A "proposer" has the option to transfer part of an endowment to a "responder." The responder passively receives any transfer. The transfer is used as a measure of pure altruism (see, e.g., Fehr and Schmidt, 2006).	Conceptual relationships: Warmth, Gregariousness, Tender-Mindedness, Hostility (opposite). Empirical relationships: Neuroticism, Agreeableness (Ashton, Paunonen, Helmes, and Jackson, 1998; Osiński, 2009; Bekkers, 2006; Ben-Ner and Kramer, 2011).
<i>Trust</i> : Willingness to make oneself vulnerable to opportunistic individuals	<i>Trust game</i> : An "investor" receives an endowment and can decide to transfer some of it to a "trustee." The amount transferred increases in value. The trustee can then decide to transfer some back to the investor but has no monetary incentive to do so. The amount the investor transfers to the trustee is used a measure of trust (see, e.g., Fehr and Schmidt, 2006). <i>Example survey question</i> : "In general, one can trust people" (see, e.g., Dohmen, Falk, Huffman, and Sunde, 2008).	Conceptual relationships: Trust. Empirical relationships: Neuroticism, Agreeableness, Openness, Conscientiousness (Dohmen, Falk, Huffman, and Sunde, 2008).

 Table 1.6 Measures of Standard Preference Parameters and Analogous Measures in the Psychology Literature—continued

Continued

Preference	Survey Questions and Experiments Used to Elicit Preference	Overview of Relationship to Personality Measures
Reciprocity: The way in which one person responds to another's actions. Positive reciprocity: Tendency to reward kind actions. Negative reciprocity: Tendency to punish others for unkind actions.	Ultimatum game: A "proposer" offers part of an endowment to a "responder." The responder can choose to accept the offer in which case both players keep the payoffs, or the responder can choose to reject the offer in which case the players receive nothing. The responder's choice is a measure of reciprocity (Fehr and Schmidt, 2006). <i>Trust game</i> : See above description. The trustee's action is used as a measure of reciprocity. <i>Gift exchange game</i> : An "employer" proposes a wage and an amount of desired effort to a potential "worker." The worker can either reject the proposal so that no one receives anything or can accept the proposal so that no one receives anything or can accept the proposal and choose any amount of effort. The employer receives a payment proportional to the worker's effort net of the wage. The workers action is used as a measure of reciprocity (see, e.g., Fehr and Schmidt, 2006). <i>Example survey question (positive reciprocity)</i> : "If someone does me a favor, I am prepared to return it." (see, e.g., Dohmen, Falk, Huffman, and Sunde, 2008). <i>Example survey question (negative reciprocity)</i> : "If I suffer a serious wrong, I will take revenge as soon as possible, no matter the cost." (see, e.g., Dohmen, Falk, Huffman, and Sunde, 2008).	Conceptual relationships: Warmth, Gregariousness, Hostility (opposite). Empirical relationships: Neuroticism, Agreeableness, Conscientiousness (Dohmen, Falk, Huffman, and Sunde, 2008).

Table 1.6 Measures of Standard Preference Parameters and Analogous Measures in the Psychology Literature—continued

likely relates to Conscientiousness, Self-control, and Consideration of Future Consequences. Risk Aversion is likely related to Openness to Experience and impulsive sensation seeking, a trait proposed by Zuckerman, Kolin, Price, and Zoob (1964), defined as "the tendency to seek novel, varied, complex, and intense sensations and experiences and the willingness to take risks for the sake of such experience."¹⁶⁰

Preferences for leisure may be related to several personality measures. The Big Five includes an Achievement Striving subscale of Conscientiousness, which describes ambition, the capacity for hard work, and an inclination toward purposeful behavior. Jackson's Personality Research Form (1974) includes an achievement scale measuring the aspiration to accomplish difficult tasks and to put forth effort to attain excellence, as well as an endurance scale, measuring willingness to work long hours and perseverance in the face of difficulty, and a play scale, measuring the inclination to participate in games, sports, and social activities "just for fun." Industriousness has been proposed as one of six facets of Conscientiousness (Roberts, Chernyshenko, Stark, and Goldberg, 2005) and is plausibly related to the preference for leisure.

Social preferences also have conceptual analogues in the personality literature. Warmth and Gregariousness are facets of Extraversion; Trust, Altruism, and Tender-Mindedness are facets of Agreeableness; and Hostility is a facet of Neuroticism.

Despite this intuitive mapping of preferences to traits, the empirical evidence supporting such mappings is weak. The few studies investigating empirical links typically report only simple regressions or correlations without discussing any underlying model. Some use survey and self-report measures similar to those used by psychologists rather than elicited preferences. The last column of Table 1.6 gives an overview of papers investigating these links.

The evidence relating personality to time preferences is mixed. Using data from an experiment involving college students, Daly, Delaney, and Harmon (2009) find that a factor that loads heavily on self-control, consideration of future consequences, elaboration of consequences, affective mindfulness, and Conscientiousness is negatively associated with the discount rate. Dohmen, Falk, Huffman, and Sunde (2010) measure time preferences experimentally. Although time preference is related to cognition, Openness to Experience is the only Big Five trait that explains any variation in time preference. Figure 1.7 reports correlations between experimental measures of time preference, Big Five factors, and measures of cognition. In that figure, only cognitive measures are correlated with time preference.¹⁶¹

Dohmen, Falk, Huffman, and Sunde (2010) find that Openness to Experience and Agreeableness are related to risk aversion. Figure 1.7 reports correlations between risk aversion, the Big Five, and measures of cognition for a sample of Germans. Of the

¹⁶⁰ See Zuckerman (1994).

¹⁶¹ Figure A2 in Section A6 of the Web Appendix displays correlations among the survey measures in the GSOEP.

Time preference								
_	Risk tolerance							
0.01	0.14***	0						
0.07	-0.07	0.11***	с					
0.07	0.08	0.42***	0.11***	Е				
0.08*	-0.09**	0.08***	0.29***	0.05	А			
0.06	-0.03	0.19***	0.06**	-0.01	-0.01	N		_
-0.12***	0.18***	0.16***	-0.15***	0.15***	-0.14***	-0.06**	Symbol test	
-0.10**	0.19***	0.15***	-0.13***	0.12***	-0.18***	-0.08**	0.40***	Word test

Figure 1.7 Pairwise Correlations between Time Preference (Impatience), Risk Tolerance, Personality, and Cognitive Ability for Males and Females from GSOEP.

Notes: *Statistically significant at the 10% level; **statistically significant at the 5% level; ***statistically significant at the 1% level. **O**, Openness to Experience; **C**, Conscientiousness; **E**, Extraversion; **A**, Agree-ableness; **N**, Neuroticism. The value in each box is the pairwise correlation. Darker-shaded boxes have lower p-values. The measures of the Big Five are based on three questions each. The measures of cognitive ability (symbol test and word test) are based on timed modules similar to the Wechsler Adult Intelligence Scale (WAIS). Time preference and risk tolerance were elicited through a real-stakes experiment. Source: The data come from Dohmen, Falk, Huffman, and Sunde (2010), available online. The calculations were conducted by the authors of this Handbook chapter.

Big Five, Openness to Experience and Agreeableness are correlated with risk aversion. There is little evidence connecting risk aversion and sensation seeking, but Eckel and Grossman (2002) include it as a control in a study of risk aversion and find no statistically significant effect. However, Bibby and Ferguson (2011) find that sensation seeking is associated with a lottery measure of risk tolerance.¹⁶² They also find that people who

¹⁶² Bibby and Ferguson report this as a measure of loss aversion, but it is more akin to a measure of risk tolerance.

are better at processing emotional information and who are less extraverted are more susceptible to framing effects when making risky decisions. Borghans, Golsteyn, Heckman, and Meijers (2009) show that risk aversion is positively associated with Neuroticism, which contains measures of fear and strong emotional responses to bad outcomes. They also find that risk aversion is negatively associated with ambition, a trait that may involve investment in uncertain opportunities. Further, Agreeableness is positively associated with risk aversion.

Figure 1.8 displays a related analysis by Anderson, Burks, DeYoung, and Rustichini (2011), who find that both cognitive ability and Agreeableness are positively associated with delay acceptance elicited from a real-stakes experiment in a sample of truck driver trainees.¹⁶³

When they separately regress delay acceptance on Neuroticism, Agreeableness, Extraversion, Conscientiousness, cognitive skill, race, marital status, age, and education, none of the personality traits are statistically significant at the 10% level. However, when they split Conscientiousness into an inhibitive side (moral scrupulousness and cautiousness) and a proactive side (the need for achievement), they find that only the inhibitive side is positively associated with delay acceptance ($\beta = 0.13$, p < 0.10). This result highlights the importance of examining facets of the Big Five when considering the relationship between preferences and personality.

As shown in Fig. 1.8, Anderson, Burks, DeYoung, and Rustichini (2011) find that of the Big Five, only Neuroticism is positively associated with risk aversion but only for lotteries over gains, not losses. In a separate regression controlling for Neuroticism, Agreeableness, Extraversion, Conscientiousness, cognitive skill, race, marital status, age, and education, risk aversion is positively associated with both Neuroticism ($\beta = 0.15$, p < 0.01) and the inhibitive side of Conscientiousness ($\beta = 0.10$, p < 0.10).

The links between social preferences and the Big Five traits are better established. Ben-Ner and Kramer (2011) find that Extraversion is associated with higher giving in a dictator game. Dohmen, Falk, Huffman, and Sunde (2008) use an experimentally validated survey measure of trust and find that Conscientiousness and Neuroticism are negatively associated with trust, whereas Agreeableness and Openness to Experience are positively associated with trust. Agreeableness and Conscientiousness are associated with more positive reciprocity and less negative reciprocity, whereas Neuroticism is associated with more negative reciprocity.

In sum, although many measures of personality and preferences seem conceptually related, the empirical associations are not uniform across studies, and often the measures of preference are uncorrelated with intuitively similar personality traits. Nevertheless, in several studies, Neuroticism is associated with risk aversion, and facets of Conscientiousness are associated with delay acceptance. Some evidence suggests that considering facets of the Big Five

¹⁶³ They do not use a measure of Openness to Experience to separate out its influence from that of cognitive ability.

Risk acceptance (gains)							
0.47***	Risk acceptance (losses)						
0.07**	0.01	Delay acceptance					
0.10***	-0.14***	0.22***	Cognitive skill				
-0.01	-0.02	0.00	-0.04	С		_	
0.01	-0.03	-0.01	0.00	0.27***	E		
-0.04	0.01	0.07**	-0.03	0.39***	0.05*	A	
-0.05*	-0.02	0.02	-0.06*	-0.40***	-0.32***	-0.36***	N

Figure 1.8 Pairwise Correlations between Risk Acceptance, Delay Acceptance, Cognitive Ability, and Personality in a Sample of US Truckers.

Notes: *Statistically significant at the 10% level; **statistically significant at the 5% level; ***statistically significant at the 1% level. **C**, Conscientiousness; **E**, Extraversion; **A**, Agreeableness; **N**, Neuroticism. The value in each box is the pairwise correlation. Darker-shaded boxes have lower p values. Delay Acceptance and Risk Acceptance for Gains and Losses come from real-stakes experiments. Cognitive Skill is the first factor from a Raven's Progressive Matrix test, a numeracy test, and the Hit 15. The facets of the Big Five were constructed from the Multidimensional Personality Questionnaire. The sample consists of 1,065 trainee truck drivers in the United States.

Source: Adapted from Anderson, Burks, DeYoung, and Rustichini (2011).

might help establish a mapping between personality and preferences. However, the empirical links between preference parameters and personality traits depend on the data used.

6.2. Mapping Preferences into Personality

Despite some plausible, empirical and conceptual links between preferences and traits, a precise mapping between the measures is not yet available. In Section 3, we argued that measured personality is generated by underlying preference parameters and constraints.

However, the preferences measured by economists are often chosen to ensure identification on particular types of data on choices and may not be sufficiently rich. Further, studies documenting relationships between preferences and traits typically only study correlations without being motivated by an underlying model. Hence, causal claims are, at this stage, largely premature. There are two main reasons for the disconnection between measures of personality and measures of preferences.

First, economists typically study marginal rates of substitution, measured over relevant ranges via observed choices. Personality psychologists typically do not study these trade-offs and often do not study choice behavior. Most approaches to measuring preferences in economics, whether observational or experimental, use some variation of revealed preference given observed choices. In contrast, psychologists typically use surveys to elicit preferences, information, or "typical" actions. Some questions elicit how respondents would feel about a given outcome, without presenting an alternative outcome. Although such questions may elicit some (unspecified) feature of preferences, it is not clear what is being measured. The difference in approach makes it intrinsically difficult to compare economic and psychological measures.

Second, traditional preference parameters may not span the entire space of human decisions measured by psychologists. Time, risk, social, and leisure preferences do not capture the only trade-offs in life. Although time preference, risk aversion, leisure preference, and social preference have analogues in psychology, many personality psychologists do not perceive self-control and delay of gratification, risk-taking behavior and sensation seeking, and motivation and ambition as the most important aspects of human decision making.

Economists typically make strong simplifying assumptions to make their models tractable and to secure identification. The estimated parameters are used to build models, evaluate policy, and create counterfactuals. The most widely used specifications of trade-offs are through parameterizations assuming separability and assume that marginal rates of substitution are summarized by one or two parameters. Personality psychologists do not have the same incentives as economists to describe behavior by simple specifications as they are often content to stop with rich descriptions and do not use their estimated relationships in policy analyses.

6.3. Do Measured Preference Parameters Predict Real-World Behavior?

One test of the stability of measured preferences is whether they predict behavior in other contexts. Several recent studies have investigated whether risk preferences predict behavior. For example, Dohmen et al. (2011) use an experimentally validated measure of risk preference in the German Socio-Economic Panel (GSOEP) and find that it predicts self-reported risky behaviors, such as holding stocks, being self-employed, participating in sports, and smoking, but it does not predict as well a survey question about "willingness to take risks in general." However, the observed

relationship might arise because both the self-reported behaviors and questions about willingness to take risk are noisy contemporaneous survey measures. Barsky, Juster, Kimball, and Shapiro (1997) measure risk tolerance, time preference, and the intertemporal elasticity of substitution and find that risk tolerance predicts smoking and drinking, holding insurance and stock, and decisions to immigrate and be selfemployed. However, measures of risk tolerance only explain a small fraction of the variation in risky behaviors.

Benz and Meier (2008) compare measures of social preferences with charitable giving in a field experiment and find that experimental measures do not predict real-life behavior well. Levitt and List (2007) and List (2009) discuss the more general discrepancy between results from the laboratory and the field and argue that this does not arise because people behave inconsistently, but because experimenters are not controlling for relevant aspects of the choice situation. Their work is a rehash of the old person-situation debate. Falk and Heckman (2009) present a different interpretation of the value of experiments. We discuss the evidence on this question below.

6.4. Integrating Traits into Economic Models

Behavioral economics has incorporated some aspects of personality psychology to investigate how standard models of preferences can be improved to better reflect reality. Behavioral economics has highlighted many so-called anomalies, ways in which standard preferences do not accurately describe human behavior. We can divide these attempts into two main approaches.

First, behavioral economists have tried to improve models of behavior by developing more flexible functional forms for preferences. Below we discuss some of the now-standard examples, such as loss aversion, hyperbolic discounting, and reciprocity. These are not anomalies with respect to rationality but are examples that challenge standard models of preferences. For example, the time-inconsistent actions induced by hyperbolic discounting (defined below) are often described as "errors," but they are not. The agent is simply optimizing nonstandard preferences.

Second, behavioral economists have introduced the concept of bounded rationality. They discuss behaviors for which there is no reasonable preference specification that can rationalize a behavior. They are called anomalies or biases relative to conventional economic choice frameworks. Examples include failure to predict the winner's curse, mental accounting, framing effects, failure to apply Bayesian updating, and default effects. We think of these as mental constraints, or traits, along the lines of the models discussed in Section 3. However, these examples are consistent with evidence reviewed below on the interaction between cognitive ability and preference parameters.

Note that while some of the nonstandard features of preferences may seem compelling, the high level of generality used in many quarters of behavioral economics tends to make it difficult to identify the parameters in the data commonly used by economists (see the discussion in Hansen, 2005).

6.4.1 Traits as Constraints

Preference measurements that do not account for all of the constraints that agents face might be biased. In the model of personality in Section 3, we describe how agents act based on both preference parameters and productive traits that embody constraints. The marginal rate of substitution is typically identified through price variation. However, the true price ratio might also depend partly on the unobserved traits of the individual. Failure to account for the traits that arise from constraints could lead to bias.

The empirical literature has focused on the interaction between cognition and preference parameters. Virtually all methods of estimating time preference assume that respondents are equally numerate, but Peters et al. (2006) show that this assumption is often untrue. Furthermore, more numerate individuals are less susceptible to framing effects and draw stronger and more precise meaning from numbers and comparisons using numbers. The confound with numeracy may explain why more intelligent (or educated) individuals often display lower discount rates when decisions require complex calculations to compare subtly different delays or reward amounts (e.g., de Wit, Flory, Acheson, McCloskey, and Manuck, 2007; Dohmen, Falk, Huffman, and Sunde, 2010), but it is less helpful in explaining why smarter individuals also have lower discount rates when choosing between relatively simple cash sums (Funder and Block, 1989) and between noncash rewards (such as smaller versus larger candy bars, as in Mischel and Metzner, 1962).¹⁶⁴ A meta-analysis by Shamosh and Gray (2007) of 24 studies in which both IQ and discount rates were measured shows that the two traits are inversely related (r = -0.23). The complexity entailed by comparing the present and future values of rewards suggests that the inverse relationship between discount rates and intelligence is not just an artifact of measurement. One explanation for this could be that cognitive ability is related to the ability to direct attention. Daly, Delaney, and Harmon (2009) find that lower discount rates are associated with cognitive mindfulness, which includes the ability to control attention. Further, an individual with poor working memory and low intelligence may not be capable of accurately calculating or even perceiving the value of a deferred reward. At the least, making such calculations is more effortful (that is, costly) for individuals of low cognitive ability. If the cost of making calculations exceeds the expected benefit of such deliberation, the individual may choose by default the immediate, certain reward. However, it is important to be aware of reverse causality, since more-patient individuals may also invest more in cognitive ability.

¹⁶⁴ Heckman (1976) shows that more-educated people have lower discount rates. More-able people are more likely to attend more years of school.

Measures of cognitive constraints also relate to measured risk preference. There is an inverse relationship between cognitive ability and risk aversion, where higher IQ people have higher risk tolerance (Dohmen, Falk, Huffman, and Sunde, 2010).¹⁶⁵ Reference dependence can lead subjects to be susceptible to framing because they will perceive two identical lotteries differentially when one is framed as a loss and the other is framed as a gain. Some evidence suggests that individuals with higher cognitive ability and education are less risk averse. Burks, Carpenter, Goette, and Rustichini (2009) find that higher-IQ individuals are more consistent in their choices between a lottery and fixed sums. They hypothesize that agents with higher cognitive ability can better translate their preferences into choices between lotteries.

Borghans, Golsteyn, Heckman, and Humphries (2011) find that while risk aversion is related to personality traits, ambiguity is not. In particular, IQ does not explain how subjects choose between a risky and an ambiguous urn.

6.4.2 Traits as Preferences

Some aspects of traits may be more naturally believed as aspects of preferences than as constraints. For example, Openness to Experience might relate to a preference for learning, and Extraversion might reflect a preference for social interactions. The distinction between preferences and constraints often seems tautological. One way of incorporating personality into preferences is by modifying functional forms, which fall into two broad and sometimes overlapping categories. First, some of the domains that are traditionally treated as fundamentally different, such as risk and time preference, social and risk preference, and leisure and time preference, may be closely related and generated from a common set of psychological traits. Second, nonseparabilities could confound measures of trade-offs. The literature on addiction presents an interesting class of nonseparable models,¹⁶⁶ as does the literature on exotic preferences in economics.¹⁶⁷

6.4.2.1 Multidimensionality

Marginal rates of substitution are often assumed to be generated by only one or two parameters, for example, the discount factor and the intertemporal elasticity of substitution. This facilitates identification given sparse data, and if it is a sensible specification of preferences, it gives a convenient description of behavior. However, one or two parameters may not describe behavior well. Conversely, some of the concepts analyzed separately in the literature may be governed by the same parameters.

¹⁶⁵ The two cognitive ability tests used by Dohmen, Falk, Huffman, and Sunde (2010) were coding speed and vocabulary tests.

¹⁶⁶ See Becker and Murphy (1988).

¹⁶⁷ See Epstein and Zin (1989), Hansen (2005), and Hansen and Sargent (2008).

In discussing the concept of time discounting, Frederick, Loewenstein, and O'Donoghue (2002) argue that time preference has three dimensions: *impulsivity*, the tendency to act spontaneously and without planning; *compulsivity*, the tendency to stick with plans; and *inhibition*, the ability to override automatic responses to urges or emotions. There are multiple interpretations of this assertion.

First, the trade-off between different time periods might be described by several parameters. Second, impulsivity, inhibition, and compulsivity might reflect constraints, that is, something that affects shadow prices of consumption in different time periods. Third, the relevant trade-off might not be between different time periods but, for example, in the case of impulsivity might be between various levels of sensation seeking, a behavior which is also related to risk seeking.

Like time preference, risk preference may depend on multiple parameters. As noted by Rabin (2000), the simple expected utility framework does not explain risk aversion over small stakes since it would imply an implausibly high curvature of the utility function. See Starmer (2000) for a review of the literature on departures from expected utility. When psychologists started measuring risk-taking behavior, they were puzzled by the large variance across domains (see the discussion of situational specificity in Section 2). More recently, Weber (2001) shows that risk preference varies by domain, and a scale that assesses risk taking in five different domains shows low correlations across these domains (Weber, Blais, and Betz, 2002). One can be quite risk averse when it comes to financial decisions but risk loving when it comes to health decisions (Hanoch, Johnson, and Wilke, 2006). Weber's risk-return model of risk taking (Weber and Milliman, 1997, and Weber and Hsee, 1998) finds that low correlations among risk-taking preference across domains can be explained by domain-specific perceptions of riskiness and return. Dohmen et al. (2011) find that a survey question on willingness to take risks within a domain predicts self-reported behaviors within each domain. Einav, Finkelstein, Pascu, and Cullen (2010) also find that there are domain-specific components of risk-taking behavior. Domain specificity might arise because sensation seeking, enjoyment of risk per se, is an important aspect of risk preferences.¹⁶⁸

Ambiguity aversion, the disutility from model uncertainty, might help explain some apparent inconsistencies. Ambiguity aversion is measured as the trade-off between lower expected return and higher model uncertainty. Ambiguity aversion explains Ellsberg's paradox: people tend to prefer an urn with a 0.5 probability of winning to an urn with an unknown probability in which they are allowed to choose which side to bet on. One version of preferences over ambiguity is due to Gilboa and Schmeidler (1989).

¹⁶⁸ Zuckerman (2007) suggests that sensation seeking is related more closely to Big Five Conscientiousness (inversely), but there is obvious conceptual overlap with excitement seeking, a facet of Big Five Extraversion on the NEO-PI-R questionnaire, as well as with Big Five Openness to Experience.

They specify max-min preferences, where the agent maximizes an expected utility function that has been minimized with respect to the prior probabilities, that is,

$$U(X_1, X_2, ..., X_K) = \min_{(\pi_1, \pi_2, ..., \pi_K)} \pi_1 u(X_1) + \pi_2 u(X_2) + \dots + \pi_K u(X_K)$$

subject to
$$\sum_{i=1}^K \pi_i = 1.$$

Borghans, Golsteyn, Heckman, and Meijers (2009) measure ambiguity aversion and risk aversion in a group of Dutch high-school students and show that this aspect of choice is distinct from risk aversion.

There is no consensus on how social preferences govern choices. Social preferences refer to any explanation for nonselfish behavior, usually as measured in a dictator game in which people have to divide a sum between themselves and another person. Typically, more than 60% of proposers give positive amounts, averaging 20% of the sum. A variation of this game is the classic ultimatum game in which a giver divides a sum between him or herself and another subject, the receiver, and the other subject can accept or decline the sum. If he declines, both will lose their money. Studies typically find that receivers decline if offered less than 20%. These results cannot be explained by pure selfishness. In the dictator game, the giver is willing to forgo his own consumption in order to increase another person's consumption, and in the ultimatum game, the receiver is willing to forgo his own consumption in order to decrease the giver's consumption if he pays him too little. Many studies seek to find deeper traits that govern these behaviors, such as preferences over the utility of one's own consumption compared to that of others, efficiency, and fairness. The notion of fairness covers various concepts, including equality and rewards in proportion to talent, effort, kindness, or intentions. For reviews of this literature, see List (2009) and Camerer and Fehr (2004).

In the linear, separable model, where each good X_i is the consumption of person *i*, we can think of the weights as caring or altruism, that captures the preference that people often care about other people's utility or consumption. See Meier (2007) for a review. Fehr and Schmidt (1999) analyze inequality aversion. People dislike inequality rather than valuing the consumption or utility of agents per se.

Caring and altruism have been shown to decrease with social distance. People typically care more about themselves than about others, and they are less altruistic the less well they know other people.

The social preference of reciprocity has been studied. Fehr and Gächter (2000) and Falk and Fischbacher (2006) present evidence on reciprocity and conditional cooperation, in which agents act in a prosocial or antisocial manner depending on the behavior of others with whom they interact. People exhibit positive reciprocity if they tend to reward others for kindness but negative reciprocity if they tend to punish others for unkindness. More precisely, they are willing to incur a cost in order to reward or

punish others. Falk and Fischbacher (2006) develop a theory of reciprocity in which utility depends on the kindness of others, which is a function not only of the outcome from another person's action but also of the perceived intentions. Reciprocity then reflects how much value a person puts on rewarding kindness. Economists could model these features by letting the person-specific weights on the subutilities depend on social distance and past actions of others. Reciprocity is often measured using a gift-giving game in which the proposer offers a wage to a responder, who then subsequently chooses a level of effort. However, List (2009) argues that the importance of fairness preferences may have been overstated in the literature and that many of the observed results are due to concerns over either reputation or scrutiny by experimenters. Several studies have shown that observed reciprocity fades over a longer time frame than the short duration of lab experiments (Gneezy and List, 2006; Hennig-Schmidt, Rockenbach, and Sadrieh, 2010; Kube, Maréchal, and Puppe, 2006). Andreoni's (1995) warm glow model of altruism suggests that people do not care about others, but value the act of giving.

Inequality aversion is distinct from caring in the sense that A's utility may be decreasing in B's consumption if it is higher than A's (see Fehr and Schmidt, 2006, for a review). Fehr and Schmidt (1999) suggest the following asymmetric specification for the utility of agent n:

$$U_n(X_1, ..., X_j, ..., X_n, ..., X_K) = X_n - \alpha_n \frac{1}{K - 1} \sum_{j \neq n} \max\{ |X_j - X_n|, 0\} - \beta_n \frac{1}{K - 1} \sum_{j \neq n} \max\{ |X_n - X_j|, 0\},\$$

where the weights satisfy $\beta_n \leq \alpha_n$ and $0 \leq \beta_n \leq 1$. People place a higher weight on own consumption compared to others', but they place value on inequality in a situation.

People appear to be more tolerant of inequality if they believe that it represents a difference earned through effort rather than from differences in exogenously given talent (see Tausch, Potters, and Riedl, 2010, for a review). This finding may be related to the notion of reciprocity. The distinction may be whether the preference is for people who have earned their reward for doing something "for me" or something admirable in general.

Some aspects of preferences appear to be multidimensional. However, many preference parameters are correlated. For example, the social preference of "trust" relates to risk aversion and reciprocity. Altmann, Dohmen, and Wibral (2008) measure trust as the willingness to give money to an investor in a trust game in which he will only be repaid if the investor decides to return the favor. In this game, one can think of trust as the belief about how own actions affect those of others. They find that trust and positive reciprocity are positively related. Using the German Socio-Economic Panel (GSOEP), Dohmen, Falk, Huffman, and Sunde (2008) find that most people exert positive reciprocity; positive reciprocity and negative reciprocity are only weakly correlated; and people who are negatively reciprocal are less willing to trust others. In situations involving trust, it seems natural that trust is closely related to risk and ambiguity aversion and that a person who is more prone to accept uncertainty is also more likely to trust others. Altmann, Dohmen, and Wibral (2008) also find that people who are less risk averse are also more willing to trust. However, they do not measure which beliefs the agents hold.

Care has to be taken in distinguishing trust from risk aversion. Kosfeld, Heinrichs, Zak, and Fehr (2005) find that people who receive oxytocin exhibit more trusting behavior in a real-stakes trust game. However, oxytocin does not make subjects more generous, suggest-ing that trust is not simply altruism. In addition, oxytocin does not affect people's decision over risky outcomes when playing against a computer rather than a human. Together, these findings suggest that there is a unique characteristic that affects willingness to trust, distinct from altruism and risk aversion. Fehr (2009) posits that this missing element might be "betrayal aversion." Using survey data from Germany, Fehr (2009) finds that risk preferences, betrayal aversion, and altruism (as expressed through volunteering) predict people's self-reported willingness to trust others.

6.4.2.2 Preference Specifications and Their Consequences

The most restrictive version of the additively separable model suggests that the marginal rate of substitution between any two goods does not depend on the consumption of other goods outside of the pair. Browning, Hansen, and Heckman (1999) present ample evidence against this assumption. Apparent inconsistencies can arise if nonseparability is ignored. Further, estimates will suffer from omitted variable bias.¹⁶⁹

The additively separable intertemporal model imposes the requirement that the intertemporal elasticity of substitution is the same as the relative risk aversion parameter. However, Barsky, Juster, Kimball, and Shapiro (1997) find no evidence that the intertemporal elasticity of substitution is correlated with risk tolerance. However, the sample on which they measure these parameters is small. Green and Myerson (2004) argue that risk and time belong to different underlying psychological processes. As evidence, they point out that the two constructs react differently to the same stimulus: for example, an increase in the size of a reward generally decreases the time discount factor but increases the discount rate when rewards are probabilistic.¹⁷⁰ This is evidence against the standard intertemporally separable model of risk aversion.

¹⁶⁹ See Section A6.4 in the Web Appendix for a discussion of additive separability and its implications.

¹⁷⁰ Further support for this disassociation comes from a cross-cultural study by Du, Green, and Myerson (2002), in which Chinese graduate students discounted delayed rewards much more steeply than Japanese students, but Japanese students discounted probabilistic rewards more steeply than did the Chinese. Barsky, Juster, Kimball, and Shapiro (1997) report that their estimates of time preference and risk tolerance are independent.

One type of nonseparability is between goods and the state or time period. The additively separable model allows for this type of dependence, represented by the subscript v on the utility function. Although exponential discounting is still the most common representation of time preferences, experiments show that people tend to put higher weight on the present than on future periods than would be predicted by exponential discounting. This is the motive for hyperbolic discounting. The most often used specification is (β, δ) -preferences, where β is the usual discount factor while δ is an additional discounting of all future periods,

$$U_{v}(X_{v}, X_{v+1}, ...) = u(X_{v}) + \delta\beta u(X_{v+1}) + \delta\beta^{2} u(X_{v+2}) + \cdots$$

The consequence of these preferences is that the trade-off between period v and period v + 1 is not evaluated the same way from the perspective of period v - 1 and period v, leading to time inconsistency.¹⁷¹ Other possibilities are that the discount rates change with age. Hyperbolic and age-dependent discounting make use of the subscript v on the utility function. We may think of an agent in multiple periods as several agents who play a game with each other. The agent today might account for what future agents might do. Further, discount rates appear to vary inversely with the size of reward and vary with the type of reward offered.¹⁷²

As previously noted, the expected utility form for risk preferences does not explain risk preferences over small stakes (Rabin, 2000). If subutility functions represent utility of lifetime wealth in different states, people should be approximately risk neutral for small stakes. However, people often avoid more-than-fair small bets. If this is explained by expected utility, then the curvature of the utility of wealth function would have to be implausibly high. Kahneman and Tversky (1979) suggest that people are loss averse, that is, that losses weigh higher than gains in the utility function. This would imply that people have state-dependent preferences, which can be expressed as

$$U_n(X_{1,}X_2,...X_n,...X_K) = \pi_1 u(X_1 - X_n) + \pi_2 u(X_2 - X_n) + \dots + \pi_K u(X_K - X_n),$$

where *n* is the current state and u'(y) is higher for negative *y* than for positive *y*. Note that this specification is very similar to that of inequality aversion discussed above. Both models share the feature that people do not have stable preferences over levels, but over differences.

The concepts of loss aversion, reference-dependence, and endowment effects (Thaler, 1980, and Kahneman and Tversky, 1979) are variations on this theme. If an agent has had an object in his possession for even a short amount of time, it affects

¹⁷¹ This specification originates in the work of Phelps and Pollak (1968).

¹⁷² See Green, Fry, and Myerson (1994); Chapman, Nelson, and Hier (1999); Kirby (1997); Chapman and Coups (1999); Estle, Green, Myerson, and Holt (2007); Bickel, Odum, and Madden (1999); and Bonato and Boland (1983).

how he trades it off against other goods. In certain experiments, List (2003) has shown that this effect disappears when agents have market experience.

Reference-dependence has also been demonstrated in dictator games. In the standard dictator game, the first player, the "dictator," is given a positive endowment while the second player receives nothing. The game has a single move. The first player can choose to transfer some or none of his endowment to a second player. At this point the game ends. The second player can take no action. Numerous studies have shown that most first players transfer a positive amount, even though they have no monetary incentive to do so. List (2007) and Bardsley (2008) modified the standard dictator game by giving the second player an endowment and allowing the first player to transfer a positive amount to the second player or to take part of the second player's endowment. With this modification, most first players did not transfer positive amounts to the second player.

Experimental measures of social preferences vary greatly across studies. Levitt and List (2007) and List (2009) argue that the degree of scrutiny in the laboratory as opposed to in the real world may make subjects behave more prosocially (Bandiera, Barankay, and Rasul, 2005).¹⁷³ Further, several studies have found that people tend to be more selfish when the stakes of the game increase (Carpenter, Verhoogen, and Burks, 2005; Slonim and Roth, 1998; Parco, Rapoport, and Stein, 2002).

There is evidence of substantial heterogeneity in preferences between and within socioeconomic groups. Marginal rates of substitution depend on other factors such as education, age, cultural values, etc.¹⁷⁴ This evidence supports the claim that people are different at a basic level, since preferences govern the choices that shape life. However, preferences may be experience dependent. Although most studies view life outcomes as the result of choices governed by exogenous preferences, and hence infer preferences from outcomes, initial conditions might determine both preferences and constraints on the available choices.

The motivation for preference specifications in economics is typically introspection, axioms about rationality, and convenience, rather than on predictive power. When measuring preferences, functional forms are chosen in an attempt to minimize approximation error subject to identification. However, economists typically consider preferences over a limited range of fundamental attributes. Time, risk, and social preferences may not be the right dimensions of choice over which parameters are stable. Each of these domains seems to be guided by multiple parameters, and some of these parameters seem to matter for each of the domains. Personality psychology may help in guiding economists as where to look for more fundamental parameters. However, the potential is largely unexploited.

¹⁷³ See, however, Falk and Heckman (2009).

¹⁷⁴ See the evidence in Browning, Hansen, and Heckman (1999).

6.5. Summary of Section 6 and Some Concluding Thoughts

Table 1.7 summarizes the main papers relating economic preference parameters to psychological measurements. The lack of any close correspondence between the traits of personality psychology and the parameters of economics suggests great opportunities for both fields. Each can learn from the other.

However, the different emphases of interest in the two fields may present a challenge to integration. The greater emphasis on prediction, intervention, and causality in economics, compared to the greater emphasis on description in personality psychology, may lead to a new choice-motivated set of psychological traits that supplant traditional Big Five measures. Developments in this direction are discussed in Ferguson, Heckman, and Corr (2011). Personality psychologists may well adopt a more choiceoriented set of parameters to supplant the description-oriented Big Five parameters. Moreover, both the conventional psychological traits and economic preference parameters may be stable manifestations of deeper parameters, connected to human motivation and goal-seeking that remain to be discovered.

7. THE PREDICTIVE POWER OF PERSONALITY TRAITS

This section discusses the empirical evidence on the power of personality in predicting life outcomes. A growing body of evidence suggests that personality measures—especially those related to Conscientiousness and, to a lesser extent, Neuroticism—predict a wide range of outcomes. The predictive power of any particular personality measure tends to be less than the predictive power of IQ, but in some cases rivals it.

For three reasons, summarizing the large literature on the predictive power of personality on outcomes is a daunting task. First, the measures of personality and cognition differ among studies. As noted in Section 5, not all psychologists use the Big Five. We attempt to cast all measures into Big Five categories. When this is not possible, we discuss the measures used and how they relate to the Big Five measures.

Second, different studies use different measures of predictive power. Many studies report only simple correlations or simple standardized regression coefficients.¹⁷⁵ Such estimated relationships do not control for other factors that may influence outcomes. This is particularly problematic for estimated relationships between personality measures and other outcomes that do not control for cognition, situation, or the effect of other personality measures. Where possible, we report both simple and partial correlations.

¹⁷⁵ Standardized regressions produce regression coefficients of outcomes divided by their standard deviations regressed on explanatory variables divided by their standard deviations. This produces correlation coefficients in bivariate regressions and partial correlation coefficients in multivariate regressions. See, e.g., Goldberger (1968).

Author(s)	Main Variable(s)	Data and Methods	Causal Evidence	Main Result(s)
Altmann, Dohmen, and Wibral (2008)	Outcome(s): <i>trust</i> —amount the first player sends in a real- stakes experimental trust game Explanatory Variable(s): <i>reciprocity</i> —amount returned by the second player in a real- stakes experimental trust game; <i>risk aversion</i> —certainty equivalent as measured by real-stakes choices over lotteries	Data: collected by authors; 240 students from the University of Bonn Methods: OLS	Controls: gender Timing of Measurements: The measures are contemporaneous. Theory: People might generally value adhering to social norms associated with trust and reciprocity.	Reciprocity and trust are positively related ($p < 0.01$). Risk aversion and trust are positively related ($p < 0.05$).
Borghans, Golsteyn, Heckman, and Meijers (2009)	Outcome(s): <i>risk aversion</i> — choices over real-stakes lotteries; <i>ambiguity aversion</i> — comparison of the willingness to bet on lotteries when the probability distribution is unknown Explanatory Variable(s): <i>gender</i> ; <i>personality</i> — self-reported measures of the Big Five, ambition, flexible thinking, and self- control	Data: collected by authors; 347 students aged 15 to 16 from a Dutch high school Methods: OLS, F-test	Controls: n/a Timing of Measurements: The measures are contemporaneous. Theory: Risk aversion and ambiguity aversion represent different preferences and might reflect different personality traits.	Men are less risk averse than women ($p < 0.001$) but more ambiguity averse ($p < 0.05$). Risk-aversion is mediated by personality ($p < 0.05$), while ambiguity aversion is not. Risk aversion is positively associated with Agreeableness and Neuroticism and is negatively associated with ambition ($p < 0.05$).

Table 1.7 Links among Personality Traits and Preferences

Continued

Author(s)	Main Variable(s)	Data and Methods	Causal Evidence	Main Result(s)
Borghans, Meijers, and ter Weel (2008)	Outcome(s): cognitive ability— number of correct answers on an IQ test; effort—time spent on each question Explanatory Variable(s): risk aversion—survey response to lotteries; time preference— survey response to trade-offs across time; leisure preference— survey response; experiment incentives—payment for correct answers to the IQ test; personality—self-reported Big Five, performance motivation, positive and negative fear of failure, locus of control, social desirability, curiosity, resilience, enjoyment of success, attitude toward work	Data: collected by authors; 128 university students from a Dutch University Methods: probit	Controls: type of cognitive test, the amount of incentive pay, and time constraints Timing of Measurements: They measured IQ both before and after providing incentives. Theory: People with different personalities and preferences might be willing to expend different amounts of mental effort during a test.	Performance motivation, fear of failure, internal locus of control, curiosity, low discount rates, and risk aversion are positively associated with more correct answers ($p < 0.05$). Negative fear of failure, Extroversion, Openness to Experience, and Agreeableness are negatively associated with answering the question correctly ($p < 0.05$). Incentives did not affect the number of questions answered correctly. Intrinsic motivation, curiosity, internal locus of control, Emotional Stability, Conscientiousness, and discount rates are negatively associated with responsiveness to incentives ($p < 0.05$). Risk aversion is negatively associated with responsiveness to incentives ($p < 0.10$). Leisure preference and Openness to Experience are positively associated with responsiveness ($p < 0.05$).

Table 1.7 Links among Personality Traits and Preferences—continued

Continued

Author(s)	Main Variable(s)	Data and Methods	Causal Evidence	Main Result(s)
Burks, Carpenter, Goette, and Rustichini (2009)	Outcome(s): risk aversion— choices over real-stakes lotteries; time discounting— choices over real-stakes payments at different times; inconsistent risk and time preference—making at least one inconsistent choice in the experiments eliciting preferences; job performance— whether a worker leaves before the end of the first year Explanatory Variable(s): cognitive ability—IQ as measured by an adaptation of Raven's Standard Progressive Matrices	Data: collected by authors, administrative data; 892 trainee truckers from a US trucking company (2005– 2006) Methods: OLS, interval regressions, linear probability model, Cox proportional hazard	Controls: race, age, age squared, education, household income, absorption, achievement, aggression, alienation, control harm avoidance, social closeness, social potency, stress reaction, traditionalism, and well-being Timing of Measurements: The measures are contemporaneous, except for job turnover that was evaluated after the experiment. Theory: People with higher IQ can better forecast the future.	An increase in IQ from the bottom quartile to the top quartile is associated with an increase in risk-taking consistency of 25 percentage points ($p < 0.001$), an increase of intertemporal consistency of 15 percentage points ($p < 0.001$), a decrease in discount rate ($p < 0.001$), and a decrease in risk aversion ($p < 0.001$). People in the lowest quartile of IQ are about twice as likely to leave the job within the first year ($p < 0.001$).
Daly, Delaney, and Harmon (2009)	Outcome(s): <i>time preference</i> — discount rate measured by a real-stakes choices over delayed payments Explanatory Variable(s): <i>health</i> —blood pressure, body fat, blood glucose, weight, height, heart rate; <i>personality</i> — questionnaire measures of the Big Five, self-control, consideration of future consequences, elaboration of potential outcomes, emotional regulation, cognitive and affective mindfulness, suppression of unwanted thoughts, experiential avoidance	Data: collected by authors; 204 students from Trinity College Dublin. Methods: factor analysis, OLS	Controls: age and sex Timing of Measurements: The measures are contemporaneous. Theory: Personality traits and health indicators might be associated with willingness to delay gratification.	Age and sex do not predict the estimated discount rate. A factor that loads heavily on self-control, consideration of future consequences, elaboration of consequences, affective mindfulness, and Conscientiousness is negatively associated with the discount rate ($p < 0.01$). A factor that loads on blood pressure is positively associated with the discount rate ($p < 0.10$).

Table 1.7	Links among	Personality	Traits and	Preferences-	-continued

Author(s)	Main Variable(s)	Data and Methods	Causal Evidence	Main Result(s)
Dohmen et al. (2011)	Outcome(s): <i>experimental risk</i> <i>measure</i> —real-stakes choices over lotteries and cash payments Explanatory Variable(s): <i>survey</i> <i>risk measure</i> —survey responses on an 11-point scale, relating to general risk preference and risk preference relating to car driving, financial matters, leisure and sports, career, and health	Data: collected by the authors; 450 adults from Germany Methods: OLS	Controls: gender, age, height, and other personal characteristics Timing of Measurements: The measures are contemporaneous. Theory: Survey and experimentally elicited risk measure the same concept.	Survey measures of general risk attitude predict incentive compatible, experimentally elicited measures of risk attitude ($p < 0.01$).
Ding, Hartog, and Sun (2010)	Outcome(s): <i>experimental risk</i> <i>measure</i> —real-stakes choices over lotteries and cash payments Explanatory Variable(s): <i>survey</i> <i>risk measure</i> —responses on an 11-point scale, relating to general risk preference and risk preference relating to car driving, financial matters, leisure and sports, career and health, survey responses to hypothetical lotteries	Data: collected by the authors; 121 students of PKU in Beijing, who participated in an experiment (2008) Methods: OLS, correlations	Controls: major, gender, family income, and class rank Timing of Measurements: The measures are contemporaneous. Theory: There could be an underlying risk parameter that applies in all situations.	The survey measures of risk explain at most 10% of the variance in the experimental measures of risk (general risk attitude and financial risk are the best). Self-assessed risk depends much on the domain or context; the highest correlation between context- based survey questions is $r =$ 0.55. Women are more risk averse than men; risk aversion decreases with parental income; and risk attitudes depend on domain (context). People view winning and losing money differently.

Table 1.7 Links among Personality Traits and Preferences—continued

We also consider a measure of predictive validity that extends traditional conceptions of variance explained. Recent work by economists relaxes the normality and linearity assumptions that underlie the use of simple partial correlations and standardized regression coefficients that are used in psychology. This method measures the predictive power of variables by the slopes of percentile changes on outcomes and not by variance explained. If outcomes are characterized by substantial measurement error, a low R^2 for a predictor may still be consistent with a substantial effect of the predictor on means and quantiles.¹⁷⁶

For example, Heckman, Stixrud, and Urzua (2006) report the effects of percentile changes in cognitive and personality measures on a variety of outcomes over the full range of estimated relationships, relaxing traditional normality or linearity assumptions and not relying directly on measures of variance explained. This approach to measuring predictive power is increasingly being applied by economists.¹⁷⁷

Third, many studies do not address the question of causality, that is, does the measured trait cause (rather than just predict) the outcome? Empirical associations are not a reliable basis for policy analysis. Problems with reverse causality are rife in personality psychology. Contemporaneous measures of personality and outcomes are especially problematic. For example, does greater Neuroticism lower earnings, is it the other way around, or do they mutually influence each other?

Few economists or psychologists working on the relationship between personality and outcomes address the issue of causality, and when they do so, it is usually by employing early measures of cognition and personality to predict later outcomes. As discussed in Section 4, using early measures of personality traits to predict later outcomes raises problems of its own. We delineate how each study surveyed addresses causality.

7.1. An Overview of the Main Findings

Before presenting a detailed survey of the effects of personality and cognition on a variety of outcomes, it is useful to have an overview of the main findings. One principal finding of our survey, consistent with the claims of the early psychologists cited in Section 2, is that Conscientiousness is the most predictive Big Five trait across many outcomes. However, other personality measures predict some outcomes.

¹⁷⁶ The slope versus variance explained distinction is an old one. However, the use of slopes as measures of "importance" is problematic in general because of the arbitrariness in the scales of the dependent and independent variables (see Goldberger, 1968). This arbitrariness is resolved in the measure used in the recent literature by mapping quantiles into quantiles. This literature is nonparametric. The measure is clear in its choice of units but the economic significance is still questionable. A better measure would relate costs of a change in the independent variable to the benefits.

¹⁷⁷ See, e.g., Piatek and Pinger (2010).

Measures of personality predict a range of educational outcomes. Of the Big Five, Conscientiousness best predicts overall attainment and achievement. Other traits, such as Openness to Experience, predict finer measures of educational attainment, such as attendance and course difficulty selected. Traits related to Neuroticism also affect educational attainment, but the relationship is not always monotonic. Conscientiousness predicts college grades to the same degree as SAT scores. Personality measures predict performance on achievement tests and, to a lesser degree, performance on intelligence tests.

Personality measures also predict a variety of labor market outcomes. Of the Big Five traits, Conscientiousness best predicts overall job performance but is less predictive than measures of intelligence. However, Conscientiousness predicts performance and wages across a broad range of occupational categories, whereas the predictive power of measures of intelligence decreases with job complexity. In addition, traits related to Neuroticism (e.g., locus of control and self-esteem) predict a variety of labor market outcomes, including job search effort. Many traits predict sorting into occupations, consistent with the economic models of comparative advantage discussed in Section 3. Personality traits are valued differentially across occupations.

All Big Five traits predict some health outcomes. However, Conscientiousness is the most predictive and better predicts longevity than intelligence or background. Personality measures predict health both through the channel of education and by improving health-related behavior, such as smoking.

The evidence on the effect of personality measures on crime suggests that traits related to Conscientiousness and Agreeableness are important predictors of criminality. These findings are consistent with the possibility that personality is related to social preferences as discussed in Section 6.

The survey presented in the text, even though extensive, is not fully comprehensive. We place additional material in the Web Appendix.

7.2. Educational Attainment and Achievement

We now turn to evidence for the predictive power of personality traits for educational outcomes, separately considering educational attainment, grades, and test scores.

7.2.1 Educational Attainment

Despite recent increases in college attendance, American high-school dropout rates remain high. About one in four American students drops out of formal schooling before receiving a high school diploma, and in recent decades, the dropout rate has increased slightly (Heckman and LaFontaine, 2010). A growing body of research finds that personality is associated with educational attainment, suggesting that further study of personality and its determinants might shed light on the recent stagnation in educational attainment. We begin by reviewing evidence about the relationship of personality measures with years of schooling and then consider specific aspects of educational achievement. Traits such as perseverance and preferences related to an interest in learning might lead people to attain more total years of schooling. Indeed some evidence suggests that this might be the case. Table 1.8 presents associations between years of schooling and the Big Five from three nationally representative samples. The studies yield different results, possibly because they control for different covariates or because they come from different countries. The first study controls for age, sex, and gender and finds that of the Big Five, Openness to Experience and Conscientiousness are most related to years of schooling attained (Goldberg, Sweeney, Merenda, and Hughes, 1998). The second study—which also controls for parental education and father's occupational status reports a strong relationship with Openness to Experience but a much weaker relationship with Conscientiousness than the first study, suggesting that parental background might mediate some of the effects of Conscientiousness (Van Eijck and De Graaf, 2004).

The first two samples lack information on cognitive ability. However, Openness to Experience is the only Big Five factor with moderate associations with general intelligence (r = 0.33 in a meta-analysis by Ackerman and Heggestad, 1997), and intelligence is associated with years of education (r = 0.55 in Neisser et al., 1996). Thus, Openness to Experience may proxy for intelligence. However, as Fig. 1.9 illustrates, controlling for measures of crystallized intelligence and fluid intelligence does not affect the coefficients on the Big Five within the third sample.¹⁷⁸ This sample differs from the others because Openness to Experience is not strongly associated with years of education unconditional on intelligence, possibly because it is based on a smaller inventory of questions. Conscientiousness is associated with years of schooling to a similar degree as intelligence. In each study, schooling and personality are measured at the same point in time so that for older individuals, personality is measured long after schooling has been completed. This complicates the interpretation of the estimated effects of personality on schooling in older samples.

Nevertheless, the components of Openness to Experience representing an intrinsic interest in ideas and learning may affect aspects of educational achievement not measured by total years of schooling such as the student's difficulty with classes and attendance. Consistent with this supposition, a longitudinal study of talented high-school students showed that when controlling for PSAT score, students who expressed more intrinsic motivation in learning took more difficult math courses 1 year later ($\beta = 0.30$, p < 0.05), 2 years later ($\beta = 0.31$, p < 0.05), and 3 years later ($\beta = 0.26$, p < 0.10) but did not have higher grades in a standardized set of courses.¹⁷⁹ Similarly, of the Big Five, Openness to Experience is most consistently associated with fewer contemporaneously measured school absences in seventh grade (r = -0.31, p < 0.01), tenth grade

¹⁷⁸ Table A7 in Section A7 of the Web Appendix presents the full results from this regression. Table A8 in Section A7 of the Web Appendix presents analogous results for high-school graduation.

¹⁷⁹ Wong and Csikszentmihalyi (1991).

Source	Sample	Timing of Measurement and Outcome	Controls	Metric	Results	
Goldberg, Sweeney, Merenda, and Hughes (1998)	Representative sample of US working adults aged $18-75$ (N = 3,629)	All the variables were measured in the same year, but years of schooling were cumulative.	Age, gender, ethnicity	Partial correlation with years of schooling (r)	Openness Conscientiousness Extraversion Agreeableness Neuroticism	0.31*** 0.12*** -0.04** -0.08*** -0.03
Van Eijck and De Graaf (2004)	Representative sample of Dutch adults aged 25–70 (N = 1,735)	All the variables were measured in the same year, but years of schooling were cumulative.	Age, gender, father's education, mother's education, father's occupational status	Standardized regression coefficient (β)	Openness Conscientiousness Extraversion Agreeableness Neuroticism	0.14*** 0.05*** -0.07*** -0.07** -0.09***
German Socio- Economic Panel GSOEP (2004– 2008), own calculations	Representative sample of Germans aged 21–94 (<i>N</i> = 2,381)	The Big Five were measured three years prior to the measurement of schooling, but years of schooling were cumulative.	Age, age ² , gender, crystallized intelligence, fluid intelligence	Standardized regression coefficient (β)	Openness Conscientiousness Extraversion Agreeableness Neuroticism	-0.03 0.18*** -0.02 -0.03 -0.09***

Table 1.8 The Relationship Between Years of Educational Attainment and Big Five Traits

 ** Statistically significant at the 5% level; *** statistically significant at the 1% level.





Notes: The figure displays standardized regression coefficients from a multivariate regression of years of school attended on the Big Five and intelligence, controlling for age and age squared. The bars represent standard errors. The Big Five coefficients are corrected for attenuation bias. The Big Five were measured in 2005. Years of schooling were measured in 2008. Intelligence was measured in 2006. The measures of intelligence were based on components of the Wechsler Adult Intelligence Scale (WAIS). The data is a representative sample of German adults between the ages 21 and 94.

Source: German Socio-Economic Panel (GSOEP), waves 2004–2008, own calculations.

(r = -0.19, p < 0.01), and twelfth grade (r = -0.27, p < 0.01; Lounsbury, Steel, Loveland, and Gibson, 2004). Still, interest in learning is not the whole story. Using prospective data, Lleras (2008) finds that controlling for cognitive ability, three behaviors associated with Conscientiousness (completing homework, working hard, arriving promptly to class) in tenth grade predicted educational attainment 10 years later, whereas relating well to others, a behavior related to Extraversion and Agreeableness, did not.

Examining discrete educational decisions, rather than total years of education, gives a more nuanced picture. The decision to obtain a GED is a particularly telling example. Many view GED certification as equivalent to earning a high-school diploma. Indeed GED recipients have the same distribution of measured achievement test scores as high-school graduates who do not attend college. However, controlling for cognitive ability, GED recipients have lower hourly wages and annual earnings and attain fewer years of education, suggesting they may "lack the abilities to think ahead, to persist in tasks, or to adapt to their environments" (Heckman and Rubinstein, 2001, p. 146). Figure 1.10, taken from Heckman, Humphries, Urzua, and Veramendi (2011), shows that GED recipients have cognitive skills similar to students who obtain high-school diplomas but do not attend college. However, GED recipients have noncognitive skills (personality traits) similar to those of high-school dropouts.¹⁸⁰

Supporting the evidence from the GED program that personality plays an important role in explaining educational attainment in adolescence, several prospective studies have shown that facets of Conscientiousness (e.g., self-control, distractibility) and facets of Neuroticism (e.g., internal locus of control) predict successful graduation from high school (Bowman and Matthews, 1960; Gough, 1964; Hathaway, Reynolds, and Monachesi, 1969; Janosz, LeBlanc, Boulerice, and Tremblay, 1997; Kelly and Veldman, 1964; Whisenton and Lorre, 1970).¹⁸¹ Table 1.9 presents findings from three more recent studies examining the relationship between locus of control, a trait related to Emotional Stability, and high-school graduation. Although the level of statistical significance varies across studies, the studies report remarkably similar estimates. When control-ling for basic demographics, a one–standard deviation increase in locus of control is associated with a 4.5–6.8% point increase in graduating from high school. Two of the studies control for cognitive ability and find that doing so reduces the association to only 1.4–1.5%. However, the measures of cognitive ability (course grades and AFQT score) are partly determined by locus of control as discussed later in this section.

Several recent studies using methods that address measurement error and reverse causality corroborate the evidence that traits related to Neuroticism affect educational attainment. For example, Heckman, Stixrud, and Urzua (2006) account for the effect of family background on test scores. They correct for the influence of schooling on

¹⁸⁰ See the discussion of the GED program in Heckman, Humphries, and Mader (2011).

¹⁸¹ See Section 5.4 for a discussion of the links between these personality facets and the Big Five traits.



Figure 1.10 Distribution of Cognitive and Noncognitive Skills by Education Group. (a) Female Cognitive Ability (no college sample). (b) Female Noncognitive Ability (no college sample). (c) Male Cognitive Ability (no college sample). (d) Male Noncognitive Ability (no college sample).

Notes: The data is from the National Longitudinal Study of Youth 1979 (no college sample, all ethnic groups). The distributions above represent noncognitive ability factors estimated using measures of early violent crime, minor crime, marijuana use, regular smoking, drinking, early sexual intercourse, and educational attainment as laid out in Hansen, Heckman, and Mullen (2004). The sample is restricted to the cross-sectional subsample for both males and females. Distributions show only those with no post-secondary educational attainment. The noncognitive ability factors are separately normalized to be mean zero standard deviation one.

Source: Reproduced from Heckman, Humphries, Urzua, and Veramendi (2011).

personality. They address measurement error in test scores. (Their estimates of the effect of schooling on these traits and on cognitive measures are discussed in Section 8.) Figure 1.11 shows that better adolescent personality traits—as measured by locus of control and self-esteem (traits related to Neuroticism)—increases the probability of graduating from, and stopping at, high school for males at the lowest quantiles of the personality distribution. However, at higher quantiles, the probability of stopping education at high-school graduation is decreasing in measured personality because such students continue on to college. As discussed in Section 3, the effects of traits on outcomes need not be monotonic, but they can be, as Fig. 1.12 shows, where both higher cognitive and personality traits have strong

Source	Sample	Timing of Measurement and Outcome	Controls	Metric	Results	
Báron and Cobb- Clark (2010)	Australians born in 1987 or 1988 (<i>N</i> = 2,065)	Contemporaneous	Welfare receipts, family structure, sex, parental education, parental immigration status, parental involvement in education, indigenous background, and born early for their grade	The effect of a standard deviation increase in locus of control on the probability of high- school graduation (<i>b</i>)	Locus of control	4.5*
Cebi (2007)	Nationally representative sample of students in the United States $(N = 1,394)$	Locus of control was measured in grades 10 or 11	(1) Race, gender, urban, parental education, family structure; (2) race, gender, urban, parental education, family structure, home life, AFQT.	The effect of a standard deviation increase in locus of control on the probability of high-school graduation (<i>b</i>).	Locus of control (1) Locus of control (2)	4.6*** 1.5
Coleman and DeLeire (2003)	Nationally representative sample of students in the United States ($N =$ (1) 13,720 and (2) 12,896).	Locus of control was measured in grade 8.	 (1) Race, gender; (2) race, gender, eighth-grade math score, eighth-grade reading score, eighth-grade GPA, parent's education, parenting controls, family structure 	The effect of a standard deviation increase in locus of control on the probability of high-school graduation (<i>b</i>).	Locus of control (1) Locus of control (2)	6.8 1.4**

Table 1.9 The Relationship between Probability of High-School Graduation and Locus of Control

Notes: The numbers in the "Controls" column indicate the controls used in different specifications. The numbers preceding the estimate reported in the "Results" column indicate the model used as defined in the "Controls" column.

*Statistically significant at 10% level; **statistically significant at 5% level; ***statistically significant at 1% level.



Figure 1.11 Probability of Being a High-School Graduate at Age 30 and Not Going on to Further Education, Males. (a) By Decile of Cognitive and Noncognitive Factors. (b) By Decile of Cognitive Factor. (c) By Decile of Noncognitive Factor.

Notes: The data are simulated from the estimates of the model and the NLSY79 sample. Higher deciles are associated with higher values of the variable. The confidence intervals are computed using bootstrapping (200 draws). Solid lines depict probability, and dashed lines, 2.5–97.5% confidence intervals. The upper curve is the joint density. The two marginal curves (ii) and (iii) are evaluated at the mean of the trait not being varied.

Source: Heckman, Stixrud, and Urzua (2006, Figure 19).

effects on graduating from a 4-year college at all deciles. Moving from the lowest decile to the highest decile in the measured personality distribution increases the probability of graduating from college more than a similar change in the cognitive trait distribution. These examples show why considering broad measures of education might obscure important relationships between skills and educational attainment and why assuming a linear—or even monotonic—relationship between skills and educational attainment might be incorrect.¹⁸²

Cunha, Heckman, and Schennach (2010) use a dynamic factor model to investigate the development of both cognitive skills and personality traits during childhood, allowing for endogenous investment in skills and dynamic complementarities. They find that

¹⁸² See the nonmonotonicity in Figure 1.11.


Figure 1.12 Probability of Being a 4-year College Graduate or Higher at Age 30, Males. (a) By Decile of Cognitive and Noncognitive Factors. (b) By Decile of Cognitive Factor. (c) By Decile of Noncognitive Factor.

Notes: The data are simulated from the estimates of the model and the NLSY79 sample. Higher deciles are associated with higher values of the variable. The confidence intervals are computed using bootstrapping (200 draws). Solid lines depict probability, and dashed lines, 2.5–97.5% confidence intervals. The upper curve is the joint density. The two marginal curves (b) and (c) are evaluated at the mean of the trait not being varied.

Source: Heckman, Stixrud, and Urzua (2006, Figure 21).

adolescent personality—as measured by a variety of behavior inventories—accounts for 12% of the variation in educational attainment, whereas adolescent cognitive ability accounts for 16% of the variation.

A separate, but related, literature examines the importance of early attention (a trait related to Conscientiousness) and aggression (a trait related to low Agreeableness) in determining graduation from high school. Some studies find that aggression is particularly important compared to attention. Duncan and Magnuson (2010) find that when controlling for measures of intelligence and demographic variables, antisocial behavior, but not attention measured in childhood, predicts high-school completion, where antisocial behavior is negatively associated with completion. Similarly, Fergusson and Horwood (1998) find that teacher and parent ratings of conduct problems at age 8 are negatively related to predicted high-school completion at age 18. In contrast, Vitaro,

Brendgen, Larose, and Tremblay (2005) examine individuals in a population-based sample of Quebec children and find that kindergarten teacher ratings of hyperactivity-inattention (inversely) predicted completion of high school better than did aggressiveness-opposition. Both attention and aggression likely play roles, but there is no consensus on their relative importance.

In sum, traits related to Big Five Openness to Experience and Conscientiousness are important in determining how many total years of education individuals complete in their lifetimes. Two traits related to Neuroticism, locus of control and self-esteem, play a particularly important role for adolescent schooling decisions. Their effects differ across schooling attainment levels, suggesting that analysts should be wary of using years of schooling attained as the outcome variable but should use the probability of attainment at different grades. Attention and early aggression, traits related to Conscientiousness and Agreeableness, are also predictive.

7.2.2 Course Grades

Conscientiousness is the most robust Big Five predictor of course grades, in terms of raw and partial correlations. Poropat (2009) conducted a meta-analysis of Big Five personality traits and course grades in primary, secondary, and post-secondary education, presented in Fig. 1.13. Associations between grades and Conscientiousness are almost as large as



Figure 1.13 Correlations of the Big Five and Intelligence with Course Grades.

Notes: All correlations are significant at the 1% level. The correlations are corrected for scale reliability and come from a meta-analysis representing a collection of studies representing samples of between N = 31,955 to N = 70,926, depending on the trait. The meta-analysis did not clearly specify when personality was measured relative to course grades. Source: Poropat (2009). those between grades and cognitive ability. Associations with grades are substantially smaller for other Big Five factors, the largest of which is Openness to Experience.

A few prospective, longitudinal studies have estimated the effect of Conscientiousness on course grades when controlling for baseline levels of grades. These studies help isolate the effects of personality on grades by reducing the potential for omitted variable bias and misleading halo effects—the propensity for teachers to favor students based on traits unrelated to academic achievement. In general, these studies support the conclusions of studies that do not account for halo effects. For instance, in a sample of American middle-school students, self-control predicts report card grades, controlling for both general intelligence and baseline grades (Duckworth and Seligman, 2005). Similarly, Duckworth, Tsukayama, and May (2010) use longitudinal hierarchical linear models to show that changes in self-control predict subsequent changes in report card grades. In a sample of Chinese primary school children, effortful control predicted report card grades when controlling for baseline grades (Zhou, Main, and Wang, 2010).

Figure 1.14 shows that associations between course grades and personality and cognitive ability and grades are generally stronger in the primary grades, a pattern consistent with censoring.¹⁸³ A notable exception to this trend is Conscientiousness, which has the same association with course grades at all levels. If censoring on cognitive and



Figure 1.14 Correlations with Course Grades by Level of Education.

Notes: The reported values for the Big Five are partial correlations, controlled for intelligence. The meta-analysis did not address when personality was measured relative to course grades. Source: Poropat (2009).

¹⁸³ The estimated predictive validity diminishes by grade due to censoring. Censoring was not accounted for in the meta-analysis in Poropat (2009), presumably because norms for variance in representative samples are generally unavailable for personality measures (Duckworth, 2009).

personality traits attenuates observed associations with course grades among students at higher grade levels, Conscientiousness might be even *more* predictive of course grades as students progress through the education system.¹⁸⁴ Consistent with this possibility, in a prospective study of an entire cohort of Belgium's medical students, the correlation (corrected for censoring) of Conscientiousness for GPA increased from r = 0.18 in the first year to r = 0.45 in the seventh and final year (Lievens, Dilchert, and Ones, 2009).¹⁸⁵

Overall, the empirical evidence suggests that Conscientiousness may be as predictive as cognitive ability in predicting and possibly causing higher course grades. Why? Even intelligent students might not enjoy the work (Wong and Csikszentmihalyi, 1991). Indeed, there is evidence that the association between Conscientiousness and course grades is mediated by positive study habits and attitudes, effort, and prosocial behavior in the classroom.¹⁸⁶

7.2.3 Standardized Achievement Test Scores

Like course grades, standardized achievement test scores reflect a student's acquired skills and knowledge. Thus, dimensions of personality that influence the acquisition of skills and knowledge should predict both outcomes. One expects, therefore, that traits related to Conscientiousness predict achievement test scores. Ample empirical evidence shows that aspects of personality predict both metrics of performance, although studies using standardized achievement tests are less common than studies using grades. As shown in Section 5, two traits related to Neuroticism, locus of control and self-esteem, explain a considerable portion of the variance of the Armed Forces Qualification Test (AFQT), an achievement test that is often used as a measure of pure intelligence in studies in economics. Similarly, Fig. 1.15 shows that in samples from three New York City middle schools, controlling for IQ, Openness to Experience is associated with Standardized Achievement Test Scores.

Roy Martin and colleagues were among the first to demonstrate that teacher and parent ratings of early childhood persistence, (low) distractibility, and (low) activity prospectively predict both course grades and standardized achievement test scores (see Martin, 1989, for a summary). Similarly, in a representative sample of Baltimore first graders, teacher ratings of attention span—restlessness in first grade—predicted both course grades and standardized achievement test scores 4 years later (Alexander, Entwisle, and Dauber, 1993).

¹⁸⁴ Cameron and Heckman (1998).

¹⁸⁵ The estimated correlation was corrected for truncation.

¹⁸⁶ Credé and Kuncel (2008); Lubbers, Van der Werf, Kuyper, and Hendriks (2010); Noftle and Robins (2007); Valiente, Lemery-Chalfant, and Castro (2007); and Valiente, Lemery-Chalfant, Swanson, and Reiser (2008).



Figure 1.15 Associations with Standardized Achievement Test Scores.

Notes: The values represent standardized regression coefficients in models including personality, IQ, gender, and ethnicity. The bars represent standard errors around the estimate. IQ is measured using Raven's Progressive Matrices. The achievement tests are based on the Comprehensive Testing Program test in the private school sample and the English/Language Arts and Mathematics standardized achievement tests in the public school sample. Source: Duckworth (2011).

More recently, in a sample of preschool children from low-income homes, parent and teacher ratings of effortful control, a facet of Conscientiousness, predicted standardized achievement test scores in kindergarten, even after controlling for general intelligence (Blair and Razza, 2007). Similarly, in a sample of kindergarteners, teacher and parent ratings of effortful control predicted performance on standardized achievement tests 6 months later when controlling for both verbal intelligence and family socioeconomic status (Valiente, Lemery-Chalfant, and Swanson, 2010). Teacher ratings of inattention at the beginning of the school year predicted standardized achievement test scores at the end of the school year in a sample of fourth graders (Finn, Pannozzo, and Voelkl, 1995).

Task measures of effortful control, a trait related to Conscientiousness, predict performance on standardized achievement tests much later in life. For instance, the number of seconds a child waits for a more preferred treat in a preschool test of delay of gratification predicts the SAT college admission test more than a decade later, with raw correlations of r = 0.42 for the verbal section and r = 0.57 for the quantitative section (Mischel, Shoda, and Rodriguez, 1989). The Head-to-Toes and Head-Toes-Knees-Shoulders tasks require young children to inhibit automatic responses, pay attention, and keep instructions in working memory (e.g., to touch their heads when the experimenter says "touch your toes"; Ponitz et al., 2008; Ponitz, McClelland, Matthews, and Morrison, 2009). Performance on this brief task predicts later performance on standardized achievement tests (McClelland et al., 2007).

Perhaps most conclusively, Duncan et al. (2007) analyzed six large, longitudinal data sets and found that school-entry attention skills, measured variously by task and questionnaire measures, prospectively predict achievement test scores, even when controlling for school-entry academic skills. In contrast, internalizing behavior (e.g., depression, anxiousness, withdrawal) and externalizing behaviors (e.g., aggression, hyperactivity, antisocial behavior) at school entry do not reliably predict standardized achievement test scores. Attention skills are related to Conscientiousness; externalizing behaviors are related to Neuroticism.

In sum, traits related to Conscientiousness play an important role in predicting achievement tests above and beyond cognitive ability. Nevertheless, as discussed in Section 6, time discounting and risk aversion also relate to test score performance, suggesting that both personality-related traits and preferences are important determinants of outcomes, consistent with the economic model presented in Section 3. In contrast to educational attainment, traits related to Emotional Stability (the opposite of Neuroticism), such as locus of control, are less important for test performance.

7.2.4 Where Course Grades and Standardized Achievement Test Scores Diverge

Course grades and standardized test scores are generally highly correlated. Each form of assessment provides reciprocal evidence on the validity of the other. Willingham, Pollack, and Lewis (2002) estimate a raw correlation of r = 0.62 (p < 0.01) between grade point average and achievement test scores.¹⁸⁷ This strong association—and the objective of each form of assessment to gauge student learning—explains why standardized achievement tests and grades are widely assumed to be "mutual surrogates; that is, measuring much the same thing, even in the face of obvious differences."¹⁸⁸ What are these differences, and how might the contribution of personality to performance vary accordingly?

Standardized achievement tests are designed to enable apples-to-apples comparisons of students from diverse contexts. To this end, standardized achievement tests are uniform in subject matter, format, administration, and grading procedure across all test takers. A course grade, on the other hand, might depend on a particular teacher's judgment.

The power of standardized achievement tests to predict later academic, and occupational outcomes is well established (Kuncel and Hezlett, 2007; Sackett, Borneman, and Connelly, 2008; Willingham, 1985). Nevertheless, cumulative high-school GPA

¹⁸⁷ The correlations were even higher when the test and grades were based on similar subject matter. They use the data from the National Education Longitudinal Study (NELS).

¹⁸⁸ Willingham, Pollack, and Lewis (2002, p. 2).

predicts graduation from college dramatically better than SAT/ACT scores do, even without adjusting for differences in high-school quality (Bowen, Chingos, and McPherson, 2009b). Similarly, high-school GPA more powerfully predicts college rank-in-class (Bowen, Chingos, and McPherson, 2009b, and Geiser and Santelices, 2007).

Perhaps more important than *which* measure of academic achievement—course grades or standardized achievement test scores—is more predictive of later outcomes is *why* these outcomes are related but not entirely interchangeable. Bowen, Chingos, and McPherson (2009b) speculate that aspects of Conscientiousness seem *differentially* essential to earning strong course grades because of what is required of students to earn them. Standardized achievement tests, in contrast to teacherdesigned quizzes, exams, homework assignments, and long-term projects, challenge students to solve *relatively* novel problems. Therefore, it is not surprising that Frey and Detterman (2004) found a correlation of r = 0.82 (p < 0.01) between SAT scores and performance on the ASVAB, an aptitude and achievement test developed for the US Army. In a separate sample, Frey and Detterman found a correlation of r = 0.72 (p < 0.01) between SAT scores and IQ when accounting for censoring. In contrast, the correlation between GPA and IQ is r = 0.23 (p < 0.01) (Poropat, 2009).

In three longitudinal, prospective studies of middle-school students, Duckworth, Quinn, and Tsukayama (2010) compare the variance explained in year-end standardized achievement test scores and GPA by self-control (a facet of Conscientiousness) and fluid intelligence measured at the beginning of the school year. For example, in a national sample of children, fourth-grade self-control was a stronger predictor of ninth-grade GPA ($\beta = 0.40$, p < 0.001) than was fourth-grade IQ $(\beta = 0.28, p < 0.001)$. In contrast, fourth-grade self-control was a weaker predictor of ninth-grade standardized test scores ($\beta = 0.11$, p < 0.05) than was fourth-grade IQ ($\beta = 0.64$, p < 0.001). These findings are consistent with those of Willingham, Pollack, and Lewis (2002), who show that conscientious classroom behaviors are more strongly associated with GPA than with standardized achievement test scores. Similarly, Oliver, Guerin, and Gottfried (2007) found that parent and self-report ratings of distractibility and persistence at age 16 predicted high-school and college GPA, but not SAT test scores. Table 1.10 presents results showing that Conscientiousness and SAT scores are similarly predictive of college GPA. However, in each of the studies, Conscientiousness was measured in college, which presents problems for a causal interpretation of this evidence due to the potential for reverse causality.

In sum, standardized achievement tests and teacher-assigned course grades both reflect students' accumulated knowledge and skill. However, they differ in important ways. The benefits of Conscientiousness, which inclines students to more productive work habits, seem greater for course grades than for test scores. This finding might explain why girls, who are higher than boys in Conscientiousness, reliably earn higher

Source	Sample	Timing of Measurement and Outcome	Controls	Metric	Results	
Conard (2005)	University students in the United States (N = 186).	College GPA and SAT were both self-reported during college. Personality was measured in college.	Class Attendance	Standardized regression coefficient (β)	SAT Total Conscientiousness	0.27** 0.30**
Noftle and Robins (2007)	University students in the United States (N = 10,472).	College GPA and SAT were both self-reported during college. Personality was measured in college.	Gender, Other Big Five Traits.	Standardized regression coefficient (β)	SAT Verbal SAT Math Conscientiousness	0.19*** 0.16*** 0.24***
Noftle and Robins (2007)	University students in the United States (N = 465).	College GPA and SAT were both self-reported during college. ¹ Personality was measured in college.	Gender, Other Big Five Traits	Standardized regression coefficient (β)	SAT Verbal SAT Math Conscientiousness	0.28*** 0.28*** 0.18***
Noftle and Robins (2007)	University students in the United States (N = 444).	College GPA and SAT were both self-reported during college. Personality was measured in college.	Gender, Other Big Five Traits	Standardized regression coefficient (β) .	SAT Verbal SAT Math Conscientiousness	0.18*** 0.25*** 0.22***
Wolfe and Johnson (1995)	University students in the United States (N = 201).	GPA and SAT were provided by the College's Record Office. Personality was measured in college.	High School GPA	Standardized regression coefficient (β).	SAT Total Conscientiousness	0.23*** 0.31***

Table 1.10 The Predictive Power of Conscientiousness Relative to SAT Scores for College GPA

Notes: (1) Self-reported SAT scores and those obtained from college records were highly correlated (r = 0.92). Self-reported GPA and that obtained from college records were highly correlated (r = 0.89).

*Statistically significant at the 10% level; **statistically significant at the 5% level; ***statistically significant at the 1% level.

grades than boys in every subject from primary school through college, but do not reliably outperform boys on either standardized achievement or intelligence tests (Duckworth and Seligman, 2006).

7.3. Labor Market Outcomes

"Eighty percent of success is showing up."

Woody Allen, as quoted in Safire (1989).

It is intuitive that personality traits affect labor market outcomes. Showing up is required for completing a task. However, precisely quantifying the direct effects of personality is more difficult.¹⁸⁹ Recently, social scientists have started to tackle the problem and, in general, find that of the Big Five, Conscientiousness and traits associated with Neuroticism (locus of control and self-esteem) play a particularly important role in determining job performance and wages.¹⁹⁰ The evidence suggests multiple channels of influence, including occupational matching, incentive scheme selection, absenteeism, turnover, and job search.

Aspects of job performance are related to academic performance. For example, both require completing work on a schedule and involve intelligence to varying degrees. Therefore, it is not surprising that as with academic performance, numerous studies and meta-analyses have found that Conscientiousness is associated with job performance and wages (Nyhus and Pons, 2005; Salgado, 1997; Hogan and Holland, 2003; Barrick and Mount, 1991). Figure 1.16 presents correlations of the Big Five and IQ with job performance. Of the Big Five, Conscientiousness is the most associated with job performance but is about half as predictive as IQ. However, Conscientiousness may play a more pervasive role than IQ. The importance of IQ increases with job complexity, defined as the information-processing requirements of the job: cognitive skills are more important for professors, scientists, and senior managers than for semiskilled or unskilled laborers (Schmidt and Hunter, 2004). In contrast, the importance of Conscientiousness does not vary much with job complexity (Barrick and Mount, 1991), suggesting that it pertains to a wider spectrum of jobs. Causality remains an open question, as it does in most of the literature in psychology. The raw correlations presented in Fig. 1.16 do not account for reverse causality, and the authors do not clearly delineate when the measures of personality were taken.

¹⁸⁹ Allen admits that his estimate is partially based on the fact that "80" has better cadence than "70" (Safire, 1989).

¹⁹⁰ Bowles, Gintis, and Osborne (2001b) discuss evidence on the association between personality traits and labor market outcomes.



Figure 1.16 Associations with Job Performance.

Notes: The values for personality are correlations that were corrected for sampling error, censoring, and measurement error. Job performance was based on performance ratings, productivity data, and training proficiency. The authors do report the timing of the measurements of personality relative to job performance. Of the Big Five, the coefficient on Conscientiousness is the only one statistically significant with a lower bound on the 90% credibility value of 0.10. The value for IQ is a raw correlation.

Source: The correlations reported for personality traits come from a meta-analysis conducted by Barrick and Mount (1991). The correlation reported for IQ and job performance come from Schmidt and Hunter (2004).

Facets related to Emotional Stability (the opposite of Neuroticism) are also important for labor market success. However, accounting for reverse causality is particularly important because strong evidence suggests that labor market participation can affect traits related to Neuroticism (see the discussion of Gottschalk, 2005, in Section 8). Several studies have addressed this problem by using measures of personality measured well before individuals enter the labor market and find that locus of control and self-esteem, two facets of Emotional Stability, predict wages (Judge and Hurst, 2007; Drago, 2008; Duncan and Dunifon, 1998). Table 1.11 presents results from the structural model of Heckman, Stixrud, and Urzua (2006), suggesting that standardized adolescent measures of locus of control and self-esteem predict adult earnings to a similar degree as cognitive ability. However, the effects vary across educational levels. In general, their measure of noncognitive ability (personality) affects wages to a similar degree across all education levels, whereas cognitive ability tends to have little effect for GED recipients, high-school dropouts, and college dropouts.

However, more recent evidence suggests that personality affects wages mostly through the channel of educational attainment. In Section 7.1, we presented evidence that personality measures (along with measurements of cognition) are strong predictors of educational attainment. Heckman, Humphries, Urzua, and Veramendi (2011) estimate

	Males		Females	
Schooling Level	Cognitive	Noncognitive	Cognitive	Noncognitive
High-school dropout	.113	.424	.322	.208
	(.076)	(.092)	(.125)	(.103)
GED	.175	.357	.020	.242
	(.107)	(.117)	(.137)	(.153)
High-school graduate	.259	.360	.341	.564
	(.041)	(.059)	(.049)	(.056)
Some college, no degree	.069	.401	.093	.569
	(.086)	(.110)	(.084)	(.116)
Two-year college degree	.039	.368	.206	.279
	(.138)	(.209)	(.096)	(.145)
Four-year college degree	.296	060	.290	.379
	(.075)	(.175)	(.066)	(.103)

Table 1.11 Estimated Coefficients of Cognitive and Noncognitive Factors for Log Hourly Wages

Notes: Standard errors are in parentheses. Sample from NLSY79 males and females at age 30. The sample excludes the oversample of blacks, Hispanics, and poor whites, the military sample, and those currently enrolled in college. The cognitive measure represents the standardized average over the raw ASVAB scores (arithmetic reasoning, word knowledge, paragraph comprehension, math knowledge, and coding speed). The noncognitive measure is computed as a standardized average of the Rosenberg Self-Esteem Scale and Rotter Internal-External Locus of Control Scale. The model also includes a set of cohort dummies, local labor market conditions (unemployment rate), and the region of residence.

Source: Heckman, Stixrud, and Urzua (2006).

a model of sequential educational choice and find that personality, as measured by participation in adolescent risky behaviors, primarily affects age 30 earnings through its effects on education. They find that given educational attainment, the effects of personality variables on a variety of outcomes are weak.¹⁹¹ Further highlighting the role of traits in explaining outcomes by education level, Fig. 1.17 shows that GED recipients—who have lower levels of noncognitive skills but comparable levels of cognitive skills (see the previous section)—have lower wages, lower total wage income, and work fewer hours relative to high-school graduates, when controlling for ability. Other studies by Heckman, Stixrud, and Urzua (2006) and Cattan (2011), using other measures of personality traits, find that the traits affect earnings above and beyond their effects on education and the effects of education on earnings. Resolving these disparate findings is an important topic for future research.

These various studies have shown that personality is associated with wages, but do not explain *why* they are associated other than suggesting that the relationship occurs through the channel of educational attainment. Other mechanisms might be absenteeism, selfemployment, and unemployment. Indeed, controlling for basic demographics, employment

¹⁹¹ See Heckman, Humphries, Urzua, and Veramendi (2011) for a discussion of why their results differ from Heckman, Stixrud, and Urzua (2006).



Figure 1.17 Ability-Adjusted Economic Gaps Relative to Dropouts: GEDs and High-School Graduates for (a) Males and (b) Females.

Notes: Regressions control for baseline AFQT scores, age, mother's highest grade completed, and dummies for urban residence at age 14, southern residence at age 14, and race. Baseline test scores are estimated using the procedure of Hansen, Heckman, and Mullen (2004) as implemented in Carneiro, Heckman, and Masterov (2005). The regressions use the cross-sectional subsample and minority oversamples of the NLSY79 data. The estimation sample is restricted to individuals who never attend college and who have not been incarcerated. Regressions for hourly wage and hours worked are restricted to those reporting more than \$1/h and less than \$100/h, and individuals working less than 4000 hours in a given year. Wage income regressions are restricted to individuals reporting wage incomes between \$1,000/year and \$100,000/year. All monetary values are in 2005 dollars. Standard errors are clustered by individual.

Source: Data come from National Longitudinal Survey of Youth 1979 (NLSY79) as analyzed by Heckman, Humphries, and Mader (2011).

history, and health, Störmer and Fahr (2010) estimate that a standard deviation increase in Emotional Stability and Agreeableness is associated with 12% (p < 0.01) and 9% (p < 0.05) fewer absent days for men and a standard deviation increase in Openness to Experience is associated with 13% (p < 0.01) more absent days for women. However, the study uses contemporaneous measures of personality and absenteeism.¹⁹²

¹⁹² All other Big Five traits were not statistically significant at the 10% level.

Personality plays a role outside of formal employer-employee relationships. Self-employed workers, with either very low or high levels of risk aversion, a trait related to dimensions of personality as discussed in Section 6, tend to remain self-employed for a shorter time, suggesting that they are less suited to self-employment (Caliendo, Fossen, and Kritikos, 2010).¹⁹³

Personality could directly affect the duration of unemployment spells. Gallo, Endrass, Bradley, Hell, and Kasl (2003) find that an internal locus of control is associated with a higher probability of reemployment. A couple of studies have explicitly incorporated locus of control into standard job search models. For example, Caliendo, Cobb-Clark, and Uhlendorff (2010) examine whether a higher locus of control increases the perceived marginal benefit of exerting search effort so that people with a greater internal locus of control will search more intensely and will have a higher reservation wage. Supporting their theory, a one-standard deviation increase in internal locus of control was associated with a 1.9% increase in the reservation wage (p < 0.01) and a 5.3% increase in the number of job applications submitted (p < 0.01), controlling for demographic characteristics and past employment history.¹⁹⁴ Although the measures were contemporaneous, the respondents became unemployed near the time that the locus of control was measured, potentially limiting the role of reverse causality. Similarly, McGee (2010) proposes a model in which people with a higher locus of control believe that search effort has a higher return. His model predicts that those with an internal locus of control search more intensely but have higher reservation wages so that the effect on the hazard rate of leaving unemployment is ambiguous. In line with his predictions, he finds that a one-standard deviation increase in locus of control, measured before market entry, is associated with a 1.3% increase in the reservation wage (p < 0.01) and a 20% increase in the time spent searching for a job per week (p = 0.14).¹⁹⁵ Those with moderate levels of locus of control have the highest hazard rates for leaving unemployment. Consistent with the interpretation that locus of control affects beliefs (not productivity), locus of control has no effect on reemployment wages when controlling for reservation wages.

Personality traits also affect occupational choice. From an economic perspective, some personality traits that reflect ability might be valued more highly in some occupations, and on the supply side, people with certain personality traits that relate to preferences might value the nonpecuniary benefits associated with particular occupations. Supporting this notion, Conscientiousness (Barrick and Mount, 1991, and Ham, Junankar, and Wells, 2009), locus of control and self-esteem (Heckman, Stixrud and Urzua, 2006) predict sorting

¹⁹³ Caliendo, Fossen, and Kritikos (2010) use measures of risk-aversion from 2004 and employment status from 2000 to 2005, assuming that risk aversion did not change over this period.

¹⁹⁴ The associations were partially mediated when controlling for the Big Five, suggesting that locus of control overlaps with the Big Five as discussed in Section 5.

¹⁹⁵ The effect on the reservation wage is higher for people looking for their first jobs.

into occupations. However, these studies use relatively broad occupational categories that might obscure more nuanced influences of personality. Analyzing 18 occupational categories, Cobb-Clark and Tan (2009) find that for men, a one-standard deviation increase in Agreeableness is associated with a 2.8% decrease in the probability of being a manager (p < 0.01) and a 2.9% decrease in being a business professional (p < 0.01). A standard-deviation increase in internal locus of control is associated with 2.8% decrease in the probability of being a manager (p < 0.01). In contrast, for women, a one-standard deviation increase in Openness to Experience is associated with a 2.5% increase in being a manager (p < 0.01).^{196,197}

Furthermore, the value of cognitive ability and personality differs by occupation just as it does by education. Cattan (2011) estimates a structural model of comparative advantage along the lines discussed in Section 3 and finds that different skills are valued differently, depending on the occupation. Accounting for selection, a one-standard deviation increase in adolescent sociability (related to Extraversion) leads to a 7% increase in the wages of managers (p < 0.01), a 4% increase in the wages of sales and service workers (p < 0.01), but leads to a 2% (p < 0.05) decrease in the wages of professionals and has no significant impact on the wages of blue-collar and clerical workers. Selfesteem and locus of control are positively valued in all occupations, but the magnitudes also depend on the occupation. The effects of traits are not uniform on wages across occupations even after controlling for schooling.

Personality affects not only the occupational selection but also the type of compensation scheme selected within an occupation. Dur, Non, and Roelfsema (2010) extend the standard principal-agent model by allowing for workers to reciprocate positive attention from managers by working harder. Their theoretical model implies that promotions, rather than monetary incentives, are more effective for eliciting effort from reciprocal workers. Workers self-select into different compensation schemes. Supporting their model, they find that a one-point increase on a seven-point reciprocity scale for workers is associated with a 5-percentage-point increase of having a job with promotion incentives (p < 0.01). They use contemporaneous measures of reciprocity and job attributes, which could be problematic if pay-for-performance schemes affect reciprocity. Similarly, Dohmen, Falk, Huffman, and Sunde (2009) find that in a German sample self-reported positive reciprocity is positively associated with income, employment, and working overtime. Negative reciprocity tends to operate in the opposite direction. As discussed in Section 6, these measures of social preference relate to personality.¹⁹⁸

¹⁹⁶ The data for occupational categories came from 2001 to 2006, whereas locus of control was measured in 2003–2004 and the Big Five were measured in 2005. Thus, these concerns about reverse causality are potentially important.

¹⁹⁷ They find other statistically significant results at the 5% and 10% levels, which we omit for brevity.

¹⁹⁸ Agreeableness and Conscientiousness are associated with greater positive reciprocity and lower negative reciprocity, whereas Neuroticism is associated with greater negative reciprocity (Dohmen, Falk, Huffman, and Sunde, 2008).

In sum, there are good theoretical reasons as well as some empirical evidence that personality affects labor market outcomes through channels other than education. Conscientiousness and Neuroticism are associated with job performance and wages to a similar but lesser degree than cognitive ability. Personality traits are more important for people with lower levels of job complexity or education level, whereas cognitive ability is more important at higher levels of job complexity. Nevertheless, some research suggests that facets related to Neuroticism might affect labor outcomes primarily through the channel of educational attainment. Other traits, such as Openness to Experience and Agreeableness, affect more specific outcomes, such as selection into particular careers or types of compensation. Table A10 in Web Appendix A7 summarizes a variety of studies that associate personality with labor market outcomes.

7.4. Personality and Health¹⁹⁹

A link between personality and health has been noted for thousands of years. Hippocrates argued that an imbalance of the four temperaments would affect both personality and physical health.²⁰⁰ Consistent with Hippocrates' ideas, recent evidence suggests that personality predicts health. The mechanisms are relatively unexplored but some empirical evidence suggests that personality affects health-related behavior, psychological responses, and social relationships (Kern and Friedman, 2010a).

A growing body of research shows that personality measures predict longevity. Roberts, Kuncel, Shiner, Caspi, and Goldberg (2007) review evidence from 34 different studies on the predictive validity of the Big Five personality traits, relative to that of cognitive ability and socioeconomic status. Most studies in their meta-analysis control for relevant background factors, including gender and severity of disease. Roberts and colleagues convert the results of each study into correlation coefficients that can be compared across studies. As shown in Fig. 1.18, Conscientiousness was a stronger predictor of longevity than any other Big Five trait and a stronger predictor than either IQ or socioeconomic status.²⁰¹ In general, traits related to Conscientiousness, Openness to Experience, and Agreeableness are associated with longer lives, whereas those related to Neuroticism are associated with shorter life spans.²⁰² However, the magnitudes of the relationships vary across studies and not all results are replicable. Although the specific channels through which personality affects longevity and health are largely unknown, several studies provide some clues.

¹⁹⁹ This section is a summary of Pietro Biroli's extensive discussion of personality and health that is presented in Web Appendix A7.1.

²⁰⁰ See Hampson and Friedman (2008) and Friedman (2007) for a brief historic review.

²⁰¹ The timing of the measurements of personality relative to the outcomes varies by study.

²⁰² See Martin, Friedman, and Schwartz (2007); Kern and Friedman (2008); Mroczek and Spiro (2007); Boyle et al. (2005); Schulz, Bookwala, Knapp, Scheier, and Williamson (1996); and Kubzansky, Sparrow, Vokonas, and Kawachi (2001).



Figure 1.18 Correlations of Mortality with Personality, IQ, and Socioeconomic Status (SES).

Notes: The figure represents results from a meta-analysis of 34 studies. Average effects (in the correlation metric) of low socioeconomic status (SES), low IQ, low Conscientiousness (C), low Extraversion/Positive Emotion (E/PE), Neuroticism (N), and low Agreeableness (A) on mortality. Error bars represent standard error. The lengths of the studies represented vary from 1 year to 71 years. Source: Roberts, Kuncel, Shiner, Caspi, and Goldberg (2007).

Personality may affect health-related behavior, such as smoking, diet, and exercise. For example, Hampson, Goldberg, Vogt, and Dubanoski (2007) find that high scores of teacher assessments of Extraversion, Agreeableness, and Conscientiousness during elementary school predict overall health behaviors during midlife (less smoking, more exercise, better self-rated health) and indirectly affect health through educational attainment.²⁰³ The effects that were statistically significant at the 5% level or less ranged from 0.06 for the effect of Extraversion on physical activity to 0.12 for the effect of Conscientiousness on self-reported health status. Both the initial level and the growth in hostility (a facet of Neuroticism) throughout elementary school predict cigarette, alcohol, and marijuana use in high school, and sociability (a trait related to Extraversion) predicts drinking but not smoking (Hampson, Tildesley, Andrews, Luyckx, and Mroczek, 2010). As Fig. 1.19 illustrates, Heckman, Stixrud and Urzua (2006) find that their personality factor affects the probability of daily smoking for males. The gradient is steepest at the high and low quantiles of the distribution.

²⁰³ Conti, Heckman, and Urzua (2010a,b) and Heckman, Humphries, Urzua, and Veramendi (2011) present evidence on the causal relationship between education and health, and also survey the previous literature on this question.



Figure 1.19 Probability of Daily Smoking by Age 18 for Males. (a) By Decile of Cognitive and Noncognitive Factors. (b) By Decile of Cognitive Factor. (c) By Decile of Noncognitive Factor.

Notes: The data are simulated from the estimates of the model of Heckman, Stixrud, and Urzua (2006) and their NLSY79 sample. They use the standard convention that higher deciles are associated with higher values of the variable. The confidence intervals are computed using bootstrapping (200 draws). Solid lines depict probability, and dashed lines, 2.5–97.5% confidence intervals. The upper curve is the joint density. The two marginal curves (ii) and (iii) are evaluated at the mean of the trait not being varied.

Source: Heckman, Stixrud, and Urzua (2006, Figure 22).

Few studies explore *how* personality affects health throughout the life cycle (Kern and Friedman, 2010b). The relationship between health and personality is complicated because health can affect personality.²⁰⁴ Some studies investigate the mechanisms by which personality affects health by considering how initial endowment of traits and health affect midlife outcomes, such as healthy behavior and education, which in turn can influence health and longevity. For example, Gale, Batty, and Deary (2008) find that a one–standard deviation increase in age-10 locus of control decreases the risk of adult obesity by 8% (p < 0.05). Similarly, Friedman, Kern, and Reynolds (2010) find that in a cohort of gifted children, Conscientiousness better predicted longevity and

²⁰⁴ Pesonen et al. (2008); Ryden et al. (2003); Sell, Tooby, and Cosmides (2009); and Hoffman, Fessler, Gneezy, and List (2010).

social interactions at age 70. They find that Neuroticism is associated with worse health for women but better health for men. Most studies do not account for the possibility that health and personality exhibit dynamic complementarities over the life cycle.

Several studies have controlled for reverse-causality by using structural models to estimate the life-cycle evolution of health. Using a structural model, Conti and Heckman (2010) estimate the causal relationship between personality traits, initial health endowments, and schooling choices and postschooling outcomes. Children sort into higher education based on cognitive ability, personality traits, and initial health endowment. Furthermore, personality and health status measured during youth explain more than 50% of the difference in poor health, depression, and obesity at age 30, observed between the educated and less educated. Figure 1.20 shows that for males, personality and health endowments are more predictive than are cognitive endowments, whereas for females, all three are roughly equally predictive. Using similar methods, Savelyev (2010) finds that both Conscientiousness (measured in youth) and higher education increase survival through age 80, but these traits serve as substitutes for each other so that effects of education are strongest at low levels of Conscientiousness.

In sum, Conscientiousness seems to be the most important Big Five trait in predicting health outcomes. Personality likely affects health through behaviors such as smoking, eating, and exercising. Studies that model the dynamic evolution of health over the life cycle find that personality affects health outcomes as much as cognitive measures or even more so in some cases.



Figure 1.20 Effects of Cognitive, Noncognitive, and Health Endowments on Self-rated Health (A Lower Number Corresponds to a Better Outcome). (a) Males. (b) Females.

Notes: Effects of endowments on fair or poor health outcomes for males (a) and females (b). The endowments and the outcomes are simulated from the estimates of the model in each panel; when the authors compute the effect of each endowment on the outcome, they integrate out the observable characteristics and fix the other two endowments at their overall means. Source: Conti and Heckman (2010).

7.5. Crime²⁰⁵

Few studies have examined the relationship between the Big Five and criminal behavior. The available evidence suggests that Big Five Conscientiousness and Agreeableness are important protective factors against criminal activity. Figure 1.21 illustrates that in a sample of at-risk youth, boys who had committed severe delinquent behaviors were more than three quarters of a standard deviation lower in Agreeableness and Conscientiousness, as measured by mother's reports at age 12 or 13, than boys who had committed minor or no delinquent behaviors up to that age (John, Caspi, Robins, and Moffitt, 1994).

Much of the literature in criminology focuses on the effects of self-control on crime. People with low self-control are "impulsive, insensitive, physical (as opposed to mental), risk taking, short sighted, and nonverbal" (Gottfredson and Hirschi, 1990, p. 90). Measures of self-control are associated with Big Five Conscientiousness (O'Gorman and Baxter, 2002). Several studies have confirmed that self-control is associated with criminal activity. In an international sample, controlling for basic demographics, a measure



Figure 1.21 Juvenile Delinquency and the Big Five.

Notes: Delinquents are those who have committed at least one of the following: breaking and entering, strong-arming, or selling drugs. Nondelinquents have committed at most one of the following: stealing at home, vandalism at home, or theft of something less than \$5. The y-axis reports mean differences in standardized scores of the Big Five measures based on mother's reports. The measures were taken at ages 12–13 and reflect cumulative delinquent behavior. Source: John, Caspi, Robins, and Moffitt (1994).

²⁰⁵ This section summarizes the more comprehensive survey of the literature on personality and crime prepared by Amanda Agan. See Web Appendix Section A7.2 for her survey.

of self-control explained between 10% and 16% of the variance in contemporaneously measured theft, assault, drug use, and vandalism (Vazsonyi, Pickering, Junger, and Hessing, 2001). Self-control relates to controlling impulsive behavior, so it is not surprising that sensation seeking and impulsivity are also positively associated with crime. In a sample of college students, partial correlations between a crime factor²⁰⁶ and sensation seeking and impulsive behavior were 0.27 and 0.13, respectively, when controlling for peer behavior and measures of risk appraisal (Horvath and Zuckerman, 1993).

Self-control might not be the entire story. Negative emotionality—a tendency toward depression likely related to Neuroticism—is associated with contemporaneously measured delinquency. Raw correlation coefficients range from r = 0.13 for whites (p < 0.05) and r = 0.20 for blacks (p < 0.05) in one sample (Caspi et al., 1994) to r = 0.22 (p < 0.01) in another sample (Agnew, Brezina, Wright, and Cullen, 2002). None of these studies control for cognitive ability or address causality.

An emerging literature investigates the causal effects of education on crime. Heckman, Stixrud, and Urzua (2006) estimate a causal model of personality and education accounting for reverse causality. They find that both cognitive traits and noncognitive traits, as captured by locus of control and self-esteem, are affected by schooling.²⁰⁷ These traits in turn are approximately equally predictive of criminal activity.²⁰⁸

Using changes in compulsory schooling laws as an instrument, Lochner and Moretti (2004) and Machin, Marie, and Vujić (2010) find that years of education are negatively associated with criminal activities in the United States and United Kingdom, respectively. In a structural model of skill production, Cunha, Heckman, and Schennach (2010) show that personality traits are relatively more important in predicting criminal activity than are cognitive traits.

8. STABILITY AND CHANGE IN PERSONALITY TRAITS AND PREFERENCES

In this section, we review empirical evidence that shows that personality and IQ change over the life cycle. We explore three channels through which personality can change. First, we discuss the contribution of ontogeny (programmed developmental processes common to all persons) and sociogeny (shared socialization processes), and show how aspects of personality, such as sensation seeking, evolve as the brain develops.

²⁰⁶ The crime factor is based on arrest for selling or buying drugs, shoplifting, driving while drunk, perjury, forging checks, and vandalizing.

²⁰⁷ We discuss this work in Section 8.

²⁰⁸ Their measure of prediction is the effect of decile improvements of cognition and personality traits on the probability of being in jail.

Second, we show how personality changes through external forces that operate through alterations in normal biology, such as brain lesions and chemical interventions. Third, and most relevant for policy, we show that education, early childhood interventions, and parental investment can affect personality throughout the life cycle. We also discuss the less-abundant evidence on the malleability of preferences.

8.1. Broad Evidence on Changes in Traits over the Life Cycle

The malleability of personality can be defined and measured in several ways: *Mean-level change* refers to change over time in absolute levels of a trait and is measured by changes in measures of a trait over time. *Rank-order change*, in contrast, refers to changes in the ordinal ranking of a trait in a population and is measured by rank correlations among longitudinal measures. One commonly held view is that rank-order, as well as mean-level, change in personality is nearly impossible after early adulthood. The speculation of James (1890) that "in most of us, by the age of thirty, the character has set like plaster, and will never soften again" (pp. 125–126) is widely touted (see Costa and McCrae, 1994; McCrae and Costa, 1990, 1994, 1996, 2003; Costa, McCrae, and Siegler, 1999). However, mounting evidence suggests that the personality-as-plaster view is not correct (Roberts, Walton, and Viechtbauer, 2006, and Roberts and Mroczek, 2008).

During the early years of life, mean-level changes in measured traits are obvious and dramatic. For example, children become much more capable of self-control as they move from infancy into toddler and preschool years (McCabe, Cunnington, and Brooks-Gunn, 2004; Mischel and Metzner, 1962; Posner and Rothbart, 2000; Vaughn, Kopp, and Krakow, 1984). But mean-level changes in measured personality are also apparent, albeit less extreme, later in life. In a 2006 meta-analysis of longitudinal studies, Roberts, Walton, and Viechtbauer (2006) examine cumulative lifetime change in Big Five Openness to Experience, Conscientiousness, Extraversion, and Agreeableness. They disaggregate Big Five "Extraversion" into social dominance (assertiveness, dominance) and social vitality (talkativeness, gregariousness, and sociability). Figure 1.22 shows that people typically become more socially dominant, conscientious, and emotionally stable (nonneurotic) across the life cycle, whereas social vitality and Openness to Experience rise early in life and then decrease in old age.²⁰⁹ Surprisingly, after childhood, the greatest mean-level change in most measured personality traits takes place not during adolescence but rather in young adulthood.

²⁰⁹ Figure A3 in Section A9 of the Web Appendix presents results for a variety of cognitive, personality, and preference parameters from a cross-sectional study based on the GSOEP data. Samples are small and standard errors are large. Many preference parameters show a surprising stability over the life cycle.



Figure 1.22 Cumulative Mean-Level Changes in Personality across the Life Cycle.

Note: Social vitality and social dominance are aspects of Big Five Extraversion. Cumulative d values represent total lifetime change in units of standard deviations ("effect sizes"). Source: Figure taken from Roberts, Walton, and Viechtbauer (2006) and Roberts and Mroczek (2008). Reprinted with permission of the authors.

In contrast, a longitudinal study of adult intellectual development shows mean-level declines in cognitive skills, particularly cognitive processing speed, after age 55 or so (Schaie, 1994). Figure 1.23a shows mean-level changes in cognitive skills using a longitudinal analysis, and Figure 1.23b shows mean-level changes using a cross-sectional analysis.²¹⁰ As

²¹⁰ Cross-sectional estimates of mean-level change are biased by cohort effects (e.g., the Flynn effect), whereas longitudinal estimates are biased by test–retest learning (when the same IQ tests are administered repeatedly to the same subjects) and by selective attrition. Thus, both estimates must be considered in conjunction as evidence for mean-level change.



Figure 1.23 (a) Longitudinal Analysis and (b) Cross-Sectional Analysis of Mean-Level Change in Cognitive Skills over the Life-Span.

Note: T-scores on the y-axis are standardized scores with a mean of 50 and a standard deviation of 10. Source: Figures taken from Schaie (1994), used with permission of the publisher.

schematically illustrated in Fig. 1.24, fluid intelligence decreases and crystallized intelligence increases over the life cycle (Horn, 1970). Accumulated skills and knowledge are important: most of us would rather use an experienced cardiac surgeon who has seen hundreds of cases just like ours to perform our surgery, rather than an exceptionally bright young surgeon with minimal experience.



Figure 1.24 Fluid Intelligence Decreases and Crystallized Intelligence Increases across the Life-Span. *Source: Figure from Horn (1970), used with permission of Elsevier.*



Figure 1.25 Rank-Order Stability of Personality over the Life Cycle.

Notes: The meta-analysis reflects test–retest correlations over, on average, 6.7-year periods. Source: Figure taken from Roberts and DelVecchio (2000). Reprinted with permission of the authors.

Rank-order stability in measured personality increases steadily over the life span. Figure 1.25 shows that 7-year test-retest stability estimates for personality plateau far from unity, at r = 0.74, which is about the same level as terminal stability estimates for IQ (Roberts and DelVecchio, 2000). However, measured personality does not reach



Figure 1.26 Rank-Order Stability of IQ across the Life Span.

Notes: The points represent 10-year, test-retest correlations over 10-year intervals. Grade level, not age, is on the x-axis.

Source: Figure reproduced from Hopkins and Bracht (1975), used with permission of the publisher.

this plateau until at least age 50, whereas IQ reaches this plateau by age six or eight (Hopkins and Bracht, 1975, and Schuerger and Witt, 1989). Figure 1.26 shows rank-order stability of IQ over broad age ranges.

8.2. Evidence on Ontogenic and Sociogenic Change

A useful dichotomy contrasts *normative change*, defined as changes that are caused either by biological programming (ontogenic) or by predictable changes in social roles (sociogenic), with *nonnormative change*, encompassing both intentional change, caused by deliberate, self-directed efforts, deliberately chosen changes in social roles and atypical life events (trauma, for example).²¹¹

If, as McCrae and colleagues have claimed, normative changes reflect genetically programmed processes, then investment should not affect change. The current literature in psychology claims that genetic factors are largely responsible for *stability* in personality in

²¹¹ "Normative" in this context refers to what most people, or average persons, experience. If most people deliberately do something that causes change, it would be normative. But that seems unlikely. Therefore, most deliberative change is nonnormative, but logically this is not necessarily true.

adulthood, whereas environmental factors are mostly responsible for *change* (Blonigen, Hicks, Krueger, Patrick, and Iacono, 2006, and Plomin and Nesselroade, 1990).^{212,213}

In a longitudinal study of twins surveyed at age 20 and then again at age 30, about 80% of the variance of the stable component of personality was attributed to genetic factors (McGue, Bacon, and Lykken, 1993). In the same study, change in measured personality was mostly attributed to environmental factors. Helson, Kwan, John, and Jones (2002), for example, document the substantial influence that social roles and cultural milieu can have on personality development. Their analysis is consistent with an economic model of investment and the response of measured traits to incentives. However, recent evidence suggests that environmental factors and, in particular, stable social roles also contribute to stability in personality and that genetic factors can contribute to change (see Roberts, Wood, and Caspi, 2008, for a review).

Research on IQ also points to the enduring effects of genes, in contrast to more transient effects of environmental influences, which depend on a multitude of unstable variables, including social roles, levels of physical maturity and decline, and historical and cultural milieu.²¹⁴ Increases in the heritability of IQ from childhood (about 40%) to adulthood (estimates range from 60-80%) are well documented in studies of behavioral genetics and possibly reflect increasing control of the individual (vs. parents) over environment (Bergen, Gardner, and Kendler, 2007; McGue, Bouchard, Iacono, and Lykken, 1993; Plomin, DeFries, Craig, and McGuffin, 2002).²¹⁵ Heritability estimates for Big Five traits are relatively stable across the life cycle at about 40-60% (Bouchard and Loehlin, 2001).²¹⁶ Behavioral genetics studies typically estimate the effect of common parental environments on adult measures of outcomes to be near zero, but Turkheimer, Haley, Waldron, D'Onofrio, and Gottesman (2003) find estimates from such studies to be biased downward by the overrepresentation of middle- and upperclass families. Among poor families, Turkheimer et al. find that 60% of the variance in IQ is accounted for by shared environment. In addition, he finds that heritability estimates are much smaller than they are for affluent families, whereas among affluent families, the contribution of heritability is much larger. Krueger and colleagues have

²¹² Plomin and the essays in the December issue of Monographs for the Society for Research in Child Development (Kovas et al., 2007) extend this analysis to childhood.

²¹³ We note that there is controversy in the literature about the validity of conventional estimates of heritability. It centers on the linearity and additivity assumptions, the assumed absence of interactions between genes and environment, and the assumption that genes do not select environments.

²¹⁴ We note here that while genes remain constant through the life cycle, the expression of genes is determined, in part, by experience.

²¹⁵ Devlin, Daniels, and Roeder (1997) suggest that traditional estimates of the heritability of IQ may be inflated because they fail to take into account the effect of the environment on conditions in the maternal womb. See also Rutter (2006b) and an emerging literature on epigenetics.

²¹⁶ Lykken (2007) suggests that heritability estimates for personality are substantially higher when situational influence and measurement error are minimized by taking multiple measures at least a few months apart.

recently demonstrated that other moderators also influence the heritability of traits (see Krueger, South, Johnson, and Iacono, 2008).²¹⁷

Genes exert their influence in part through the selection and evocation of environments that are compatible with one's genotype—a phenomenon sometimes referred to as "gene–environment correlation" or "nature via nurture" (see Rutter, 2006a). As individuals move from childhood to adulthood, they have more control over their environments, and thus, gene–environment correlation becomes more important because shared environments become less common.²¹⁸

Substantial but temporary influence of environment is a basic assumption of the Dickens-Flynn model reconciling the high heritability of IQ and massive gains of IQ between generations (Dickens and Flynn, 2001).²¹⁹ The relatively short half-life of common environmental influences may also explain why adopted children resemble their biological parents more and more and their adopted parents less and less as they grow older (Scarr, Weinberg, and Waldman, 1993).²²⁰

It is important to note that the family studies of genetic influence measure only the effects of *shared* environments, which become less similar as children age. Thus, even identical twins may be motivated to seek out different environments over time (Rutter, 2006a). Recent evidence that first-born children grow up, on average, to have three points higher IQ than their younger siblings reinforces the point that parents do not necessarily provide identical environments in childhood (Kristensen and Bjerkedal, 2007). Lizzeri and Siniscalchi (2008) develop an economic model of differential parenting of siblings.

As mentioned earlier, genes could affect not only the base level of personality but also how personality changes over the life cycle. Just as people grow taller throughout childhood, people's personalities might naturally develop, even without investment. Steinberg (2008) speculates that typical biological (ontogenic) development explains the surge of risk taking in adolescence followed by the decline in adulthood. Figure 1.27 illustrates

²¹⁷ It is important to note that shared environment is not the same as environment. Children may be treated individually by parents.

²¹⁸ Gene-environment interactions are another means by which genes and environment jointly influence traits. The effects of the environment depend on the genes and vice versa (see Caspi et al., 2003; Moffitt, Caspi, and Rutter, 2005; and Caspi et al., 2002).

²¹⁹ A second crucial assumption is that environmental influence can be amplified by a "social multiplier" effect: smarter individuals create for one another an enriched environment, which in turn increases intelligence, e.g. Some caution must be taken in relying on the claims in this literature. Blair, Gamson, Thorne, and Baker (2005) attribute the Flynn effect to increasing access to formal schooling early in the twentieth century and, from the mid-century onward, to increasing fluid cognitive demand of mathematics curricula. Flynn (2007) concurs about the former but believes that the latter had negligible impact.

²²⁰ The literature establishes that shared environments become less important as children age. This literature does *not* say that environments do not matter. This effect can arise because genetically similar children (or their parents) choose different environments to distinguish themselves or because of parental investment (Lizzeri and Siniscalchi, 2008).



Figure 1.27 Proportion of Individuals in Each Age Group Scoring at or Above the Mean for 26- to 30-Year Olds on Indices of Intellectual and Psychosocial Maturity.

Source: From Steinberg, Cauffman, Woolard, Graham, and Banich (2009) submitted for publication.

his conjecture about how basic intellectual ability and psychosocial maturity (related to, e.g., impulsivity, risk perception, sensation seeking, future orientation) evolve over the life cycle.²²¹ He argues that intellectual ability matures more rapidly than psychosocial maturity. In his model, increases in adolescent risk taking are due to a restructuring of the brain's dopaminergic system (responsible for the brain's reward processing) in such a way that immediate or novel experiences yield higher rewards, especially in the presence of peers. He attributes declines in risk taking to the development of the brain's cognitive control system, specifically improvements in the prefrontal cortex that promote aspects of executive function such as response inhibition, planning ahead, weighing risks and rewards, and the simultaneous consideration of multiple information sources. Interestingly, even in his model, sensation seeking partially depends on the presence of peers, which corresponds to aspects of the situation (*h* in the framework of Section 3). This example highlights the difficulty in disentangling situational and biological changes in personality.

What factors other than preprogrammed genetic influences might account for meanlevel changes in personality? Personality change in adulthood may be precipitated by major shifts in social roles (e.g., getting a job for the first time or becoming a parent). If social role changes are experienced by most people in a population at the same time, we will observe the effects as mean-level changes in measured personality. If, on the other hand, these social roles are not assumed synchronously, we will observe rank-order changes.

One difficulty with many of the studies that address this question is the problem of reverse causality. Changes in personality may drive social role changes rather than the other way around.

²²¹ Spear (2000a,b) also finds that sensation seeking reaches its peak in adolescence.

8.3. External Changes to Biology

The previous subsection discusses the difficulty in disentangling biological changes in personality from environmental or situational effects. In this subsection, we provide some evidence on causal changes in personality due to external forces that either damage parts of the brain or abruptly alter the chemistry of the brain.

8.3.1 Brain Lesion Studies

Brain lesion studies provide the most dramatic and convincing evidence that personality can change. The most famous example is Phineas Gage, a construction foreman whose head was impaled by a metal spike and who subsequently changed from being polite and dependable to rude and unreliable but preserved his problemsolving abilities (Damasio, Grabowski, Frank, Galaburda, and Damasio, 2005). Since then, there have been many more case studies of patients with brain damage. For example, Mataró et al. (2001) describe the behavior of a Spanish patient whose head was impaled by an iron spike, injuring both frontal lobes. Like Phineas Gage, his behavior changed. After the accident, he had difficulty planning, became more irritable, and had problems regulating emotions. Unlike Phineas, he was cheerful and did not display antisocial behavior, suggesting that personality is malleable in different dimensions, even through brain damage. The effects of brain damage are persistent. After 5 years, patients who suffered traumatic head injuries have social impairments, such as anger control, even when their performance on cognitive tasks returns to the normal range (Lezak, 1987).

Using more advanced methods, neuroscientists have delved deeper into the inner workings of the brain. Some recent studies have investigated how two parts of the brain, the amygdala and ventromedial prefrontal cortex (VMPC), affect personality by regulating emotion. Bechara (2005) discusses how emotion might allow people to assign and store value to particular outcomes in a way that is useful for decision making. The amygdala is believed to signal impulsive emotional responses to immediate environmental stimuli, such as reacting quickly to a snake. In contrast, the VMPC is believed to signal reflective emotional responses to memories and knowledge. These two parts of the brain conflict with each other when people make decisions: signals from the amygdala induce behavior that implicitly values immediate outcomes, whereas signals from the VMPC reflect long-run considerations. The stronger signal dictates the resultant behavior. People with damage to these parts of the brain exhibit changes in personality. For example, people with damage to the VMPC, the part that regulates reflective emotion, tend to act impulsively and seem to overvalue short-term outcomes in a way that leads to long-term financial loss and loss of friendships, despite having relatively normal levels of intellectual capacity. These findings are consistent with McClure, Laibson, Loewenstein, and Cohen's β/δ system that describes hyperbolic discounting (McClure, Laibson, Loewenstein, and Cohen, 2004). However, some recent research in neuroscience challenges this theory and presents empirical evidence that contradicts β/δ theory (Monterosso and Luo, 2010).

Further experiments involving these parts of the brain highlight why attempts to separate cognitive and noncognitive traits might be futile. For example, Bechara and Damasio (2005) study the performance of patients with lesions in the VMPC in a seemingly cognitive task. Participants in their experiment were given the Iowa Gambling Task, in which they repeatedly chose between four decks of cards that represented lotteries of different value, unknown to the participant at the onset. Throughout the experiment, the authors also measured skin conductance responses (SCRs), a known physiological reflection of emotion. By trial and error, participants without lesions learned to choose the "better" decks of cards with lower short-term payoffs but higher average payoffs. The normal participants also showed emotional activity both when picking their card and when receiving the rewards or penalties. In contrast, people with lesions never learned to pick the better decks, seemingly because they could not develop emotional responses. Patients with damage to the amygdala never showed emotional response to rewards or penalties, suggesting they never learned to value the outcomes. Patients with damage to the VMPC showed emotional response only when receiving the reward or penalty but not when selecting decks, suggesting that they might not have reflective emotional responses crucial in considering future consequences. These findings suggest that emotion helps to guide decision making. Numerous other studies show the role of the amygdala in signaling emotions and its relationship to cognition and behavior (Phelps, 2006).

8.3.2 Chemical and Laboratory Interventions

A few recent studies show that it is possible to alter preferences and personality through experiments that change the brain's chemistry. For example, magnetic disruption of the left lateral prefrontal cortex can increase experimentally elicited discount rates (Figner et al., 2010). Similarly, nasal sprays of oxytocin increase trust (distinct from altruism or ability to assess probabilities) in a game-theoretic experiment (Kosfeld, Heinrichs, Zak, and Fehr, 2005). As discussed in Section 5, the Big Five traits are linked to personality disorders. Therefore, it is not surprising that administering paroxetine, a drug for treating depression, decreases Neuroticism and increases Extraversion. More surprising is that the drug affects personality above and beyond its direct effects on depression. Furthermore, patients who become less neurotic are also less likely to relapse even after treatment, suggesting that paroxetine might have a long-lasting impact through a biochemical change in the brain (Tang et al., 2009). Similarly, Knutson et al. (1998) find evidence that paroxetine can diminish hostile behavior through a decrease in general negative effect.

8.4. The Evidence on the Causal Effects of Parental Investment, Education, and Interventions

Even though brain lesion studies and laboratory experiments provide convincing causal evidence that personality can be changed, they are not viable mechanisms for large-scale policy interventions. A growing body of evidence suggests that education, parental investment, and interventions can causally affect personality traits. More than just ontogenic and sociogenic processes are at work. A major contribution of economics to the literature in psychology is to develop and apply a framework to investigate how investment, including education, work experience, and self help, changes traits. We discuss the evidence on trait changes through these mechanisms, using the theoretical framework introduced in Section 3.8 as a guide. In all of the models considered in this subsection, the development of traits arises from purposive actions of agents and not just from exogenous biological processes.

The empirical literature has not estimated the investment model (1.16) in Section 3.8 in its full generality. It focuses on estimating productivity functions (1.1) specified in terms of traits θ . Due to data limitations, there is no empirical work yet to report that standardizes for effort or for situation. To simplify the notation, we keep h implicit.

Denote the productivity traits at age v by θ^{v} . Substituting for actions in terms of their determinants, the performance on task *j* at age *v* is

$$P_{i}^{\mathsf{v}} = \phi_{i}^{\mathsf{v}}(\theta^{\mathsf{v}}, e_{i}^{\mathsf{v}}), \, j \in \{1, \dots, J\}, \, \mathsf{v} \in \mathcal{V}$$
(1.22)

where $e_j^{\mathbf{v}}$ is effort devoted to task *j* at time **v**. For simplicity, break $\theta^{\mathbf{v}}$ into cognitive, μ , and personality, π , components:

$$\theta^{\mathsf{v}} = (\theta^{\mathsf{v}}_{\mu}, \theta^{\mathsf{v}}_{\pi}),$$

using the notation of Section 3.²²² e_j^v depends on preferences, rewards and information.

The vector of productivity traits evolves via a simplified version of (1.16):

$$\theta^{\nu+1} = \eta^{\nu}(\theta^{\nu}, IN^{\nu}, h^{\nu}), \ \nu \in \mathcal{V}.$$
(1.23)

 IN^{ν} is interpreted very broadly to include investment by parents, schools, work experience, and interventions. θ^0 is a vector of initial endowments. Some components of effort may be included in investment.

The productivity of investment can depend on the age at which it is made. A crucial feature of the technology that helps to explain many findings in the literature on

²²² See Eq. (1.15).

skill formation (see Cunha and Heckman, 2007, 2009) is complementarity of traits with investment:

$$\frac{\partial^2 \eta^{\mathsf{v}}(\theta^{\mathsf{v}}, IN^{\mathsf{v}}, h^{\mathsf{v}})}{\partial \theta^{\mathsf{v}} \partial (IN^{\mathsf{v}})^{'}} \ge 0.$$
(1.24)

Technology (1.23) is characterized by *static complementarity* between period v traits and period v investment. The higher the θ^v , the higher the productivity of the investment. There is also *dynamic complementarity* if the technology determines period v + 1traits (θ^{v+1}). This generates complementarity between investment in period v + 1 and investment in period s, s > v + 1. Higher investment in period v raises θ^{v+1} because technology is increasing in IN^v , which in turn raises θ^s because the technology is increasing in θ^v , between v and s. This, in turn, increases $\frac{\partial \eta^s(\cdot)}{\partial IN^s}$ because θ^s and IN^s are complements, as a consequence of (1.24).

Dynamic complementarity explains the evidence that early nurturing environments affect the ability of animals and humans to learn. It explains why investments in disadvantaged young children are so productive (see Knudsen, Heckman, Cameron, and Shonkoff, 2006). Early investments enhance the productivity of later investments. Dynamic complementarity also explains why investment in low-ability adults often has such low returns—because the stock of θ^{v} is low.²²³ Using dynamic complementarity, one can define *critical* and *sensitive* periods for investment. If $\frac{\partial \eta^{v}(\cdot)}{\partial IN^{v}} = 0$ for $v \neq v^{*}$, v^{*} is a critical period for that investment. If $\frac{\partial \eta^{v}(\cdot)}{\partial IN^{v}}$ for all $v' \neq v$, v is a sensitive period.²²⁴ The technology of skill formation is consistent with a body of evidence that shows critical and sensitive periods in human development for a variety of traits.²²⁵

Figure 1.28 shows how adult outcomes are shaped by *sequences* of investments over the life cycle. The importance of the early years depends on how easy it is to compensate for adverse early effects with later investment. The literature shows that resilience and remediation are possible, but are more costly later on.²²⁶ The accumulation of investments over the life cycle of the child determines adult outcomes and the choices people will make when they become adults. To capture these interactive effects requires nonlinear models.

For the purposes of policy analysis, it is important to know at which stage of the life cycle interventions are the most effective and to move beyond the correlations between early life and later life events to understand the mechanisms of skill formation. Cunha

²²³ See the evidence in Cunha and Heckman (2007); Heckman (2007); Heckman (2008b); and in Cunha, Heckman, Lochner, and Masterov (2006).

²²⁴ This expression is evaluated at common levels of the inputs on both sides of the expression.

²²⁵ See the evidence summarized in Heckman (2008b); Cunha and Heckman (2009); and Cunha, Heckman, Lochner, and Masterov (2006).

²²⁶ See Cunha and Heckman (2007, 2009); Cunha, Heckman, and Schennach (2010); and Heckman (2008b).



Figure 1.28 A Life Cycle Framework for Organizing Studies and Integrating Evidence: V + 1 Period Life Cycle. θ^{v} : Capacities at *v*. IN^{v} : Investment at *v*. h^{v} : environments at time *v*. $\theta^{v+1} = \eta^{v}$ (θ^{v} , IN^{v} , h^{v}).

and Heckman (2008) and Cunha, Heckman, and Schennach (2010) estimate technologies of skill formation to understand how the skills of children evolve in response to the stock of skills children have already accumulated, the investments made by their parents and the stock of skills accumulated by the parents.

The most general empirical specification of the technology to date is that of Cunha, Heckman, and Schennach (2010). They allow for Q + 1 different developmental stages in the life of the child: $q \in \{0,...,Q\}$. Developmental stages may be defined over specific ranges of ages, $v \in \{0,...,V\}$, so $Q \leq V$. They assume that each component of θ^v and IN^v can be represented by a scalar as can environment h^{v} .²²⁷ Letting IN_k^v be investment in trait k at age v, they estimate a CES, stage-specific version of (1.23) for trait k at stage q:

$$\theta_{k}^{\mathbf{v}+1} = \left[\gamma_{\mu,k}^{q} (\theta_{\mu}^{\mathbf{v}})^{\sigma_{k}^{q}} + \gamma_{\pi,k}^{q} (\theta_{\pi}^{\mathbf{v}})^{\sigma_{k}^{q}} + \gamma_{IN,k}^{q} (IN_{k}^{\mathbf{v}})^{\sigma_{k}^{q}} + \gamma_{E,k}^{q} (h^{\mathbf{v}})^{\sigma_{k}^{q}} \right]^{\overline{\sigma_{k}^{q}}},$$

$$\gamma_{m,k}^{q} \ge 0, \quad \sum_{m \in (\mu,\pi,IN,h)} \gamma_{m,k}^{q} = 1 \text{ for all } k \in \{\mu,\pi\} \text{ and } q \in \{0,...,Q\}.$$
(1.25)

²²⁷ For them, environment is parental environment.

A main finding of Cunha, Heckman, and Schennach (2010) is that the elasticity of substitution σ_{μ}^{q} governing the acquisition of cognitive traits decreases with q. This is consistent with other evidence that shows the declining malleability of cognition with age, that is, cognitive deficits are easier to remedy at early ages than at later ages. At the same time, σ_{π}^{q} , associated with personality, stays roughly constant over q. This is consistent with evidence on the emergence of psychological maturity, as shown in Figure 27.²²⁸

Adjoined with measurement systems for productivity on tasks in period v (Eq. (1.22)), the econometric model is a "state space" model that accounts for errors in measurements and endogeneity of inputs. Cunha and Heckman (2008) and Cunha, Heckman, and Schennach (2010) estimate these models on panel data on the growth dynamics of individuals and show that accounting for measurement error and endogeneity is empirically important.

Cunha, Heckman, and Schennach (2010) estimate technology (1.25) using longitudinal data on the development of children with rich measures of parental investment and of child traits. They examine the estimated substitution parameters to examine the issue of the cost of remediating early disadvantage at later ages. Their findings shed light on the dynamic process of capability formation in a way that raw correlations do not. They find that self-productivity becomes stronger as children become older, for both cognitive and noncognitive capability formation. The elasticity of substitution for cognitive inputs is *smaller* in second-stage production, so that it is more difficult to compensate for the effects of adverse environments on cognitive endowments at later ages than it is at earlier ages. This is consistent with the high rank stability of cognition over ages past 10–12.

This finding helps to explain the evidence on ineffective cognitive remediation strategies for disadvantaged adolescents documented in Cunha, Heckman, Lochner, and Masterov (2006); Knudsen, Heckman, Cameron, and Shonkoff (2006); and Cunha and Heckman (2007). Personality traits foster the development of cognition but not vice versa. It is equally easy to substitute at both stages for socioemotional skills over the life cycle. Overall, 16% of the variation in educational attainment is explained by adolescent cognitive traits, 12% is due to adolescent personality (socioemotional traits), and 15% is due to measured parental investments.

8.4.1 Evidence of Change in Traits from Other Studies of Parental Investment

Cunha, Heckman, Lochner, and Masterov (2006) summarize a large literature on child development. Evidence from a substantial literature suggests that for intelligence, the enduring effects of environment are greater earlier in life. Duyme, Dumaret, and Tomkiewicz (1999) studied children with IQs below 86, who were adopted between the ages of four and six into stable homes. As measured in their adolescent years, children adopted into high-SES homes gained an average of 19.5 IQ points; children

²²⁸ Cunha and Heckman (2008) estimate a linear version of the technology. Their specification rules out interaction and assumes that, over the feasible range, investment can perfectly substitute for skill deficits.

adopted into low-SES homes showed an average gain of 7.7 IQ points. In a study of Romanian children taken from impoverished orphanages and placed into middle-class British homes, the long-term salutary effects of adoption on cognitive ability were dramatic when infants were placed before they reached 6 months and markedly less so when adoption was delayed until later ages (Beckett et al., 2006). Notably, children adopted at different ages between six to 42 months did not differ at age 11 from each other in the terms of cognitive ability, with all children demonstrating an average deficit of 15 IQ points relative to children who had been adopted earlier in life. Low nutrition had no effect on cognitive outcomes at age 11, suggesting a prominent role for psychological deprivation. As Beckett and colleagues point out, these findings are consistent with the existence of a very early critical or sensitive period for intellectual development in which particular environmental stimuli are necessary for normative axonal rewiring (see Uylings, 2006, and Rutter, 2006b, for reviews).²²⁹

8.4.2 The Effects of Schooling on Cognitive and Personality Traits

Despite a large literature on the effects of schooling on shaping preferences (see Bowles and Gintis, 1976, and the literature it spawned), there is surprisingly little direct evidence on the effect of schooling on cognitive and personality traits. An exception is the analysis of Heckman, Stixrud, and Urzua (2006). The authors formulate and estimate an economic model that identifies the effect of cognitive and personality traits on schooling and a variety of other outcomes. The model controls for the effect of schooling in boosting both cognitive and personality measures and thus controls for reverse causality. They estimate their model on the National Longitudinal Survey of Youth 1979 (NLSY79), which has measures on the components of the Armed Services Vocational Battery (ASVAB) that are used to create the Armed Forces Qualifying Test (AFQT), a widely used measure of cognition. In addition, the NLSY79 has two measures of personality. The Rotter Locus of Control Scale, discussed in Section 5, is designed to capture the extent to which individuals believe that they have control over their lives through self-motivation or self-determination as opposed to the extent that the environment controls their lives (Rotter, 1966). The NLSY79 data also contain the Rosenberg Self-Esteem Scale, which attempts to assess the degree of approval or disapproval of oneself (Rosenberg, 1965). The relationship between these measures and the Big Five traits of Neuroticism is discussed in Section 5.

Different traits might be more responsive to investment at different ages. Figure 1.29 shows the causal effects of years of schooling attained on five components of the Armed Forces Qualifying Test (AFQT). Schooling in the high-school years has moderate but positive effects on the measures of cognition, consistent with previous research by

²²⁹ However, the data are also consistent with alternative explanations such as extreme stress permanently damaging brain structures.



Figure 1.29 Causal Effect of Schooling on ASVAB Measures of Cognition. (a) Arithmetic Reasoning. (b) Word Knowledge. (c) Paragraph Comprehension. (d) Math Knowledge. (e) Coding Speed.

Notes: Effect of schooling on components of the ASVAB. The first four components are averaged to create males with average ability. We standardize the test scores to have within-sample mean zero and variance one. The model is estimated using the NLSY79 sample. Solid lines depict average test scores, and dashed lines, 2.5–97.5% confidence intervals. Source: Heckman, Stixrud, and Urzua (2006, Figure 4).
Hansen, Heckman, and Mullen (2004); Neal and Johnson (1996); and Winship and Korenman (1997). The most dramatic causal effects on cognition arise from college attendance. Figure 1.30 shows the causal effects of years of schooling attained on locus of control and self esteem. In contrast, locus of control is primarily affected by high-school attendance but not college attendance. On measures of self-esteem, an additional year of high school and college play powerful roles.²³⁰

Some other evidence supports the possibility that school can affect measures of intelligence. Cahan and Cohen (1989) use a quasi-experimental paradigm comparing children who differ in both age and schooling to show that schooling increases intelligence test scores independently of age. Schooler and his colleagues show that complex (i.e., cognitively demanding) work increases intellectual functioning among adults and vice versa (Schooler, Mulatu, and Oates, 1999, and Kohn and Schooler, 1978).

8.4.3 Evidence from Interventions

As noted in the introduction, the Perry Preschool Program, did not have a lasting improvement on cognitive ability but did improve important later-life outcomes through personality (Heckman, Malofeeva, Pinto, and Savelyev, first draft 2008, revised 2011). The Perry Preschool Program enriched the lives of low-income black children with initial IQs below 85 at age 3. In addition, there were home visits to promote parent-child interactions. The program ended after 2 years of enrollment and both treatments and controls entered the same school. Participants were taught social skills in a "plan-do-review" sequence in which students planned a task, executed it, and then reviewed it with teachers and fellow students. They learned to work with others when problems arose.²³¹ The program was evaluated by the method of random assignment.

The program had strong effects for both boys and girls, although the effects differ by age and outcomes. The program had a statistically significant rate of return of around 6–10% per annum for both boys and girls. These returns are above the post–World War II, pre-2008 meltdown, stock market returns to equity in US labor market that are estimated to be 5.8% per annum.²³² The Perry Preschool Program worked primarily through socioemotional channels. Figure 1.31 shows that the program improved scores on the California Achievement Test (CAT). However, the program did not have a lasting effect on IQ scores. This evidence is consistent with the discussions in Sections 5 and 7 that show that achievement test results are strongly dependent on personality traits (see Borghans, Duckworth, Heckman, and ter Weel, 2008, and Borghans, Golsteyn, Heckman, and Meijers, 2009). Indeed the personalities of participants improved. Participants had better direct measures of personal behavior (a weighted average of "absences

²³⁰ See Hansen, Heckman, and Mullen (2004) for additional estimates of the causal effect of schooling on AFQT.

²³¹ Sylva (1997) describes the Perry Program as a Vygotskian program fostering personality traits.

²³² See DeLong and Magin (2009).





Notes: Effect of schooling on socioemotional scales for males with average ability, with 95% confidence bands. The locus of control scale is based on the four-item abbreviated version of the Rotter Internal-External Locus of Control Scale. This scale is designed to measure the extent to which individuals believe that they have control over their lives through self-motivation or self-determination (internal control) as opposed to the extent to which individuals believe that the environment controls their lives (external control). The self-esteem scale is based on the 10-item Rosenberg Self-Esteem Scale. This scale describes a degree of approval or disapproval toward oneself. In both cases, we standardize the test scores to have within-sample mean zero and variance one, after taking averages over the respective sets of scales. The model is estimated using the NLSY79 sample. Solid lines depict average test scores, and dashed lines, 2.5–97.5% confidence intervals.

Source: Heckman, Stixrud, and Urzua (2006, Figure 5).



Figure 1.31 Perry Age 14 Total CAT Scores, by Treatment Group. CAT, California Achievement Test. Treatment: N = 49; Control: N = 46. Statistically Significant Effect for Males and Females (*p*-values 0.009 and 0.021, respectively).

Source: Heckman, Malofeeva, Pinto, and Savelyev (first draft 2008, revised 2011).

and truancies," "lying and cheating," "stealing," and "swears or uses obscene words" measured by teachers in the elementary school years). Participants of both genders improved their externalizing behavior, which, as noted in Section 5, is related to Agreeableness and Conscientiousness. For girls, Openness to Experience (proxied by academic motivation) was also improved. Heckman, Malofeeva, Pinto, and Savelyev (first draft 2008, revised 2011) decompose the treatment effects of the Perry Program into components due to experientially induced changes of cognition (IQ), measures of personality, and residual factors. Improvements in personality led to beneficial changes in life outcomes.

Analyses of data from Project STAR, a program that randomly assigned kindergarteners and teachers to classes of different sizes, yields results similar to the Perry Program. Using data from Project STAR, Dee and West (2008) find that assignment to a small class is associated with positive changes in personality. In a follow-up reanalysis, Chetty et al. (2010) examine the Project STAR program and find that students placed in higher quality kindergarten classes—as measured by their peer's average performance on a Stanford Achievement Test—tend to have higher test scores at the end of kindergarten. The effect fades out over time; by eighth grade, students in better kindergarten classes perform no differently on tests. However, as with the Perry Program, the benefits re-emerge later in life. People in better kindergarten classrooms had significantly higher earnings in early adulthood. Furthermore, kindergarten classroom quality also predicted better fourth- and eighth-grade behavior as measured by teacher-assessed effort, initiative, interest in the class, and disruptive behavior.²³³ In turn, behavior predicted earnings in adulthood, suggesting that personality is the channel through which better kindergarten classrooms improve earnings.

The Perry Program and Project STAR did not primarily focus on improving personality traits, but a few programs did. The Promoting Alternative Thinking Strategies (PATHS) curriculum teaches self-control, emotional awareness, and social problemsolving skills and is aimed at elementary school children (see Bierman et al., 2010). A recent random-assignment, longitudinal study demonstrates that the PATHS curriculum reduces teacher and peer ratings of aggression, improves teacher and peer ratings of prosocial behavior, and improves teacher ratings of academic engagement.²³⁴ PATHS is an exemplar of school-based social and emotional learning (SEL) programs, whose impact on both course grades (d = 0.33), where d is measured in units of standard deviations ("effect sizes") and standardized achievement test scores (d = 0.27), was recently documented in a meta-analysis of controlled studies involving over 270,000 children in kindergarten through college (Durlak, Weissberg, Dymnicki, Taylor, and Schellinger, 2011).²³⁵ Similarly, a random assignment evaluation of Tools of the Mind, a preschool and early primary school curriculum, shows that in short-term follow-ups it improves classroom behavior as well as executive function, defined as higher level cognitive skills including inhibitory control, working memory, and cognitive flexibility (Barnett et al., 2008; Barnett, Yarosz, Thomas, and Hornbeck, 2006; Bodrova and Leong, 2001; Bodrova and Leong, 2007; Diamond, Barnett, Thomas, and Munro, 2007). Similar findings are reported for the Montessori preschool curriculum (Lillard and Else-Quest, 2006).

There is also evidence that targeted intervention efforts can improve aspects of Conscientiousness. These studies are typically more short term and, in contrast to the multifaceted curricula described above, are designed to isolate a particular mechanism for behavioral change. For instance, Rueda, Rothbart, McCandliss, Saccomanno, and Posner (2005) designed a set of computer exercises to train attention in children between 4 and 6 years of age. Children in the intervention group improved in performance on computer tasks of attention relative to children who instead watched interactive videos for a comparable amount of time. Similarly, Stevens, Fanning, Coch, Sanders, and Neville (2008) designed a 6-week computerized intervention and showed that it can improve selective auditory attention (i.e., the ability to attend to a target auditory signal in the face of an irrelevant, distracting auditory signal). Again, all of these programs have short-term follow-ups.

 ²³³ These scales are based on more-detailed questionnaires. Only a subset of the sample has their behavioral measures.
²³⁴ Bierman et al. (2010)

²³⁵ Note, however, that the largest federal study to date on character education programs, including PATHS, failed to find evidence for improvements in behavior or academic performance (see Social and Character Development Research Consortium, 2010).

Several studies suggest that personality can be remediated in adolescence. Martins (2010) analyzes data from EPSIS, a program developed to improve student achievement of 13- to 15-year olds in Portugal by increasing motivation, self-esteem, and study skills. The program consists of one-on-one meetings with a trained staff member or meetings in small groups. The intervention was tailored to each participant's individual skill deficit. Overall, the program was successful and cost-effective, decreasing grade retention by 10% points. Bloom, Gardenhire-Crooks, and Mandsager (2009) analyze the data from the National Guard Youth ChalleNGe program, a 17-month intervention for youth who have dropped out of high school. Although the program does not require enrollment in the military, it stresses aspects of military discipline. The program features a 2-week assessment period, a 20-week residential program often conducted at a military base, and a 1-year mentoring program. Nine months after entry, participants in the program were 12% more likely to obtain a high-school diploma or GED, 9% more likely to be working full time, and less likely to be arrested. Furthermore, participants had higher levels of self-efficacy (a trait related to Emotional Stability), suggesting that personality change might have helped with the improvements. However, the 9-month follow-up period is too short to know if the program has long-lasting effects. Although these studies show that adolescent personality can be improved through intervention, a couple of other studies show less-promising results (Rodríguez-Planas, 2010, and Holmlund and Silva, 2009).

Behncke (2009) provides some experimental evidence that short-term exogenous shocks to noncognitive skills affect test performance. She finds that giving words of encouragement before a diagnostic math test (an intervention that might boost short-term self-efficacy or self-esteem), was associated with 2.5% higher scores among all students (p < 0.05) and 8% higher scores among those with self-reported difficulties with math (p < 0.01). The result suggests that noncognitive skills can be shaped, at least in the very short term.

The evidence for adults corroborates the finding of Cunha, Heckman, and Schennach (2010) for children. Personality is malleable throughout the life cycle. For example, Gottschalk (2005) shows evidence from a randomized control trial that working at a job can improve locus of control. He uses data from the Self-Sufficiency Project (SSP) in which some welfare recipients were randomly offered substantial subsidies to work. The subsidy more than doubled the earnings of a minimum wage worker, and people in the experimental group worked about 1/3 more hours than those in the control group. After 36 months, those who received the subsidy were more likely to have an improved locus of control.

Several other studies find similar results. Clausen and Gilens (1990) claim that female labor force participation increases self-confidence. Roberts (1997) reports an increase of social dominance, and Roberts and Chapman (2000) report a decrease in Neuroticism for working women. Others show that marital and family experiences shape personality (Helson and Picano, 1990, and Roberts, Helson, and Klohnen, 2002). However, these studies are all correlational in nature. None of these studies have the random assignment features of the Gottschalk study.

Personality may even be malleable at the end of life. Jackson, Hill, Payne, Roberts, and Stine-Morrow (2010) investigate causal mechanisms behind the association between Openness to Experience and IQ, using data from a 16-week intervention designed to boost inductive reasoning for elderly people. The intervention consisted of laboratory training for how to recognize novel patterns and around 10 hours a week of solving crossword, Sudoku, and logic puzzles. Controlling for inductive reasoning, self-reported Openness to Experience increased for participants during the training program relative to those in a wait-listed control group. However, the elderly people were not followed after the program to determine whether the change was long lasting or whether important outcomes, like life expectancy, improved.

Table 1.12 summarizes the evidence on the effects of interventions that is discussed in this subsection. The evidence is consistent with effects of interventions, but there are woefully few causal studies with long-term follow-up.

8.4.4 Evidence from Psychotherapy

The accomplishments of psychotherapy also support the possibility of intentional, mean-level, and rank-order change. In a 1980 meta-analysis, Smith, Glass, and Miller (1980) summarized 475 controlled studies, concluding that individuals who undergo psychotherapy are about 0.85 standard deviations better on outcome measures than those who do not. The large benefits of therapy are not permanent, however: the effect of psychotherapy over control conditions falls to about half a standard deviation 2 years after therapy is concluded. Moreover, it is not clear that the effects of psychotherapy on individuals who seek change generalize to individuals who are not actively seeking treatment for a condition that causes them distress.²³⁶

8.5. Stability of Economic Preference Parameters

Less is known about the stability of economic preferences. To our knowledge, no longitudinal study has measured the mean-level or rank-order stability of time preference over the life cycle (Frederick, Loewenstein, and O'Donoghue, 2002). A handful of cross-sectional studies using relatively small samples have examined mean-level stability, and their findings are mixed. Green, Fry, and Myerson (1994) and Harrison,

²³⁶ Some evidence that further intervention can produce enduring change in nonclinical populations comes from Gillham and Reivich (1999) who show that children taught to make more optimistic causal attributions about negative events maintain this optimistic outlook 2 years postintervention.

Author(s)	Main Variable(s)	Data and Methods	Causal Evidence	Main Result(s)
Barnett et al. (2008)	Outcome(s): <i>internalizing</i> <i>and externalizing behavior</i> — teacher-assessed Problem Behaviors Scale of the Social Skills Rating System (SSRS) Intervention: participation in a year-long Tools of the Mind preschool program compared to a generic curriculum.	Data: collected by authors; 210 children aged 3 and 4 Methods: Students were randomly assigned to classrooms within the same school after parental consent was obtained. Teachers were randomly assigned to control and treatment classrooms.	Control Variables: n/a. Timing of Measurements: <i>Baseline</i> —Behavior measures were taken prior to the program in October–November of 2002. <i>Posttreatment</i> —Behavioral measures taken immediately after the program in May–June of 2003.	Participants in the program had a 0.47 standard deviation lower score for the behavioral problems index (p < 0.05).
Behncke (2009)	Outcome(s): <i>cognitive</i> <i>ability</i> —performance on a diagnostic math test for a college economics class. Intervention: verbal encouragement before the test.	Data: collected by author; 440 students from a Swiss University Methods: The treatment was randomly assigned to already-established classroom sections. Students were unaware they were in an experiment.	Control Variables: n/a Timing of Measurements: <i>Post treatment</i> : The diagnostic test was given immediately after the treatment.	Verbal encouragement raised test scores by 2.5% among all students (p < 0.05) and by 8% among students who reported difficulties with math $(p < 0.01)$.

Table 1.12 The Effect of Interventions on Personality

Author(s)	Main Variable(s)	Data and Methods	Causal Evidence	Main Result(s)
Bierman et al. (2010)	Outcome(s): teacher- assessed behavior—Social Health Profile (SHP) including authority acceptance, cognitive concentration, and social competence; peer-assessed behavior—survey questions about behavior labeled as aggressive, prosocial, and hyperactive. Intervention: participation in a 3-year- long Fast Track PATHS program focused on improving self-control and positive social behavior.	Data: 2,937 children (grades 1–3) Methods: School administrators were offered participation in the experiment, knowing the school would receive treatment with a 50% probability.	Control Variables: time, time squared, individual baseline, school baseline, city fixed effects, poverty level, interactions of intervention with time, time squared, individual baseline, poverty, and poverty and time. Timing of Measurements: <i>Baseline</i> —Behavioral measures were taken prior to the program in the fall of first grade. <i>Posttreatment</i> —Behavioral measures were taken again in the spring of third grade around the end of the program.	Immediately after the 3-year program, participation was associated with a 0.24 standard deviation increase in authority acceptance ($p < 0.001$), a 0.12 standard deviation increase in cognitive concentration ($p < 0.001$), and a 0.34 standard deviation increase in social competence ($p < 0.0001$) compared with the control group. The effects were stronger in more- disadvantaged schools. Similar but weaker results apply for the peer-assessed measures.
Bloom, Gardenhire- Crooks, and Mandsager (2009)	Outcome(s): <i>educational</i> <i>attainment</i> —high-school diploma; <i>labor force</i> <i>participation</i> —whether working at a job; <i>personality</i> —self-efficacy and social adjustment Intervention: Participation in the ChalleNGe program consisting of a 2-week assessment period, 20- week residential program often conducted at a military base, and a 1-year mentoring program.	Data: 1,018 young people between the ages 16 and 18, who have dropped out of school. Methods: The control group was constructed out of applicants who qualified for the program but were not taken due to lack of space.	Control Variables: sample member characteristics Timing of Measurements: <i>During treatment</i> — Outcomes were measured approximately 9 months after entering the study.	Participants in the program were 12.0 percentage points more likely to earn a high- school diploma (p < 0.01), 9.1 percentage points more likely to be working $(p < 0.01)$, and 9.6 percentage points less likely to report a self- efficacy and social adjustment score one standard deviation below the mean $(p < 0.01)$. The program also improved measures of criminality and health.

Table 1.12 The Effect of Interventions on Personality—continued

Author(s)	Main Variable(s)	Data and Methods	Causal Evidence	Main Result(s)
Chetty et al. (2010)	Outcome(s): <i>noncognitive</i> <i>skills</i> —an index based on the teacher's observations of the students Intervention: randomly assigned kindergarten class quality as measured by difference in percentiles of the mean end-of-year test scores of the students' classmates and the scores of the other kindergarteners at the same school.	Data: Project STAR; 1,671 fourth-grade students and 1,780 eighth-grade students Methods: Students and teachers were randomly assigned in kindergarten to classrooms of different sizes. The students were assigned to the same size classroom through third grade.	Control Variables: wave fixed effects, student gender, free-lunch status, age, race, a quartic in the parent's household income interacted with parent's marital status, mother's age at child's birth, whether the parents own a home and whether the parents made a 401(k) contribution between 1996 and 2008. Timing of Measurements: <i>During treatment</i> — Age-relevant SAT tests were administered in kindergarten and grades 1–3. <i>Posttreatment</i> —Age- relevant SAT tests and behavioral surveys were given in fourth and eighth grades. College quality and attendance was measured at age 19. Earnings were measured at age 27.	A 1-percentile improvement in kindergarten class quality increases an index of noncognitive skills by 0.15 percentiles in fourth grade ($p < 0.05$) and 0.13 percentiles in eighth grade ($p < 0.05$). Better classrooms were also associated with better life outcomes.

Table 1.12 The Effect of Interventions on Personality-continued

Author(s)	Main Variable(s)	Data and Methods	Causal Evidence	Main Result(s)
Diamond, Barnett, Thomas, and Munro (2007)	Outcome(s): <i>Executive</i> <i>Function</i> —Dots-Mixed task, Reverse-Flanker task. Intervention: participation in a Tools of the Mind program instead of the regular school district's balanced literacy program.	Data: 147 preschoolers Methods: Teachers and students were randomly assigned to classrooms within the same school.	Control Variables: age, gender, years in program Timing of Measurements: <i>Posttreatment</i> —The tasks were given at the end of the second year of the program.	84% of students in Tools were successful in the Reverse Flanker task compared with 65% in the control group. Almost twice as many students in the Tools program achieved greater than 75% accuracy on the Dots-Mixed task compared with the control group.
Durlak, Weissberg, Dymnicki, Taylor, and Schellinger (2011)	Outcome(s): social and emotional learning skills, attitudes, positive social behavior, conduct problems, emotional distress, academic performance. Intervention: Meta- analysis of school-based, universal social and emotional learning program.	Data: 270,034 kindergarten through high-school students Methods: All studies include a control group.	Control Variables: n/a Timing of Measurements: All studies contained follow-up data at least 6 months after the intervention.	The mean difference in standard deviations between the treatment and control groups are as follows: social and emotional learning skills = $0.57 \ (p < 0.05)$; attitudes $= 0.23 \ (p < 0.05)$; positive social behavior $= 0.24$ (p < 0.05); conduct problems $= 0.22 \ (p < 0.05)$; emotional distress $=$ $0.24 \ (p < 0.05)$; academic performance $= 0.27$ (p < 0.05). All variables are coded so that positive numbers reflect better outcomes.

Table 1.12 The Effect of Interventions on Personality—continued

Table 1.12 The Effect of Interventions on Personality—cont	inued
--	-------

Author(s)	Main Variable(s)	Data and Methods	Causal Evidence	Main Result(s)
Gottschalk (2005)	Outcome(s): <i>Personality</i> — four measures of locus of control based on whether the respondent agrees strongly, agrees, disagrees, or strongly disagrees with statements Intervention: A subsidy for full-time work during a 36-month period.	Data: Self-Sufficiency Project; 4,958 single parents over the age of 19 in New Brunswick and British Columbia Methods: The subsidy was randomly offered to a population of people receiving Income Assistance (IA).	Control Variables: age, age squared, region, gender, speaks French, number of children Timing of Measurements: <i>Baseline</i> —Locus of control was measured before the program. <i>During treatment</i> —Locus of control was measured again 18 and 36 months after the baseline.	Using whether the participant received the subsidy as an instrument for hours worked, the authors find that working tends to improve locus of control by the 36-month re-interview.
Heckman, Malofeeva, Pinto, and Savelyev (first draft 2008, revised 2011)	Outcome(s): <i>externalizing</i> <i>behavior, internalizing</i> <i>behavior</i> —measured using Pupil Behavior Inventory (PBI) of teacher reports Intervention: participation in the Perry Preschool Program, an intervention that lasted 2 years and enriched the lives of low-income black children.	Data: Perry Preschool Program; 123 preschool students Methods: The students were randomly assigned to treatment.	Control Variables: n/a Timing of Measurements: <i>Post treatment</i> —The measure of externalizing and internalizing behavior are taken ages 7–9 (2–4 years after treatment). Other life outcomes were measured at ages 19, 27, and 40.	The intervention improved mean externalizing behavior for both males and females (p < 0.05). It improved Openness to Experience, as measured by academic motivation, for females (p < 0.10) but not for males. The program also generated a wide range of later-life outcomes primarily through improvements in noncognitive skills.

Author(s)	Main Variable(s)	Data and Methods	Causal Evidence	Main Result(s)
Holmlund and Silva (2009)	Outcome(s): <i>academic</i> <i>petformance</i> —average of standardized test scores in English, Math, and Science. Intervention: participation in the "xl programme" targeting the noncognitive skills of secondary school students aged 14.	Data: "xl club programme," National Pupil Database (NPD), Pupil Level Annual Schools Census (PLASC); 2,333 and 259,189 treated and control students aged 14 in England (2004) Methods: logit, propensity score matching, OLS, difference-in-difference, double differences, random-growth model.	Control Variables: sex, language, eligibility for school meals, special needs status, and race Timing of Measurements: <i>Baseline</i> —Standardized exams were taken at age 11 and age 14 before the start of the program. <i>Posttreatment</i> — Standardized national exams were taken again at age 16 at the end of the program (2 years after the beginning of the program).	Unconditional on observables, the performance of the students in the xl club is 1.2-1.4 standard deviations lower than the control subjects ($p <$ 0.01). Using OLS, the effect is -0.17 . The propensity score estimates are -0.13 and -0.15 . For the difference-in- difference models estimated using OLS and propensity score matching, there is no longer a significant effect of the program in either direction. Overall the program had little effect.
Social and Character Development Research Consortium (2010)	Outcome(s): Social and Emotional Competence— self-efficacy for peer interaction, normative beliefs about aggression, empathy; Behavior— altruistic behavior, positive social behavior, problem behavior, ADHD-related behavior; Academics—engagement with learning, academic competence and	Data: Social and Character Development (SACD) Research Program; around 6,000 elementary school students Methods: Schools were first asked to participate in the program and were then randomly assigned one of the seven SACD programs or left with their traditional	Control Variables: gender, race, parental education, family structure, household income, measures of poverty, parental labor force participation, teacher race, teacher experience. (Note: the specific set depended on the outcome of interest.) Timing of Measurements: <i>Baseline</i> —Initial measures	Fall 2003 to Spring 2005: Of the 20 outcomes, the only significant effects were that participation in any program was associated with a 0.07 standard deviation higher primary caregiver- reported altruistic behavior ($p < 0.10$), a 0.06 standard deviation lower child-reported altruistic behavior

Table 1.12 The Effect of Interventions on Personality-continued

Author(s)	Main Variable(s)	Data and Methods	Causal Evidence	Main Result(s)
	motivation; <i>Perceptions of</i> <i>School Climate</i> —positive school orientation, negative school orientation, student afraid at school, victimization at school, feelings of safety, student support for teachers. Intervention: seven different programs (ABC, CSP, LBW, PA, PATHS, 4Rs, SS) aimed to build Social and Character Development (SACD) compared to the "standard practice" programs at nontreated schools.	curriculum. The data were analyzed using hierarchical linear modeling (HLM).	were collected near the start of the program in the fall of 2004. <i>During treatment</i> —Data were collected in the spring of 2005, the fall of 2005, and the spring of 2006. <i>Posttreatment</i> —Data were collected near the end of the program in the spring of 2007.	(p < 0.10), and a 0.12 standard deviation higher teacher-reported student support for teachers (p < 0.05). Fall 2003 to Spring 2006: Of the 20 outcomes, the only significant effects were that participation in any program was associated with a 0.07 standard deviation lower child-reported self- efficacy for peer interactions $(p < 0.10)$ and 0.16 standard deviation higher teacher- reported student support for teachers $(p < 0.05)$. Fall 2003 to Spring 2007: There were no statistically different effects of participating in any program. Other Analyses: The results were similar when analyzing each of the programs separately and when using growth curves. There is some evidence that programs were beneficial for high- risk students.

Table 1.12 The Effect of Interventions on Personality-continued

Author(s)	Main Variable(s)	Data and Methods	Causal Evidence	Main Result(s)
Jackson, Hill, Payne, Roberts, and Stine- Morrow (2010)	Outcome(s): <i>Personality</i> — Openness to Experience Intervention: participation in a 16- week inductive reasoning training program coupled with 10 hours of puzzle solving per week.	Data: collected by the authors; 183 adults aged 60–94. Methods: Participants were randomly assigned to treatment and control groups after deciding to participate in the experiment.	Control Variables: n/a Timing of Measurements: Baseline—Openness to Experience was measured pre-treatment. During treatment— Openness to Experience was measured at week 5 and at week 10. Post treatment—Openness to Experience was measured at the end of the program in week 16.	On average, participants in the program were 0.39 standard deviations higher in Openness to Experience after the program relative to people in the control group ($p < 0.05$).
Martins (2010)	Outcome(s): <i>Educational</i> <i>attainment</i> —grade retention Intervention: participation in the EPIS program that boosts noncognitive skills including motivation, self-esteem, and study skills.	Data: EPIS database; 15,307 students of grade 7–9 in Portugal Methods: linear probability model, quasi- randomization.	Control Variables: student-fixed effects, time-fixed effects Timing of Measurements: <i>Baseline</i> —Measures of academic achievement were taken before the intervention in grades 7 and 8. <i>During treatment</i> — Measures were taken each quarter that the students participate in the program through seven academic quarters after the beginning of the program (students entered the program at different times and remained in treatment for different lengths of time but were followed if they left treatment).	The program reduced annual grade retention by at least 10.1 percentage points (<i>p</i> < 0.001).

Table 1	1.12	The Effect	of In	terventions	on	Personality—cc	ntinued
i abic i		THE ENCE	O 1 111	cerverneroris		reisonancy ce	maca

148

Author(s)	Main Variable(s)	Data and Methods	Causal Evidence	Main Result(s)
Rodríguez-Planas (2010)	Outcome(s): educational attainment—high-school completion and post- secondary education; academic achievement— math test score percentile, reading test score percentile, GPA; labor market success—earnings during the last year of the program, 3 years after the program, and five years after the program. Intervention: participation in the Quantum Opportunity Program (QOP) that was available for 5 years, centered around mentoring, developing social skills, community service, and providing incentives for academic success for ninth graders.	Data: Quantum Opportunity Program (QOP); 1,069 students from seven large US cities Methods: Students in schools participating in the program were randomly assigned to treatment or control groups.	Control Variables: n/a Timing of Measurements: <i>Post treatment</i> —Interviews were conducted during the last year of the program, 3 years after the program, and 5 years after the program.	During last year of the program: Participation in the program was associated with a 7–percentage-point increase in the probability of graduating high school ($p < 0.10$) and 6–percentage-point increase in the probability of attending college ($p < 0.10$). There were no differences in academic achievement. 3 years after the program: Participation in the program was associated with a 7–percentage-point increase in the probability of ever attending college ($p < 0.10$), 9–percentage- point increase in the probability of attending college ($p < 0.05$), and a 7–percentage-point decrease in the probability of having a job ($p < 0.10$). 5 years after the program: There are no significant differences 5 years after the program. Findings for subpopulations: The program benefited people who were 14 or less upon entering high school significantly more than older students. It also tended to benefit girls more than boys.

Author(s)	Main Variable(s)	Data and Methods	Causal Evidence	Main Result(s)
Stevens, Fanning, Coch, Sanders, and Neville (2008)	Outcome(s): <i>attention</i> — ERP index of selective auditory attention; <i>language skills</i> —Clinical Evaluation of Language Fundamentals-3. Intervention: Participation in a six- week (100 min/day) computerized training program for boosting language skills (Fast ForWord program).	Data: collected by the authors; 33 children aged 7 on average. Methods: The students who received treatment were compared to a control group who did not.	Control Variables: Test scores were normalized by age. Timing of Measurements: <i>Baseline</i> —Measures were taken right before the start of the program. <i>Post treatment</i> —Measures were taken again at the end of the program (6 weeks after the start).	The increase in the attention score was 0.81 standard deviations higher for the participants than for the nonparticipants $(p < 0.01)$. The increase in the receptive language scores was 0.91 standard deviations higher in the participants than for the control group $(p < 0.01)$. There was no significant effect on expressive language scores between the participants and the control group.

Table 1.12 The Effect of Interventions on Personality—continued

Lau, and Williams (2002) find that discount rates are lower among older individuals. On the other hand, Chesson and Viscusi (2000) claim to find that older adults have *higher* discount rates than younger adults. Chao, Szrek, Sousa Pereira, and Pauly (2007), de Wit, Flory, Acheson, McCloskey, and Manuck (2007), and Coller and Williams (1999) find no relationship between age and discount rate. Finally, Read and Read (2004) find a curvilinear relationship in which older people discount more than younger people, and middle-aged people discount less than either group. Sahm (2007) shows that risk aversion increases with age. Table 1.13 below summarizes (but not exhaustively) the findings from a variety of recent economic studies on the heritability, malleability, and stability of preferences and personality.

8.6. Summary of Section 8

We have reviewed the evidence on change in personality over the life cycle. The evidence is strong that personality changes over the life cycle, both in terms of mean-level and rank-order change. The evidence on the source of the change is less clear-cut. Three competing visions of the source of change are discussed: (1) The ontogenic and sociogenic model that describes how biology and socialization produce changes in average traits. This approach does not explain why individuals develop with different trajectories; (2) The biological and pharmacological model that describes how alterations in the biology of the person can explain variations in personality and its evolution; and (3) The intervention/family influence model that describes how investment and environments generate changes. No study considers all three sources of development at the same time, largely due to data limitations. The evidence from the intervention and family influence studies suggests that interventions that target personality may be effective but much further evidence is required to specify the exact mechanisms through which the interventions work.

9. SUMMARY AND CONCLUSIONS

We summarize this chapter by providing provisional answers to the eight questions posed in Section 1.

1. How can we fit psychological constructs of personality into an economic framework? Can conventional models of preferences in economics characterize the main theories in personality psychology?

We have defined personality as a response function of agents that depends on situations (including incentives), endowments of traits, information, and resources within a conventional economic model. Psychologists analyze a richer class of actions than economists normally consider. We show how to integrate these actions into

Table 1.13 The Heritability, Malleability, and Stability of Preferences and Personality

Author(s)	Main Variable(s)	Data and Methods	Causal Evidence	Main Result(s)
Booth and Nolen (2009)	Outcome(s): <i>risk aversion</i> - choice whether to accept a real-stakes lottery versus a certain payment. Explanatory Variable(s): <i>short-term gender</i> <i>environment</i> -whether the student was assigned to a co-ed or single-sex group during the experiment; <i>long-term gender</i> <i>environment</i> -whether the student attends a co-ed or single-sex school.	Data: Collected by the authors; 260 students in grades 10 and 11 from eight publicly funded schools in England (2007) Methods: probit.	Controls: n/a Timing of Measurements: the measures are contemporaneous. Theory: growing up in an environment with males might cause girls to act more "feminine" and take fewer risks. Similarly, boys in co-ed environments might exhibit more risk-taking in co-ed environments to try to impress girls.	Girls from co-ed high schools in England were 36% ($p < 0.01$) less likely to accept a real-stakes lottery. Girls assigned to experimental group with all girls were 12% ($p < 0.10$) more likely to accept the lottery than girls in co-ed experimental groups.

Author(s)	Main Variable(s)	Data and Methods	Causal Evidence	Main Result(s)
Burks, Carpenter, Goette, and Rustichini (2010)	Outcome(s): demand for information—whether people request the results of their IQ and numeracy tests; overconfidence—the difference between ex- ante estimate of quintile in the IQ distribution and the true quintile in the IQ distribution. Explanatory Variable(s): self-assessment—before and after test assessments of the quintile of performance on the tests; personality—a self- reported measure of personality traits, social closeness, social potency, and stress reaction.	Data: Collected by authors, administrative data from a human resources department; 1,063 trainee truckers from a US trucking company. Methods: probit, ordered probit, linear spline.	Controls: actual test performance, harm avoidance, education levels, ethnicity, sex, age, age squared, household income, before-test belief, and post-test belief Timing of Measurements: People are asked about their expected performance on the IQ and numeracy tests before and after they take the test. Later, they are given the option to receive the results of their tests. Theory: People misjudge their own ability due to systematically biased noisy signals; value their self-assessed ability and avoid updating when the assessment is positive; or subconsciously misrepresent their own ability for strategic advantage.	Demand for information: A one-quintile increase in a person's post-test belief about their test performance is positively associated with 3.0 percentage point higher probability of demanding information about the IQ test ($p < 0.01$) and a 3.9 percentage point higher probability of demanding information about the numeracy test ($p < 0.01$). Overconfidence: Harm avoidance and stress reaction are negatively correlated with overconfidence on the IQ test ($p < 0.01$, p < 0.05). Social potency is positively linked to overconfidence on the IQ test ($p < 0.01$). Stress reaction is negatively associated with overconfidence on the numeracy test ($p < 0.01$). Social potency is positively associated with overconfidence on the numeracy test ($p < 0.05$).

Table 1.13 The Heritability, Malleability, and Stability of Preferences and Personality—continued

Author(s)	Main Variable(s)	Data and Methods	Causal Evidence	Main Result(s)
Dohmen et al. (2011)	Outcome(s): <i>risk preference</i> — survey responses on an 11-point scale, relating to general risk preference and risk preference relating to car driving, financial matters, leisure and sports, career and health. Explanatory Variable(s): (see controls).	Data: Collected by the authors/German Socio- Economic Panel (GSOEP); 450 adults from Germany/22,019 people living in Germany Methods: interval regression, probit.	Controls: sex, age, height, parental education, 2002 household wealth, 2003 household income Timing of Measurements: The measures are contemporaneous. Theory: People might have a stable, underlying preference for risk cross contexts.	Determinants of risk attitude: being female and age are negatively associated with willingness to take risks (p < 0.01). Height is positively associated with a general willingness to take risks $(p < 0.01)$. Mother and father's education is positively associated with willingness to take risks $(p < 0.01)$. Stability of risk: The 6 measures of contextualized risk aversion are correlated with each other, ranging from 0.456–0.609.
Einav, Finkelstein, Pascu, and Cullen (2010)	Outcome(s): <i>risk</i> <i>preference</i> —order rankings of observed decisions to insurance for health, prescription drugs, dental, short-term and long-term disability and 401(k) plans. Explanatory Variable(s): <i>predictable risk</i> — predictions from modeling risk based on observables; <i>idiosyncratic</i> <i>risk</i> —realization of risk in the next period.	Data: Administrative data; 12,752 employees of Alcoa, Inc. (2004) Methods: Spearman correlations, OLS.	Controls: the menu of benefits the employee faced, predictable and idiosyncratic risk Timing of Measurements: Most of the financial decisions were made in the same year. Theory: There is an underlying preference for risk that applies across many contexts.	The average correlations between the various domains are 0.164. Empirical results hold up under several robustness checks.

Table 1.13 The Heritability, Malleability, and Stability of Preferences and Personality—continued

Author(s)	Main Variable(s)	Data and Methods	Causal Evidence	Main Result(s)
Kosfeld, Heinrichs, Zak, and Fehr (2005)	Outcome(s): <i>trust</i> — willingness to "invest" in a real-stakes two-player trust game by the investee (the first player); <i>risk</i> <i>preference</i> —real-stakes trust game played against a computer that randomly gives payoffs; <i>altruism</i> —the amount transferred back by the investee in the trust game. Explanatory Variable(s): <i>biological determinant</i> <i>of trust</i> —nasal spray of oxytocin	Data: Experiment conducted by the authors; 194 male university students in Germany. Methods: Mann- Whitney U-test, RCT.	Controls: RCT Timing of Measurements: The measures were contemporaneous. Theory: There is a notion of "trust" distinct from altruism and risk preference.	People who receive the oxytocin nasal spray invest on average 17% more than those who do not ($p < 0.05$). Risk behavior does not differ between the two groups. Trustees (the second players in the trust game) do not show more altruistic behavior when given oxytocin.
Le, Miller, Slutske, and Martin (2010)	Outcome(s): <i>risk preference</i> — response to a 10-point survey question about willingness to take risks in general, response to a 10-point survey questions about how conservative the subject is in making decisions to spend money. Explanatory Variable(s): <i>genetic makeup</i> — differences in outcomes between monozygotic and dyzygotic twins.	Data: Australian Twin Study of Gambling; 1,875 complete twin pairs. Methods: OLS.	Controls: gender, age, education, and marital status. Timing of Measurements: The measures are contemporaneous.	Heritability of the risk measure is 0.192 ($p < 0.01$). Heritability of the conservative measure is 0.134 ($p < 0.01$).

Table 1.13 The Heritability, Malleability, and Stability of Preferences and Personality-continued

Author(s)	Main Variable(s)	Data and Methods	Causal Evidence	Main Result(s)
Sutter et al. (2010)	Outcome(s): social preferences (selfish—the agent maximizes their own payoff, regardless of the other person's; efficient—maximizing the sum of the payoffs; maximin—maximizes the minimum of the two payoffs; FS inequality— values own payoff plus a weighted average of the difference between own payoff to the others' payoffs; ERC inequality— disutility if payoff deviates from the group average)— choices of allocating resources between peers in a real-stakes experiment. Explanatory Variable(s): n/a.	Data: Collected by the authors; 883 students aged 8–17 living in Australia (2008). Methods: maximum likelihood error-rate analysis.	Controls: n/a Timing of Measurements: The measures are contemporaneous. Theory: Social preferences might change with age and maturity.	20% of girls and boys behave selfishly. An increase in one year of age is associated with a 0.044 increase in the probability of having efficiency preferences for males ($p < 0.05$), but has no effect for females.

Table 1.13 The Heritability, Malleability, and Stability of Preferences and Personality-continued

economic theory. The leading models of personality psychology are special cases of our model.²³⁷

2. What are the main measurement systems used in psychology for representing personality and personality traits, and how are they validated? How are different systems related to each other? What is the relationship between standard measures of personality and measures of psychopathology and child temperament?

In Section 5, we exposit the main systems for measuring personality, focusing primarily, but not exclusively, on the Big Five model. We consider the strengths and limitations of the systems and the relationships among competing systems. We show how measures of psychopathology are extreme manifestations of personality traits and how child temperament is related to adult traits. We link specific diagnoses of pathology with conventional measures of psychological traits.

3. What is the relationship between economic preference parameters and psychological measurements?

We review an emerging body of research that relates economic preference parameters (risk aversion, time preference, ambiguity aversion, social preferences) to the Big Five traits and to measures of self-esteem and personal control that are linked to the Big Five traits. Time preference is negatively correlated with IQ and the ability to control attention. Risk preference is negatively correlated with IQ and other measures of cognition. Higher-IQ people are more consistent in their choices under uncertainty. Although risk aversion is related to personality traits, the available evidence suggests that marginal ambiguity aversion is not. Social preferences are predicted by measured personality traits, but the evidence on this question is not strong.

4. How stable across situations and over the life cycle are preference parameters and personality traits?

We review the history of the person-situation debate between the social psychologists who maintain the primacy of the situation in determining behavior and the traits theorists who maintain the primacy of traits in explaining behavior. Behavioral economists, as a group, have adopted the situationist point of view. Extreme advocates of the situationist point of view claim that there is no stable personality construct. The issue hinges on the nonlinearity of action, effort, and productivity functions. In the presence of such nonlinearities, measured traits (e.g., actions) depend on situations and tasks.

A large body of evidence suggests that nonlinearity is an empirically important phenomenon. Nonetheless, evidence suggests that there are stable personality traits

²³⁷ Freudian models of the unconscious would make the traits that govern behavior, and especially ψ , unknown to agents but nonetheless governing choices. A pure model of behaviorism would feature the effects of constraints on choices. Borghans, Duckworth, Heckman, and ter Weel (2008) develop such a model. We review it in the Web Appendix.

that predict a variety of behaviors in different situations. Personality is neither an ephemeral creation of situations nor is its manifestation invariant across situations. Moreover, personality traits are not set in stone. They change over the life cycle. The evidence on the stability of preference parameters across situations and over the life cycle is less ample. There is evidence that standard separable models of preferences are inadequate descriptions of choice behavior. There is little evidence on the stability of preference parameters over the life cycle.

5. What is the evidence on the predictive power of cognitive and personality traits?

We present a large body of evidence that shows strong associations between personality traits and educational, labor market, health, and criminal outcomes.

6. What is the evidence on the causal power of personality on behavioral outcomes?

Few of the correlational studies relating personality to outcomes have a firm causal basis. Personality psychologists have not yet attempted to establish the causal status of personality. However, there are a few experimental manipulations that establish the causal effect of personality. Recent studies in economics establish causal status of certain personality traits on outcomes for observational studies invoking assumptions that are inevitably subject to debate. Research in this area is likely to flourish in the coming years.

7. Can personality be altered across the life cycle? Are interventions that change personality traits likely fruitful avenues for policy?

There is a small but growing body of intervention studies that establish that personality traits can be altered over long periods of time in response to interventions. Some of the major effects of early childhood intervention programs appear to operate through their lasting effects on personality. Family investment decisions also change personality. The evidence to date suggests that interventions that boost personality traits can be effective in promoting adult success.

8. Do the findings from psychology suggest that conventional economic theory should be enriched? The evidence from psychology enriches economics by providing a more nuanced interpretation of human choice and actions. It promises to provide a deeper understanding of conventional economic preference parameters and how they arise. Unfortunately, at the time of this writing, this promise remains unfulfilled. Given the current state of evidence against conventional economic preference specifications (see, e.g., Starmer (2000) and the evidence in Section 6), this line of research is very promising.

While personality psychology can enrich economics, the flow will likely be both ways. Economists have recently supplemented their traditional menu of preference parameters to account for a richer array of choices and actions. Personality psychologists have begun to use these new parameters. See Ferguson, Heckman, and Corr (2011). As personality psychologists shift their emphasis from the task

of description to study issues related to causality, policy, and prediction, the Big Five may be replaced.

Also, economists and personality psychologists might eventually derive both the psychological traits and the economic preference parameters from a deeper set of motivation-oriented parameters. Conventional psychological traits and preference parameters may be manifestations of as yet undiscovered parameters rooted in human biology.

REFERENCES

- Abbring, J.H., Heckman, J.J., 2007. Econometric Evaluation of Social Programs, Part III: Distributional Treatment Effects, Dynamic Treatment Effects, Dynamic Discrete Choice, and General Equilibrium Policy Evaluation. In: Heckman, J., Leamer, E., (Ed.), Handbook of Econometrics. Elsevier, Amsterdam, 6B, 5145–5303.
- Achenbach, T.M., McConaughy, S.H., Howell, C.T., 1987. Child/Adolescent Behavioral and Emotional Problems: Implications of Cross-Informant Correlations for Situational Specificity. Psychol. Bull. 101 (2), 213–232.
- Ackerman, P.L., Heggestad, E.D., 1997. Intelligence, Personality, and Interests: Evidence for Overlapping Traits. Psychol. Bull. 121 (2), 219–245.
- ACT, Incorporated, 2007. The ACT Technical Manual. ACT, Inc., Iowa City, IA.
- Agnew, R., Brezina, T., Wright, J.P., Cullen, F.T., 2002. Strain, Personality Traits, and Delinquency: Extending General Strain Theory. Criminology 40 (1), 43–72.
- Aigner, D.J., Hsiao, C., Kapteyn, A., Wansbeek, T., 1984. Latent Variable Models in Econometrics. In: Griliches, Z., Intriligator, M.D., (Ed.), Handbook of Econometrics. Elsevier, Amsterdam, 2, 1321–1393.
- Akerlof, G.A., Dickens, W.T., 1982. The Economic Consequences of Cognitive Dissonance. Am. Econ. Rev. 72 (3), 307–319.
- Alexander, K.L., Entwisle, D.R., Dauber, S.L., 1993. First-Grade Classroom Behavior: Its Short- and Long-Term Consequences for School Performance. Child Dev. 64 (3), 801–814.
- Allport, G.W., Odbert, H.S., 1936. Traitnames. A Psycho-Lexical Study. Psychol. Monogr. 47, 171.
- Altmann, S., Dohmen, T., Wibral, M., 2008. Do the Reciprocal Trust Less? Econ. Lett. 99 (3), 454-457.
- American Council on Education, 2009. Technical Manual: 2002 Series Ged Tests. American Council on Education, GED Testing Service, Washington, DC.
- American Psychological Association, 2007. APA Dictionary of Psychology. American Psychological Association, Washington, DC.
- Anderson, J., Burks, S., DeYoung, C., Rustichini, A., 2011. Toward the Integration of Personality Theory and Decision Theory in the Explanation of Economic Behavior. Unpublished manuscript. Presented at the IZA Workshop: Cognitive and Non-Cognitive Skills, January 27, 2011. http://www.iza.org/ conference_files/CoNoCoSk2011/rustichini_a3566.pdf.
- Anderson, T.W., Rubin, H., 1956. Statistical Inference in Factor Analysis. In: Neyman, J., (Ed.), Proceedings of the Third Symposium on Mathematical Statistics and Probability, 5. University of California Press, Berkeley, pp. 111–150.
- Andreoni, J., 1995. Warm-Glow Versus Cold-Prickle: The Effects of Positive and Negative Framing on Cooperation in Experiments. Q. J. Econ. 110 (1), 1–21.
- Ardila, A., Pineda, D., Rosselli, M., 2000. Correlation between Intelligence Test Scores and Executive Function Measures. Arch. Clin. Neuropsychol. 15 (1), 31–36.
- Aristotle, 1956. The Nicomachean Ethics. Harvard University Press, Cambridge, MA.
- Ashton, M.C., Lee, K., Perugini, M., Szarota, P., de Vries, R.E., Di Blas, L., Boies, K., De Raad, B., 2004. A Six-Factor Structure of Personality-Descriptive Adjectives: Solutions from Psycholexical Studies in Seven Languages. J. Pers. Soc. Psychol. 86 (2), 356–366.

- Ashton, M.C., Paunonen, S.V., Helmes, E., Jackson, D.N., 1998. Kin Altruism, Reciprocal Altruism, and the Big Five Personality Factors. Evol. Hum. Behav. 19 (1), 243–255.
- Ayllon, T., Kelly, K., 1972. Effects of Reinforcement on Standardized Test Performance. J. Appl. Behav. Anal. 5 (4), 477–484.
- Bagby, R.M., Bindseil, K.D., Schuller, D.R., Rector, N.A., Young, L.T., Cooke, R.G., Seeman, M.V., McCay, E.A., Joffe, R.T., 1997. Relationship between the Five-Factor Model of Personality and Unipolar, Bipolar, and Schizophrenic Patients. Psychiatry Res. 70 (2), 83–94.
- Bandiera, O., Barankay, I., Rasul, I., 2005. Social Preferences and the Response to Incentives: Evidence from Personnel Data. Q. J. Econ. 120 (3), 917–962.
- Bardsley, N., 2008. Dictator Game Giving: Altruism or Artefact? Exp. Econ. 11 (2), 122-133.
- Barenbaum, N.B., Winter, D.G., 2008. History of Modern Personality Theory and Research. In: John, O.P., Robins, R.W., Pervin, L.A., (Ed.), Handbook of Personality: Theory and Research. Guilford Press, New York, pp. 3–26.
- Barnett, W.S., Jung, K., Yarosz, D.J., Thomas, J., Hornbeck, A., Stechuk, R., Burns, S., 2008. Educational Effects of the Tools of the Mind Curriculum: A Randomized Trial. Early Child. Res. Q. 23 (3), 299–313.
- Barnett, W.S., Yarosz, D.J., Thomas, J., Hornbeck, A., 2006. Educational Effectiveness of a Vygotskian Approach to Preschool Education: A Randomized Trial. National Institute for Early Education Research, Rutgers, NJ.
- Báron, J.D., Cobb-Clark, D.A., 2010. Are Young People's Educational Outcomes Linked to Their Sense of Control? IZA Discussion Paper No. 4907.
- Barrick, M.R., Mount, M.K., 1991. The Big Five Personality Dimensions and Job Performance: A Meta-Analysis. Pers. Psychol. 44 (1), 1–26.
- Barsky, R.B., Juster, F.T., Kimball, M.S., Shapiro, M.D., 1997. Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Study. Q. J. Econ. 112 (2), 537–579.
- Baumeister, R.F., Bratslavsky, E., Muraven, M., Tice, D.M., 1998. Ego Depletion: Is the Active Self a Limited Resource? J. Pers. Soc. Psychol. 74 (5), 1252–1265.
- Bechara, A., 2005. Decision Making, Impulse Control and Loss of Willpower to Resist Drugs: A Neurocognitive Perspective. Nat. Neurosci. 8 (11), 1458–1463.
- Bechara, A., Damasio, A.R., 2005. The Somatic Marker Hypothesis: A Neural Theory of Economic Decision. Games Econ. Behav. 52 (2), 336–372.
- Becker, G.S., 1964. Human Capital: A Theoretical and Empirical Analysis, with Special Reference to Education. National Bureau of Economic Research, distributed by Columbia University Press, New York. Becker, G.S., Murphy, K.M., 1988. A Theory of Rational Addiction. J. Polit. Econ. 96 (4), 675.
- Beckett, C., Maughan, B., Rutter, M., Castle, J., Colvert, E., Groothues, C., Kreppner, J., Stevens, S., O'Connor, T.G., Sonuga-Barke, E.J.S., 2006. Do the Effects of Early Severe Deprivation on Cognition Persist into Early Adolescence? Findings from the English and Romanian Adoptees Study. Child Dev. 77 (3), 696–711.
- Behncke, S., 2009. How Do Shocks to Non-Cognitive Skills Affect Test Scores? IZA Discussion Paper No. 4222.
- Bekkers, R., 2006. Traditional and Health-Related Philanthropy: The Role of Resources and Personality. Soc. Psychol. Q. 69 (4), 349–366.
- Ben-Ner, A., Kramer, A., 2011. Personality and Altruism in the Dictator Game: Relationship to Giving to Kin, Collaborators, Competitors, and Neutrals. Pers. Indiv. Differ. In press, Corrected Proof.
- Benz, M., Meier, S., 2008. Do People Behave in Experiments as in the Field?—Evidence from Donations. Exp. Econ. 11 (3), 268–281.
- Bergen, S.E., Gardner, C.O., Kendler, K.S., 2007. Age-Related Changes in Heritability of Behavioral Phenotypes over Adolescence and Young Adulthood: A Meta-Analysis. Twin Res. Hum. Genet. 10 (3), 423–433.
- Bibby, P.A., Ferguson, E., 2011. The Ability to Process Emotional Information Predicts Loss Aversion. Pers. Indiv. Differ. In Press, Corrected Proof.
- Bickel, W.K., Odum, A.L., Madden, G.J., 1999. Impulsivity and Cigarette Smoking: Delay Discounting in Current, Never, and Ex-Smokers. Psychopharmacology 146 (4), 447–454.

- Bierman, K.L., Coie, J.D., Dodge, K.A., Greenberg, M.T., Lochman, J.E., McMahon, R.J., Pinderhughes, E., 2010. The Effects of a Multiyear Universal Social-Emotional Learning Program: The Role of Student and School Characteristics. J. Consult. Clin. Psychol. 78 (2), 156–168.
- Binet, A., Simon, T., 1916. The Development of Intelligence in Children (the Binet-Simon Scale). Williams & Wilkins Co, Baltimore, MD.
- Blair, C., 2006. How Similar Are Fluid Cognition and General Intelligence? A Developmental Neuroscience Perspective on Fluid Cognition as an Aspect of Human Cognitive Ability. Behav. Brain Sci. 29 (2), 109–160.
- Blair, C., Gamson, D., Thorne, S., Baker, D., 2005. Rising Mean IQ: Cognitive Demand of Mathematics Education for Young Children, Population Exposure to Formal Schooling, and the Neurobiology of the Prefrontal Cortex. Intelligence 33 (1), 93–106.
- Blair, C., Razza, R.P., 2007. Relating Effortful Control, Executive Function, and False Belief Understanding to Emerging Math and Literacy Ability in Kindergarten. Child Dev. 78 (2), 647–663.
- Blais, A.R., Weber, E.U., 2006. A Domain-Specific Risk-Taking (DOSPERT) Scale for Adult Populations. Judgm. Decis. Mak. 1 (1), 33–47.
- Block, J., 1995. A Contrarian View of the Five-Factor Approach to Personality Description. Psychol. Bull. 117 (2), 187–215.
- Blonigen, D.M., Hicks, B.M., Krueger, R.F., Patrick, C.J., Iacono, W.G., 2006. Continuity and Change in Psychopathic Traits as Measured Via Normal-Range Personality: A Longitudinal-Biometric Study. J. Abnorm. Psychol. 115 (1), 85–95.
- Bloom, D., Gardenhire-Crooks, A., Mandsager, C.L., 2009. Reengaging High School Dropouts: Early Results of the National Guard Youth ChalleNGe Program Evaluation. MDRC Report, New York.
- Boake, C., 2002. From the Binet-Simon to the Wechsler-Bellevue: Tracing the History of Intelligence Testing. J. Clin. Exp. Neuropsychol. 24 (3), 383–405.
- Bodrova, E., Leong, D.J., 2001. Tools of the Mind: A Case Study of Implementing the Vygotskian Approach in American Early Childhood and Primary Classrooms. International Bureau of Education, UNESCO, Geneva.
- Bodrova, E., Leong, D.J., 2007. Tools of the Mind: The Vygotskian Approach to Early Childhood Education. Pearson Education, Inc, Upper Saddle River.
- Bonato, D.P., Boland, F.J., 1983. Delay of Gratification in Obese Children. Addict. Behav. 8 (1), 71-74.
- Booth, A.L., Nolen, P.J., 2009. Gender Differences in Risk Behaviour: Does Nurture Matter? IZA Discussion Paper No. 4026.
- Borghans, L., Duckworth, A.L., Heckman, J.J., ter Weel, B., 2008. The Economics and Psychology of Personality Traits. J. Hum. Resour. 43 (4), 972–1059.
- Borghans, L., Golsteyn, B.H.H., Heckman, J.J., and Humphries, J.E., 2011. Identification Problems in Personality Psychology, *Personality and Individual Differences* 51 (Special Issue on Personality and Economics), 315–320.
- Borghans, L., Golsteyn, B.H.H., Heckman, J.J., Meijers, H., 2009. Gender Differences in Risk Aversion and Ambiguity Aversion. J. Eur. Econ. Assoc. 7 (2–3), 649–658.
- Borghans, L., Meijers, H., ter Weel, B., 2008. The Role of Noncognitive Skills in Explaining Cognitive Test Scores. Econ. Inquiry 46 (1), 2–12.
- Borkenau, P., Mauer, N., Riemann, R., Spinath, F.M., Angleitner, A., 2004. Thin Slices of Behavior as Cues of Personality and Intelligence. J. Personal. Soc. Psychol. 86 (4), 599–614.
- Borsboom, D., Mellenbergh, G.J., van Heerden, J., 2003. The Theoretical Status of Latent Variables. Psychol. Rev. 110 (2), 203–219.
- Bouchard, T.J., Loehlin, J.C., 2001. Genes, Evolution and Personality. Behav. Genet. 31 (3), 243-273.
- Bowen, W.G., Chingos, M.M., McPherson, M.S., 2009a. Educational Attainment: Overall Trends, Disparities, and the Public Universities We Study. In: Crossing the Finish Line: Completing College at America's Public Universities. Princeton University Press, Princeton, NJ, pp. 1–19.
- Bowen, W.G., Chingos, M.M., McPherson, M.S., 2009b. Test Scores and High School Grades as Predictors. In: Crossing the Finish Line: Completing College at America's Public Universities. Princeton University Press, Princeton, NJ, pp. 112–133.

- Bowles, S., Gintis, H., 1976. Schooling in Capitalist America: Educational Reform and the Contradictions of Economic Life. Basic Books, New York.
- Bowles, S., Gintis, H., Osborne, M., 2001a. The Determinants of Earnings: A Behavioral Approach. J. Econ. Lit. 39 (4), 1137–1176.
- Bowles, S., Gintis, H., Osborne, M., 2001b. Incentive-Enhancing Preferences: Personality, Behavior, and Earnings. Am. Econ. Rev. 91 (2), 155–158.
- Bowman, P.C., Matthews, C.V., 1960. Motivations of Youth for Leaving School. University of Chicago Development Project, Chicago IL.
- Boyle, S.H., Williams, R.B., Mark, D.B., Brummett, B.H., Siegler, I.C., Barefoot, J.C., 2005. Hostility, Age, and Mortality in a Sample of Cardiac Patients. Am. J. Cardiol. 96 (1), 64–66.
- Breuning, S.E., Zella, W.F., 1978. Effects of Individualized Incentives on Norm-Referenced IQ Test Performance of High School Students in Special Education Classes. J. Sch. Psychol. 16 (3), 220.
- Bridgman, P.W., 1959. The Way Things Are. Harvard University Press, Cambridge.
- Browning, M., Hansen, L.P., Heckman, J.J., 1999. Micro Data and General Equilibrium Models. In: Taylor, J.B., Woodford, M., (Ed.), Handbook of Macroeconomics. Elsevier, Amsterdam, 1A, 543–633.
- Buck, R., 2003. Advanced Calculus. Waveland Press, Inc., Long Grove, IL.
- Burks, S.V., Carpenter, J.P., Goette, L., Rustichini, A., 2009. Cognitive Skills Affect Economic Preferences, Strategic Behavior, and Job Attachment. Proc. Natl. Acad. Sci. 106 (19), 7745–7750.
- Burks, S.V., Carpenter, J., Goette, L., Rustichini, A., 2010. Overconfidence Is a Social Signaling Bias. IZA Discussion Paper No. 4840.
- Cahan, S., Cohen, N., 1989. Age Versus Schooling Effects on Intelligence Development. Child Dev. 60 (5), 1239–1249.
- Caliendo, M., Cobb-Clark, D., Uhlendorff, A., 2010. Locus of Control and Job Search Strategies. IZA Discussion Paper No. 4750.
- Caliendo, M., Fossen, F., Kritikos, A., 2010. The Impact of Risk Attitudes on Entrepreneurial Survival. Journal of Economic Behavior & Organization, 76 (1), 45–63.
- Camerer, C.F., Fehr, E., 2004. Measuring Social Norms and Preferences Using Experimental Games: A Guide for Social Scientists. In: Henrich, J., Boyd, R., Bowles, S., Camerer, C., Fehr, E., Gintis, H., (Ed.), Foundations of Human Sociality: Experimental and Ethnographic Evidence from 15 Small-Scale Societies. Oxford University Press, New York, pp. 55–95.
- Cameron, S.V., Heckman, J.J., 1998. Life Cycle Schooling and Dynamic Selection Bias: Models and Evidence for Five Cohorts of American Males. J. Polit. Econ. 106 (2), 262–333.
- Canli, T., 2006. Biology of Personality and Individual Differences. Guilford Press, New York.
- Caplin, A., Leahy, J., 2001. Psychological Expected Utility Theory and Anticipatory Feelings. Q. J. Econ. 116 (1), 55–79.
- Carneiro, P., Heckman, J.J., Masterov, D.V., 2005. Labor Market Discrimination and Racial Differences in Pre-Market Factors. J. Law Econ. 48 (1), 1–39.
- Carpenter, P.A., Just, M.A., Shell, P., 1990. What One Intelligence Test Measures: A Theoretical Account of the Processing in the Raven Progressive Matrices Test. Psychol. Rev. 97 (3), 404–431.
- Carpenter, J., Verhoogen, E., Burks, S., 2005. The Effect of Stakes in Distribution Experiments. Econ. Lett. 86 (3), 393–398.
- Carroll, J.B., 1993. Human Cognitive Abilities: A Survey of Factor-Analytic Studies. Cambridge University Press, New York.
- Casey, B.J., Tottenham, N., Liston, C., Durston, S., 2005. Imaging the Developing Brain: What Have We Learned About Cognitive Development? Trends Cogn. Sci. 9 (3), 104–110.
- Caspi, A., 2000. The Child Is Father of the Man: Personality Continuities from Childhood to Adulthood. J. Personal. Soc. Psychol. 78 (1), 158–172.
- Caspi, A., McClay, J., Moffitt, T.E., Mill, J., Martin, J., Craig, I.W., Taylor, A., Poulton, R., 2002. Role of Genotype in the Cycle of Violence in Maltreated Children. Science 297 (5582), 851–854.
- Caspi, A., Moffit, T.E., Silva, P.A., Stouthamer-Loeber, M., Krueger, R.F., Schmutte, P.S., 1994. Are Some People Crime-Prone? Replications of the Personality-Crime Relationship across Countries, Genders, Races, and Methods. Criminology 32 (2), 163–195.

- Caspi, A., Shiner, R.L., 2006. Personality Development. In: Damon, W., Lerner, R.M., (Ed.), Handbook of Child Psychology. John Wiley & Sons, Inc., Hoboken, NJ, pp. 300–365.
- Caspi, A., Sugden, K., Moffitt, T.E., Taylor, A., Craig, I.W., Harrington, H., McClay, J., Mill, J., Martin, J., Braithwaite, A., Poulton, R., 2003. Influence of Life Stress on Depression: Moderation by a Polymorphism in the 5-HTT Gene. Science 301 (5631), 386–389.
- Cattan, S., 2011. Heterogeneity and Selection in the Labor Market. PhD Thesis. Economics Department, University of Chicago.
- Cattell, J.M., 1890. Mental Tests and Measurements. Mind 15 (59), 373.
- Cattell, R.B., 1971. Abilities: Their Structure, Growth, and Action. Houghton Mifflin, Boston.
- Cattell, R.B., 1987. Intelligence: Its Structure, Growth and Action. Elsevier, Amsterdam.
- Cebi, M., 2007. Locus of Control and Human Capital Investment Revisited. J. Hum. Resour. 42 (4), 919–932.
- Cervone, D., 2004. The Architecture of Personality. Psychol. Rev. 111 (1), 183-204.
- Cervone, D., 2005. Personality Architecture: Within-Person Structures and Processes. Ann. Rev. Psychol. 56, 423–452.
- Cervone, D., Pervin, L.A., 2009. Personality: Theory and Research. John Wiley and Sons, Inc., Hoboken, NJ.
- Chamorro-Premuzic, T., Furnham, A., 2005. Personality and Intellectual Competence. L. Erlbaum Associates, Mahwah, NJ.
- Chao, L.-W., Szrek, H., Sousa Pereira, N., Pauly, M.V., 2007. Time Preference and Its Relationship with Age, Health, and Longevity Expectations. Universidade do Porto Faculdade de Economia do Porto 0706.
- Chapman, G.B., Coups, E.J., 1999. Time Preferences and Preventive Health Behavior: Acceptance of the Influenza Vaccine. Med. Decis. Mak. 19 (3), 307.
- Chapman, G.B., Nelson, R., Hier, D.B., 1999. Familiarity and Time Preferences: Decision Making About Treatments for Migraine Headaches and Crohn's Disease. J. Exp. Psychol. Appl. 5 (1), 17.
- Chesson, H., Viscusi, W.K., 2000. The Heterogeneity of Time-Risk Tradeoffs. J. Behav. Decis. Mak. 13 (2), 251–258.
- Chetty, R., Friedman, J.N., Hilger, N., Saez, E., Diane, S.W., Yagan, D., 2010. How Does Your Kindergarten Classroom Affect Your Earnings? Evidence from Project STAR. NBER Working Paper No. 16381.
- Clausen, J.A., Gilens, M., 1990. Personality and Labor Force Participation across the Life Course: A Longitudinal Study of Women's Careers. Sociol. Forum 5 (4), 595–618.
- Clingman, J., Fowler, R.L., 1976. The Effects of Primary Reward on the I.Q. Performance of Grade-School Children as a Function of Initial I.Q. Level. J. Appl. Behav. Anal. 9 (1), 19–23.
- Cloninger, C.R., 1987. A Systematic Method for Clinical Description and Classification of Personality Variants: A Proposal. Arch. Gen. Psychiatry 44 (6), 573–588.
- Cloninger, C.R., Svrakic, D.M., Bayon, C., Przybeck, T.R., 1999. Measurement of Psychopathology as Variants of Personality. In: Cloninger, C.R., (Ed.), Personality and Psychopathology. American Psychiatric Press, Washington, DC.
- Cobb-Clark, D., Tan, M., 2009. Noncognitive Skills, Occupational Attainment, and Relative Wages. Forthcoming, Labour Econ.
- Coleman, M., DeLeire, T., 2003. An Economic Model of Locus of Control and the Human Capital Investment Decision. J. Hum. Resour. 38 (3), 701–721.
- Coller, M., Williams, M.B., 1999. Eliciting Individual Discount Rates. Exp. Econ. 2 (2), 107-127.
- Conard, M.A., 2005. Aptitude Is Not Enough: How Personality and Behavior Predict Academic Performance. J. Res. Pers. 40 (2006), 339–346.
- Conti, G., Heckman, J.J., 2010. Understanding the Early Origins of the Education-Health Gradient: A Framework That Can Also Be Applied to Analyze Gene-Environment Interactions. Perspect. Psychol. Sci. 5 (5), 585–605.
- Conti, G., Heckman, J.J., Lopes, H.F., Piatek, R., 2010. Constructing Economically Justified Aggregates: An Application of the Early Origins of Health. Unpublished manuscript. Department of Economics, University of Chicago.
- Conti, G., Heckman, J.J., Urzua, S., 2010a. Early Endowments, Education, and Health. Unpublished manuscript, University of Chicago, Department of Economics.

- Conti, G., Heckman, J.J., Urzua, S., 2010b. The Education-Health Gradient. American Economic Review. 100 (2), 1–5.
- Conti, G., Pudney, S., 2007. The Dynamics of Cognitive Development, Unpublished manuscript, Department of Economics, University of Essex.

Costa, P.T., McCrae, R.R., 1992a. Four Ways Five Factors Are Basic. Pers. Individ. Dif. 13 (6), 653-665.

- Costa, P.T., McCrae, R.R., 1992b. Revised NEO Personality Inventory (NEO PI-R), NEO Five-Factor Inventory (NEO-FFI). Psychological Assessment Resources, Odessa, FL.
- Costa, P.T., McCrae, R.R., 1994. Set Like Plaster? Evidence for the Stability of Adult Personality. In: Heatherton, T.F., Weinberger, J.L., (Ed.), Can Personality Change? American Psychological Association, Washington, DC, pp. 21–40.
- Costa, P.T., McCrae, R.R., Siegler, I.C., 1999. Continuity and Change over the Adult Life Cycle: Personality and Personality Disorders. In: Cloninger, C.R., (Ed.), Personality and Psychopathology. American Psychiatric Press, Washington, DC.
- Cragg, J.G., Donald, S.G., 1997. Inferring the Rank of a Matrix. J. Econom. 76 (1-2), 223-250.
- Credé, M., Kuncel, N.R., 2008. Study Habits, Skills, and Attitudes: The Third Pillar Supporting Collegiate Academic Performance. Perspect. Psychol. Sci. 3 (6), 425–453.
- Csikszentmihalyi, M., 1996. Creativity: Flow and the Psychology of Discovery and Invention. Harper Collins Publishers, New York.
- Cudeck, R., MacCullum, R.C., 2007. Factor Analysis at 100: Historical Developments and Future Directions. Lawrence Erlibaum Associates, Mahwah, NJ.
- Cunha, F., Heckman, J.J., 2007. The Technology of Skill Formation. Am. Econ. Rev. 97 (2), 31-47.
- Cunha, F., Heckman, J.J., 2008. Formulating, Identifying and Estimating the Technology of Cognitive and Noncognitive Skill Formation. J. Hum. Resour. 43 (4), 738–782.
- Cunha, F., Heckman, J.J., 2009. The Economics and Psychology of Inequality and Human Development. J. Eur. Econ. Assoc. 7 (2–3), 320–364.
- Cunha, F., Heckman, J.J., Lochner, L.J., Masterov, D.V., 2006. Interpreting the Evidence on Life Cycle Skill Formation. In: Hanushek, E.A., Welch, F., (Ed.), Handbook of the Economics of Education. North-Holland, Amsterdam, 1, 697–812.
- Cunha, F., Heckman, J.J., Navarro, S., 2005. Separating Uncertainty from Heterogeneity in Life Cycle Earnings, the 2004 Hicks Lecture. Oxf. Econ. Pap. 57 (2), 191–261.
- Cunha, F., Heckman, J.J., Schennach, S.M., 2010. Estimating the Technology of Cognitive and Noncognitive Skill Formation. Econometrica 78 (3), 883–931.
- Daly, M., Delaney, L., Harmon, C.P., 2009. Psychological and Biological Foundations of Time Preferences. J. Eur. Econ. Assoc. 7 (2–3), 659–669.
- Damasio, H., Grabowski, T., Frank, R., Galaburda, A.M., Damasio, A.R., 2005. The Return of Phineas Gage: Clues About the Brain from the Skull of a Famous Patient. In: Cacioppo, J.T., Berntson, G.G., (Ed.), Social Neuroscience: Key Readings. Key Reading in Social Psychology. Psychology Press, New York, pp. 21–28.
- de Wit, H., Flory, J.D., Acheson, A., McCloskey, M., Manuck, S.B., 2007. IQ and Nonplanning Impulsivity Are Independently Associated with Delay Discounting in Middle-Aged Adults. Pers. Individ. Dif. 42 (1), 111–121.
- Dee, T., West, M., 2008. The Non-Cognitive Returns to Class Size. NBER Working Paper No. 13994.
- DeLong, J.B., Magin, K., 2009. The U.S. Equity Return Premium: Past, Present and Future. J. Econ. Perspect. 23 (1), 193–208.
- Devlin, B., Daniels, M., Roeder, K., 1997. The Heritability of IQ. Nature 388 (6641), 468-471.
- DeYoung, C.G., Hirsh, J.B., Shane, M.S., Papademetris, X., Rajeevan, N., Gray, J.R., 2010. Testing Predictions from Personality Neuroscience: Brain Structure and the Big Five. Psychol. Sci. 21 (6), 820–828.
- DeYoung, C.G., Quilty, L.C., Peterson, J.B., 2007. Between Facets and Domains: Ten Aspects of the Big Five. J. Pers. Soc. Psychol. 93 (5), 880–896.
- Diamond, A., Barnett, S., Thomas, J., Munro, S., 2007. Preschool Program Improves Cognitive Control. Science 318 (5855), 1387–1388.
- Dickens, W.T., Flynn, J.R., 2001. Heritability Estimates Versus Large Environmental Effects: The IQ Paradox Resolved. Psychol. Rev. 108 (2), 346–369.

Digman, J.M., 1997. Higher-Order Factors of the Big Five. J. Pers. Soc. Psychol. 73 (6), 1246–1256.

- Ding, X., Hartog, J., Sun, Y., 2010. Can We Measure Individual Risk Attitudes in a Survey? IZA Discussion Paper, 4807.
- Dohmen, T., Falk, A., Huffman, D., Sunde, U., 2008. Representative Trust and Reciprocity: Prevalence and Determinants. Econ. Inquiry 46 (1), 84–90.
- Dohmen, T., Falk, A., Huffman, D., Sunde, U., 2009. Homo Reciprocans: Survey Evidence on Behavioural Outcomes. Econ. J. 119 (536), 592–612.
- Dohmen, T., Falk, A., Huffman, D., Sunde, U., 2010. Are Risk Aversion and Impatience Related to Cognitive Ability? Am. Econ. Rev. 100 (3), 1238–1260.
- Dohmen, T., Falk, A., Huffman, D., Sunde, U., Schupp, J., Wagner, G.G., 2011. Individual Risk Attitudes: Measurement, Determinants and Behavioral Consequences. J. Eur. Econ. Assoc. 9 (3), 522–550.
- Drago, F., 2008. Self-Esteem and Earnings. IZA Discussion Paper, 3577.
- Du, W., Green, L., Myerson, J., 2002. Cross-Cultural Comparisons of Discounting Delayed and Probabilistic Rewards. Psychol. Rec. 52 (4), 479.
- Duckworth, A.L., 2007. Unpublished Dataset. Department of Psychology, University of Pennsylvania.
- Duckworth, A.L., 2009. (Over and) Beyond High-Stakes Testing. Am. Psychol. 64 (4), 279-280.
- Duckworth, A.L., 2011. Unpublished. Dataset on Personality as a Predictor of Grade Point Average and Achievement Test Scores. University of Pennsylvania, Department of Psychology.
- Duckworth, A.L., Quinn, P.D., Tsukayama, E., 2010. What No Child Left Behind Leaves Behind: The Roles of IQ and Self-Control in Predicting Standardized Achievement Test Scores and Report Card Grades. Unpublished manuscript. Department of Psychology, University of Pennsylvania.
- Duckworth, A.L., Schulze, R. 2009. Jingle Jangle: A Meta-Analysis of Convergent Validity Evidence for Self-Control Measures. Manuscript. Department of Psychology, University of Pennsylvania.
- Duckworth, A.L., Seligman, M.E.P., 2005. Self-Discipline Outdoes IQ in Predicting Academic Performance of Adolescents. Psychol. Sci. 16 (12), 939–944.
- Duckworth, A.L., Seligman, M.E.P., 2006. Self-Discipline Gives Girls the Edge: Gender in Self-Discipline, Grades, and Achievement Scores. J. Educ. Psychol. 98 (1), 198–208.
- Duckworth, A.L., Tsukayama, E., May, H., 2010. Establishing Causality Using Longitudinal Hierarchical Linear Modeling: An Illustration Predicting Achievement from Self-Control. Soc. Psychol. Personal. Sci. 1 (4), 311–317.
- Duncan, G.J., Dowsett, C.J., Claessens, A., Mugnuson, K., Huston, A.C., Klebanov, P., Pagani, L.S., Feinstein, L., Engel, M., Brooks-Gunn, J., Sexton, H., Duckworth, K., Japel, C., 2007. School Readiness and Later Achievement. Dev. Psychol. 43 (6), 1428–1446.
- Duncan, G.J., Dunifon, R., 1998. Soft-Skills and Long-Run Labor Market Success. In: Polachek, S.W., (Ed.), Research in Labor Economics. JAI Press, Stamford, CT, 17, 123–149.
- Duncan, G.J., Magnuson, K., 2010. The Nature and Impact of Early Achievement Skills, Attention Skills, and Behavior Problems. Working paper 2010 at the Department of Education, UC Irvine.
- Dur, R., Non, A., Roelfsema, H., 2010. Reciprocity and Incentive Pay in the Workplace. J. Econ. Psychol. 31 (4), 676–686.
- Durlak, J.A., Weissberg, R.P., Dymnicki, A.B., Taylor, R.D., Schellinger, K.B., 2011. The Impact of Enhancing Students' Social and Emotional Learning: A Meta-Analysis of School-Based Universal Interventions. Child Dev 82 (1), 405–432.
- Duyme, M., Dumaret, A.-C., Tomkiewicz, S., 1999. How Can We Boost IQs Of "Dull Children?: A Late Adoption Study. Proc. Natl. Acad. Sci. 96 (15), 8790–8794.
- Eckel, C.C., Grossman, P.J., 2002. Sex Differences and Statistical Stereotyping in Attitudes toward Financial Risk. Evol. Hum. Behav. 23 (4), 281–295.
- Edlund, C.V., 1972. The Effect on the Behavior of Children, as Reflected in the IQ Scores, When Reinforced after Each Correct Response. J. Appl. Behav. Anal. 5 (3), 317–319.
- Einav, L., Finkelstein, A., Pascu, I., Cullen, M.R., 2010. How General Are Risk Preferences? Choices under Uncertainty in Different Domains. NBER Working Paper No. 15686.
- Endler, N., Magnusson, D., (Ed.), 1976. Interactional Psychology and Personality. Hemisphere Publishing Corp, New York, NY.

- Epstein, S., 1979. The Stability of Behavior: I. On Predicting Most of the People Much of the Time. J. Personal. Soc. Psychol. 37 (7), 1097–1126.
- Epstein, L.G., Schneider, M., 2003. Recursive Multiple-Priors. J. Econ. Theory 113 (1), 1-31.
- Epstein, L.G., Zin, S.E., 1989. Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework. Econometrica 57 (4), 937–969.
- Estle, S.J., Green, L., Myerson, J., Holt, D.D., 2007. Discounting of Monetary and Directly Consumable Rewards. Psychol. Sci. 18 (1), 58.
- Eysenck, H.J., 1991. Dimensions of Personality: 16, 5 or 3? Criteria for a Taxonomic Paradigm. Pers. Individ. Dif. 12 (8), 773–790.
- Falk, A., Fischbacher, U., 2006. A Theory of Reciprocity. Games Econ. Behav. 54 (2), 293-315.
- Falk, A., Heckman, J.J., 2009. Lab Experiments Are a Major Source of Knowledge in the Social Sciences. Science 326 (5952), 535–538.
- Fehr, E., 2009. On the Economics and Biology of Trust. J. Eur. Econ. Assoc. 7 (2-3), 235-266.
- Fehr, E., Gächter, S., 2000. Fairness and Retaliation: The Economics of Reciprocity. J. Econ. Perspect. 14 (3), 159–181.
- Fehr, E., Schmidt, K.M., 1999. A Theory of Fairness, Competition, and Cooperation. Q. J. Econ. 114 (3), 817–868.
- Fehr, E., Schmidt, K.M., 2006. The Economics of Fairness, Reciprocity and Altruism-Experimental Evidence and New Theories. In: Kolm, S.-C., Ythier, J.M., (Ed.), Handbook of the Economics of Giving, Altruism and Reciprocity. Elsevier, Amsterdam, 1, 615–691.
- Feingold, A., 1982. The Validity of the Information and Vocabulary Subtests of the WAIS. J. Clin. Psychol. 38 (1), 169–174.
- Ferguson, E., Heckman, J.J., Corr, P., 2011. Editorial: Personality and Economics: Overview and Proposed Framework. Personality and Individual Differences, 51 (Special Issue on Personality and Economics), 201–209.
- Fergusson, D.M., Horwood, L.J., 1998. Early Conduct Problems and Later Life Opportunities. J. Child Psychol. Psychiatry 39 (8), 1097–1108.
- Figner, B., Knoch, D., Johnson, E.J., Krosch, A.R., Lisanby, S.H., Fehr, E., Weber, E.U. 2010. Lateral Prefrontal Cortex and Self-Control in Intertemporal Choice. Nat. Neurosci. 13 (5), 538–539.
- Finn, J.D., Pannozzo, G.M., Voelkl, K.E., 1995. Disruptive and Inattentive-Withdrawn Behavior and Achievement among Fourth Graders. Elem. Sch. J. 95 (5), 421–434.
- Fleeson, W., 2001. Toward a Structure- and Process-Integrated View of Personality: Traits as Density Distributions of States. J. Personal. Soc. Psychol. 80 (6), 1011–1027.
- Fleeson, W., Noftle, E., 2008. The End of the Person-Situation Debate: An Emerging Synthesis in the Answer to the Consistency Question. Soc. Personal. Psychol. Compass 2 (4), 1667–1684.
- Flynn, J.R., 2007. What Is Intelligence? Beyond the Flynn Effect. Cambridge University Press, New York. Foster, J.E., Sen, A.K., 1997. On Economic Inequality. Oxford University Press, New York.
- Frederick, S., Loewenstein, G., O'Donoghue, T., 2002. Time Discounting and Time Preference: A Critical Review. J. Econ. Lit. 40 (2), 351–401.
- Freud, S., 1909. Five Lectures on Psycho-Analysis (the Standard Edition). W. W. Norton & Company, New York.
- Frey, M.C., Detterman, D.K., 2004. Scholastic Assessment or g? The Relationship between the Scholastic Assessment Test and General Cognitive Ability. Psychol. Sci. 15 (6), 373–378.
- Friedman, H.S., 2007. Personality, Disease, and Self-Healing. In: Friedman, H.S., Silvers, R.C., (Ed.), Foundations of Health Psychology. Oxford University Press, New York, pp. 172–199.
- Friedman, H.S., Kern, M.L., Reynolds, C.A., 2010. Personality and Health, Subjective Well-Being, and Longevity. J. Personal. 78 (1), 179–216.
- Friedman, N.P., Miyake, A., Corley, R.P., Young, S.E., DeFries, J.C., Hewitt, J.K., 2006. Not All Executive Functions Are Related to Intelligence. Psychol. Sci. 17 (2), 172–179.
- Funder, D.C., 2008. Persons, Situations, and Person-Situation Interactions. In: John, O.P., Robins, R.W., Pervin, L.A., (Ed.), Handbook of Personality: Theory and Research. Guilford Press, New York, pp. 568–580.

- Funder, D.C., 2009. Persons, Behaviors and Situations: An Agenda for Personality Psychology in the Postwar Era. J. Res. Personal. 43 (2), 120–126.
- Funder, D.C., Block, J., 1989. The Role of Ego-Control, Ego-Resiliency, and IQ in Delay of Gratification in Adolescence. J. Personal. Soc. Psychol. 57 (6), 1041–1050.
- Gale, C.R., Batty, G.D., Deary, I.J., 2008. Locus of Control at Age 10 Years and Health Outcomes and Behaviors at Age 30 Years: The 1970 British Cohort Study. Psychosom. Med. 70 (4), 397–403.
- Gallo, W.T., Endrass, J., Bradley, E.H., Hell, D., Kasl, S.V., 2003. The Influence of Internal Control on the Employment Status of German Workers. Schmollers Jahrbuch 123 (1), 71–81.
- Galton, F., 1883. Inquiries into Human Faculty and Its Development. Macmillan and Company, London.
- Gardner, H., 2004. Frames of Mind: The Theory of Multiple Intelligences. Basic Books, New York, NY.
- Garon, N., Bryson, S.E., Smith, I.M., 2008. Executive Function in Preschoolers: A Review Using an Integrative Framework. Psychol. Bull. 134 (1), 31–60.
- Geiser, S., Santelices, M.V., 2007. Validity of High School Grades in Predicting Student Success Beyond the Freshman Year: High-School Record Vs. Standardized Tests as Indicators of Four-Year College Outcomes. Center for Studies in Higher Education at the University of California, Berkeley CSHE.6.07.
- Gilboa, I., Schmeidler, D., 1989. Maxmin Expected Utility with Non-Unique Prior. J. Math. Econ. 18 (2), 141–153.
- Gillham, J.E., Reivich, K.J., 1999. Prevention of Depressive Symptoms in School Children: A Research Update. Psychol. Sci. 10 (5), 461–462.
- Gneezy, U., List, J.A., 2006. Putting Behavioral Economics to Work: Testing for Gift Exchange in Labor Markets Using Field Experiments. Econometrica 74 (5), 1365–1384.
- Goldberg, L.R., 1993. The Structure of Phenotypic Personality Traits. Am. Psychol. 48 (1), 26-34.
- Goldberg, L.R., Sweeney, D., Merenda, P.F., Hughes, Jr. J.E., 1998. Demographic Variables and Personality: The Effects of Gender, Age, Education, and Ethnic/Racial Status on Self-Descriptions of Personality Attributes. Pers. Individ. Dif. 24 (3), 393–403.
- Goldberger, A.S., 1968. Topics in Regression Analysis. Macmillan, New York.
- Goldsmith, H.H., Buss, A.H., Plomin, R., Rothbart, M.K., Thomas, A., Chess, S., Hinde, R.A., McCall, R.B., 1987. Roundtable: What Is Temperament? Four Approaches. Child Dev. 58 (2), 505–529.
- Gorsuch, R., 1983. Factor Analysis. Lawrence Erlbaum Associates, Hillsdale, NJ.
- Gottfredson, L.S., 1998. The General Intelligence Factor. Sci. Am. Presents 9 (4), 24-29.
- Gottfredson, L.S., 2002. g: Highly General and Highly Practical. In: Sternberg, R.J., Grigorenko, E.L., (Ed.), The General Factor of Intelligence: How General Is It? Lawrence Erlbaum Associates Publishers, Mahwah, NJ, pp. 331–380.
- Gottfredson, M.R., Hirschi, T., 1990. A General Theory of Crime. Stanford University Press, Palo Alto, CA.
- Gottschalk, P., 2005. Can Work Alter Welfare Recipients' Beliefs? J. Policy Anal. Manage 24 (3), 485–498.
- Gough, H.G., 1964. Graduation from High School as Predicted from the California Psychological Inventory. Psychol. Sch. 3 (3), 208–216.
- Gough, H.G., Heilbrun, A.B., 1983. The Adjective Check List Manual. Consulting Psychologists Press, Palo Alto, CA.
- Gray, J.R., 2004. Integration of Emotion and Cognitive Control. Curr. Dir. Psychol. Sci. 13 (2), 46-48.
- Green, L., Fry, A.F., Myerson, J., 1994. Discounting of Delayed Rewards: A Life-Span Comparison. Psychol. Sci. 5 (1), 33–36.
- Green, L., Myerson, J., 2004. A Discounting Framework for Choice with Delayed and Probabilistic Rewards. Psychol. Bull. 130 (5), 769.
- Greene, A.C., Sapp, G.L., Chissom, B., 1990. Validation of the Stanford-Binet Intelligence Scale: Fourth Edition with Exceptional Black Male Students. Psychol. Sch. 27 (1), 35–41.
- Griliches, Z., 1977. Estimating the Returns to Schooling: Some Econometric Problems. Econometrica 45 (1), 1–22.
- Groth-Marnat, G., 2009. Handbook of Psychological Assessment. John Wiley and Sons, Hoboken, NJ.
- Gul, F., Pesendorfer, W., 2004. Self-Control and the Theory of Consumption. Econometrica 72 (1), 119–158.

- Gunderson, J.G., Triebwasser, J., Phillips, K.A., Sullivan, C.N., 1999. Personality and Vulnerability to Affective Disorders. In: Cloninger, C.R., (Ed.), Personality and Psychopathology. American Psychiatric Press, Washington, DC.
- Gustafsson, J.-E., 1988. Hierarchical Models of Individual Differences in Cognitive Abilities. In: Sternberg, R.J., (Ed.), Advances in the Psychology of Human Intelligence. Lawrence Erlbaum Associates, Hillsdale, NJ, p. 4.
- Ham, R., Junankar, P.N., Wells, R., 2009. Occupational Choice: Personality Matters. IZA Discussion Paper No. 4105.
- Hampson, S.E., Friedman, H.S., 2008. Personality and Health: A Lifespan Perspective. In: John, O.P., Robins, R., Pervin, L., (Ed.), The Handbook of Personality: Theory and Research. Guilford, New York, pp. 770–794.
- Hampson, S.E., Goldberg, L.R., Vogt, T.M., Dubanoski, J.P., 2007. Mechanisms by Which Childhood Personality Traits Influence Adult Health Status: Educational Attainment and Healthy Behaviors. Health Psychol. 26 (1), 121–125.
- Hampson, S.E., Tildesley, E., Andrews, J.A., Luyckx, K., Mroczek, D.K., 2010. The Relation of Change in Hostility and Sociability During Childhood to Substance Use in Mid Adolescence. J. Res. Personal. 44 (1), 103–114.
- Hanoch, Y., Johnson, J.G., Wilke, A., 2006. Domain Specificity in Experimental Measures and Participant Recruitment: An Application to Risk-Taking Behavior. Psychol. Sci. 17 (4), 300.
- Hansen, L.P., 2005. Comment On "Exotic Preferences for Macroeconomics." By D.K Backus, B.R. Routledge, and, S.E. Zin, In: Gertler, M. Rogoff, K., (Ed.), NBER Macroeconomics Annual. MIT Press, Cambridge, MA, pp. 391–405.
- Hansen, K.T., Heckman, J.J., Mullen, K.J., 2004. The Effect of Schooling and Ability on Achievement Test Scores. J. Econom. 121 (1–2), 39–98.
- Hansen, L.P., Sargent, T.J., 2008. Robustness. Princeton University Press, Princeton, NJ.
- Harman, H.H., 1976. Modern Factor Analysis, third ed. University of Chicago Press, Chicago, IL.
- Harrison, G.W., Lau, M.I., Williams, M.B., 2002. Estimating Individual Discount Rates in Denmark: A Field Experiment. Am. Econ. Rev. 92 (5), 1606–1617.
- Hartigan, J.A., Wigdor, A.K., 1989. Fairness in Employment Testing: Validity Generalization, Minority Issues, and the General Aptitude Test Battery. National Academy Press, Washington, DC, p. 354.
- Hartlage, L.C., Steele, C.T., 1977. WISC and WISC-R Correlates of Academic Achievement. Psychol. Sch. 14 (1), 15–18.
- Hartog, J., 1980. Earnings and Capability Requirements. Rev. Econ. Stat. 62 (2), 230-240.
- Hartog, J., 2001. On Human Capital and Individual Capabilities. Rev. Income Wealth 47 (4), 515-540.
- Hathaway, S.R., Reynolds, P.C., Monachesi, E.D., 1969. Follow-up of the Later Careers and Lives of 1,000 Boys Who Dropped out of High School. J. Consult. Clin. Psychol. 33 (3), 370–380.
- Heckman, J.J., 1976. A Life-Cycle Model of Earnings, Learning, and Consumption. J. Polit. Econ. 84 (4, Part 2), S11–S44.
- Heckman, J.J., 2007. The Economics, Technology and Neuroscience of Human Capability Formation. Proc. Natl. Acad. Sci. 104 (3), 13250–13255.
- Heckman, J.J., 2008a. Econometric Causality. Int. Stat. Rev. 76 (1), 1–27.
- Heckman, J.J., 2008b. Schools, Skills and Synapses. Econom. Inquiry 46 (3), 289-324.
- Heckman, J.J., Honoré, B.E., 1990. The Empirical Content of the Roy Model. Econometrica 58 (5), 1121–1149.
- Heckman, J.J., Humphries, J.E., Mader, N., 2011. The GED. In: Hanushek, E.A., Machin, S., Wößmann, L., (Ed.), Handbook of the Economics of Education. Elsevier, Amsterdam, pp. 423–484.
- Heckman, J.J., Humphries, J.E., Urzua, S., Veramendi, G., 2011. The Effects of Schooling on Labor Market and Health Outcomes. Unpublished Manuscript. Department of Economics, University of Chicago.
- Heckman, J.J., LaFontaine, P.A., 2010. The American High School Graduation Rate: Trends and Levels. Rev. Econom. Stat. 92 (2), 244–262.

- Heckman, J.J., Malofeeva, L., Pinto, R., Savelyev, P.A., 2011. Understanding the Mechanisms through Which an Influential Early Childhood Program Boosted Adult Outcomes. Unpublished manuscript. Department of Economics, University of Chicago, First draft, 2008.
- Heckman, J.J., Masterov, D.V., 2007. The Productivity Argument for Investing in Young Children. Rev. Agric. Econ. 29 (3), 446–493.
- Heckman, J.J., Moon, S.H., Pinto, R., Savelyev, P.A., Yavitz, A.Q., 2010a. Analyzing Social Experiments as Implemented: A Reexamination of the Evidence from the Highscope Perry Preschool Program. Quant. Econom. 1 (1), 1–46.
- Heckman, J.J., Moon, S.H., Pinto, R., Savelyev, P.A., Yavitz, A.Q., 2010b. The Rate of Return to the Highscope Perry Preschool Program. J. Public Econom. 94 (1–2), 114–128.
- Heckman, J.J., Navarro, S., 2007. Dynamic Discrete Choice and Dynamic Treatment Effects. J. Econom. 136 (2), 341–396.
- Heckman, J.J., Rubinstein, Y., 2001. The Importance of Noncognitive Skills: Lessons from the GED Testing Program. Am. Econom. Rev. 91 (2), 145–149.
- Heckman, J.J., Sedlacek, G.L., 1985. Heterogeneity, Aggregation, and Market Wage Functions: An Empirical Model of Self-Selection in the Labor Market. J. Polit. Econ. 93 (6), 1077–1125.
- Heckman, J.J., Stixrud, J., Urzua, S., 2006. The Effects of Cognitive and Noncognitive Abilities on Labor Market Outcomes and Social Behavior. J. Labor Econ. 24 (3), 411–482.
- Heckman, J.J., Williams, B., 2011. Estimating Treatment Effects with Proxies for Controls. Unpublished manuscript. Department of Economics, University of Chicago.
- Heitz, R.P., Unsworth, N., Engle, R.W., 2005. Working Memory Capacity, Attention Control, and Fluid Intelligence. In: Wilhelm., O., Engle, R.W., (Ed.), Handbook of Understanding and Measuring Intelligence. Sage Publications Inc., Thousand Oaks, CA, pp. 61–77.
- Helson, R., Kwan, V.S.Y., John, O.P., Jones, C., 2002. The Growing Evidence for Personality Change in Adulthood: Findings from Research with Personality Inventories. J. Res. Personal. 36 (4), 287–306.
- Helson, R., Picano, J., 1990. Is the Traditional Role Bad for Women? J. Personal. Soc. Psychol. 59 (2), 311–320.
- Hembree, R., 1988. Correlates, Causes, Effects, and Treatment of Test Anxiety. Rev. Educ. Res. 58 (1), 47–77.
- Hennig-Schmidt, H., Rockenbach, B., Sadrieh, A., 2010. In Search of Workers' Real Effort Reciprocity a Field and a Laboratory Experiment. J. Eur. Econ. Assoc. 8 (4), 817–837.
- Herrnstein, R.J., Murray, C.A., 1994. The Bell Curve: Intelligence and Class Structure in American Life. Free Press, New York.
- Hicks, J.R., 1946. Value and Capital: An Inquiry into Some Fundamental Principles of Economic Theory. Clarendon Press, Oxford.
- Hoffman, M., Fessler, D., Gneezy, U., List, J., 2010. Height and Competitiveness. Unpublished manuscript, Department of Economics, University of Chicago.
- Hofstee, W.K., de Raad, B., Goldberg, L.R., 1992. Integration of the Big Five and Circumplex Approaches to Trait Structure. J. Personal. Soc. Psychol. 63 (1), 146–163.
- Hogan, R., Hogan, J., 2007. Hogan Personality Inventory Manual. Hogan Assessment Systems, Tulsa, OK.
- Hogan, R., Hogan, J., Roberts, B.W., 1996. Personality Measurement and Employment Decisions: Questions and Answers. Am. Psychol. 51 (5), 469–477.
- Hogan, J., Holland, B., 2003. Using Theory to Evaluate Personality and Job-Performance Relations: A Socioanalytic Perspective. J. Appl. Psychol. 88 (1), 100–112.
- Hogan, R., Roberts, B.W., 2001. Personality Psychology in the Workplace. American Psychological Association, Washington, DC.
- Holmlund, H., Silva, O., 2009. Targeting Non-Cognitive Skills to Improve Cognitive Outcomes: Evidence from a Remedial Education Intervention. IZA Discussion Paper No. 4476.
- Holt, M.M., Hobbs, T.R., 1979. The Effects of Token Reinforcement, Feedback and Response Cost on Standardized Test Performance. Behav. Res. Ther. 17 (1), 81–83.

- Hopkins, K.D., Bracht, G.H., 1975. Ten-Year Stability of Verbal and Nonverbal IQ Scores. Am. Educ. Res. J. 12 (4), 469–477.
- Horn, J.L., 1970. Organization of Data on Life-Span Development of Human Abilities. In: Goulet, L.R., Baltes, P.B., (Ed.), Life-Span Developmental Psychology: Research and Theory. Academic Press, New York, pp. 423–466.
- Horn, J.L., McArdle, J.J., 2007. Understanding Human Intelligence since Spearman. In: Cudek, R., MacCallum, R.C., (Ed.), Factor Analysis at 100: Historical Developments and Future Directions. Lawrence Erlbaum Associates, Mahwah, NJ, pp. 205–249.
- Horvath, P., Zuckerman, M., 1993. Sensation Seeking, Risk Appraisal, and Risky Behavior. Pers. Individ. Dif. 14 (1), 41–52.
- Hough, L.M., 1992. The "Big Five" Personality Variables–Construct Confusion: Description Versus Prediction. Hum. Perf. 5 (1–2), 139–155.
- Hough, L.M., Eaton, N.K., Dunnette, M.D., Kamp, J.D., McCloy, R.A., 1990. Criterion-Related Validities of Personality Constructs and the Effect of Response Distortion on Those Validities. J. Appl. Psychol. 75 (5), 581–595.
- Hough, L.M., Ones, D.S., 2002. The Structure, Measurement, Validity, and Use of Personality Variables in Industrial, Work, and Organizational Psychology. In: Anderson, N., Ones, D.S., Sinangil, H.K., Viswesvaran, C., (Ed.), Handbook of Industrial, Work and Organizational Psychology, Volume 1: Personnel Psychology. Sage Publications Inc., Thousand Oaks, CA, pp. 233–277.
- Hough, L.M., Oswald, F.L., 2000. Personnel Selection: Looking toward the Future–Remembering the Past. Annu. Rev. Psychol. 51, 631–664.
- Hsiao, C., 2003. Analysis of Panel Data. Cambridge University Press, New York.
- Hunter, J.E., 1986. Cognitive Ability, Cognitive Aptitudes, Job Knowledge, and Job Performance. J. Vocat. Behav. 29 (3), 340–362.
- Jackson, D.N., 1974. Personality Research Form Manual. Research Psychologists Press, Goshen, NY.
- Jackson, J.J., Hill, P.L., Payne, B.R., Roberts, B.W., Stine-Morrow, E.A.L., 2010. Can an Old Dog Learn (and Want to Experience) New Tricks? Cognitive Training Increases Openness in Older Adults. Unpublished Manuscript. Department of Psychology, University of Illinois.
- James, W., 1890. The Principles of Psychology. H. Holt and Company, New York.
- Janosz, M., LeBlanc, M., Boulerice, B., Tremblay, R.E., 1997. Disentangling the Weight of School Dropout Predictors: A Test on Two Longitudinal Samples. J. Youth Adolesc. 26 (6), 733–762.
- Jensen, A.R., 1998. The g Factor: The Science of Mental Ability. Praeger, Westport, CT.
- John, O.P., 1990. The "Big Five" Factor Taxonomy: Dimensions of Personality in the Natural Language and in Questionnaires. In: Pervin, L.A., John, O.P., (Ed.), Handbook of Personality: Theory and Research. Guilford Press, New York, pp. 66–100.
- John, O.P., Caspi, A., Robins, R.W., Moffitt, T.E., 1994. The "Little Five": Exploring the Nomological Network of the Five-Factor Model of Personality in Adolescent Boys. Child Dev. 65 (1), 160–178.
- John, O.P., Robins, R.W., Pervin, L.A., 2008. Handbook of Personality: Theory and Research. The Guilford Press, New York.
- John, O.P., Srivastava, S., 1999. The Big Five Trait Taxonomy: History, Measurement and Theoretical Perspectives. In: Pervin, L.A., John, O.P., (Ed.), Handbook of Personality: Theory and Research. The Guilford Press, New York, 102–138.
- Johnson, W., Bouchard, Jr., T.J., Krueger, R.F., McGue, M., Gottesman, I.I., 2004. Just One g: Consistent Results from Three Test Batteries. Intelligence 32 (1), 95–107.
- Johnson, J., Wilke, A., Weber, E.U., 2004. Beyond a Trait View of Risk Taking: A Domain-Specific Scale Measuring Risk Perceptions, Expected Benefits, and Perceived-Risk Attitudes in German-Speaking Populations. Pol. Psychol. Bull. 35 (3), 153–172.
- Judge, T.A., Bono, J.E., Erez, A., Locke, E.A., 2005. Core Self-Evaluations and Job and Life Satisfaction: The Role of Self-Concordance and Goal Attainment. J. Appl. Psychol. 90 (2), 257–268.
- Judge, T.A., Erez, A., Bono, J.E., Thoresen, C.J., 2002. Are Measures of Self-Esteem, Neuroticism, Locus of Control, and Generalized Self-Efficacy Indicators of a Common Core Construct? J. Personal. Soc. Psychol. 83 (3), 693–710.
- Judge, T.A., Hurst, C., 2007. Capitalizing on One's Advantages: Role of Core Self-Evaluations. J. Appl. Psychol. 92 (5), 1212–1227.
- Kahneman, D., Tversky, A., 1979. Prospect Theory: An Analysis of Decision under Risk. Econometrica 47 (2), 263–291.
- Kelly, F.J., Veldman, D.J., 1964. Delinquency and School Dropout Behavior as a Function of Impulsivity and Nondominant Values. J. Abnorm. Soc. Psychol. 69 (2), 190–194.
- Kern, M.L., Friedman, H.S., 2008. Do Conscientious Individuals Live Longer? Health Psychol. 27 (5), 505–512.
- Kern, M.L., Friedman, H.S., 2010a. Personality and Pathways of Influence on Physical Health. Soc. Personal. Psychol. Compass 5 (1), 76–87.
- Kern, M.L., Friedman, H.S., 2010b. Why Do Some People Thrive While Other Succumb to Disease and Stagnation? Personality, Social Relations, and Resilience. In: Fry, P.S., Keyes, C.L.M., (Ed.), New Frontiers in Resilient Aging: Life-Strengths and Well-Being in Late Life. Cambridge University Press, Cambridge, pp. 162–184.
- Kilburn, M.R., Hanser, L.M., Klerman, J.A., 1998. Information in the NELS and Approaches to Imputing Test Scores. Estimating the AFQT Scores for National Educational Longitudinal Study (NELS) Respondents. National Defense Research Institute, Arroya Center: Rand Corp, Santa Monica, CA.
- Kirby, K.N., 1997. Bidding on the Future: Evidence against Normative Discounting of Delayed Rewards. J. Exp. Psychol. Gen. 126 (1), 54.
- Kirby, K.N., Herrnstein, R.J., 1995. Preference Reversals Due to Myopic Discounting of Delayed Reward. Psychol. Sci. 6 (2), 83–89.
- Knudsen, E.I., Heckman, J.J., Cameron, J.L., Shonkoff, J.P., 2006. Economic, Neurobiological, and Behavioral Perspectives on Building America's Future Workforce. Proc. Natl. Acad. Sci. 103 (27), 10155–10162.
- Knutson, B., Wolkowitz, O.M., Cole, S.W., Chan, T., Moore, E.A., Johnson, R.C., Terpstra, J., Turner, R.A., Reus, V.I., 1998. Selective Alteration of Personality and Social Behavior by Serotonergic Intervention. Am. J. Psychiatry 155 (3), 373–379.
- Kobrin, J.L., Patterson, B.F., Shaw, E.J., Mattern, K.D., Barbuti, S.M., 2008. Validity of the SAT for Predicting First-Year College Grade Point Average. Coll. Board 5:1–10.
- Kogan, N., Wallach, M.A., 1967. Risk Taking as a Function of the Situation, the Person, and the Group. In: Mandler, G., Mussen, P., Kogan, N., Wallach, M., (Ed.), New Directions in Psychology III. Holt, Rinehart & Winston, New York, pp. 111–278.
- Kohn, M.L., Schooler, C., 1978. The Reciprocal Effects of the Substantive Complexity of Work and Intellectual Flexibility: A Longitudinal Assessment. Am. J. Soc. 84 (1), 24–52.
- Kosfeld, M., Heinrichs, M., Zak, P.J., Fischbacher, U., Fehr, E., 2005. Oxytocin Increases Trust in Humans. Nature 435 (7042), 673–676.
- Köszegi, B., 2006. Ego Utility, Overconfidence, and Task Choice. J. Eur. Econ. Assoc. 4 (4), 673-707.
- Kovas, Y., Haworth, C.M.A., Philip, S.D., Plomin, R., Weinberg, R.A., Thomson, J.M., Fischer, K.W., 2007. The Genetic and Environmental Origins of Learning Abilities and Disabilities in the Early School Years. Monogr. Soc. Res. Child Dev. 72 (3), vii–144.
- Kristensen, P., Bjerkedal, T., 2007. Explaining the Relation between Birth Order and Intelligence. Science 316 (5832), 1717.
- Krueger, R.F., Eaton, N., 2010. Personality Traits and the Classification of Mental Disorders: Toward a More Complete Integration in DSM 5 and an Empirical Model of Psychopathology. Personal. Disord. Theory Res. Treat. 1 (2), 97–118.
- Krueger, R.F., Johnson, W., 2008. Behavioral Genetics and Personality: A New Look at the Integration of Nature and Nurture. In: John, O.P., Robins, R.W., Pervin, L.A., (Ed.), Handbook of Personality: Theory and Research. Guilford Press, New York, pp. 287–310.
- Krueger, R.F., South, S., Johnson, W., Iacono, W.G., 2008. The Heritability of Personality Is Not Always 50%: Gene-Environment Interactions and Correlations between Personality and Parenting. J. Personal. 76 (6), 1485–1522.
- Kube, S., Maréchal, M.A., Puppe, C., 2006. Putting Reciprocity to Work–Positive Versus Negative Responses in the Field. University of St. Gallen Working Paper, 2006–27.

- Kubzansky, L.D., Sparrow, D., Vokonas, P., Kawachi, I., 2001. Is the Glass Half Empty or Half Full? A Prospective Study of Optimism and Coronary Heart Disease in the Normative Aging Study. Psychosom. Med. 63 (6), 910.
- Kuncel, N.R., Hezlett, S.A., 2007. Standardized Tests Predict Graduate Students' Success. Science 315 (5815), 1080–1081.
- Kuncel, N.R., Ones, D.S., Sackett, P.R., 2010. Individual Differences as Predictors of Work, Educational, and Broad Life Outcomes. Pers. Individ. Dif. 49 (4), 331–336.
- Kvist, A.V., Gustafsson, J.-E., 2008. The Relation between Fluid Intelligence and the General Factor as a Function of Cultural Background: A Test of Cattell's Investment Theory. Intelligence 36 (5), 422–436.
- Kyllonen, P.C., Walters, A.M., Kaufman, J.C., 2005. Noncognitive Constructs and Their Assessment in Graduate Education: A Review. Educ. Assess. 10 (3), 153–184.
- Larson, G.E., Saccuzzo, D.P., Brown, J., 1994. Motivation: Cause or Confound in Information Processing/ Intelligence Correlations? Acta Psychol 85 (1), 25–37.
- Layard, R., 2005. Happiness: Lessons from a New Science. Penguin, London.
- Le, A.T., Miller, P.W., Slutske, W.S., Martin, N.G., 2010. Are Attitudes Towards Economic Risk Heritable? Analysis Using the Australian Twin Study of Gambling. IZA Discussion Paper No. 4859.
- Lear, G.R., 2004. Happy Lives and the Highest Good: Aristotle's Nicomachean Ethics. Princeton University Press, Princeton.
- Lejuez, C.W., Aklin, W.M., Zvolensky, M.J., Pedulla, C.M., 2003. Evaluation of the Balloon Analogue Risk Task (BART) as a Predictor of Adolescent Real-World Risk-Taking Behaviours. J. Adolesc. 26 (4), 475–479.
- Lejuez, C.W., Read, J.P., Kahler, C.W., Richards, J.B., Ramsey, S.E., Stuart, G.L., Strong, D.R., Brown, R.A., 2002. Evaluation of a Behavioral Measure of Risk Taking: The Balloon Analogue Risk Task (BART). J. Exp. Psychol. Appl. 8 (2), 75–84.
- Levitt, S.D., List, J.A., 2007. What Do Laboratory Experiments Measuring Social Preferences Reveal About the Real World? J. Econ. Perspect. 21 (2), 153–174.
- Lezak, M.D., 1987. Relationships between Personality Disorders, Social Disturbances, and Physical Disability Following Traumatic Brain Injury. J. Head Trauma Rehabil. 2 (1), 57–69.
- Lievens, F., Dilchert, S., Ones, D.S., 2009. Personality Scale Validities Increase Throughout Medical School. J. Appl. Psychol. 94 (6), 1514–1535.
- Lillard, A., Else-Quest, N., 2006. The Early Years: Evaluating Montessori. Science 313 (5795), 1893–1894.
- List, J.A., 2003. Does Market Experience Eliminate Market Anomalies? Q. J. Econom. 118 (1), 41-71.
- List, J.A., 2006. The Behavioralist Meets the Market: Measuring Social Preferences and Reputation Effects in Actual Transactions. J. Polit. Econ. 114 (1), 1–37.
- List, J.A., 2007. On the Interpretation of Giving in Dictator Games. J. Polit. Econ. 115 (3), 482.
- List, J.A., 2009. Social Preferences: Some Thoughts from the Field. Ann. Rev. Econom. 1, 563-583.
- Lizzeri, A., Siniscalchi, M., 2008. Parental Guidance and Supervised Learning. Q. J. Econom. 123 (3), 1161–1195.
- Lleras, C., 2008. Do Skills and Behaviors in High School Matter? The Contribution of Noncognitive Factors in Explaining Differences in Educational Attainment and Earnings. Soc. Sci. Res. 37 (3), 888–902.
- Lochner, L.J., Moretti, E., 2004. The Effect of Education on Crime: Evidence from Prison Inmates, Arrests, and Self-Reports. Am. Econom. Rev. 94 (1), 155–189.
- Lounsbury, J.W., Steel, R.P., Loveland, J.M., Gibson, L.W., 2004. An Investigation of Personality Traits in Relation to Adolescent School Absenteeism. J. Youth Adolesc. 33 (5), 457–466.
- Lubbers, M.J., Van der Werf, M.P.C., Kuyper, H., Hendriks, A.A.J., 2010. Does Homework Behavior Mediate the Relation between Personality and Academic Performance? Learn. Individ. Differ. 20 (3), 203–208.
- Lubinski, D., 2004. Introduction to the Special Section on Cognitive Abilities: 100 Years after Spearman's (1904) "General Intelligence," Objectively Determined and Measured'. J. Personal. Soc. Psychol. 86 (1), 96.
- Lucas, R.E., Donnellan, M.B., 2009. If the Person-Situation Debate Is Really over, Why Does It Still Generate So Much Negative Affect? J. Res. Personal. 43 (2), 146–149.

Lykken, D.T., 2007. A More Accurate Estimate of Heritability. Twin Res. Hum. Genet. 10 (1), 168–173. Machin, S., Marie, O., Vujić, S., 2010. The Crime Reducing Effect of Education. CEP Discussion Paper No. 979.

- Mandelbrot, B.B., 1962. Paretian Distributions and Income Maximization. Q. J. Econom. 76 (1), 57-85.
- Martin, R.P., 1989. Activity Level, Distractibility, and Persistence: Critical Characteristics in Early Schooling. In: Kohnstamm, G.A., Bates, J.E., Rothbart, M.K., (Ed.), Temperament in Childhood. John Wiley and Sons Ltd., Chichester, England, pp. 451–461.
- Martins, P.S., 2010. Can Targeted, Non-Cognitive Skills Programs Improve Achievement? IZA Discussion Paper No. 5266.
- Martin, L.R., Friedman, H.S., Schwartz, J.E., 2007. Personality and Mortality Risk across the Life Span: The Importance of Conscientiousness as a Biopsychosocial Attribute. Health Psychol. 26 (4), 428–436.
- Mataró, M., Jurado, A., García-Sánchez, C., Barraquer, L., Costa-Jussá, F.R., Junqué, C., 2001. Long-Term Effects of Bilateral Frontal Brain Lesion: 60 Years after Injury with an Iron Bar. Arch. Neurol. 58 (7), 1139–1142.
- Matsumoto, K., Tanaka, K., 2004. Conflict and Cognitive Control. Science 303 (5660), 969-970.
- Mayer, J.D., Salovey, P., 1997. What Is Emotional Intelligence? In: Salovey, P., Sluyter, D.J., (Ed.), Emotional Development and Emotional Intelligence: Educational Implications. Basic Books, New York, p. 3.
- McAdams, D.P., 2006. The Person: A New Introduction to Personality Psychology. J. Wiley and Sons, Hoboken, NJ.
- McArdle, J.J., Hamagami, F., Meredith, W., Bradway, K.P., 2000. Modeling the Dynamic Hypotheses of Gf-Gc Theory Using Longitudinal Life-Span Data. Learn. Individ. Differ. 12 (1), 53–79.
- McCabe, L.A., Cunnington, M., Brooks-Gunn, J., 2004. The Development of Self-Regulation in Young Children: Individual Characteristics and Environmental Contexts. In: Baumeister, R.F., Vohs, K.D., (Ed.), Handbook of Self-Regulation: Research, Theory, and Applications. Guilford Press, New York, pp. 340–356.
- McClelland, M.M., Cameron, C.E., Connor, C.M., Farris, C.L., Jewkes, A.M., Morrison, F.J., 2007. Links between Behavioral Regulation and Preschoolers' Literacy, Vocabulary, and Math Skills. Dev. Psychol. 43 (4), 947–959.
- McClelland, D.C., Koestner, R., Weinberger, J., 1989. How Do Self-Attributed and Implicit Motives Differ? Psychol. Rev. 96 (4), 690–702.
- McClure, S.M., Laibson, D.I., Loewenstein, G.F., Cohen, J.D., 2004. Separate Neural Systems Value Immediate and Delayed Monetary Rewards. Science 306 (5695), 503–507.
- McCrae, R.R., Costa, P.T., 1990. Personality in Adulthood. Guilford Press, New York.
- McCrae, R.R., Costa, P.T., 1994. The Stability of Personality: Observation and Evaluations. Curr Dir. Psychol. Sci. 3 (6), 173–175.
- McCrae, R.R., Costa, P.T., 1996. Toward a New Generation of Personality Theories: Theoretical Contexts for the Five-Factor Model. In: Wiggins, J.S. (Ed.), The Five-Factor Model of Personality: Theoretical Perspectives. Guilford Press, New York, pp. 51–87.
- McCrae, R.R., Costa, P.T., 1997a. Conceptions and Correlates of Openness to Experience. In: Hogan, R., Johnson, J.A., Briggs, S.R., (Ed.), Handbook of Personality Psychology. Academic Press, San Diego, CA, pp. 825–847.
- McCrae, R.R., Costa, P.T., 1997b. Personality Trait Structure as a Human Universal. Am. Psychol. 52 (5), 509–516.
- McCrae, R.R., Costa, P.T., 2003. Personality in Adulthood: A Five-Factor Theory Perspective. Guilford Press, New York.
- McCrae, R.R., Costa, Jr., P.T., 2008. The Five-Factor Theory of Personality. In: John, O.P., Robins, R.W., Pervin, L.A., (Ed.), Handbook of Personality: Theory and Research. The Guilford Press, New York, pp. 159–181.
- McGee, A., 2010. How the Perception of Control Influences Unemployed Job Search. Unpublished manuscript. Department of Economics, Ohio State University.
- McGue, M., Bacon, S., Lykken, D.T., 1993. Personality Stability and Change in Early Adulthood: A Behavioral Genetic Analysis. Dev. Psychol. 29 (1), 96–109.

- McGue, M., Bouchard, T.J., Iacono, W.G., Lykken, D.T., 1993. Behavioral Genetics of Cognitive Ability: A Life-Span Perspective. In: Plomin, R., McClearn, G.E., (Ed.), Nature, Nurture and Psychology. American Psychological Association, Washington, DC, pp. 59–76.
- McHenry, J.J., Hough, L.M., Toquam, J.L., Hanson, M.A., Ashworth, S., 1990. Project a Validity Results: The Relationship between Predictor and Criterion Domains. Pers. Psychol. 43 (2), 335–354.
- Meier, S., 2007. A Survey of Economic Theories and Field Evidence on Pro-Social Behavior. In: Frey, B.S., Stutzer, A., (Ed.), Economics and Psychology: A Promising New Cross-Disciplinary Field. MIT Press, Cambridge, MA, pp. 51–88.
- Mershon, B., Gorsuch, R.L., 1988. Number of Factors in the Personality Sphere: Does Increase in Factors Increase Predictability of Real-Life Criteria? J. Personal. Soc. Psychol. 55 (4), 675–680.
- Michalko, K.T., Saklofske, D.H., 1996. A Psychometric Investigation of the Wechsler Individual Achievement Test with a Sample of Saskatchewan Schoolchildren. Can. J. Sch. Psychol. 12 (1), 44–54.
- Miller, E.K., 2000. The Prefrontal Cortex: No Simple Matter. NeuroImage 11 (5), 447-450.
- Miller, E.K., Cohen, J.D., 2001. An Integrative Theory of Prefrontal Cortex Function. Ann. Rev. Neurosci. 1 (1), 167–202.
- Mischel, W., 1968. Personality and Assessment. Wiley, New York.
- Mischel, W., 2004. Toward an Integrative Science of the Person. Ann. Rev. Psychol. 55: 1-22.
- Mischel, W., 2009. From Personality and Assessment (1968) to Personality Science, 2009. J. Res. Personal. 43 (2), 282–290.
- Mischel, W., Ayduk, O., Berman, M.G., Casey, B.J., Gotlib, I.H., Jonides, J., Kross, E., Teslovich, T., Wilson, N.L., Zayas, V., Shoda, Y., 2010. 'Willpower' over the Life Span: Decomposing Self-Regulation. Soc. Cogn. Affect. NeuroSci. http://scan.oxfordjournals.org/content/early/2010/09/18/scan.nsq081.full.pdf+html.
- Mischel, W., Metzner, R., 1962. Preference for Delayed Reward as a Function of Age, Intelligence, and Length of Delay Interval. J. Abnorm. Soc. Psychol. 64 (6), 425–431.
- Mischel, W., Shoda, Y., 1995. A Cognitive-Affective System Theory of Personality: Reconceptualizing Situations, Dispositions, Dynamics, and Invariance in Personality Structure. Psychol. Rev. 102 (2), 246–268.
- Mischel, W., Shoda, Y., 2008. Toward a Unifying Theory of Personality: Integrating Dispositions and Processing Dynamics within the Cognitive-Affective Processing System. In: John, O.P., Robins, R.W., Pervin, L.A., (Ed.), Handbook of Personality Psychology. Guilford Press, New York, pp. 209–241.
- Mischel, W., Shoda, Y., Rodriguez, M.L., 1989. Delay of Gratification in Children. Science 244 (4907), 933–938.
- Miyake, A., Friedman, N.P., Emerson, M.J., Witzki, A.H., Howerter, A., 2000. The Unity and Diversity of Executive Functions and Their Contributions to Complex "Frontal Lobe" Tasks: A Latent Variable Analysis. Cogn. Psychol. 41 (1), 49–100.
- Möbius, M.M., Niederle, M., Niehaus, P., Rosenblat, T.S., 2010. Managing Self-Confidence: Theory and Experimental Evidence. Unpublished manuscript. Department of Economics, Iowa State University.
- Moffitt, T.E., Caspi, A., Rutter, M., 2005. Strategy for Investigating Interactions between Measured Genes and Measured Environments. Arch. Gen. Psychiatry 62 (5), 473–481.
- Monterosso, J.R., Luo, S., 2010. An Argument against Dual Valuation System Competition: Cognitive Capacities Supporting Future Orientation Mediate Rather Than Compete with Visceral Motivations. J. Neurosci. Psychol. Econom. 3 (1), 1–14.
- Morgeson, F.P., Campion, M.A., Dipboye, R.L., Hollenbeck, J.R., Murphy, K., Schmitt, N., 2007. Reconsidering the Use of Personality Tests in Personnel Selection Contexts. Pers. Psychol. 60 (3), 683–729.
- Moskowitz, D.S., 1982. Coherence and Cross-Situational Generality in Personality: A New Analysis of Old Problems. J. Personal. Soc. Psychol. 43 (4), 754–768.
- Moutafi, J., Furnham, A., Tsaousis, I., 2006. Is the Relationship between Intelligence and Trait Neuroticism Mediated by Test Anxiety? Pers. Individ. Dif. 40 (3), 587.
- Mroczek, D.K., Spiro, A., 2007. Personality Change Influences Mortality in Older Men. Psychol. Sci. 18 (5), 371–376.

- Mueller, G., Plug, E., 2006. Estimating the Effect of Personality on Male and Female Earnings. Ind. Labor Relat. Rev. 60 (1), 3–22.
- Mueser, P.R., 1979. The Effects of Noncognitive Traits. In: Jencks, C., Bartlett, S., Corcoran, M., Crouse, J., Eaglesfield, D., Jackson, G., McClelland, K., Mueser, P., Olneck, M., Schwartz, J., Ward, S., Williams, J., (Ed.), Who Gets Ahead? The Determinants of Economic Success in America. Basic Books, New York, pp. 122–158.
- Mundlak, Y., 1978. On the Pooling of Time Series and Cross Section Data. Econometrica 46 (1), 69-85.
- Murray, H.A., 1938. Explorations in Personality: A Clinical and Experimental Study of Fifty Men of College Age. Oxford University Press, New York.
- Neal, D.A., Johnson, W.R., 1996. The Role of Premarket Factors in Black-White Wage Differences. J. Polit. Econ. 104 (5), 869–895.
- Neisser, U., Boodoo, G., Bouchard, Jr., T.J., Boykin, A.W., Brody, N., Ceci, S.J., Halpern, D.F., Loehlin, J.C., Perloff, R., Sternberg, R.J., Urbina, S., 1996. Intelligence: Knowns and Unknowns. Am. Psychol. 51 (2), 77–101.
- Newman, D.L., Caspi, A., Moffitt, T.E., Silva, P.A., 1997. Antecedents of Adult Interpersonal Functioning: Effects of Individual Differences in Age 3 Temperament. Dev. Psychol. 33 (2), 206–217.
- Noftle, E.E., Robins, R.W., 2007. Personality Predictors of Academic Outcomes: Big Five Correlates of GPA and SAT Scores. J. Personal. Soc. Psychol. 93 (1), 116–130.
- Nyhus, E.K., Pons, E., 2005. The Effects of Personality on Earnings. J. Econom. Psychol. 26 (3), 363-384.
- O'Gorman, J.G., Baxter, E., 2002. Self-Control as a Personality Measure. Pers. Individ. Dif. 32 (3), 533–539.
- O'Leary, U.-M., Rusch, K.M., Guastello, S.J., 1991. Estimating Age-Stratified WAIS-R IQs from Scores on the Raven's Standard Progressive Matrices. J. Clin. Psychol. 47 (2), 277–284.
- Oliver, P.H., Guerin, D.W., Gottfried, A.W., 2007. Temperamental Task Orientation: Relation to High School and College Educational Accomplishments. Learn. Individ. Differ. 17 (3), 220–230.
- Omizo, M.M., 1980. The Differential Aptitude Tests as Predictors of Success in a High School for Engineering Program. Educ. Psychol. Meas. 40 (1), 197–203.
- Ones, D.S., Viswesvaran, C., 1998. The Effects of Social Desirability and Faking on Personality and Integrity Assessment for Personnel Selection. Hum. Perf. 11 (2–3):245–269.
- Osiński, J., 2009. Kin Altruism, Reciprocal Altruism and Social Discounting. Pers. Individ. Dif. 47 (4), 374–378.
- Pailing, P.E., Segalowitz, S.J., 2004. The Error-Related Negativity as a State and Trait Measure: Motivation, Personality, and Erps in Response to Errors. Psychophysiology 41 (1), 84–95.
- Parco, J.E., Rapoport, A., Stein, W.E., 2002. Effects of Financial Incentives on the Breakdown of Mutual Trust. Psychol. Sci. 13 (3), 292–297.
- Paulhus, D.L., 1984. Two-Component Models of Socially Desirable Responding. J. Personal. Soc. Psychol. 46 (3), 598–609.
- Paunonen, S.V., Ashton, M.C., 2001. Big Five Factors and Facets and the Prediction of Behavior. J. Personal. Soc. Psychol. 81 (3), 524–539.
- Perkins, D.N., Tishman, S., 2001. Dispositional Aspects of Intelligence. In: Collis, J.M., Messick, S., (Ed.), Intelligence and Personality: Bridging the Gap in Theory and Measurement. Lawrence Erlbaum Associates, Publishers, Mahwah, NJ, pp. 233–257.
- Pesonen, A.-K., Räikkönen, K., Heinonen, K., Andersson, S., Hovi, P., Järvenpää, A.-L., Eriksson, J.G., Kajantie, E., 2008. Personality of Young Adults Born Prematurely: Helsinki Study of Very Low Birth Weight Adults. J. Child Psychol. Psychiatry 49 (6), 609–617.
- Peters, E., Vastfjall, D., Slovic, P., Mertz, C.K., Mazzocco, K., Dickert, S., 2006. Numeracy and Decision Making. Psychol. Sci. 17 (5), 407–413.
- Phelps, E.A., 2006. Emotion and Cognition: Insights from Studies of the Human Amygdala. Ann. Rev. Psychol. 57, 27–53.
- Phelps, E.S., Pollak, R.A., 1968. On Second-Best National Saving and Game-Equilibrium Growth. Rev. Econom. Stud. 35 (2), 185–199.

- Piatek, R., Pinger, P., 2010. Maintaining (Locus of) Control? Assessing the Impact of Locus of Control on Education Decisions and Wages. Institute for the Study of Labor (IZA), Discussion Paper No. 5289.
- Pinker, S., 2002. The Blank Slate: The Modern Denial of Human Nature. Penguin Books, New York.
- Plomin, R., DeFries, J.C., Craig, I.W., McGuffin, P., 2002. Behavioral Genetics in the Postgenomic Era. American Psychological Association, Washington, DC.
- Plomin, R., Nesselroade, J.R., 1990. Behavioral Genetics and Personality Change. J. Personal. 58 (1), 191–220.
- Ponitz, C.C., McClelland, M.M., Jewkes, A.M., Connor, C.M., Farris, C.L., Morrison, F.J., 2008. Touch Your Toes! Developing a Direct Measure of Behavioral Regulation in Early Childhood. Early Child. Res. Q. 23 (2), 141–158.
- Ponitz, C.C., McClelland, M.M., Matthews, J.S., Morrison, F.J., 2009. A Structured Observation of Behavioral Self-Regulation and Its Contribution to Kindergarten Outcomes. Dev. Psychol. 45 (3), 605–619.
- Poropat, A.E., 2009. A Meta-Analysis of the Five-Factor Model of Personality and Academic Performance. Psychol. Bull. 135 (2), 322–338.
- Posner, M.I., Rothbart, M.K., 2000. Developing Mechanisms of Self-Regulation. Dev. Psychopathol. 12 (3), 427–441.
- Project Head Start, 1969. The Impact of Head Start: An Evaluation of the Effects of Head Start on Children's Cognitive and Affective Development. Westinghouse Learning Corporation, Bladensburg, MD.
- Putnam, S.P., Ellis, L.K., Rothbart, M.K., 2001. The Structure of Temperament from Infancy through Adolescence. In: Eliasz, A., Angleitner, A., (Ed.), Advances in Research on Temperament. Pabst Science Publishers, Lengerich, Germany, pp. 165–182.
- Rabin, M., 2000. Diminishing Marginal Utility of Wealth Cannot Explain Risk Aversion. In: Kahneman, D., Tversky, A., (Ed.), Choices, Values, and Frames. Cambridge University Press, New York, pp. 202–208.
- Raudenbush, S.W., Bryk, A.S., 2001. Hierarchical Linear Models: Applications and Data Analysis Methods (Advanced Quantitative Techniques in the Social Sciences). Sage Publications, Inc, Thousand Oaks, CA.
- Raven, J.C., 1962. Advanced Progressive Matrices: Sets I and II. Lewis, London.
- Raven, J., Raven, J.C., Court, J.H., 1988. Manual for Raven's Progressive Matrices and Vocabulary Scales. Harcourt Assessment, San Antonio, TX.
- Raver, C.C., Zigler, E.F., 1997. Social Competence: An Untapped Dimension in Evaluating Head Start's Success. Early Child. Res. Q. 12 (4), 363–385.
- Read, D., Read, N.L., 2004. Time Discounting over the Lifespan. Organ. Behav. Hum. Decis. Process. 94 (1), 22–32.
- Revelle, W., 1997. Extraversion and Impulsivity: The Lost Dimension? In: Nyborg, H., (Ed.), The Scientific Study of Human Nature: Tribute to Hans J. Eysenck at Eighty. Pergamon/Elsevier Science, Amsterdam, p. 189.
- Revelle, W., Wilt, J., Condon, D.M., 2011. Individual Differences and Differential Psychology: A Brief History and Prospect. In: Chamorro-Premuzic, T., Furnham, A., von Stumm, S., (Ed.), The Wiley-Blackwell Handbook of Individual Differences, Wiley-Blackwell, Hoboken, NJ.
- Richard, F.D., Bond, C.F., Stokes-Zoota, J.J., 2003. One Hundred Years of Social Psychology Quantitatively Described. Rev. Gen. Psychol. 7 (4), 331–363.
- Rindermann, H., 2007. The g-Factor of International Cognitive Ability Comparisons: The Homogeneity of Results in PISA, TIMSS, PIRLS and IQ-Tests across Nations. Eur. J. Personal. 21 (5), 667.
- Roberts, B.W., 1997. Plaster or Plasticity: Are Adult Work Experiences Associated with Personality Change in Women? J. Personal. 65 (2), 205–232.
- Roberts, B.W., 2006. Personality Development and Organizational Behavior. In: Staw, B., 27, (Ed.), Research on Organizational Behavior. Elsevier Science/JAI Press, Oxford, 1–41.
- Roberts, B.W., 2007. Contextualizing Personality Psychology. J. Personal. 75 (6), 1071-1082.
- Roberts, B.W., 2009. Back to the Future: *Personality and Assessment* and Personality Development. J. Res. Personal. 43 (2), 137–145.

- Roberts, B.W., Chapman, C.N., 2000. Change in Dispositional Well-Being and Its Relation to Role Quality: A 30-Year Longitudinal Study. J. Res. Personal. 34 (1), 26–41.
- Roberts, B.W., Chernyshenko, O.S., Stark, S., Goldberg, L.R., 2005. The Structure of Conscientiousness: An Empirical Investigation Based on Seven Major Personality Questionnaires. Pers. Psychol. 58 (1), 103–139.
- Roberts, B.W., DelVecchio, W.F., 2000. The Rank-Order Consistency of Personality Traits from Childhood to Old Age: A Quantitative Review of Longitudinal Studies. Psychol. Bull. 126 (1), 3–25.
- Roberts, R.D., Goff, G.N., Anjoul, F., Kyllonen, P.C., Pallier, G., Stankov, L., 2000. The Armed Services Vocational Aptitude Battery (ASVAB): Little More Than Acculturated Learning (Gc)!? Learn. Individ. Differ. 12 (1), 81–103.
- Roberts, B.W., Harms, P., Smith, J.L., Wood, D., Webb, M., 2006. Using Multiple Methods in Personality Psychology. In: Eid, M., Diener, E., (Ed.), Handbook of Multimethod Measurement in Psychology. American Psychological Association, Washington, DC, pp. 321–335.
- Roberts, B.W., Helson, R., Klohnen, E.C., 2002. Personality Development and Growth in Women across 30 Years: Three Perspectives. J. Personal. 70 (1), 79–102.
- Roberts, B.W., Jackson, J.J., 2008. Sociogenomic Personality Psychology. J. Personal. 76 (6), 1523–1544.
- Roberts, B.W., Kuncel, N.R., Shiner, R., Caspi, A., Goldberg, L.R., 2007. The Power of Personality: The Comparative Validity of Personality Traits, Socioeconomic Status, and Cognitive Ability for Predicting Important Life Outcomes. Perspect. Psychol. Sci. 2 (4), 313–345.
- Roberts, R.D., Markham, P.M., Matthews, G., Zeidner, M., 2005. Assessing Intelligence: Past, Present, and Future. In: Wilhelm, O., Engle, R.W., (Ed.), Handbook of Understanding and Measuring Intelligence. Sage Publications Inc., Thousand Oaks, CA, pp. 333–360.
- Roberts, B.W., Mroczek, D., 2008. Personality Trait Change in Adulthood. Curr. Dir. Psychol. Sci. 17 (1), 31–35.
- Roberts, B.W., Walton, K.E., Viechtbauer, W., 2006. Patterns of Mean-Level Change in Personality Traits across the Life Course: A Meta-Analysis of Longitudinal Studies. Psychol. Bull. 132 (1), 1–25.
- Roberts, B.W., Wood, D., Caspi, A., 2008. Personality Development. In: John, O.P., Robins, R.W., Pervin, L.A., (Ed.), Handbook of Personality: Theory and Research. Guilford Press, New York, pp. 375–398.
- Robson, A.J., 1996. A Biological Basis for Expected and Non-Expected Utility. J. Econom. Theory 68, 397–424.
- Robson, A.J., 2001. Why Would Nature Give Individuals Utility Functions? J. Polit. Econ. 109 (4), 900–914.
- Robson, A.J., Samuelson, L., 2007. The Evolution of Intertemporal Preferences. Am. Econ. Rev. 97 (2), 496–500.
- Robson, A.J., Samuelson, L., 2009. The Evolution of Time Preference with Aggregate Uncertainty. Am. Econ. Rev. 99 (5), 1925–1953.
- Rodríguez-Planas, N., 2010. Mentoring, Educational Services, and Economic Incentives: Longer-Term Evidence on Risky Behaviors from a Randomized Trial. IZA Discussion Paper No. 4968.
- Rosenberg, M., 1965. Society and the Adolescent Self-Image. Princeton University Press, Princeton, NJ.
- Rosenberg, M., 1989. Society and the Adolescent Self-Image (Rev. Ed.). Wesleyan University Press, Middletown, CT.
- Ross, L., Nisbett, R.E., 1991. The Person and the Situation: Perspectives of Social Psychology. McGraw-Hill, New York.
- Rothbart, M.K., Ahadi, S.A., Evans, D.E., 2000. Temperament and Personality: Origins and Outcomes. J. Personal. Soc. Psychol. 78 (1), 122–135.
- Rothlisberg, B., 1987. Comparing the Stanford-Binet, Fourth Edition to the WISC-R: A Concurrent Validity Study. J. Sch. Psychol. 25 (2), 193–196.
- Rotter, J.B., 1966. Generalized Expectancies for Internal Versus External Control of Reinforcement. Psychol. Monogr. 80 (1), 1–28.
- Roy, A.D., 1951. Some Thoughts on the Distribution of Earnings. Oxf. Econ. Pap. 3 (2), 135-146.

- Rueda, M.R., Rothbart, M.K., McCandliss, B.D., Saccomanno, L., Posner, M.I., 2005. Training, Maturation, and Genetic Influences on the Development of Executive Attention. Proc. Natl. Acad. Sci. 102 (41), 14931–14936.
- Rust, J., 2008. Dynamic Programming. In: Durlauf, S.N., Blume, L.E. (Ed.), The New Palgrave Dictionary of Economics. Palgrave Macmillan, Ltd, London.
- Rutter, M., 2006a. Genes and Behavior: Nature-Nurture Interplay Explained. Blackwell Publishers, Oxford.
- Rutter, M., 2006b. Implications of Resilience Concepts for Scientific Understanding. Ann. N.Y. Acad. Sci. 1094 (1), 1–12.
- Ryden, A., Sullivan, M., Torgerson, J.S., Karlsson, J., Lindroos, A.K., Taft, C., 2003. Severe Obesity and Personality: A Comparative Controlled Study of Personality Traits. Int. J. Obes. 27 (12), 1534–1540.
- Sackett, P.R., Borneman, M.J., Connelly, B.S., 2008. High Stakes Testing in Higher Education and Employment: Appraising the Evidence for Validity and Fairness. Am. Psychol. 63 (4), 215–227.
- Sackett, P.R., Hardison, C.M., Cullen, M.J., 2004. On Interpreting Stereotype Threat as Accounting for African American-White Differences on Cognitive Tests. Am. Psychol. 59 (1), 7–13.
- Safire, W., 1989. On Language; the Elision Fields. N.Y. Times 13 August.
- Sahm, C.R., 2007. Does Risk Tolerance Change? Unpublished manuscript, University of Michigan Institute for Social Research.
- Salgado, J.F., 1997. The Five Factor Model of Personality and Job Performance in the European Community. J. Appl. Psychol. 82 (1), 30–43.
- Samuel, D.B., Widiger, T.A., 2008. A Meta-Analytic Review of the Relationships between the Five-Factor Model, DSM-IV-Tr Personality Disorders: A Facet Level Analysis. Clin. Psychol. Rev. 28 (8), 1326–1342.
- Samuelson, P.A., 1956. Social Indifference Curves. Q. J. Econom. 70 (1), 1-22.
- Saucier, G., 2003. An Alternative Multi-Language Structure for Personality Attributes. Eur. J. Pers. 17 (3), 179–205.
- Savage, L.J., 1954. The Foundations of Statistics. Wiley, New York.
- Savelyev, P.A., 2010. Conscientiousness, Education, and Longevity of High-Ability Individuals. Unpublished manuscript. Department of Economics, University of Chicago.
- Scarr, S., Weinberg, R.A., Waldman, I.D., 1993. IQ Correlations in Transracial Adoptive Families. Intelligence 17 (4), 541–555.
- Schaie, K.W., 1994. The Course of Adult Intellectual Development. Am. Psychol. 49 (4), 304-313.
- Schmeichel, B.J., Vohs, K.D., Baumeister, R.F., 2003. Intellectual Performance and Ego Depletion: Role of the Self in Logical Reasoning and Other Information Processing. J. Personal. Soc. Psychol. 85 (1), 33–46.
- Schmidt, F.L., Hunter, J.E., 1983. Individual Differences in Productivity: An Empirical Test of Estimates Derived from Studies of Selection Procedure Utility. J. Appl. Psychol. 68 (3), 407–414.
- Schmidt, F.L., Hunter, J.E., 1998. The Validity and Utility of Selection Methods in Personnel Psychology: Practical and Theoretical Implications of 85 Years of Research Findings. Psychol. Bull. 124 (2), 262–274.
- Schmidt, F.L., Hunter, J., 2004. General Mental Ability in the World of Work: Occupational Attainment and Job Performance. J. Personal. Soc. Psychol. 86 (1), 162–173.
- Schmidt, F.L., Hunter, J.E., Larson, M., 1988. General Cognitive Ability Versus General and Specific Aptitudes in the Prediction of Training Performance: Some Preliminary Findings. Paper presented at the 96th Annual Convention of the American Psychological Association, Atlanta, GA.
- Schooler, C., Mulatu, M.S., Oates, G., 1999. The Continuing Effects of Substantively Complex Work on the Intellectual Functioning of Older Workers. Psychol. Aging 14 (3), 483–506.
- Schuerger, J.M., Witt, A.C., 1989. The Temporal Stability of Individually Tested Intelligence. J. Clin. Psychol. 45 (2), 294–302.
- Schulz, R., Bookwala, J., Knapp, J.E., Scheier, M., Williamson, G.M., 1996. Pessimism, Age, and Cancer Mortality. Psychol. Aging 11, 304–309.
- Schumpeter, J.A., 1954. History of Economic Analysis. Oxford University Press, New York.

- Schwartz, S.H., 1992. The Universal Content and Structure of Values: Theoretical Advances and Empirical Tests in 20 Countries. In: Zanna, M.P., (Ed.), Advances in Experimental Social Psychology, Academic Press, Inc., San Diego, CA. pp. 1–65.
- Schweinhart, L.J., Montie, J., Xiang, Z., Barnett, W.S., Belfield, C.R., Nores, M., 2005. Lifetime Effects: The High/Scope Perry Preschool Study through Age 40. High/Scope Press, Ypsilanti, MI.
- Segal, C., 2008. Motivation, Test Scores, and Economic Success. Department of Economics and Business, Universitat Pompeu Fabra, Working Paper No. 1124.
- Sell, A., Tooby, J., Cosmides, L., 2009. Formidability and the Logic of Human Anger. Proc. Natl. Acad. Sci. 106 (35), 15073–15078.
- Shamosh, N., Gray, J.R., 2007. Delay Discounting and Intelligence: A Meta-Analysis. Department of Psychology, Yale University, Unpublished manuscript.
- Shiner, R.L., 1998. How Shall We Speak of Children's Personalities in Middle Childhood? A Preliminary Taxonomy. Psychol. Bull. 124 (3), 308–332.
- Shiner, R., Caspi, A., 2003. Personality Differences in Childhood and Adolescence: Measurement, Development, and Consequences. J. Child Psychol. Psychiatry 44 (1), 2–32.
- Siegler, R.S., 1992. The Other Alfred Binet. Dev. Psychol. 28 (2), 179.
- Sijtsma, K., 2009. On the Use, the Misuse, and the Very Limited Usefulness of Cronbach's Alpha. Psychometrika 74 (1), 107–120.
- Skiadas, C., 1998. Recursive Utility and Preferences for Information. Econom. Theory 12 (2), 293-312.
- Slonim, R., Roth, A.E., 1998. Learning in High Stakes Ultimatum Games: An Experiment in the Slovak Republic. Econometrica 66 (3), 569–596.
- Slovic, P., 1962. Convergent Validation of Risk Taking Measures. J. Abnorm. Soc. Psychol. 65 (1), 68-71.
- Slovic, P., 1966. Risk-Taking in Children: Age and Sex Differences. Child Dev. 37 (1), 169–176.
- Slovic, P., 1972. Psychological Study of Human Judgment: Implications for Investment Decision Making. J. Finance 27 (4), 779–799.
- Smith, M.L., Glass, G.V., Miller, T.I., 1980. The Benefits of Psychotherapy. Johns Hopkins University Press, Baltimore.
- Social and Character Development Research Consortium, 2010. Efficacy of Schoolwide Programs to Promote Social and Character Development and Reduce Problem Behavior in Elementary School Children (NCER 2011-2001). National Center for Education Research, Institute of Education Sciences, U.S. Department of Education, Washington, DC.
- Spear, L.P., 2000a. The Adolescent Brain and Age-Related Behavioral Manifestations. Neurosci. Biobehav. Rev. 24 (4), 417.
- Spear, L.P., 2000b. Neurobehavioral Changes in Adolescence. Curr. Dir. Psychol. Sci. 9 (4), 111.
- Spearman, C., 1904. "General Intelligence," Objectively Determined and Measured. Am. J. Psychol. 15 (2), 201–293.
- Spearman, C., 1927. The Abilities of Man; Their Nature and Measurement. The Macmillan Company, New York.
- Stankov, L., 2005. g Factor: Issues of Design and Interpretation. In: Wilhelm, O., Engle, R.W., (Ed.), Handbook of Understanding and Measuring Intelligence. Sage Publications Inc., Thousand Oaks, CA, 279–293.
- Stanovich, K.E., 1999. Who Is Rational?: Studies of Individual Differences in Reasoning. Lawrence Erlbaum Associates, Mahwah, NJ.
- Starmer, C., 2000. Developments in Non-Expected Utility Theory: The Hunt for a Descriptive Theory of Choice under Risk. J. Econom. Lit. 38 (2), 332–382.
- Steele, C.M., Aronson, J., 1998. Stereotype Threat and the Test Performance of Academically Successful African Americans. In: Jencks, C., Phillips, M., (Ed.), The Black-White Test Score Gap. Brookings Institution Press, Washington, DC, pp. 401–427.
- Steinberg, L., 2008. A Social Neuroscience Perspective on Adolescent Risk-Taking. Dev. Rev. 28 (1), 78–106.
- Steinberg, L., Cauffman, E., Woolard, J., Graham, S., Banich, M., 2009. Are Adolescents Less Mature Than Adults? Minors' Access to Abortion, the Juvenile Death Penalty, and the Alleged APA 'Flip-Flop'. Am. Psychol. 64 (7), 583–594.

- Sternberg, R.J., 2001. Successful Intelligence: Understanding What Spearman Had Rather Than What He Studied. In: Collis, J.M., Messick, S., Schiefele, U., (Ed.), Intelligence and Personality: Bridging the Gap in Theory and Measurement. Lawrence Erlbaum Associates, Mahwah, NJ, pp. 347–373.
- Sternberg, R.J., Forsythe, G.B., Hedlund, J., Horvath, J.A., Wagner, R.K., Williams, W.M., Snook, S.A., Grigorenko, E., 2000. Practical Intelligence in Everyday Life. Cambridge University Press, New York.
- Stevens, C., Fanning, J., Coch, D., Sanders, L., Neville, H., 2008. Neural Mechanisms of Selective Auditory Attention Are Enhanced by Computerized Training: Electrophysiological Evidence from Language-Impaired and Typically Developing Children. Brain Res. 1205, 55–69.
- Sticht, T.G., 1995. The Military Experience and Workplace Literacy: A Review and Synthesis for Policy and Practice. National Center on Adult Literacy, University of Pennsylvania Working paper, TR94-01.
- Sticht, T.G., Hooke, L.R., Caylor, J.S., 1982. Literacy, Oracy, and Vocational Aptitude as Predictors of Attrition and Promotion in the Armed Services. Human Resources Research Organization, Washington, DC.
- Störmer, S., Fahr, R., 2010. Individual Determinants of Work Attendance: Evidence on the Role of Personality. IZA Discussion Paper No. 4927.
- Sutter, M., Feri, F., Kocher, M.G., Martinsson, P., Nordblom, K., Rützler, D., 2010. Social Preferences in Childhood and Adolescence: A Large-Scale Experiment. IZA Discussion Paper No. 5016.
- Sylva, K., 1997. The Quest for Quality in Curriculum. In: Schweinhart, L.J., Weikart, D.P., (Ed.), Lasting Differences: The High/Scope Preschool Curriculum Comparison Study through Age 23. High/Scope Press, Ypsilanti, pp. 89–93.
- Tang, T.Z., DeRubeis, R.J., Hollon, S.D., Amsterdam, J., Shelton, R., Schalet, B., 2009. Personality Change During Depression Treatment: A Placebo-Controlled Trial. Arch. Gen. Psychiatry 66 (12), 1322–1330.
- Tausch, F., Potters, J., Riedl, A., 2010. Preferences for Redistribution and Pensions: What Can We Learn from Experiments? CESifo Working Paper No. 3156.
- Tellegen, A., 1985. Structures of Mood and Personality and Their Relevance to Assessing Anxiety, with an Emphasis on Self-Report. In: Tuma, A.H., Maser, J.D., (Ed.), Anxiety and the Anxiety Disorders. Lawrence Erlbaum Associates, Hillsdale, NJ, pp. 681–706.
- Terman, L.M., 1916. The Measurement of Intelligence: An Explanation of and a Complete Guide for the Use of the Stanford Revision and Extension of the Binet-Simon Intelligence Scale. Houghton Mifflin Co, Boston.
- Terman, L.M., Merrill, M.A., 1960. Stanford-Binet Intelligence Scale: Manual for the Third Revision Form, L-M. Houghton Mifflin, Boston.
- Thaler, R.H., 1980. Toward a Positive Theory of Consumer Choice. J. Econ. Behav. Organ. 1 (1), 39-60.
- Thaler, R.H., 2008, July 25–27. A Short Course in Behavioral Economics. Edge Master Class 2008 Retrieved November 4, 2010, http://www.edge.org/3rd_culture/thaler_sendhil08/class6.html.
- Turkheimer, E., Haley, A., Waldron, M., D'Onofrio, B., Gottesman, I.I., 2003. Socioeconomic Status Modifies Heritability of IQ in Young Children. Psychol. Sci. 14 (6), 623–628.
- Uylings, H.B.M., 2006. Development of the Human Cortex and the Concept Of "Critical" Or "Sensitive" Periods. Lang. Learn. 56 (Suppl 1), 59–90.
- Valiente, C., Lemery-Chalfant, K., Castro, K.S., 2007. Children's Effortful Control and Academic Competence. Merrill Palmer Q. 53 (1), 1–25.
- Valiente, C., Lemery-Chalfant, K., Swanson, J., 2010. Prediction of Kindergartners' Academic Achievement from Their Effortful Control and Emotionality: Evidence for Direct and Moderated Relations. J. Educ. Psychol. 102 (3), 550–560.
- Valiente, C., Lemery-Chalfant, K., Swanson, J., Reiser, M., 2008. Prediction of Children's Academic Competence from Their Effortful Control, Relationships, and Classroom Participation. J. Educ. Psychol. 100 (1), 67–77.
- van Eijck, K., de Graaf, P.M., 2004. The Big Five at School: The Impact of Personality on Educational Attainment. Neth. J. Soc. Sci. 40 (1), 24–40.
- van Praag, B.M.S., 1985. Linking Economics with Psychology: An Economist's View. J. Econom. Psychol. 6 (3), 289–311.

- van Praag, B.M.S., van Weeren, J., 1988. Memory and Anticipation Processes and Their Significance for Social Security and Income Inequality. In: Maital, S., (Ed.), Applied Behavioral Economics Volume 1. New York University Press, New York, pp. 731–751.
- Vaughn, B.E., Kopp, C.B., Krakow, J.B., 1984. The Emergence and Consolidation of Self-Control from Eighteen to Thirty Months of Age: Normative Trends and Individual Differences. Child Dev. 55 (3), 990–1004.
- Vazsonyi, A.T., Pickering, L.E., Junger, M., Hessing, D., 2001. An Empirical Test of a General Theory of Crime: A Four-Nation Comparative Study of Self-Control and the Prediction of Deviance. J. Res. Crime Deling. 38, 91–131.
- Viswesvaran, C., Ones, D.S., 1999. Meta-Analyses of Fakability Estimates: Implications for Personality Measurement. Educ. Psychol. Meas. 59 (2), 197–210.
- Vitaro, F., Brendgen, M., Larose, S., Tremblay, R.E., 2005. Kindergarten Disruptive Behaviors, Protective Factors, and Educational Achievement by Early Adulthood. J. Educ. Psychol. 97 (4), 617–629.
- Wang, L., 1993. The Differential Aptitude Test: A Review and Critique. Annual Meeting of the Southwest Educational Research Association, Austin, TX.
- Watson, D., Clark, L.A., Chmielewski, M., 2008. Structures of Personality and Their Relevance to Psychopathology: II. Further Articulation of a Comprehensive Unified Trait Structure. Journal of Personality, 76 (6), 1545–1586.
- Webb, E., 1915. Character and Intelligence. Br. J. Psychol. Suppl. 1 (3), 1-99.
- Weber, E.U., 2001. Personality and Risk Taking. In: Smelser, N.J., Baltes, P.B., (Ed.), International Encyclopedia of the Social and Behavioral Sciences. Pergamon, Oxford, p. 11274.
- Weber, E.U., Blais, A.-R., Betz, N.E., 2002. A Domain-Specific Risk-Attitude Scale: Measuring Risk Perceptions and Risk Behaviors. J. Behav. Decis. Mak. 15 (4), 263–290.
- Weber, E.U., Hsee, C., 1998. Cross-Cultural Differences in Risk Perception, but Cross-Cultural Similarities in Attitudes Towards Perceived Risk. Manag. Sci. 44 (9), 1205.
- Weber, E.U., Milliman, R.A., 1997. Perceived Risk Attitudes: Relating Risk Perception to Risky Choice. Manag. Sci. 43 (2), 123.
- Wechsler, D., 1939. The Measurement of Adult Intelligence. Williams & Wilkins Co., Baltimore, MD.
- Weikart, D.P., Epstein, A.S., Schweinhart, L., Bond, J.T., 1978. The Ypsilanti Preschool Curriculum Demonstration Project: Preschool Years and Longitudinal Results. In: Monographs of the High/Scope Educational Research Foundation. High/Scope Press, Ypsilanti, MI, p. 4.
- Whisenton, J.T., Lorre, M.R., 1970. A Comparison of the Values, Needs, and Aspirations of School Leavers with Those of Non-School Leavers. J. Negro Educ. 39 (4), 325–332.
- Widiger, T.A., Costa, Jr., P.T., 2002. Five-Factor Model Personality Disorder Research. In: Costa, Jr., P.T., Widiger, T.A., (Ed.), Personality Disorders and the Five-Factor Model of Personality. American Psychological Association, Washington, DC.
- Widiger, T.A., Trull, T.J., Clarkin, J.F., Sanderson, C., Costa, Jr., P.T., 2002. A Description of the DSM-IV Personality Disorders with the Five-Factor Model of Personality. In: Costa, Jr., P.T., Widiger, T.A. (Ed.), Personality Disorders and the Five-Factor Model of Personality. American Psychological Association, Washington, DC.
- Wilhelm, O., Engle, R.W., 2005. Intelligence: A Diva and a Workhorse. In: Wilhelm, O., Engle, R.W. (Ed.), Handbook of Understanding and Measuring Intelligence. Sage Publications, Thousand Oaks, CA.
- Willingham, W.W., 1985. Success in College: The Role of Personal Qualities and Academic Ability. College Entrance Examination Board, New York.
- Willingham, W.W., Pollack, J.M., Lewis, C., 2002. Grades and Test Scores: Accounting for Observed Differences. J. Educ. Meas. 39 (1), 1–37.
- Willis, R.J., Rosen, S., 1979. Education and Self-Selection. J. Polit. Econ. 87 (5, Part 2), S7-S36.
- Winship, C., Korenman, S., 1997. Does Staying in School Make You Smarter? The Effect of Education on IQ in the Bell Curve. In: Devlin, B., Feinberg, S., Resnick, D., Roeder, K., (Ed.), Intelligence, Genes, and Success: Scientists Respond to the Bell Curve. Springer, Copernicus, New York, pp. 215–234.
- Wolfe, R.N., Johnson, S.D., 1995. Personality as a Predictor of College Performance. Educ. Psychol. Meas. 55 (2), 177–185.

- Wong, M.M., Csikszentmihalyi, M., 1991. Motivation and Academic Achievement: The Effects of Personality Traits and the Duality of Experience. J. Personal. 59 (3), 539–574.
- Wood, D., Roberts, B.W., 2006. Cross-Sectional and Longitudinal Tests of the Personality and Role Identity Structural Model (PRISM). J. Personal. 74 (3), 779–809.
- Young, J.W., Kobrin, J.L., 2001. Differential Validity, Differential Prediction, and College Admission Testing: A Comprehensive Review and Analysis. College Entrance Examination Board, New York, NY.
- Zentner, M., Bates, J.E., 2008. Child Temperament: An Integrative Review of Concepts, Research Programs, and Measures. Eur. J. Dev. Sci. 2 (1–2), 7–37.
- Zhou, Q., Main, A., Wang, Y., 2010. The Relations of Temperamental Effortful Control and Anger/Frustration to Chinese Children's Academic Achievement and Social Adjustment: A Longitudinal Study. J. Educ. Psychol. 102 (1), 180–196.
- Zigler, E., Butterfield, E.C., 1968. Motivational Aspects of Changes in IQ Test Performance of Culturally Deprived Nursery School Children. Child Dev. 39 (1), 1–14.
- Zuckerman, M., 1994. Behavioral Expressions and Biosocial Bases of Sensation Seeking. Cambridge University Press, New York.
- Zuckerman, M., 2007. Sensation Seeking and Risky Behavior. American Psychological Association, Washington, DC.
- Zuckerman, M., Kolin, E.A., Price, L., Zoob, I., 1964. Development of a Sensation-Seeking Scale. J. Consult. Psychol. 28 (6), 477–482.



Nonproduction Benefits of Education: Crime, Health, and Good Citizenship

Lance Lochner

University of Western Ontario

Contents

1.	Introduction	184
2.	Education's Effects on Crime	184
	2.1. The Economics of Education and Crime	186
	2.1.1 A Dynamic Model of Schooling, Work, and Crime	186
	2.1.2 Other Ways in Which Education May Affect Crime	192
	2.1.3 School Attendance and Contemporaneous Crime	192
	2.2. Evidence on Education and Crime	193
	2.2.1 Educational Attainment and Crime	193
	2.2.2 School Quality and Crime	202
	2.2.3 Contemporaneous Schooling and Crime	205
	2.3. Evidence on Human Capital-Based Interventions from Birth to Young Adulthood	207
	2.4. Policy Lessons	216
	2.4.1 Valuing the Social Benefits from Crime Reduction	216
	2.4.2 Where Are the Big Returns?	218
	2.4.3 Additional Policy Lessons	219
	2.5. Interpreting the Evidence on Crime	220
3.	Education's Effects on Health and Mortality	224
	3.1. A Simple Model of Education and Mortality	225
	3.1.1 Valuing Mortality Benefits from Education	228
	3.1.2 Allocative Efficiency	229
	3.2. Evidence on the Effects of Education on Health and Mortality	230
	3.2.1 Does Education Affect Health and Mortality?	230
	3.2.2 Why Does Education Affect Health and Mortality?	247
	3.3. Implications for Education and Health Policy	258
	3.4. Suggestions for Future Research on Education and Health	260
4.	Education's Effect on Citizenship, Political Participation, and Democracy	261
	4.1. Macro Evidence on Education and Democracy	262
	4.2. Estimating "Causal" Effects of Education on Citizenship	264
5.	Conclusions	269
Appendix: Estimating Mediating Effects		271
Acknowledgments		274
References		274

Abstract

A growing body of work suggests that education offers a wide range of benefits that extend beyond increases in labor market productivity. Improvements in education can lower crime, improve health, and increase voting and democratic participation. This chapter reviews recent developments on these "nonproduction" benefits of education.

Keywords

Education Crime Health Mortality Democracy Citizenship Social Benefits

1. INTRODUCTION

Economists have long recognized and measured the effect of education on an individual's own lifetime earnings.¹ More recently, attention has been paid to various "nonproduction" benefits of education.² A growing body of work suggests that education offers a wide range of personal and social benefits that extend beyond increases in labor market productivity. This chapter reviews recent developments on the impacts of education on criminal behavior, health and mortality, and voting and democratic participation.³ The primary focus is on new empirical studies (largely by economists); however, this chapter also discusses simple economic models that can be used to interpret the evidence. Important implications for education-based policies are also discussed.

2. EDUCATION'S EFFECTS ON CRIME

In 1997, over two-thirds of all prison inmates in the United States were high school dropouts (Harlow (2003)). Despite promising evidence that education-based policies and early childhood interventions can play an important role in helping to reduce

¹ See, for example, Card (1999) or Heckman, Lochner, and Todd (2006, 2008).

² Other studies attempted to estimate production externalities associated with education (see, e.g., Acemoglu and Angrist (2001); Heckman and Klenow (1997); Moretti (2004a, 2004b); and Ciccone and Peri (2006)). Lange and Topel (2006) offered a recent review of these studies.

³ Of course, education impacts individuals and society in other ways as well. For example, see McMahon (2009) for a recent attempt to value a more comprehensive range of benefits associated with education.

crime, evidence is still limited and sometimes mixed. The link between schooling and crime is more complicated than simple prison statistics suggest. This section reviews evidence in this rapidly growing area and develops a human capital-based theory for interpreting much of this evidence.

We first discuss the relationship between education and crime from an economic perspective, developing a simple life-cycle model that sheds light on key ways in which early childhood programs and policies that encourage schooling may affect both juvenile and adult crime. The model is grounded in human capital theory and paints with a broad brush. It emphasizes the choice individuals face between legitimate work and criminal activity, with its associated punishments. By altering the relative rewards of work and crime, educational investments affect decisions to engage in crime. While the model does not incorporate all avenues through which education may affect crime, it serves as a useful point of reference.

We next discuss evidence on the impacts of educational attainment and school quality/choice on adult crime. The evidence from studies of educational attainment on crime is largely consistent with a human capital-based theory of crime, suggesting that increases in schooling reduce most types of adult crime (e.g., Lochner (2004) and Lochner and Moretti (2004)). Studies of school choice and increases in school quality paint a more nuanced picture: sizeable improvements in school quality produce minor (at best) improvements in student achievement and educational attainment, while they appear to substantially reduce crime during late adolescence and early adulthood (Cullen, Jacob, and Levitt (2006) and Deming (2009a)). We will then discuss the contemporaneous relationship between school attendance and crime. Using exogenous policy changes and other events that effectively force students to stay in school or take extra days off (e.g., changes in compulsory schooling laws, teacher in-service days and strikes), a few recent studies have shown that school attendance affects crime in rich and complex ways.

We also review a number of recent studies that examine the long-run impacts of early childhood, school-based, and young adult training interventions on juvenile and adult crime. While a few early preschool programs have produced sizeable long-run reductions in crime—most famously, Perry Preschool—other quite similar programs have not. School-based programs focused on improving social development among "risky" children have been shown to reduce crime through early adulthood. Finally, job training for young adults (e.g., Job Corps) appears to reduce self-reported arrests and convictions during the period of intensive training, but it yields negligible lasting effects on crime. Altogether, the evidence suggests that reductions in crime can be achieved by a wide range of human capital-based intervention strategies.

Finally, we discuss a number of policy issues related to education and its potential role as a crime-fighting strategy.

2.1. The Economics of Education and Crime

Why does education reduce crime, and does its impact vary across different types of crimes? How might education and human capital policies help reduce crime? We begin with a simple economic model that formalizes a number of key channels through which education may affect crime. We then briefly discuss other factors that may help determine the relationship between education and crime.

2.1.1 A Dynamic Model of Schooling, Work, and Crime

We develop a dynamic time allocation model of crime, work, and education in which individuals are assumed to maximize lifetime earnings from work and crime. The model is based on Lochner (2004); however, it abstracts from random shocks to the returns from work and crime and simplifies certain life-cycle aspects of the problem to focus on the role of education. It emphasizes the role of education as a human capital investment that increases future legitimate work opportunities, which discourages participation in crime.⁴ This is consistent with numerous recent studies that show that higher wages reduce crime (e.g., Grogger (1998); Machin and Meghir (2004); and Gould, Mustard, and Weinberg (2002)) and decades of research in labor economics showing that education increases wage rates (see, e.g., Card (1999)).

We consider two stages of life: a "schooling stage" followed by a "work stage." During the schooling stage (ages $t = 1, ..., t_s$), individuals decide how much time to allocate to work, crime, and human capital investment at school. During the "work stage" ($t = t_s + 1, ...$), individuals only decide between work and crime. We assume that the "work stage" lasts forever.

Time spent at work $L_t \ge 0$ returns income wH_tL_t , where w reflects the rental rate on human capital and H_t reflects an individual's legitimate market human capital. Time spent on crime $k_t \ge 0$ offers a reward of $N(k_t, H_t, \theta)$, where $N(\cdot)$ is strictly increasing and concave in k_t , strictly increasing in θ , and weakly increasing in H_t . The parameter θ reflects an individual's criminal "ability," and it is assumed that $\frac{\partial^2 N}{\partial k \partial \theta} \ge 0$. While we use the term "criminal ability," θ should be interpreted broadly, including anything that affects the rewards from crime (e.g., local neighborhood conditions, the availability of good "targets"). Our assumptions allow for the fact that labor market skills may also be rewarded in the criminal sector, an issue we discuss further below. During the "schooling stage," time and effort spent in school $I_t \ge 0$ offer no current rewards, but it increases future market skill levels according to:

$$H_{t+1} = H_t + Af(I_t, H_t),$$
(2.1)

⁴ Fella and Gallipoli (2009) embedded a similar life-cycle model of crime in a general equilibrium environment.

where $f(\cdot)$ is increasing and concave in each of its arguments, and A > 0 reflects an individual's "learning ability." Individuals are endowed with an initial skill level H_0 . Although human capital grows during the "schooling stage" for those who spend some time acquiring skills through education (or training), it is assumed to remain constant during the "work stage." We normalize an individual's time endowment to one at all ages, so $L_t + k_t + I_t = 1$ during the "schooling stage" and $L_t + k_t = 1$ during the "work stage."

Individuals who engage in crime at any age may be caught and imprisoned for J years with probability πk_t . We assume that the "incarceration rate" $\pi \in (0, 1)$, so there is some probability that individuals escape punishment even if they spend all of their time committing crime. During prison, individuals earn nothing and their skills remain unchanged when they return to the labor market.⁵ For simplicity, we further assume that youth imprisoned during the "schooling stage" of life leave prison in the "work stage" of life.

In setting up the model, it is convenient to begin with the "work stage." Since this stage lasts forever, we drop the *t* subscripts and formulate the problem in recursive form. For a time discount rate of $\beta \in (0, 1)$, the value function for someone with human capital *H* who is in the "work stage" but not currently in prison is

$$W(H) = \max_{k \in [0,1]} \{ wH(1-k) + N(k, H, \theta) + \beta [\pi k \beta^J + (1-\pi k)] W(H) \},$$
(2.2)

which reflects the fact that the value associated with imprisonment for J years is simply $\beta^J W(H)$ since individuals leave prison in the same state they entered. The interior first-order condition for time spent committing crime is

$$wH + \pi\beta(1 - \beta^J)W(H) = \frac{\partial N(k, H, \theta)}{\partial k},$$
(2.3)

where the left-hand side reflects the marginal cost of spending more time on crime, while the right-hand side reflects the marginal return. Note that crime is costly in terms of current foregone earnings (*wH*) as well as expected losses due to time spent in prison (since $\beta^{J} < 1$).

⁵ Most estimates suggest that earnings and employment decline following an arrest or prison term; although, there has been some debate about the magnitude and duration of any effects (e.g., see Grogger (1995); Kling (2003); Nagin and Waldfogel (1995); and Waldfogel (1994)). This suggests that losses in earnings power due to stigma effects and/or human capital depreciation outweigh any gains from new skills that may be acquired in prison. We abstract from these features to focus on other issues.

188 Lance Lochner

Before discussing the effects of human capital or criminal ability on criminal choices, it is worth discussing the marginal value of human capital, W'(H). The envelope condition implies that

$$W'(H) = w(1-k) + \frac{\partial N(k, H, \theta)}{\partial H} + \beta [\pi k \beta^{J} + (1-\pi k)] W'(H)$$

$$= \frac{w(1-k) + \frac{\partial N(k, H, \theta)}{\partial H}}{1-\beta + \pi k \beta (1-\beta^{J})}$$

$$\geq 0,$$
(2.4)

where the final inequality is strict if k < 1 or $\frac{\partial N}{\partial H} > 0$. Since $\frac{w}{1-\beta}$ reflects the marginal value of human capital for someone who never commits any crime, we observe three ways in which crime affects the marginal value of human capital during adulthood: (1) by reducing labor market work, it lowers the payoff from human capital by wk; (2) the rewards from crime may be increasing in human capital $\left(\frac{\partial N}{\partial H} > 0\right)$, which increases the marginal value of skills in the criminal sector; and (3) the likelihood of arrest and imprisonment reduces the expected years human capital can be utilized in both the legitimate and criminal sectors. When human capital has negligible effects on criminal earnings (i.e., $\frac{\partial N}{\partial H} \approx 0$), engagement in crime will generally lower the marginal value of human capital due to (1) and (3).

Now, we can discuss the effects of human capital and criminal ability on crime. Equation (2.3) makes clear that crime is costly for individuals with higher levels of human capital. Totally differentiating this equation with respect to H yields

$$\frac{\mathrm{d}k}{\mathrm{d}H} = \frac{w - \frac{\partial^2 N}{\partial k \partial H} + \pi \beta (1 - \beta^J) W'(H)}{\frac{\partial^2 N}{\partial k^2}}$$

If human capital raises the marginal returns from work more than crime (i.e., $w \ge \frac{\partial^2 N}{\partial k \partial H}$), it is clear that human capital reduces crime since $W'(H) \ge 0$. Crime will be increasing in human capital (or schooling) only if human capital raises the marginal returns from crime sufficiently more than the returns from legitimate work to compensate for any expected losses from imprisonment associated with crime. From this, it is clear that policies that increase schooling (or the efficiency of schooling) will reduce most types of street crime among adults; although, certain types of white-collar crime (e.g., embezzlement, fraud) may increase with education if they sufficiently reward skills learned in school.

One can also differentiate Eq. (2.3) with respect to θ to determine the effects of "criminal ability" on the extent of criminal activity. A higher criminal ability raises the current payoff from crime, but it also raises the value of staying out of prison in the future, which discourages crime. As a result, the effects of θ on crime are generally ambiguous; however, if the probability of punishment is sufficiently low or if the length of incarceration is sufficiently short, the effect of θ on current criminal payoffs will dominate and more able criminals will commit more crime.⁶ Given the relatively low probabilities of arrest and subsequent incarceration for most crimes, this is the more empirically relevant case. In the discussion below, we assume that crime is increasing in θ .

Now, consider the "schooling stage." At ages $t \le t_s$, someone not currently in prison solves the following problem

$$S_{t}(H_{t}) = \max_{k_{t}, I_{t}} \{ wH_{t}(1 - I_{t} - k_{t}) + N(k_{t}, H_{t}, \theta) + \beta [\pi k_{t} \beta^{J} W(H_{t+1}) + (1 - \pi k_{t}) S_{t+1}(H_{t+1})] \}$$
(2.5)

subject to the time constraints and human capital accumulation Eq. (2.1).⁷ At an interior, the first-order conditions for investment and crime are

$$wH_{t} = \beta [\pi k_{t} \beta^{J} W'(H_{t+1}) + (1 - \pi k_{t}) S'_{t+1}(H_{t+1})] A \frac{\partial f(I_{t}, H_{t})}{\partial I_{t}}$$
(2.6)

$$wH_{t} = \frac{\partial N(k_{t}, H_{t}, \theta)}{\partial k_{t}} - \beta \pi [S_{t+1}(H_{t+1}) - \beta^{J} W(H_{t+1})].$$
(2.7)

As long as youth are engaged in all three activities (work, school, and crime), they will equate the marginal values of each.⁸ Since the marginal return on work is fixed at any age (wH_t) , individuals simply spend their time studying and committing crime up to the point where the net marginal returns on these activities equal their current period wage rate wH_t . Investment offers future returns, while the costs (foregone earnings) are paid up front. Crime is the opposite, providing current returns with

$$\frac{\mathrm{d}k}{\mathrm{d}\theta} = \frac{\pi\beta(1-\beta^J)\frac{\partial W}{\partial\theta} - \frac{\partial^2 N}{\partial k\partial\theta}}{\frac{\partial^2 N}{\partial k^2}}$$

where $\frac{\partial W}{\partial \theta} = \frac{\partial N/\partial \theta}{1 - \beta + \pi k \beta (1 - \beta^J)}$ and $\frac{\partial^2 N}{\partial k \partial \theta}$ are both positive.

⁷ Notice $S_{t,+1}(H) = W(H)$, since individuals enter the "work stage" of life at that time.

⁶ Recognizing that $W(\cdot)$ depends on θ as well as H and totally differentiating (2.3) with respect to θ yields

⁸ Youth who do not work at all will equate the marginal return to crime with the marginal return to schooling investment (i.e., they equate the right-hand side of Eq. (2.6) with the right-hand side of (2.7)).

costs paid up front (foregone labor market earnings) and in the future (expected imprisonment).⁹ The marginal value of human capital during the schooling period is as follows:

$$S'_{t}(H_{t}) = w(1 - I_{t} - k_{t}) + \frac{\partial N(k_{t}, H_{t}, \theta)}{\partial H_{t}} + \beta [\pi k_{t} \beta^{J} W'(H_{t+1}) + (1 - \pi k_{t}) S'_{t+1}(H_{t+1})] \left[1 + A \frac{\partial f(I_{t}, H_{t})}{\partial H_{t}} \right].$$

Notice that both current and future criminal activities affect the marginal value of human capital and incentives to invest in skills. Both current and future crimes raise the probability of incarceration, which lowers the utilization of skills. Future crime also affects the marginal value of human capital, since crime and work will not generally reward human capital the same. For example, if $\frac{\partial N}{\partial H} = 0$, then future time spent on crime reduces time spent working and utilizing human capital in the same way that time spent in prison does.

On average, street criminals are quite young, have low IQ levels (Kandel et al., (1988) and White, Moffitt, and Silva (1989)), and acquire little formal education. This suggests that the returns to traditional market skills are substantially lower for common street crimes (e.g., larceny, robbery, auto theft, drug dealing) than in the legitimate labor market. On the other hand, the returns to white-collar crimes like forgery, fraud, and embezzlement would seem to depend heavily on market skill levels. It is, therefore, worth distinguishing between *unskilled crimes*, for which we assume $\frac{\partial N}{\partial H} = 0$, and *white-collar crimes*, for which we expect that $\frac{\partial N}{\partial H} > 0$ and $\frac{\partial^2 N}{\partial H \partial k} > 0$.

How do individual endowments (A, H_0, θ) influence decisions about investment and crime? Learning ability, A, largely influences crime through past investments and their effects on current skill levels. More able youth will generally invest more in their skills and accumulate more skills per unit of investment. As a result, they will commit fewer unskilled crimes at older ages than their less able counterparts. Because incarceration costs are increasing in A (due to the higher productivity of foregone learning opportunities), more able individuals will also commit less crime (than the less able) during their schooling years. For low incarceration probabilities, the latter effect is relatively unimportant and differences in crime by ability A will grow with age as human capital profiles diverge. Because differences in skill levels will tend to persist over time for a cohort, individuals who begin life with more human capital (H_0) will tend to commit fewer unskilled crime at all ages. By contrast, individuals with higher criminal returns (θ) are more likely to engage in crime at any age. Furthermore, criminal ability

⁹ It is clear that $S_t(H) \ge W(H)$, since the latter is the value associated with a constrained problem (I = 0) that is otherwise the same as the former.

(in unskilled crimes) also has a reinforcing indirect effect on crime at older ages through its effect on early human capital investment decisions—by directly encouraging crime over work, θ indirectly reduces the returns on skill investments causing high θ individuals to accumulate less human capital. Altogether, factors that reflect high A and H_0 or low θ should be negatively correlated with unskilled crime. All of these relationships are likely to be muted (or even perverse) for more skill-intensive white-collar crimes. A high A and low θ are also likely to be associated with greater schooling investments.

Now, consider the relationship between human capital and crime. This relationship is most easily analyzed for unskilled crimes, which offer no reward for human capital. Conditional on ability and other permanent characteristics, adults with high current skill levels face high opportunity costs of crime and should, on average, commit less crime. As such, past investments and education should be negatively correlated with current unskilled crime rates conditional on ability and initial skill levels. Of course, investments and education are endogenous. Individuals who are expected to commit a lot of unskilled crime in the future regardless of their current investments in skill (e.g., individuals with a high θ) have little incentive to make any investment. Thus, variation in criminal opportunities can also generate a negative correlation between educational attainment and unskilled crime.

Since a higher A and a lower θ cause individuals to invest more and commit less unskilled crime, their joint distribution is important for determining the correlation between educational attainment and unskilled crime in the population. Holding either of these endowments constant in the population, we should expect a negative correlation between schooling and unskilled crime.¹⁰ A positive correlation between education and unskilled crime is, therefore, only likely to arise if A and θ are strongly positively correlated, in which case the best criminals are also the brightest. Given the lower than average IQ levels of most criminals (Kandel et al., (1988) and White, Moffitt, and Silva (1989)), this seems unlikely.

The correlation between white-collar crime and education will typically be less negative (and may even be positive). This is because a high A leads to greater investment, but that investment need not lead to less white-collar crime. And, while a high θ encourages white-collar crime, it need not cause individuals to invest less in their skills if those skills pay off in the criminal sector. Thus, two important forces generating a negative correlation between unskilled crime and educational attainment are less powerful for white-collar crime.¹¹

¹⁰ While H_0 should be negatively correlated with unskilled crime, it is more neutral with respect to investment since human capital increases both the returns and costs of investing.

¹¹ Additionally, punishments for white-collar crimes tend to be skewed more heavily toward fines relative to incarceration. This generates similar disincentives to engage in white-collar crime for all education backgrounds, since fines are largely independent of the offender's labor market skills. In fact, the most educated individual may be deterred the least by fines due to diminishing marginal utility of wealth. By contrast, incarceration is more costly for the most educated.

Although individual endowments are important determinants of investment, work, and crime decisions, they are not the only factors affecting the education-crime relationship. Although not modeled here (see Lochner (2004)), shocks to wages and criminal returns can influence both current and future decisions. A strong positive wage or crime shock during the teenage years may cause some individuals to drop out of school, which should then affect subsequent decisions about crime. Absent the accumulation of criminal skills or habit persistence, these exogenous shocks only affect subsequent crime through differences in accumulated human capital levels. Variation in the costs of or tastes for schooling may also affect the education-crime relationship among adults through accumulated skill levels. Finally, policies and government interventions may affect schooling decisions. Policies that encourage schooling investments should reduce crime rates among youth as they substitute time from crime to school. Furthermore, increases in human capital acquired through additional schooling should reduce subsequent crime rates as well. The empirical literature has primarily focused on the latter; however, a few studies have attempted to examine the contemporaneous relationship between education policies and crime. We review some of these studies below.

2.1.2 Other Ways in Which Education May Affect Crime

Education may also teach individuals to be more patient (Becker and Mulligan (1997)). This would discourage crime, since forward-looking individuals place greater weight on any expected future punishment associated with their criminal activities. To the extent that time preferences are affected by schooling, crimes associated with long prison sentences (or other long-term consequences) should be most affected. Education may also affect preferences toward risk. If schooling makes individuals more risk averse, it should discourage crime with its greatest effects on offenses that entail considerable uncertainty in returns or punishment. Finally, schooling may affect the set of people individuals interact with on a daily basis in school, work, or their neighborhoods. Assuming more educated people interact more with other educated people who are less inclined to engage in crime, this is likely to compound any reductions in crime associated with schooling.¹² In most cases, mechanisms related to changes in preferences or social interactions suggest that educational attainment is likely to reduce most types of crime among adults.

2.1.3 School Attendance and Contemporaneous Crime

It is useful to distinguish between the effects of educational attainment on subsequent criminal activity and the way in which school attendance itself affects contemporaneous crime. The latter relationship is likely to be driven by three mechanisms (the first two being central to the time allocation problem analyzed above). First, school may have an incapacitation effect—youth cannot be in two places at once and many criminal

¹² See Glaeser, Sacerdote, and Scheinkman (1996) for a model of crime where social interactions are important.

opportunities are more limited in school than on the streets. Of course, school does not last all day, so this effect depends, in part, on the ease with which youth can engage in crime during nonschool hours. Second, longer periods of school attendance should increase human capital levels and improve future employment prospects. This, in turn, may make juvenile arrests and long periods of detention more costly, reducing incentives to engage in crime while enrolled in school. Third, schools bring hundreds of adolescents together for the day and then free them all at the same time. The social interaction effects from doing this are far from obvious, but it is quite possible that this leads to altercations and more general group-based delinquency. The incapacitation and human capital effects are likely to imply negative effects of school attendance on crime, while the social interaction effect could be positive or negative.

2.2. Evidence on Education and Crime

We now discuss evidence on the effects of educational attainment and school quality and choice on subsequent criminal outcomes. We also review empirical studies that analyze the relationship between school attendance and contemporaneous crime.

2.2.1 Educational Attainment and Crime

We have discussed four primary reasons schooling might affect subsequent crime: (1) education raises wage rates, which raises the opportunity costs of crime; (2) education may directly affect the financial or "psychic" rewards from crime; (3) education may alter preferences for risk taking or patience; and (4) schooling may affect the social networks or peers of individuals. For most crimes (except, possibly, white-collar crimes), one would expect these forces to induce a negative effect of schooling on adult crime.

Empirically, there is a strong negative correlation between educational attainment and official measures of crime. In 1997, 75% of state and 59% of federal prison inmates in the United States did not have a high school diploma (Harlow (2003)).¹³ After controlling for age, state of birth, state of residence, year of birth, and year effects, Lochner and Moretti (2004) still found significant effects of schooling (especially, high school completion) on the probability of incarceration in the United States as reported in Fig. 2.1.¹⁴ In 2001, more than 75% of convicted persons in Italy had not completed

¹³ These figures exclude those who received a General Educational Development (GED) diploma. As shown in Cameron and Heckman (1993) and Heckman and LaFontaine (2006), individuals with a GED perform like high school dropouts rather than graduates in the labor market. Roughly, 35% of state inmates and 33% of federal inmates completed their GED with more than two-thirds of these inmates earning their GED while incarcerated. A small percentage of those who did not receive a high school diploma had participated in some vocational or postsecondary courses. See Harlow (2003).

¹⁴ These figures report the coefficients on indicators for different years of completed schooling from the 1960, 1970, and 1980 censuses for white and black men ages 20–60.



Figure 2.1 Regression-Adjusted Probability of Incarceration by Education (Men Ages 20–60). (a) Whites. (b) Blacks.

Notes: From 1960, 1970, and 1980 US censuses. Regressions control for age, state of birth, state of residence, cohort of birth, state, and year effects. Source: Lochner and Moretti (2004).

high school (Buonanno and Leonida (2006)), while incarceration rates among men ages 21–25 in the United Kingdom were more than eight times higher for those without an education qualification (i.e., dropouts) relative to those with a qualification (Machin, Marie, and Vujic (2011)).

Differences by education are also apparent in self-reported survey measures of crime. For example, in the 1980 wave of the National Longitudinal Survey of Youth (NLSY), 34% of American men aged 20–23 with 11 or 12 years of completed schooling selfreported earning some income from crime, compared with 24% of those with 12 years of school, and only 17% of those with more than 12 years. The effect of education is magnified if we consider more active criminal engagement: 4.2% of 20–23-year-old NLSY men completing 10 or 11 years of school reported earning more than half their income from crime, compared with 1.4% of those with 12 years and 0.7% of those with at least some college education. Similar patterns are observed for violent crime in the NLSY. See Lochner (2004) for further details.

Western (2006) showed that while United States differences in incarceration rates by education existed in the 1970s, they have grown considerably since. Among black men born during 1945–1949, roughly 4% of those who dropped out of high school would spend some time in prison by age 35, while only 1% of those with a high school degree or GED faced incarceration. By contrast, black men born during 1965–1969 faced a 59% incarceration rate by age 35 if they had not completed high school compared with an 18% rate among those who had. Trends in the incarceration differential by education among white men are qualitatively similar. As we will discuss below, these trends are predictable responses to the rapid growth in the probability of imprisonment and prison sentence lengths, as well as the rising wage returns to schooling over the 1980s and 1990s.

Early studies of the relationship between education and crime focused on their correlation conditional on measured individual and family characteristics using standard regression methods.¹⁵ For example, Witte and Tauchen (1994) found no significant relationship between educational attainment and crime after controlling for a number of individual characteristics. Grogger (1998) estimated a significant negative effect of wages on crime, but he found no relationship between years of completed schooling and crime after controlling for individual wage rates. Of course, increased wages and earnings are important consequences of schooling. Thus, this study suggests that education may indirectly reduce crime through increased wage rates.

These earlier studies must be interpreted with caution. A negative cross-sectional correlation between education and crime, even after controlling for measured family back-ground and neighborhood characteristics, does not necessarily imply that education reduces crime. Standard regression studies are unlikely to estimate the causal effect of education on crime (i.e., the effect of increasing someone's schooling on his/her criminal activity) for a number of reasons. First, unobserved individual characteristics like patience or risk aversion are likely to directly affect both schooling and criminal decisions. Individuals who choose more schooling (even after conditioning on observable characteristics) might also choose less crime regardless of their education level, in which case regression-based estimates do not identify a causal effect. Second, using variation in crime and education across states or local communities may also produce biased estimates. Governments may face a choice between funding police or good public schools, which would tend to produce a spurious positive correlation between education and crime. Alternatively, unobserved characteristics

¹⁵ Ehrlich (1975) provided an early empirical exploration of predicted effects of education on crime from a human capital perspective. See Witte (1997) for a survey of the early empirical literature on education and crime.

about communities or their residents may directly affect the costs or benefits of both education and crime. For example, communities with few job opportunities that reward schooling may also be faced with severe gang problems. While it is often possible to account for permanent unobserved differences across communities by examining the relationship between changes in schooling and crime over time, this approach does not account for the effects of changing unobserved community characteristics. Third, reverse causality is another important concern, for reasons discussed earlier. Individuals who plan to heavily engage in crime (e.g., because they are particularly good at it, enjoy it, or live in areas with plenty of illicit opportunities) are likely to choose to leave school at a young age. Arrests or incarceration associated with juvenile crime may also cause some youth to drop out of school early (Hjalmarsson (2008)).

Recently, economists have attempted to address these difficult issues and to estimate the causal effects of schooling on crime using instrumental variable (IV) methods. In the context of estimating the effect of educational attainment on crime, an instrument is valid if it induces variation in schooling but is uncorrelated with other factors that directly affect criminal proclivity (e.g., individual preferences or abilities, local law enforcement). Intuitively, this approach exploits differences in educational attainment across individuals that arise in response to factors that have no direct impact on criminal decisions. An ideal instrument would randomly assign some youth to drop out of high school and others to finish. Then, comparing the differences in crime rates across these groups would identify the causal effect of high school completion on crime. In practice, we typically do not observe such perfect experiments, but researchers can sometimes come close.

Because crime itself is difficult to measure, researchers are often forced to use measures of arrest or incarceration rather than actual crimes committed. It is possible that education reduces the probability of arrest and incarceration or the sentence lengths administered by judges, in which case estimates based on measures of arrest or incarceration incorporate these effects in addition to any effects of education on actual crime. While there is little direct evidence on these issues, Mustard (2001) found negligible effects of defendant education levels on the sentence lengths they receive. Furthermore, results using self-reported measures of crime in the NLSY support the case that education reduces actual violent and property crime and not just the probability of arrest or incarceration conditional on crime (Lochner (2004) and Lochner and Moretti (2004)).¹⁶

¹⁶ There has been considerable debate among criminologists on the merits of self-reported measures of crime versus official measures of arrest. Most studies have found a reasonably high correlation between the two; however, it is generally agreed that the two measures offer distinct and complementary information about criminal activity. Comparisons of self-reported arrests versus official arrests tend to find a stronger correlation, with agreement increasing further for self-reported versus official measures of criminal convictions. A number of studies report greater underreporting of crimes and arrests by blacks; however, studies vary considerably on this. See the classic Hindelang, Hirschi, and Weis (1981) for comprehensive treatment of the issue or Thornberry and Krohn (2000) for a more recent survey of this literature.

Many recent empirical studies analyzed crime aggregated at some geographic level, exploring the effects of average educational attainment on crime, arrest, conviction, or incarceration rates. To address concerns with endogeneity or unobserved heterogeneity, researchers have typically turned to instrumental variables estimation or a differences-in-differences strategy using changes in state or national rules that affect schooling decisions. An aggregate-level regression is often specified as follows:

$$\gamma_{calt} = E_{alt}\beta + X_{alt}\gamma + d_{lt} + d_{cl} + d_{al} + d_{ct} + d_{at} + d_{ca} + \varepsilon_{calt}$$
(2.8)

where y_{calt} is a measure of the crime, arrest, or incarceration rate for some offense type c, age group a, in location l, in year t. In some cases, only a single measure of crime is used (e.g., incarceration or total arrests), in which case the c subscript is extraneous. E_{alt} is an aggregate measure of educational attainment for age group a in location l at time t (e.g., average schooling attainment or high school completion rates). X_{alt} is a set of observable characteristics that may vary across age, location, and time (e.g., racial composition of an area). The d's represent indicator variables that account for unobserved differences by age/cohort, location, year, and criminal offense types. The term d_{it} allows for location-specific time effects, which account for time-varying unobserved location-specific differences that may reflect differences in local public spending, economic conditions, or law enforcement. The inclusion of d_d allows the average distribution of crime or arrest types to differ across areas. For example, some states may focus arrests more heavily on one type of crime, while others focus on other types. Or, some areas may be more amenable to certain crimes while others are not. Similarly, the age distribution of crime or arrests need not be the same across areas—some age groups may be more crime-prone in some areas or arrest policies with respect to age may differ across areas. The term d_{al} absorbs long-run differences in age-arrest patterns across locations. Crime-specific and age-specific time trends in arrest common to all areas are accounted for by d_{ct} and d_{at} , respectively. Finally, d_{ca} accounts for long-term differences in age-crime profiles across different types of criminal offenses. Given these fixed effects, identification of the effect of education on crime is achieved through time variation in cohort educational attainment levels across different locations. The absence of d_{alt} indicator variables in Eq. (2.8) is, therefore, central to identification.

Lochner and Moretti (2004) examined state-level male arrest rates by criminal offense and age (5-year age categories beginning at ages 20–24 through 55–59) from the FBI's Uniform Crime Reports (UCR) for the United States in 1960, 1970, 1980, and 1990. These data are linked to 1960–1990 decennial US census data on educational attainment and race to estimate Eq. (2.8), where γ_{calt} represents log arrest rates for a specific offense, age category, state, and census year. They specifically analyzed arrest rates for murder, rape, assault, robbery, burglary, larceny, auto theft, and arson. In using log arrest rates, the effect of education is assumed to be the same in percentage terms for

each type of crime included in the regression. They explored the effects of both average years of schooling and high school completion rates at the cohort-level (cohorts are defined by year of birth given year t and age a) in state l as of time t (i.e., E_{alt}). In addition to including all the d fixed effects in Eq. (2.8), they also included the percent of males who are black in age group a living in state l at time t.

The main methodological contribution of Lochner and Moretti (2004) is the use of changes in state-specific compulsory schooling laws over time as instrumental variables for schooling.¹⁷ Intuitively, this strategy measures the extent to which an increase in a state's compulsory schooling age leads to an immediate increase in educational attainment and reduction in subsequent crime rates for affected cohorts. Because the laws only affect schooling at low levels (mainly grades 8–12), their instrumental variable (IV) estimates reflect the impact of an additional year of high school on crime.¹⁸ The impacts of schooling at other margins may differ as discussed below.

Lochner and Moretti (2004) estimated Eq. (2.8) using both ordinary least squares (OLS) and IV estimation. Using OLS, they found that a 1-year increase in average education levels in a state reduces state-level arrest rates by 11%. IV estimates suggest slightly larger effects, although they are not statistically different. These estimated effects are very similar to the predicted effects derived from multiplying the estimated increase in wages associated with an additional year of school by the estimated effects of higher wage rates on crime (from Gould, Mustard, and Weinberg (2002)). This suggests that much of the effect of schooling on crime may come through increased wage rates and opportunity costs as emphasized in the above model. Given the strong relationship between high school completion and incarceration apparent in Fig. 2.1, Lochner and Moretti (2004) also estimated the specifications using the high school completion rate as a measure of schooling. OLS estimates suggest that a ten percentage point increase in high school graduation rates would reduce arrest rates by 7%, while IV estimates suggest a slightly larger impact of 9%.

Lochner and Moretti (2004) also estimated separate effects of education for different types of crimes using OLS (including interactions of criminal offense type with education in Eq. (2.8)). These results suggest similar effects across the broad categories

¹⁷ The relevant compulsory schooling age is based on the state law that applied when a cohort was age 14. Lochner and Moretti's (2004) analysis suggests that changes in compulsory schooling laws are exogenous and not related to prior trends in schooling or state expenditures on law enforcement, so it appears to be a valid instrument. Other studies reached similar conclusions about the exogenous nature of changes in compulsory schooling laws in other contexts (e.g., Acemoglu and Angrist (2001) and Lleras-Muney (2002)).

¹⁸ More specifically, they reflect the impact of an additional year of high school on crime among those youth, who are affected by changes in the schooling age laws. See Angrist and Imbens (1994, 1995) for the interpretation of IV estimators when using discrete instrumental variables.

of violent (murder, rape, robbery, and assault) and property (burglary, larceny, motor vehicle theft, and arson) crime—a 1-year increase in average years of schooling reduces both property and violent crime by about 11–12%. However, the effects vary considerably within these categories. A 1-year increase in average years of schooling reduces murder and assault by almost 30%, motor vehicle theft by 20%, arson by 13%, and burglary and larceny by about 6%. Estimated effects on robbery are negligible, while those for rape are significantly positive. Additional specifications suggest quantitatively similar effects for a 10–20 percentage point increase in high school graduation rates. Their results for rape are surprising and not easily explained by standard economic models of crime.¹⁹

Lochner (2004) followed a similar approach using the same UCR data from 1960 to 1980; however, he also examines white-collar crime. OLS estimation of Eq. (2.8) produces *positive*, though statistically insignificant, effects of schooling on arrest rates for white-collar crimes (forgery and counterfeiting, fraud, and embezzlement). Estimates for violent and property crime are negative and similar to those of Lochner and Moretti (2004).

Lochner and Moretti (2004) also used individual-level data on incarceration and schooling from the 1960, 1970, and 1980 US censuses to estimate the effects of educational attainment on the probability of imprisonment separately for black and white men (ages 20-60). Their estimates control for age of the respondent (3-year age categories), state of birth, state of residence, cohort of birth, and state-specific year effects. Most importantly, controlling for state-specific year effects allows for the possibility that different states may have different time trends for law enforcement policies or may simply exhibit different trends in aggregate criminal activity. Analogous to their analysis of state-level arrest rates, they use state-level changes in compulsory schooling ages as an instrument for educational attainment. Although this analysis uses individual-based measures of incarceration and schooling, variation in schooling laws at the state-year level effectively identifies the effect of education on crime. As with the estimates for aggregate arrest rates, identification comes from the fact that in any given state and year, different age cohorts faced different compulsory schooling laws during their high school years, causing them to acquire different levels of schooling and to commit crime at different rates. Again, both OLS and IV estimates are very similar and suggest that, on average, an extra year of education reduces the probability of imprisonment by slightly more than .1 percentage point for whites and by about .4 percentage points for blacks. In their sample, the probability of incarceration for male whites (blacks) without a high school degree averaged .83% (3.6%), which translates into a 10-15% reduction in

¹⁹ However, the results are consistent with some specifications in Gould, Mustard, and Weinberg (2002), which suggests that local wage rates are positively correlated with local crime rates for rape.

incarceration rates for both white and black males associated with an extra year of completed schooling. These estimated effects are comparable to those for arrest rates described earlier. OLS results suggest that completion of the 12th grade causes the greatest drop in incarceration, while there is little effect of schooling beyond high school (see Fig. 2.1).

Oreopolous and Salvanes (2009) reproduced the Lochner and Moretti (2004) IV results for black males using the same estimation strategy with a slightly different specification and an expanded sample that includes men ages 25–64 from the 1950–1980 US censuses.²⁰ Their estimate suggests that an additional year of completed schooling reduces incarceration rates among black men by about 20%.

Machin, Marie, and Vujic (2011) exploited a 1972–1973 increase in the minimum schooling age (from age 15 to 16) in England and Wales to estimate the effects of schooling on criminal convictions for property and violent crimes over the period 1972–1996. Using both IV and regression discontinuity methods, identification effectively comes from cohort-level changes in schooling attainment and crime for cohorts turning 15 immediately before and after the law change.²¹ Among men, they estimated that a 1-year increase in average schooling levels reduces conviction rates for property crime by 20–30% and violent crime by roughly one-third to one-half as much.²² Compared with estimates for the United States by Lochner and Moretti (2004), the impacts of education on property crime appear to be greater in the United Kingdom, while the effects on violent crime are weaker.

Buonanno and Leonida (2006) estimated the effects of educational attainment on crime rates in Italy using regional panel data from 1980 to 1995. Their unit of observation is a region-year (they examined 20 Italian regions), and they estimated a restricted form of Eq. (2.8) using OLS. Specifically, they control for region and time-fixed effects (d_l and d_l), along with region-specific quadratic time trends (assuming $d_{lt} = \delta_{1l}t + \delta_{2l}t^2$), and a rich set of time-varying region-specific covariates.²³ These estimates are identified from the relationship between changes in regional education levels and crime rates (around smooth regional time trends). Their estimates suggest that a 10 percentage

- ²² Estimated effects on male property crime are statistically significant, while effects on male violent crime are not. Estimated effects for women are, unfortunately, very imprecise.
- ²³ Covariates included employment rates, GDP per capita, GDP growth rates, average wage rates, the fraction of crimes without an arrest, police per capita, and the length of time in the judicial process.

²⁰ Most notably, they do not include state- and state-specific year effects in their specification. They also removed individuals with greater than 12 years of schooling from their sample. Their measures of compulsory schooling ages differ as well, incorporating the fact that some states allow for exceptions to the dropout age under certain conditions.

²¹ They estimated models aggregated to the year-age level for individuals ages 18–40 from 1972–1996. To alleviate concerns that other important economic or social factors may have changed at the same time the compulsory schooling age increased, they included a rich set of covariates: year and age indicators, fraction British-born, fraction employed, fraction nonwhite, and fraction living in London.

point increase in high school graduation rates would reduce property crime rates by 4% and total crime rates by about 3%. (Effects on property crime are statistically significant, while effects on total crime are not.) They found no evidence to suggest that university completion reduces crime.²⁴

Merlo and Wolpin (2009) took a very different approach to estimate the relationship between schooling and subsequent crime. Using individual-level panel data on black males aged 13–22 from the NLSY, they estimated a discrete choice vector autoregression model in which individuals can choose to engage in crime, attend school, and/or work each year.²⁵ These decisions are allowed to depend on unobserved individualspecific returns to each activity, as well as crime, schooling, and work choices the previous year. Using estimates for their model, Merlo and Wolpin simulated the effects of changing youth schooling status at age 16 on subsequent outcomes. Their estimates suggest that, on average, attending school at age 16 reduces the probability a black male ever commits a crime over ages 19–22 by 13 percentage points and the probability of an arrest over those ages by 6 percentage points. These represent 42 and 23% reductions in self-reported crime and arrest rates, respectively, for black males not in school at age 16.

A final study worth mentioning examined the effects of an explicit education subsidy on youth burglary rates in England. Between 1999 and 2002, England piloted Educational Maintenance Allowances (EMA), which provided subsidies of up to $\pounds 40$ per week (plus bonuses for completion of coursework) for low-income 16-18-yearold youth to attend school. The program was administered in 15 local areas with low schooling participation rates. During the same time period, the Reducing Burglary Initiative (RBI) funded 63 different local burglary reduction schemes as a separate pilot project. Roughly half of all EMA pilot areas were also selected for the RBI. Sabates and Feinstein (2008) used a differences-in-differences strategy to identify the effects of each pilot program as well as the combination of the two on burglary. Specifically, they compared changes in burglary conviction rates before and after the introduction of RBI, EMA, or both against a set of comparison areas. While baseline burglary conviction rates were much higher in EMA and EMA-RBI combined areas relative to the comparison areas, annual growth rates in burglary conviction rates prior to the programs were quite similar across all three classifications. To reduce concerns about differences between the treated and untreated areas, Sabates and Feinstein control for a number

²⁴ Buonanno and Leonida (2006) also generalized their econometric specification to allow for an effect of lagged crime rates on current crime rates, estimating this using a generalized method of moments estimator to account for the endogeneity of lagged crime rates. This specification produces similar estimated effects of schooling on crime.

²⁵ Crime, work, and school are not mutually exclusive activities in this framework—individuals can do any combination of these three activities in each year. An individual is said to have engaged in crime in any year if he/she self-reported any of the following offenses: theft, other property crime, sold drugs, or assault.

of time-varying area-specific factors likely to affect crime and limit their sample of comparison areas to those that best "match" the distribution of demographic characteristics in the pilot areas.²⁶ Their findings suggest that the combination of both the EMA and RBI significantly reduced burglary rates by 1.3 per 1000 youth (about 5.5%) relative to the "matched" comparison areas. Effects of the EMA alone were slightly lower but still significant. While there are obvious concerns about the extent to which time-varying determinants of burglary are the same for treated and comparison areas, Sabates and Feinstein (2008) show that estimated effects on burglary rates for 19–21-year-olds (who were not offered the education subsidy) were much lower and statistically insignificant.

2.2.2 School Quality and Crime

If human capital acquisition, socialization, or preference modifications are important mechanisms determining the impacts of educational attainment on crime, then it is likely that the school quality and the type of schools students attend also affect criminal behavior. Although there are no studies which directly estimate the effects of measured school quality on crime, three recent studies on school choice and desegregation provide some useful insights.

Cullen, Jacob, and Levitt (2006) and Deming (2009a) examined the importance of school choice in large urban US school districts (Chicago and Charlotte-Mecklenburg, respectively) on a variety of student outcomes, including delinquency and crime. Both studies examined the effects of "winning" a randomized lottery for admission to schools children selectively apply to.²⁷ By comparing the outcomes for youth who win versus lose a particular school admission lottery, they estimated the effects of being offered admission to that school relative to the preferred alternative. This reflects the "intention to treat" (i.e., the effects of being offered the opportunity to attend better schools) and not necessarily the effects of actually attending that school, since many students did not ultimately enroll in schools for which they were admitted by lottery. However, both studies found that "winning" a lottery does significantly increase enrollment in that school. Since many students applying outside their assigned local school are from disadvantaged backgrounds and neighborhoods, on average, lottery winners end up attending

²⁶ Their regressions control for unemployment rates for individuals under 25, proportion of students eligible for free school meals, number of qualified teachers, pupil-teacher ratios, and the number of supplementary staff for ethnic minorities, percent of youth with no schooling qualifications as of age 16 (i.e., dropouts), and the percent of unauthorized half-days missed in secondary school. We discussed results based on the "matched" sample of comparison areas.

²⁷ In both cases, students could always choose to attend their neighborhood school. If any additional positions were available in a school, an open enrollment lottery was run based on all other students who applied to that school/ program. Lotteries were random within population subgroups (e.g., by race or family income).

better quality schools, as measured by such things as student achievement scores, value added (i.e., growth in achievement), student behavioral problems, or teacher quality. In this sense, these studies offer an opportunity to examine the effects of school quality, broadly defined, on delinquency and crime.

Cullen, Jacob, and Levitt (2006) found that winning a high school lottery in Chicago significantly raises peer graduation rates by 6% and the share of peers who test above national norms by about 14%; however, lottery winners appear to be placed in lower tracked classes within the better schools. Interestingly, they found no evidence that lottery winners perform better on a wide range of academic measures (e.g., math and reading tests, enrollment, days absent) and some evidence that they are more likely to drop out of high school. The latter may be due to a mismatch between student ability and school demands. Despite the disappointing findings regarding academic outcomes, students who won lotteries to high achievement Chicago public schools reported nearly 60% fewer arrests on a ninth-grade student survey. These winners also reported getting into less trouble at school, and school administrative data suggest that they had lower incarceration rates during school ages. Of course, it is possible that schools themselves affect student arrest and incarceration rates through differential disciplinary policies (or criminal opportunities), so it is important to study whether these reductions in arrests/incarceration persist beyond high school.

To this end, Deming (2009a) examined the impacts of open enrollment lotteries (for middle and high schools) on adult criminal outcomes 7 years after random assignment.²⁸ Given his interest in the effects of school choice on crime, he categorized males based on their likelihood of arrest, which he estimated as a function of demographic characteristics, earlier math and reading test scores, and other school-related behaviors at young ages. For his entire sample of middle and high school lottery participants, "high-risk" youth (defined as those in the top quintile of predicted arrest probability) have seven times more felony arrests (7 years after random assignment) than the average student from the bottom four quintiles combined.

Like Cullen, Jacob, and Levitt (2006), Deming (2009a) estimated significant effects of winning a school lottery on the quality of school attended, especially among "high-risk" youth, but no effects on achievement tests. There appears to be some effect on student enrollment during high school years, but there is no evidence that "high-risk" lottery winners are more likely to graduate from high school.²⁹ Among high school lottery winners in the high-risk category, Deming (2009a) estimated a significant 0.35 (roughly 45%) reduction in the number of adult felony arrests (cumulative as of 7 years

²⁸ He merges Charlotte-Mecklenburg school district data with data on adult (ages 16+) arrests and incarceration from Mecklenburg County and the North Carolina Department of Corrections.

²⁹ There is more evidence of effects on high school graduation and college attendance among the lower risk quintiles.

after the lottery) with an associated savings in victimization costs of \$4,600–16,600.³⁰ Because many crimes do not lead to an arrest, the total benefits to potential victims and society are likely to be much larger. His estimates suggest that winning middle school lotteries also reduces crime among high-risk youth with most effects of a similar order of magnitude.

Court-ordered school desegregation policies enacted since Brown v. Board of Education of Topeka in 1954 dramatically altered the types of schools blacks attended in many American districts. In most cases, the resources and average student achievement of schools attended by blacks would have improved markedly.³¹ Guryan (2004) estimated that these desegregation efforts significantly increased high school graduation rates among blacks by 2-3 percentage points but had no effect on white graduation rates. Weiner, Lutz, and Ludwig (2009) examined whether these changes affected county-level homicide rates.³² Their estimates suggest that homicide deaths among blacks aged 15-19 declined by 17% in the first 5 years after court-ordered desegregation, while homicide deaths among white 15-19-year-olds declined by about 23%. Homicide deaths among slightly older whites and blacks also declined. In looking at offenders, they found that arrest rates for homicide declined by one-third for blacks aged 15-19, while there was no decline for young whites. Combining Guryan's (2004) estimated effect on high school graduation rates with the estimated effects of schooling on crime from Lochner and Moretti (2004), they argue that much of the effect may be coming from the increased schooling among blacks associated with desegregation.

For some perspective, it is interesting to compare these findings with those from the Moving-to-Opportunity (MTO) experiment, which provided housing vouchers to low-income families to move out of high-poverty neighborhoods. Evaluations of MTO report that families receiving the housing vouchers moved into neighborhoods with about 25% lower poverty rates; however, these moves only led to modest improvements in the quality of schools youth attended and no improvements in their cognitive achievement (Sanbonmatsu et al. (2006)). Kling, Ludwig, and Katz (2005) reported that the MTO housing vouchers led to lasting reductions in arrests for both violent and property offenses among young females, short-term reductions in violent crime arrests for males, and delayed *increases* in property crime arrests for males. Overall, reductions in crime were modest at best.

³⁰ These victimization costs (in year 2009) assign costs based on the type of offense using cost estimates from Miller, Cohen, and Wiersema (1996). They do not include justice system or enforcement costs. The larger figure uses a cost of \$4.3 million for murder, while the smaller uses a value of \$125,000 (twice the cost of rape).

³¹ For example, Reber (2007) shows that integration efforts in Louisiana from 1965–1970 were accompanied by large increases in per pupil funding for black students.

³² They used data on homicide death by year and county over the period 1958–1988 from vital statistics and data on homicide victims and arrestees from the Supplemental Homicide Report from 1976–2003.

Taken together, these studies suggest that simply improving the schools attended by disadvantaged youth appears to be much more successful in reducing their criminal activity (though not necessarily in improving academic outcomes) than changing neighborhoods. Given the mixed findings on educational attainment levels (with modest positive effects at best), the impacts of better schools on crime appear to be driven largely by school quality and not "quantity." Whether it is the quality of teachers and instruction or of student peers is less obvious. The fact that test scores did not improve among lottery winners suggests that the main effects of attending "better" schools on delinquency and crime are likely to be attributed to better socialization, better peer interactions, improvements in noncognitive skills, or changes in preferences.³³ It is, therefore, interesting that substantial improvements in "neighborhood peers" do not yield the same benefits in terms of crime reduction.

2.2.3 Contemporaneous Schooling and Crime

We now consider the relationship between contemporaneous schooling and crime. As noted earlier, there are three main ways in which altering youth's schooling attendance is likely to affect their contemporaneous engagement in crime: (1) incapacitation, (2) raising the costs of future punishment through human capital accumulation, and (3) social interactions facilitated by bringing youth together. The incapacitation and human capital effects of schooling on crime are likely to be negative, while the sign of the social interaction effect is theoretically ambiguous.

Three relatively recent studies shed light on these effects by estimating the impacts of different "interventions" that directly affect youth schooling attendance.³⁴ Anderson (2009) examined the effects of increasing state compulsory schooling ages on crime among affected youth (i.e., forcing some youth to stay in school), while Jacob and Lefgren (2003) and Luallen (2006) estimated the effects of extra days off from school due to teacher in-service days or teacher strikes (i.e., keeping youth out of school). The policies analyzed by these studies differ in two important respects. First, increases in compulsory schooling ages typically "require" students to stay in school at least one additional year and sometimes more, whereas teacher in-service days release all students from school, changes in compulsory schooling laws typically affect a small set of marginal students. All three potential effects of school attendance on crime are likely to be relevant to changes in compulsory schooling, while the effects of in-service days

³³ See Heckman, Stixrud, and Urzua (2006) for evidence on the importance of "cognitive" versus "noncognitive" skills for crime and in the labor market.

³⁴ Using individual-level data, earlier studies by Gottfredson (1985), Farrington, Gallagher, Morley, St. Ledger, and West (1986), and Witte and Tauchen (1994) explored the cross-sectional relationship between time spent in school and contemporaneous crime, concluding that time spent in school significantly "reduces" criminal activity. Unfortunately, these findings are difficult to interpret given the simultaneous nature of the crime and schooling choices.

and teacher strikes are likely to be limited to incapacitation and social interactions. Any social interaction effects are likely to be magnified in the latter cases due to the universal nature of the policy.

Rather than use changes in compulsory schooling laws as instruments for educational attainment, Anderson (2009) estimated the direct effect of these laws on contemporaneous county-level arrest rates (from the UCR) from 1980 to 2006 among affected youth aged 16–18. Specifically, his estimates are identified from within-county fluctuations in arrests (around county-specific trends) for 16–18-year-olds (relative to 13–15-year-olds) over time as state compulsory schooling ages change.

Anderson's estimates for total arrest rates imply that a compulsory schooling age of 17 years significantly reduces age 17 arrests by about 8% (5.4 arrests per 1000 youth) compared with a compulsory schooling age of 16 or less. Similarly, a compulsory schooling age of 18 significantly reduces arrests by 9.7–11.5% at ages 16–18. Separating arrests by type of offense, he estimates that compulsory schooling laws significantly reduce both property and violent arrests for 16–18-year-olds. Although, estimated effects of schooling age laws on drug-related crimes are sizeable, the effects are typically not statistically significant. Overall, the estimates generally suggest that forcing youth to spend an extra year or two in high school significantly reduces their arrest rates over that period.

Jacob and Lefgren (2003) examined the effects of single day changes in school-wide attendance on juvenile crime and arrest rates in 29 large American cities from 1995 to 1999. Exploiting teacher in-service days across jurisdictions over time as an exogenous source of variation in school days, they essentially compared local juvenile crime rates on days when school is not in session to those when it is.³⁵ Their findings suggest that an additional day of school reduces serious juvenile property crime by about 14% that day, while it *increases* serious juvenile violent crime by 28%. These results are consistent with an "incapacitation effect" of school that limits participation in property crime. However, the increased level of interaction among adolescents facilitated through schools may raise the likelihood of violent conflicts (and other minor delinquency) after school. Interestingly, they found no evidence to suggest that school days simply shift crime to other days without changing overall crime rates.

Luallen (2006) followed a similar approach, using teacher strikes (typically lasting about 5 days) rather than in-service days as an exogenous source of school days. Using data from the state of Washington from 1980 to 2001, Luallen (2006) estimated that an extra day of school reduces arrests for property crimes by about 29% while increasing arrests for violent crimes by about 32% in urban areas. The effect on property crime is roughly double the effect estimated in Jacob and Lefgren (2003), while the effect on violent crime is quite similar. In rural and suburban areas, Luallen found insignificant effects

³⁵ Their main specification includes controls to account for the possibility that crime may be higher on certain days of the week or that different cities may experience different monthly crime cycles.
on both violent and property crime arrests. Thus, the incapacitation and social interaction effects appear to be particularly strong in urban areas and negligible elsewhere.

2.3. Evidence on Human Capital-Based Interventions from Birth to Young Adulthood

A growing body of evidence suggests that early childhood and school-age interventions can reduce adult crime rates. Most famously, the High/Scope Perry Preschool Program substantially lowered arrest rates through age 40 for a sample of low-income minority children in Ypsilanti, Michigan. Several other early childhood interventions have produced similar effects on delinquency; however, others have not. We briefly review studies of early childhood and school-age interventions that have analyzed educational and criminal/delinquency outcomes during late adolescence or adulthood.³⁶ We then discuss a few programs aimed at improving school participation among adolescents or that directly provide training to adolescents and young adults.

Table 2.1 summarizes four small-scale early childhood interventions (Abecedarian Project, Chicago Child Parent Center [CPC], High/Scope Perry Preschool, and Infant Health and Development Program [IHDP]), their target populations, study methodology, and estimated effects on educational attainment and crime at ages 18 or older. All of the programs included a preschool component, ranging from full-time full-year care from birth to kindergarten (Abecedarian) to half-day preschool at ages 3 and 4 (Chicago CPC and Perry Preschool). Perry Preschool and IHDP also included regular home visits at preschool ages as part of their curriculums.³⁷ All of the programs targeted youth facing some form of disadvantage. Abecedarian and Perry Preschool specifically targeted children at risk of having problems developing normally in school. Children enrolling in the Chicago CPC were all minorities selected from families with low socioeconomic status (SES). IHDP drew from a more heterogeneous population, targeting pre term children born of low birth weight (less than 2500 g). Overall, these studies cover a reasonably broad range of potential preschool-based interventions and target populations. (The findings for Head Start are discussed below.)

Youth from all four of these programs were followed until at least age 18, enough time to determine whether the programs have medium-term effects on the education and criminal behavior of participants. Only the Chicago CPC was not evaluated using randomized trials; however, Reynolds, Temple, Robertson, and Mann (2001) used a strong design of matching treated children with other comparison children based on age of kindergarten

³⁶ See Karoly et al. (1998) or Blau and Currie (2006) for more comprehensive surveys of early childhood programs.

³⁷ All of the programs typically provided other additional services to families and children (e.g., nutritional and health services). While a subsample of the Abecedarian participants received an extended school-age intervention for the first few years of school, we focus on the preschool component of the program. The additional school-age services did not substantially impact the educational attainment or crime outcomes discussed here.

Program	Program Description	Program Population	Methodology	Education Effects	Crime Effects
Abecedarian Project	Full-time full-year preschool from infancy to kindergarten	Developmentally at-risk children, Chapel Hill, NC	Random assign	Increased high school graduation rate by 0.03 (0.70 vs. 0.67) and enrollment in 4-yr college by 0.22* (0.36 vs. 0.14)	No significant effects by age 21
Chicago Child Parent Center	Half-day preschool (school year) ages 3 and 4	Low-income minority children, Chicago, IL	Matched sample	Increased high school completion rate by 0.09 (0.57 vs 0.48) for females and 0.14* (0.43 vs. 0.29) for males	By age 18, reduced fraction arrested by 0.08* (0.17 vs. 0.25)
High/Scope Perry Preschool	Half-day preschool (school year) ages 3 and 4, bi-weekly home visits	Low-income black children at risk of school failure, Ypislanti, MI	Random assign	Increased high school graduation rates by 0.52* (0.84 vs. 0.32) for females and reduced graduation rates by 0.04 (0.50 vs. 0.54) for males	By age 40, reduced fraction arrested 5 or more times by 0.10* (0.24 vs34) for females and 0.24* (0.45 vs. 0.69) for males
Infant Health & Development Program (IHDP)	Weekly/Bi-weekly home visits from 0–36 months, full- time full-year preschool 12–36 months	Low-birthweight pre-term infants, 8 sites	Random assign	No significant effect on high school dropout (approximately 10% dropout rate)	No significant effects on arrests by age 18

Table 2.1 Effects of Selected Early Childhood Programs on Educational Attainment and Adult Crime

Notes: Effects for Abecedarian Project taken from Campbell, Ramey, Pungello, Sparling, and Miller-Johnson (2002) and Clarke and Campbell (1998). Effects for Chicago Child Parent Center taken from Reynolds, Temple, Robertson, and Mann (2001). Effects for Perry Preschool taken from Schweinhart et al. (2005). Effects for IHDP taken from McCormick et al. (2006).

*Denotes difference is statistically significant at 0.05 level.

entry, eligibility for and participation in government-funded programs, and neighborhood and family poverty. Children from the matched comparison sample would also have been eligible for the program had they lived in a neighborhood with a center. Sample sizes range from around 100 children for Perry Preschool to 1300 for Chicago CPC.

Both Chicago CPC and Perry Preschool significantly increased high school completion rates overall; however, the Chicago CPC had more sizeable effects on male graduation rates, while Perry Preschool only raised female graduation rates (Reynolds, Temple, Robertson, and Mann (2001); Schweinhart et al. (2005)). The IHDP had no effect on high school dropout rates by age 18, while Abecedarian increased college attendance but not high school completion (McCormick et al. (2006); Campbell, Ramey, Pungello, Sparling, and Miller-Johnson (2002)). These programs typically produced short-term gains in achievement scores and sometimes generated lasting gains.

The final column of Table 2.1 reports estimated effects of these programs on late juvenile and adult crime. As alluded to above, Perry Preschool had significant effects on lifetime crime measured as of age 40 (Schweinhart et al. (2005)). Reductions in the fraction arrested five or more times were substantial for both males and females. Both showed reductions of about one-third; however, the size of the effect in absolute terms is much larger for males given their higher baseline crime rate. Reductions in crime for Perry Preschool students were observed across a broad range of crimes (e.g., drug, property, and violent crimes) and were apparent even at younger ages. The Chicago CPC also reduced arrest rates (by age 18) by about one-third (Reynolds, Temple, Robertson, and Mann (2001)). Another widely cited family support and preschool program, the Syracuse University Family Development Research Program, showed significant reductions in juvenile delinquency measured at a slightly earlier age: 6% of preschool participants had been placed under probation services by age 15 compared with 22% of controls (Lally, Mangione, and Honig (1988)).³⁸

The estimated savings in reduced criminal justice expenditures and victimization costs resulting from the crime reductions of Perry Preschool and Chicago CPC are sizeable. Using a 3% discount rate, Belfield, Nores, Barnett, and Schweinhart (2006) estimated that the Perry Preschool produced a social benefit of over \$150,000 (year 2000 dollars) per child from crime reduction alone.³⁹ Reynolds et al. (2002) estimated that

³⁸ The Elmira Nurse Home Visitation Program provided home visits by nurses to first-time mothers who were young, unmarried, or of low SES. Nurses visited homes for randomly assigned mothers during pregnancy and for the first 2 years of the child's life. Olds et al. (1998) reported mixed but encouraging effects of the program on delinquency at age 15: treated youth were more likely to self-report being stopped by the police but had fewer incidences of arrests and convictions.

³⁹ This figure is for benefits through age 40. Using a 7% discount rate produces a social benefit from crime of about \$67,000 (Belfield, Nores, Barnett, and Schweinhart (2006)). Heckman, Moon, Pinto, Savelyev, and Yavitz (2009) reported that savings from crime reduction account for about 40–65% of the benefit-cost ratio for Perry Preschool, depending on assumptions about discount rates (0–7%) and the cost of murder.

reductions in juvenile crime through age 18 associated with the Chicago CPC saved society roughly \$8000. Findings like these, especially those for Perry Preschool, led Donohue and Siegelman (1998) to conclude that small, rigorous early intervention programs may pay for themselves through reduced crime rates alone, if they can be targeted to high-crime groups.

Not all early childhood programs in Table 2.1 yield reductions in crime. Although modest reductions in self-reported convictions and incarceration through age 21 are observed for Abecedarian, none of these effects is statistically significant (Campbell, Ramey, Pungello, Sparling, and Miller-Johnson (2002)). Based on administrative records of adult criminal charges in North Carolina, Clarke and Campbell (1998) reported nearly identical rates of arrests and criminal charges (as of age 21, on average) for treatment and control children in the Abecedarian study. Similarly, IHDP produced no significant effects on crime through age 18 (McCormick et al. (2006)).

Why didn't Abecedarian and IHDP produce significant reductions in crime like Perry Preschool and Chicago CPC? It is difficult to point to any particular curriculum difference; although, not all preschools are alike. Abecedarian began preschool at infancy and continued through kindergarten—the longest of any program. It was also full-day year-round, unlike Perry Preschool or Chicago CPC. Like Perry Preschool, it showed sizeable gains in achievement and IQ, so it is difficult to attribute its lack of effects on crime to inadequate intervention. The only obvious program difference between Abecedarian and Perry Preschool or Chicago CPC that might explain the absence of any impact on crime is its lack of a "home visit" component, but IHDP included home visits by nurses from birth through 3 years of age. IHDP began early but also ended when Perry Preschool and Chicago CPC began (age 3), so it is possible that the early "home visit" combined with later preschool care is a key combination of services necessary for long-term impacts on delinquency and crime.

An alternative hypothesis is that the environments more than the specifics of the programs were important in determining impacts on crime. Chapel Hill is a mid-sized mostly white and relatively affluent university city in the South, while Ypsilanti is a smaller industrial city with a sizeable minority population. Chicago CPC sites were in low-income neighborhoods in a large urban mid-western city. (IHDP had sites throughout the United States.) It seems quite possible that the same program might have different effects in each city. As noted by Barnett and Masse (2007), crime rates were 70% higher in Ypsilanti than Chapel Hill when the respective program participants would have been age 15. They speculated that there may have been little crime to prevent among the Abecedarian sample; however, Clarke and Campbell (1998) reported that the two control samples (Perry and Abecedarian) had very similar arrest rates (around 40%) by their early 20s. McCormick et al. (2006) reported that juvenile arrest rates among controls were similar for the IHDP and Chicago CPC as well. So, among the target populations for these programs, crime rates were fairly similar even

if local crime rates were quite different. Of course, it is possible that the long-term effects of early childhood programs depend as much on the environment in which participants grow up as on individual and family characteristics of the participants themselves. If so, it is important to exercise caution in extrapolating benefits from any single program or community to the wider population.

Despite the fact that children targeted by all programs were disadvantaged, there is a sizeable difference in baseline educational attainment levels between Abecedarian and IHDP on the one hand and Chicago CPC and Perry Preschool on the other hand. High school graduation rates were 70% among Abecedarian controls; dropout rates (as of age 18) were only 10% among the IHDP controls. These both compare quite favorably with Chicago CPC and Perry Preschool controls who had high school completion rates ranging from 30–50%. Neither IHDP nor Abecedarian increased high school graduation rates. Although Abecedarian improved college attendance rates, this does not appear to be an important margin for crime (see Fig. 2.1). Given the tight link between high school dropout and crime discussed earlier, it may not be particularly surprising that Abecedarian and IHDP did not reduce crime given their negligible effects on high school completion. Yet, Perry Preschool substantially reduced male crime rates without raising educational attainment among males. Clearly, early interventions may reduce delinquency and criminal behavior without significantly improving final school-ing outcomes.

In the end, there is no easy explanation for the different findings across studies. Although the results from these studies are individually powerful given their research designs (most are based on random assignment), it is difficult to draw strong conclusions overall about the efficacy of early childhood interventions as a national crime-fighting strategy. The fact that sample sizes are quite modest and that program populations are not necessarily representative of the United States raises additional questions. This itself may explain some of the variation in findings across studies. It is natural to ask how these programs would affect other populations. Questions about scalability have also been raised: can these programs and their effects be reproduced at a larger scale? These issues have led a number of researchers to analyze the largest early childhood program in the United States: Head Start. This program targets children from low-income families usually living in low-income communities and has served hundreds of thousands of children throughout the United States since its inception in 1967.

Because no large-scale long-term random assignment studies of Head Start are available, researchers have employed nonexperimental methods. These studies generally examine impacts on national samples of individuals served by Head Start, using data from the Panel Survey of Income and Dynamics (PSID) or Children of the National Longitudinal Survey of Youth (CNLSY). We next discuss those studies that examine the impacts of Head Start on behavioral problems, delinquency, or measures of adult crime.

Garces, Thomas, and Currie (2002) and Deming (2009b) used a family-fixed effects approach to estimate the effects of Head Start on a variety of long-term outcomes. By comparing siblings who did and did not attend a Head Start program at ages 3-5, they address important concerns about permanent or long-run differences across families that may affect decisions about preschool or Head Start enrollment.⁴⁰ Garces, Thomas, and Currie (2002) used data from the PSID, examining adult outcomes for individuals born between 1964 and 1977, while Deming (2009b) used data from the CNLSY and examined outcomes for individuals born in the late 1970s and early 1980s. Despite using the same empirical approach, these two studies found quite different patterns for Head Start impacts on educational attainment and criminal behavior. Garces, Thomas, and Currie (2002) estimated significant increases in high school completion (by 20 percentage points) and college attendance (by 28 percentage points) for whites only, while Deming (2009b) estimated an 11 percentage point increase in high school completion rates and a 14 percentage point increase in college attendance for blacks only. Excluding GED recipients, Deming (2009b) estimated a smaller (7 percentage points) and statistically insignificant effect on high school completion for blacks, suggesting that much of their apparent improvement in high school completion is due to increases in the GED.⁴¹ Regarding crime, estimates by Garces, Thomas, and Currie (2002) suggest that Head Start reduces the probability of being booked or charged with a crime by about 12 percentage points among blacks, with no effect on whites. Deming (2009b) found no significant effects of Head Start on crime for blacks or whites.42

Carneiro and Ginja (2008) used a regression discontinuity design to estimate the effects of Head Start on adolescent outcomes, including the probability that someone is sentenced for a crime. Their approach exploits the fact that Head Start imposes strict eligibility criteria related to family income and structure: children aged 3–5 are eligible if family income is below the federal poverty guidelines or if the family is eligible for public assistance. Since these criteria vary across states and time, the income thresholds vary across these dimensions as well. They exploited this exogenous variation in eligibility, assuming the effects of family income (when children are ages 3–5) on subsequent outcomes are continuous. Using data from the CNLSY on youth who would have enrolled in Head Start during the 1980s and 1990s, they

⁴⁰ Of course, they leave unanswered the question as to why some siblings enroll in Head Start while others from the same family do not and, more importantly, whether different enrollment decisions are related to underlying differences in child abilities or other factors that may affect outcomes later in life.

⁴¹ The substitution between high school degrees and GED receipt is less relevant for the earlier cohort studied by Garces, Thomas, and Currie (2002), since the GED was much less common in the 1980s relative to more recent years.

⁴² His measure of crime is an indicator equal to one if the respondent reports having been convicted of a crime, been on probation, sentenced by a judge, or is in prison at the time of the interview.

estimated that participation in Head Start at ages 3–5 significantly reduces the probability (by 31 percentage points) a 16–17-year-old male is sentenced for a crime (based on self-reports). They estimated similar effects for a sample of blacks only. These estimates measure the effect of Head Start on children who were at the margin of eligibility for the program and, therefore, represent the effects we might expect with modest expansions of the program.

Altogether, the nonexperimental evidence on Head Start appears to suggest some long-term effects on education and crime, but findings vary in important ways across studies.⁴³ The strongest effects on crime appear to exist for blacks; although, Deming (2009) found no effect on crime for either blacks or whites. Combined with the evidence from smaller scale programs evaluated by randomized trials, there is limited but important evidence that early childhood interventions can reduce crime later in life for youth from disadvantaged backgrounds.

A recent program, Fast Track, introduced in four sites around the United States, provided group- and individual-based services to children from grades 1 through 10. The program specifically targets children from high crime and poverty neighborhoods, who exhibit conduct problems in kindergarten, with the primary aim of preventing antisocial behavior and psychiatric disorders. The program focuses on three elements of development: social and cognitive skills, peer relationships, and parenting. During early grades, parents were offered training and home visits to help improve parenting skills, while children were engaged in group activities to foster friendships and tutoring sessions in reading. As children aged, more individualized services were provided, along with group sessions aimed at dealing with the transition to middle school, resistance to drugs, and so on. The program also incorporated a classroom intervention during grades 1-5 at schools with program children. Teachers implemented 2-3 sessions per week designed to promote social and emotional competence and to reduce aggression. Experimental estimates based on random assignment suggest that the program produced sustained improvements in conduct disorders and antisocial behavior over grades 3-9 (Conduct Problems Prevention Research Group, CPPRG (2007)). As of grade 9, high-risk youth (those from the top 3% of conduct problems in kindergarten) receiving

⁴³ While Head Start may affect juvenile and adult crime even if it has no effect on educational attainment (as with males in the Perry Preschool program), one might speculate that any increases in schooling (especially high school years) associated with Head Start should lead to reductions in crime as estimated by Lochner and Moretti (2004). Under this assumption, estimates from Ludwig and Miller (2007), which suggest that roughly doubling Head Start spending (per capita) increases high school completion rates by as much as four percentage points, imply that this policy should also reduce arrest rates by up to 3–4%. Of course, multiplying the Garces, Thomas, and Currie (2002) estimated effects of Head Start on schooling attainment among whites by Lochner and Moretti's (2004) estimated effects of education on crime suggests that Head Start attendance should significantly reduce incarceration rates among whites, while analogous estimates from Deming (2009b) suggest that Head Start should reduce crime among blacks. Yet, these studies estimated no effect of Head Start attendance on self-reported measures of arrest, conviction or incarceration rates for these populations.

the Fast Track program showed significant reductions in self-reported delinquency and criminal behavior; however, no significant effects on antisocial behavior were found for other youth.⁴⁴ Two recent follow-up studies (CPPRG (2010, 2011)) suggest that the reductions in crime and conduct problems extend at least 2 years beyond the conclusion of the program (last measured at grade 12/age 19) and continue to be focused on youth that were initially "high risk." Effects on juvenile conduct disorders did not appear to decline after the program, while effects on crime showed some fade-out.

Experimental evaluations of two earlier, more limited elementary school-age, interventions are worth commenting on, since they also focused largely on social development among "high-risk" children. The Montreal Longitudinal Experimental Study provided social skills training to first- and second-grade children, along with teacher and parent training over those same years. Boisjoli, Vitaro, Lacourse, Barker, and Tremblay (2007) reported that by age 24, children receiving the intervention (compared with control children) were twice as likely to have completed high school and only half as likely to have a criminal record. The Seattle Social Development Project intervened over a longer period (grades 1–6); however, it only provided teacher and parent training (aimed at improving child social and emotional skill development). As of age 21, Hawkins, Catalano, Kosterman, Abbott, and Hill (2005) estimated that the 6-year intervention had increased high school graduation rates from 81 to 91% and significantly reduced self-reported crime and official lifetime court charges (from 53 to 42%).

Altogether, the evidence from Fast Track, the Montreal Longitudinal Experimental Study, and the Seattle Social Development Project suggests that comprehensive schoolage programs designed to improve social development can produce lasting impacts on educational attainment, conduct disorders, and criminal behavior. In many ways, these programs emphasized social over cognitive development relative to the preschool programs summarized in Table 2.1. Of course, both sets of programs were broad based and yielded improvements in both domains.⁴⁵

Programs targeted to older adolescents and young adults have shown mixed results. The Quantum Opportunity Program provided entering high school students with a mentor/tutor, who aided them in schoolwork and community activities for 4 years. Financial incentives designed to encourage high school graduation and college enrollment were provided for educational, service, and developmental activities. A recent random assignment evaluation of the program reported no significant improvements

⁴⁴ Results for antisocial behavior are based on an index created from self-reports of serious delinquent/criminal actions like stealing something worth more than \$100, assault, selling heroin or LSD, and sexual assault.

⁴⁵ This is largely consistent with recent estimates of skill production functions for both cognitive and 'noncognitive' skills (e.g., see Cunha and Heckman (2008)).

in schooling or reductions in crime 6 years after scheduled high school graduation (Schirm, Stuart, and McKie (2006)). In part, this may be due to the relatively low participation by youth in program activities.⁴⁶

The Job Corps provides intensive basic educational and vocational training for economically disadvantaged youth and young adults ages 16–24 throughout the United States. The program also offers a wide range of other services (e.g., counseling, social skills training, health education, job placement services). The average participant is enrolled for about 8 months, with most living in residence at training sites. The program's primary goal is to improve employment and earnings prospects. Based on a recent random assignment evaluation, Schochet, Burghardt, and Glazerman (2001) concluded that the program produced modest positive impacts on post-program employment and earnings. The program also reduced self-reported arrest rates by about 30% during the first year after random assignment, when most youth would have been enrolled. Reductions in subsequent years were smaller and statistically insignificant. The program also significantly reduced conviction rates by about 17% during the 4 years following random assignment.⁴⁷ Conclusions from the less-expensive and nonresidential JOBSTART program are largely consistent with these findings (Cave et al. (1993)).⁴⁸

Collectively, these studies indicate that human capital-based interventions from early childhood to early adulthood can reduce juvenile and adult crime, at least for some populations. To understand why, it is useful to return briefly to the model developed earlier. The model suggests that effective interventions may reduce juvenile and adult crime by improving child learning productivity, A, increasing adolescent human capital levels, H_0 , or by socializing children (i.e., lowering θ). Although preschool programs highlighted in Table 2.1 may raise learning abilities, achievement gains are generally short-lived and limited to primary school ages. Evidence of reduced criminal activity among adolescents attributed to early intervention programs suggests that these programs raise adolescent market skills (H_0) and/or reduce criminal returns (θ) through socialization. School-based programs for high-risk youth like Fast Track emphasized social development

⁴⁶ On average, youth spent only 177 hours per year on educational, community, and developmental activities. Roughly one-in-four spent no time at all in these activities by the fourth year of the program.

⁴⁷ An earlier study by Long, Mallar, Thornton (1981) estimated that the social benefits from reduced criminal activity among Job Corps participants amounted to over \$7000 (in 2008) per participant—almost 30% of the total social benefit of the program.

⁴⁸ JOBSTART offered many of the same basic components of the Job Corps to a similar population. Cave et al. (1993) found modest (and statistically insignificant) positive effects on earnings 3–4 years after random assignment for the full sample; however, earnings increased roughly 25% (in years 3 and 4) for male participants with a prior arrest (i.e., had an arrest since age 16 but prior to random assignment). Among male participants with no prior arrest, the program significantly reduced self-reports of an arrest (6.4 percentage points or 36%) during the first year after random assignment (i.e., the training year) but did not reduce the fraction arrested in subsequent years. Among males with a prior arrest, the program (insignificantly) reduced the fraction reporting an arrest over the first 4 years after random assignment by about 8% and had negligible effects on arrests during the first year. There were no significant effects on arrests for female participants.

(i.e., lowering θ) over cognitive achievement; yet, they also likely improved adolescent human capital levels H_0 . Despite the difference in emphasis between the two types of programs, both have shown the ability to significantly reduce juvenile and adult crime. Job training programs for adolescents and young adults directly operate on the incentives to invest in human capital and have led to modest reductions in crime during periods of heavy training. These programs have produced only modest increases in earnings and negligible long-run effects on crime, however, suggesting that simply training low-skilled adolescents does not provide the same promise as earlier interventions that act on individual endowments.

2.4. Policy Lessons

We next discuss a number of important policy lessons regarding human capital policies and crime. First, we summarize evidence on the social savings from crime reduction that we might expect from policies that increase educational attainment or enrollment, improve school choice and quality, or expand access to early childhood interventions. Second, we highlight a few subpopulations and schooling margins that are likely to yield the greatest social gain from crime reduction. Finally, we discuss a few other lessons based on the evidence.

2.4.1 Valuing the Social Benefits from Crime Reduction

Lochner and Moretti (2004) estimated that increasing educational attainment levels in the population yields sizeable social benefits. Specifically, they calculated the social savings from crime reduction that would result from a one percentage point increase in high school graduation rates in the United States. Table 2.2 summarizes their exercise, translating all dollar values into 2008 dollars using the Consumer Price Index for All Urban Consumers (CPI-U). Column 1 reports total costs per crime associated with murder, rape, robbery, assault, burglary, larceny/theft, motor vehicle theft, and arson.⁴⁹ Column 2 reports the predicted change in total United States arrests based on the Lochner and Moretti (2004) offense-specific arrest estimates discussed earlier and the total number of arrests in the 1990 Uniform Crime Reports. Column 3 adjusts the arrest effect in column 2 by the number of crimes per arrest. In total, nearly 100,000 fewer crimes would have taken place in 1990 if high school graduation rates had been one percentage point higher. The implied social savings from reduced crime are shown in column 4. Savings from murder alone are as high as \$1.7 billion. Savings from reduced assaults amount to nearly \$550 million. Because the estimates suggest that schooling increases rape and robbery

⁴⁹ These costs include incarceration and victim costs. See notes to Table 2.2 or Lochner and Moretti (2004) for details.

	Total Cost	Estimated Change	Estimated Change	
	per Crime	in Arrests	in Crimes	Social Benefit
Violent crimes				
Murder	4,506,253	-373	-373	\$1,683,083,243
Rape	132,938	347	1559	-\$207,270,899
Robbery	13,984	134	918	-\$12,839,495
Assault	14,776	-7798	-37,135	\$548,690,721
Property crimes				
Burglary	1471	-653	-9467	\$13,920,409
Larceny/theft	295	-1983	-35,105	\$10,347,853
Motor vehicle	1855	-1355	-14,238	\$26,414,558
Theft				
Arson	58,171	-69	-469	\$27,302,131
Total		-11,750	-94,310	\$20,898,648,519

Table 2.2 Social Benefits of Increasing High School Completion Rates by 1%

Source: Lochner and Moretti (2004).

Notes: These costs reflect incarceration and victim costs. Victim costs are taken from Miller, Cohen, and Wiersema (1996). Incarceration costs per crime equal the incarceration cost per inmate multiplied by the incarceration rate for that crime (approximately \$25,000). Incarceration rates by offense type are calculated as the total number of individuals in jail or prison (from US Department of Justice (1994)) divided by the total number of offenses that year (where the number of offenses is adjusted for nonreporting to the police). Incarceration costs per inmate are taken from US Department of Justice (1999). All dollar figures are translated into 2008 dollars using the CPI-U.

offenses, increased costs associated with these crimes partially offset the benefits from reductions in other crimes.

The final row reports the total savings from reductions in all eight types of crimes. Because these figures only include a partial list of crimes (e.g., nearly 25% of all prisoners in 1991 were incarcerated for drug offenses according to the US Department of Justice (1994)) and do not include all costs associated with each crime (e.g., private security measures are omitted), these amounts are likely to underestimate the true social benefit associated with increasing high school graduation rates. Still, the savings are substantial: the social benefits of a one percentage point increase in male US high school graduation rates (from reduced crime alone) in 1990 would have amounted to more than \$2 billion.⁵⁰ This represents more than \$3000 in annual savings *per additional male graduate*.

Machin, Marie, and Vujic (2011) also estimated the social savings from crime reduction associated with increasing the population of individuals with an education qualification (similar to high school completion in the United States). Accounting only for benefits from property crime reduction (estimated effects on violent crime are

⁵⁰ These calculations are partial equilibrium in nature, implicitly assuming that the skill prices of workers do not change in response to a policy that increases national high school completion rates. Fella and Gallipoli (2009) showed that the social benefits of such a policy would be even larger if general equilibrium effects on the wages of high school dropouts and graduates were incorporated.

statistically insignificant), their estimates suggest a savings of over $\pounds 10,000$ per additional student qualification.

Open school enrollment lotteries and desegregation efforts appear to reduce crime rates by improving school quality. Deming (2009a) estimated that reductions in arrests associated with offering better quality school options to a high-risk youth produces a roughly \$16,000 social savings to victims over the next 7 years. Because better schools are also likely to have reduced crimes that never led to an arrest, total victimization savings are likely to be substantially higher. Total social savings should be still larger once savings on prisons and other crime prevention costs are factored in.

The effects of school attendance on contemporaneous juvenile crime rates are more complicated. Studies estimating the effects of day-to-day changes in attendance suggest that in urban communities additional school days reduce property crime while increasing violent crime (Jacob and Lefgren (2003) and Luallen (2006)). Overall, the social costs associated with increased violence are likely to dominate the benefits from reduced property crime. On the other hand, Anderson (2009) estimated reductions in both violent and property juvenile crime associated with increases in compulsory schooling ages. Thus, his findings suggest an overall social savings from juvenile crime reduction, although he does not attempt to put a dollar value on the effects.

Evidence on the effects of early childhood and school-age interventions are mixed. Long-run impacts on juvenile delinquency and adult crime can be substantial for disadvantaged youth. For example, estimates suggest that Perry Preschool produced a social benefit from crime reduction of roughly \$150,000 per child (through age 40). On the other hand, Abecedarian produced no significant impacts on crime.

2.4.2 Where Are the Big Returns?

Given that the most sizeable reductions in crime appear to result from the final years of high school, policies that encourage high school completion would seem to be most promising in terms of their impacts on crime.⁵¹ Because crime rates are already quite low among high school graduates, policies that encourage college attendance or completion are likely to yield much smaller social benefits from crime reduction; although, they may be desirable on other grounds.⁵²

In general, policies designed to encourage schooling among more crime-prone groups are likely to produce the greatest benefits from crime reduction. Consistent with this, the school-age Fast Track program appears to have reduced juvenile crime only among very high-risk children, showing little impact on even moderately high-risk

⁵¹ See Hanushek and Lindseth (2009), Jacob and Ludwig (2008), or Murnane (2008) for recent discussions of policies to improve schooling outcomes in the United States.

⁵² The fact that crime declines substantially with high school completion but not college attendance suggests that net expected returns from crime for most individuals lie somewhere between the wages of high school dropouts and graduates. See Freeman (1999) for a summary of evidence regarding criminal wages and earnings.

children (CPPRG (2007, 2010)). Similarly, Deming (2009a) estimated that improved school choice for middle and high school students leads to significant reductions in arrests for high-risk youth but not for others. As Donohue and Siegelman (1998) concluded, the overall efficiency of early childhood programs as a crime-fighting strategy is likely to depend heavily on the ability to target high-risk children at very young ages. The same is likely to be true for school-age interventions.

Social benefits from crime reduction also vary across gender and race. Men commit much more crime than women, on average. Thus, it is not surprising that crime-related benefits from education policies and interventions are typically much smaller for females than males (e.g., Perry Preschool, Job Corps). This is true even though programs sometimes reduce female and male crime rates by similar amounts in percentage terms. Among men, Lochner and Moretti (2004) estimated much larger effects of additional schooling on incarceration rates among blacks relative to whites. Garces, Currie, and Thomas (2002) estimated that Head Start significantly reduced crime for blacks but not whites; however, Deming (2009b) estimated no effect on crime for either group, while Carneiro and Ginja (2008) estimated similar large effects on both. Because crime rates are much higher among blacks than whites, on average, policies would generally need to produce much larger proportional reductions in white crime rates to achieve similar absolute reductions in crime. None of the evidence surveyed here suggests that this is the case.

2.4.3 Additional Policy Lessons

A few other useful lessons can be drawn from the studies surveyed here.

First, education policies can reduce property crime and violent crime. In the United States, the estimated effects of educational attainment or school enrollment on property and violent offenses appear to be quite similar in percentage terms (Lochner and Moretti (2004) and Anderson (2009)).⁵³ Even murder appears to be quite responsive to changes in educational attainment and school quality (Lochner and Moretti (2004) and Weiner, Lutz, and Ludwig (2009)).

Second, higher wages increase the opportunity costs (including work foregone while incarcerated) of both property and violent crime. Lochner and Moretti (2004) showed that the estimated effects of educational attainment on crime can largely be accounted for by the effects of schooling on wages and the effects of wages on crime. This is important, since it suggests that policymakers can reduce crime simply by increasing labor market skills; they need not alter individual preferences or otherwise socialize youth.⁵⁴ Of course, as evidence from the Job Corps and other training programs suggests, this is not

⁵³ Estimates from Machin, Marie, and Vujic (2011) suggest that education reduces property crime more than violent crime in the United Kingdom.

⁵⁴ As Heckman, Stixrud, and Urzua (2006) showed, both "cognitive" and "noncognitive" skills are acquired in school, are rewarded in the labor market, and affect crime.

necessarily an easy task. Training programs targeted at low-skill adolescents and young adults have modest (at best) effects on earnings and crime. On the other hand, encouraging youth to finish high school (e.g., through compulsory schooling laws) appears to substantially increase earnings and reduce crime. Preventing early school dropout is likely to be more successful than trying to compensate for dropout a few years later.

Third, education-based policies need not increase educational attainment to reduce crime. Studies on school choice lotteries (Cullen, Jacob, and Levitt (2006) and Deming (2009a)) suggest that providing disadvantaged urban youth with better schools can substantially reduce juvenile and adult crime, even if it has small effects on traditional education outcomes. Perry Preschool had no effect on male schooling levels but sub-stantially reduced male crime rates through age 40 (Schweinhart et al. (2005)).

Fourth, evidence that violent crime is higher on school days than nonschool days in urban districts suggests that social interaction effects are particularly important for juvenile violent crime (Jacob and Lefgren (2003) and Luallen (2006)). Smart policing efforts may be able to help address some of the problems associated with schools releasing lots of adolescents at the same time. For example, an increased police presence immediately after school or other major adolescent congregations let out may be warranted. Or, on nonschool days, it may be wise for police to focus more on targets or areas of juvenile property crime, worrying less about violent crime. The "hot spot" or "problemoriented policing" literature in criminology suggests that informed targeting of police efforts to high-crime areas (and, by extension, times) can be effective at reducing overall crime rates.⁵⁵ Alternatively, it may be useful to consider ways of designing after-school youth programs or other weekend activities to minimize violent behavior afterwards.

2.5. Interpreting the Evidence on Crime

Much of the empirical evidence is broadly consistent with a human capital-based model of crime and work premised on the notion that schooling raises legitimate wage rates much more than the returns to most types of crimes. Indeed, Lochner and Moretti (2004) argued that the reductions in violent and property crime associated with increased schooling in the United States is roughly equivalent to the effect of education on wages multiplied by the effect of increased wages on crime.⁵⁶ Thus, most of the effect of education on violent and property crime may come from increased wages.

⁵⁵ For a recent survey of this literature, see Braga (2005).

⁵⁶ Machin and Meghir (2004) estimated the elasticity of property crime to low-skill wages to be a little over one in the United Kingdom. Grenet (2010) estimated that an extra year of school (induced by the United Kingdom's 1973 increase in compulsory schooling) increases wages by about 7%. Combining these figures, the effect of an extra year of school through increased wages should be to lower property crime by 7–10% in the United Kingdom. Estimates from Machin, Marie, and Vujic (2011) suggest that an extra year of school reduces overall property crime by 10–15% (20–30% among men and very little among women). Altogether, these calculations suggest that increased wages explain anywhere from one-half to all of the effect of education on property crime in the United Kingdom.

By contrast, education may increase the returns to white-collar crime more than the returns to work. Consistent with this, Lochner (2004) found that arrest rates for white-collar crime increase when education levels rise.

The human capital-based model of crime is further supported by recent trends in incarceration differentials by education in the United States as documented by Western (2006). Since the early 1980s, the United States has witnessed a dramatic increase in the wage returns to schooling, the probability of incarceration, and prison sentence lengths. All of these trends have served to increase the opportunity costs of crime relatively more for the most educated. As predicted by the model, incarceration rates have become more strongly correlated with education over this period.

Education-based programs may also socialize youth, reducing personal or psychic rewards from crime. Heckman and Rubinstein (2001) argued that "noncognitive" skills may be more important determinants of many life outcomes than "cognitive" skills among those at the bottom of the education distribution. Emphasizing social and emotional development, school-age programs like Fast Track have shown the ability to significantly reduce later conduct disorders and crime (among high-risk children). These programs also improved educational outcomes, which may explain some of their impacts on crime. Perry Preschool reduced male (and female) crime rates without affecting male schooling outcomes. Thus, the program appears to have improved social development or increased early skill levels (without noticeably affecting subsequent schooling investments). Evidence from school choice lotteries suggests that improvements in school and peer quality can lead to reductions in crime without raising student achievement or educational attainment. The most likely explanation for the reduction in crime is that higher quality schools better socialize youth or provide them with a better set of peers. Yet, evidence from the MTO experiment suggests that moving families to lower poverty neighborhoods does not produce the same reductions in crime, complicating any explanation related to peer effects or social networks.

Education may also increase patience or alter preferences for risk; however, neither seems to be central to the estimated impacts on crime. Property crimes are generally associated with less than 1 month of expected time in jail or prison conditional on being sentenced (see Table 2.3), hardly enough time for modest changes in patience to play much of a role. Property crimes also have very low expected probabilities of arrest (typically less than 10% chance) and even lower probabilities of incarceration (typically less than 1%), so there is little actual uncertainty in outcomes associated with these crimes (see Table 2.3). Yet, estimated impacts of schooling on property crime are similar to those for violent crime, which entails much longer and more uncertain prison sentences.

Altogether, the evidence suggests that while efforts to socialize youth can be effective, simply providing them with valuable market skills can discourage them from choosing a life of crime. In terms of crime reduction, human capital-based policies that

Crime	Probability of Arrest	Probability of Conviction Conditional on Arrest	Probability of Incarceration Conditional on Conviction	Unconditional Probability of Incarceration	Estimated Months Served if Incarcerated	Expected Days Served per Crime Committed
Violent crimes	0.25	0.22	0.79	0.043	91	119.4
Murder	0.85	0.67	0.95	0.544	248	4102.4
Forcible rape	0.15	0.39	0.90	0.051	136	212.2
Robbery	0.15	0.36	0.89	0.047	94	134.8
Aggravated assault	0.30	0.17	0.71	0.035	59	63.2
Property crimes	0.06	0.11	0.68	0.004	24	3.2
Burglary	0.07	0.27	0.76	0.015	29	13.2
Larceny-theft	0.05	0.08	0.61	0.002	20	1.4
Motor vehicle theft	0.10	0.08	0.73	0.006	17	3.1

Table 2.3 Expected Punishment Associated with Incarceration (Uniform Crime Reports)

Source: Lochner (2004).

Notes: Probability of arrest computed from crimes and arrests in the United States (from the 2000 Uniform Crime Reports) adjusted for nonreporting to the police (from the 2000 National Criminal Victimization Survey). It is assumed that all murders are reported to the police. Probability of conviction conditional on arrest divides total arrests in the United States by total State and Federal convictions for 2000. Probability of incarceration conditional on conviction is based on reporting of State courts. Estimated months served if incarcerated applies to State prisoners and is estimated by the US Department of Justice based on sentence lengths handed out that year and the average percent of sentences served by prisoners released that year. Unless otherwise noted, all criminal justice figures are for 2000 and are taken from Durose and Langan (2003).

target the most disadvantaged (and crime-prone) are likely to be the most efficient while also promoting a more equitable society. To that end, increasing high school graduation rates and improving our nation's worst inner city schools are likely to yield the greatest social return.

Although policies that increase school attendance for a year or more (e.g., increased compulsory schooling ages) appear to reduce both violent and property crime (Anderson (2009)), a few extra days off from school may actually lead to reductions in violent crime, especially in urban areas (Jacob and Lefgren (2003) and Luallen (2006)). From a human capital perspective, the increased opportunities that open up for youth attending an additional year of schooling should raise the future costs of incarceration associated with juvenile crime. This may serve as an important additional criminal deterrent that does not exist for day-to-day changes in the school calendar. In general, the effects of longer periods of attendance on contemporaneous juvenile crime are consistent with the subsequent effects of additional schooling on adult crime.

There are many ways by which early childhood interventions may affect juvenile and adult crime. The human capital approach favored in this chapter highlights the potential effects of these programs on learning abilities, adolescent skill levels, and socialization or tastes for crime. These programs may also affect childhood preferences, including risk aversion, patience, or self-control. Although a few early childhood programs have produced sizeable reductions in both juvenile and adult crime-most famously, Perry Preschool-other quite similar programs have not. School-age interventions focused on developing social and emotional skills have proven successful at reducing later conduct disorders and crime, especially among very high-risk children. The benefits from reduced crime associated with successful programs certainly warrant the attention they have received; yet, we still need to know much more about why other programs have not produced the same effects. Two things are clear. First, preschool and school-age programs have substantially reduced crime for some disadvantaged high-risk populations. Even if these gains cannot be expected in all cases, they are large enough to warrant careful consideration on a broader scale. Second, successful programs did not always increase educational attainment, even when they significantly reduced juvenile and adult crime rates. Thus, disappointing achievement or educational outcomes need not imply the absence of benefits from crime reduction.

Given current evidence, it is difficult to draw strong conclusions about the relative benefits of trying to target and "treat" children at very young ages versus intervening at later ages to keep adolescents from dropping out of high school. Of course, we need not choose one or the other. Indeed, both are likely to be important components of a broad-based national crime-fighting agenda. Calculations by Lochner and Moretti (2004) and Donohue and Seigelman (1998) suggest that both human capital-oriented policies are competitive with more traditional law enforcement and incarceration efforts when all benefits are considered.

3. EDUCATION'S EFFECTS ON HEALTH AND MORTALITY

Health and mortality gaps by education are large and have been growing for decades (see, e.g., Pappas, Queen, Hadden, and Fisher (1993); Goesling (2007); and Meara, Richards, and Cutler (2008)). In 2000, white males with at least some college education could expect to live 6.2 years more than their less educated counterparts (Meara, Richards, and Cutler (2008)). As highlighted by recent surveys of the literature on education and health (e.g., Grossman and Kaestner (1997) and Grossman (2000, 2006)), education is more strongly correlated with health than is income or occupation. As discussed below, the literature has identified many reasons education may improve health and reduce mortality.

We address three questions in this section: (1) Does education actually improve health and mortality outcomes (i.e., is there a causal effect of education)? (2) If so, why? and (3) What does this mean for education or health policy?

We first develop a simple model of education and mortality to shed light on a few key mechanisms at play. In particular, the model incorporates three often-discussed channels: First, education may directly increase health production by raising the marginal productivity of health inputs or behaviors, sometimes referred to as "productive efficiency" (Grossman (1972a)). Second, education may enhance one's ability to acquire and process health information or to follow more complicated treatments. As a result, education may improve the "allocative efficiency" of health inputs (Rosenzweig and Schultz (1982)). Third, education generally increases earnings, which makes costly health care and insurance purchases more affordable. An increase in income also raises the demand for health and longevity by increasing consumption opportunities (Grossman (1972a)). Overall, the model clarifies how education affects mortality directly and indirectly through the choice of health inputs.

We next discuss recent evidence on whether education has a causal effect on health and mortality.⁵⁷ In particular, we focus on the growing number of studies that attempt to address well-known concerns about the endogeneity of schooling and unobserved heterogeneity in preferences using instrumental variable (IV) or regression discontinuity (RD) techniques. We discuss both the credibility of these studies and their findings, concluding that the most convincing of these studies find that additional schooling leads to modest improvements in health and small reductions in mortality.

Much of the most recent literature on education and health has attempted to explain why schooling may improve health and reduce mortality. Most of these studies take one of two approaches. Some studies attempt to empirically decompose the education–health gradient by regressing various health outcomes or behaviors on education and different subsets of potential mediators (e.g., income, ability, preferences, information about

⁵⁷ For surveys of earlier work, see Grossman and Kaestner (1997) or Grossman (2000, 2006).

treatments or diseases) to see how much of the education-health gradient can be "explained by" these mediators. Other studies focus on a single potential mechanism, theory, or explanation. Most studies (by economists) in this vein have attempted to determine whether differences in health/mortality by education are driven by differences in specific health knowledge or behavior. We briefly review both types of studies and discuss the strengths and weaknesses of each approach.

Finally, we discuss a few important policy issues. First, we argue that it is important to determine whether individuals implicitly or explicitly pay for any health benefits associated with greater education (e.g., better health insurance or more expensive treatments). This is crucial for determining the net return on education and the value of education as a health-improving policy goal, since additional costs should be netted out. Second, it is important to determine whether health benefits accrue exclusively to the individual who becomes educated, to other family members, or to broader social networks and society at large. Evidence of important health externalities may justify public expenditures on schooling. Third, it is reasonable to ask whether individuals are aware of future health benefits associated with additional schooling. If health benefits are effectively free and individuals are unaware of those benefits (or fail to incorporate them) when making their education decisions, governments may wish to subsidize education even in the absence of externalities.

3.1. A Simple Model of Education and Mortality

Grossman (1972a, 1972b, 1975) developed a dynamic model of health investment related to the traditional human capital investment models of Becker (1967) and Ben-Porath (1967).⁵⁸ In his framework, health affects the time endowment for workers, while education may affect health and returns in the labor market. Grossman (2006) discussed key aspects of the model in a simpler static environment.

In this section, we develop a simple model in which education and "health inputs" affect mortality. The model helps formalize and clarify three mechanisms: productive efficiency, allocative efficiency, and differential demand for health by income. It also facilitates a "back of the envelope" calculation offered below on the potential benefits of education through improved mortality, an important and objective measure of health. In many ways, the model developed here is similar to that of Becker and Mulligan (1997), which allows individuals to affect their rate of time preference through costly up-front expenditures.⁵⁹

⁵⁸ Also, see Ehrlich and Chuma (1990).

⁵⁹ There are some important differences. Becker and Mulligan (1997) assumed a fixed resource endowment, the cost of increasing patience (survival probability in our case) is paid up front with a one-time expenditure, and individuals can borrow and save to smooth consumption optimally over time. Given the emphasis on mortality here, our model assumes that individuals receive income each period they remain alive and must pay each period to reduce their mortality rate. For simplicity, we assume that individuals cannot borrow or save.

Individuals wish to maximize discounted lifetime utility from consumption *c*, but they face some probability of death each period. Schooling $s \in \{\underline{s}, \overline{s}\}$ may directly improve the survival probability ("productive efficiency") or individuals may purchase costly inputs *x* at price *p* to increase this probability. While we model these inputs as goods with price *p*, one could easily introduce other costs more directly related to foregone opportunities (e.g., giving up smoking or drinking). We assume that the survival probability at any age, $\pi(x, s)$, is strictly increasing and concave in inputs $x \ge 0$ and increasing in schooling *s*. We also assume complementarity between health inputs and schooling: $\frac{\partial^2 \pi}{\partial x \partial s} \ge 0$. Furthermore, $\pi(0, \underline{s}) \ge 0$ and $\lim_{x\to\infty} \pi(x, \overline{s}) = \overline{\pi} < 1$, ensuring that people have some probability of surviving each period but do not live forever. Since we are not concerned with life-cycle issues here, we assume that $\pi(x, s)$ is a time invariant function and that individuals must choose the same *x* at all ages. Naturally, one would want to relax these assumptions to study how health or mortality investments vary over the life-cycle.

Schooling also increases individual earnings $\gamma(s)$ so $\gamma'(s) \ge 0$. However, we ignore any life-cycle effects on income in assuming that postschool earnings are constant over time. We also abstract from savings/borrowing decisions and assume consumption $c = \gamma(s) - px$. While alive, individuals receive utility u(c) from consumption at each date and discount utility from future consumption at rate $\tilde{\beta} \in (0, 1)$. It is also useful to define the "effective discount rate" $\beta(x, s) \equiv \tilde{\beta}\pi(x, s)$. Without loss of generality, we effectively normalize the utility of death at zero.⁶⁰

We begin by considering optimal health input decisions given schooling attainment *s*. Given our assumptions, individuals who have completed *s* years of school solve the following problem:

$$V(s) = \max_{x:x \ge 0} \left\{ \sum_{t=0}^{\infty} [\widetilde{\beta}\pi(x,s)]^{t} u[\gamma(s) - px] \right\}$$
$$= \max_{x:x \ge 0} \left\{ \frac{u[\gamma(s) - px]}{1 - \beta(x,s)} \right\}.$$

Optimal health input x decisions depend on the marginal benefits less the marginal costs. These net marginal benefits are reflected in the following function:

$$N(x,s) \equiv \frac{\partial \beta(x,s)}{\partial x} u(c) - [1 - \beta(x,s)] p u'(c), \qquad (2.9)$$

where the first term reflects the marginal benefits of extending life from an increase in health expenditures, while the second term reflects the loss in utility from foregone

⁶⁰ The "bad" state in our model could be something other than death. The key assumption is that utility in this state must be independent of schooling and health expenditures.

consumption. If inputs x are not sufficiently beneficial, some individuals may choose not to purchase them at all (i.e., individuals with N(0, s) < 0). Given our assumption that $\frac{\partial^2 \pi}{\partial x \partial s} \ge 0$, there exists a threshold schooling level \hat{s} , above which individuals invest $x = x^*(s) > 0$ (satisfying the interior first-order condition $N(x^*(s), s) = 0$), and below which individuals invest $x = 0.^{61}$ (In general, we would not expect all health inputs to be strictly zero for any individuals. Allowing for multiple inputs would generally yield separate thresholds for each input, such that more educated individuals would utilize a wider range of inputs.)

Among individuals with $s > \hat{s}$ who purchase positive health inputs, it is straightforward to show that those inputs are increasing in completed schooling:

$$\frac{\mathrm{d}x^*(s)}{\mathrm{d}s} = \frac{\left[\frac{\partial^2\beta}{\partial x\partial s} + \left(\frac{1}{1-\beta}\right)\frac{\partial\beta}{\partial s}\frac{\partial\beta}{\partial x}\right]u(c) + \left[\frac{\partial\beta}{\partial x}u'(c) - (1-\beta)pu''(c)\right]\gamma'(s)}{-\frac{\partial^2\beta}{\partial x^2}u(c) - p^2(1-\beta)u''(c)} > 0.$$

This reveals two main effects of past schooling on health inputs. The first term in the numerator reflects the fact that schooling and health inputs are complementary, both directly and indirectly. The indirect complementarity arises because the marginal value of extending life today is greater if future survival rates are also high. Thus, by increasing future survival rates, schooling raises the marginal value of investing in health today. The second term in the numerator reflects simple income effects due to the market returns from schooling. By increasing earnings, schooling increases the demand for longevity and, therefore, health inputs. Overall, health inputs may be zero for a low range of schooling levels ($s \le \hat{s}$) and will be strictly increasing in schooling among those with $s > \hat{s}$.

Since health inputs are increasing in schooling, it is clear that

$$\frac{\mathrm{d}\pi}{\mathrm{d}s} = \frac{\partial\pi}{\partial s} + \frac{\partial\pi}{\partial x}\frac{\mathrm{d}x}{\mathrm{d}s} > \frac{\partial\pi}{\partial s} \ge 0$$

and that schooling increases survival rates. Among individuals with low schooling levels $(s \leq \hat{s})$, the total effect of schooling equals the partial effect (i.e., $\frac{d\pi}{ds} = \frac{\partial \pi}{\partial s} \geq 0$). Among more educated individuals with $s > \hat{s}$, the total effect of schooling $\frac{d\pi}{ds}$ on mortality is greater than its partial effect $\frac{\partial \pi}{\partial s}$ due to the responsiveness of health inputs. Indirect effects through increased health expenditures will be large when (1) schooling and health inputs exhibit strong direct complementarity, (2) schooling raises income a lot, and (3) mortality is very responsive to health inputs and/or schooling. This suggests that

⁶¹ The complementarity assumption $\frac{\partial^2 \pi}{\partial x \partial s} \ge 0$ ensures that N(0, s) is increasing in *s*, yielding the threshold \hat{s} as the solution to $N(0, \hat{s}) = 0$. Second-order conditions for a maximum are satisfied by the assumptions on $\pi(x, s)$ and u(c).

productive efficiency may lead to a stronger education-health/mortality gradient at the top end of the education distribution (i.e., $s > \hat{s}$), where health input decisions play a more important role.

It is interesting to consider how a change in the price of health inputs affects health decisions. It is straightforward to show that input purchases are decreasing in price $\left(\frac{dx^*}{dp} < 0\right)$, while \hat{s} is increasing in prices. Thus, innovations that lower the price of health inputs will cause highly educated individuals to increase their health investments and some less-educated individuals to begin purchasing the cheaper input. The latter effect will tend to shrink health and mortality differences by education; however, it is more difficult to sign the effect of a price change $\frac{dx^*}{ds}$ on those already using the input. While a decline in the price of an input should unambiguously reduce education differences in the fraction of individuals using that input, it could actually widen education differences in the quantity of inputs used (among initial users).

What happens if new information arises suggesting that some particular health input (or behavior) is good for one's health? For example, suppose $\frac{\partial \beta}{\partial x}$ was previously thought to be zero but is now thought to improve health, that is, $\frac{\partial \beta}{\partial x} > 0$. This new information should cause individuals with $s > \hat{s}$ to begin using the new input, with the most educated increasing their use the most. By contrast, lower-educated individuals may not respond at all to the new information. The higher the price of the new input, the narrower the response will be among the most educated. In general, new information that some input or behavior is beneficial for health is likely to increase education—health/ mortality gradients even if that information is distributed and interpreted equally by everyone.⁶² We would expect a similar response to the introduction of new health inputs.

3.1.1 Valuing Mortality Benefits from Education

Using the envelope condition, it is straightforward to derive the marginal lifetime benefit of schooling:

$$V'(s) = \left[\frac{u'(c)}{1-\beta}\right] \left[\left(\frac{\widetilde{\beta}}{1-\beta} \frac{u(c)}{u'(c)}\right) \frac{\partial \pi}{\partial s} + \gamma'(s) \right]$$

= $\left[\frac{u'(c)}{1-\beta}\right] \left[\text{VSL} \times \frac{\partial \pi}{\partial s} + \gamma'(s) \right],$ (2.10)

The marginal value of schooling can be decomposed into a "mortality effect" and a standard lifetime earnings effect. Interestingly, the mortality effect can be calculated from

⁶² As with price changes, it is more difficult to generally sign the response of $\frac{dx^*}{dp}$ to marginal increases in the productivity of inputs when some of the population is already using *x*.

standard estimates of the value of a statistical life (VSL) and the estimated (partial) effects of schooling on survival rates.⁶³

Below, we calculate the "mortality effect" using plausible values for the effect of schooling on mortality and measures of the value of a statistical life. Empirically, it is difficult to distinguish between the total effect of schooling on mortality $\frac{d\pi}{ds}$ and the partial effect $\frac{\partial \pi}{\partial s}$ observed in Eq. (2.10). Most recent studies that attempt to estimate the causal effects of schooling on health estimate the total effect rather than the partial effect; however, as this analysis and our discussion below suggests, the partial effect is crucial for education and health policy.

3.1.2 Allocative Efficiency

The notion that education may affect the allocation of health inputs is often discussed by economists. See Rosenzweig and Schultz (1982) for an early treatment of the issue. Health economists have discussed a number of reasons more educated individuals may make different health-related choices. For example, more educated individuals may be better at acquiring or processing health information or they may better understand and follow complicated directions. Education may also raise the demand more for some health inputs than others. We briefly discuss some of these ideas within the context of a generalized version of the above model.

To incorporate the idea of "allocative efficiency," consider two potential inputs x_1 and x_2 with prices p_1 and p_2 , respectively. To keep things very simple, assume that these inputs are perfect substitutes in production (e.g., two different forms of medicine for the same disease), so we can write the survival probability as $\pi(a_1x_1 + a_2x_2, s)$.

Except when $\frac{a_1}{a_2} = \frac{p_1}{p_2}$, perfect substitutability implies that individuals will never purchase both inputs. They will purchase input x_1 if $\frac{a_1}{a_2} > \frac{p_1}{p_2}$; otherwise, they will choose x_2 . The least-educated individual may choose to purchase neither input. The characterization of the above input decisions carries over here to the desired input. As such, the introduction of a new input x_2 that is equally productive as x_1 (i.e., $a_1 = a_2$) but which has a lower price should cause individuals to switch to the new input, choosing a higher input level. The effects on health are equivalent to those observed for a price reduction in the more simple previously discussed single-input model.

$$VSL \equiv \frac{d\gamma_t}{d\pi_t} \mid_{V_t = \overline{V_t}} = \frac{\widetilde{\beta}}{1 - \beta} \frac{u(c)}{u'(c)}.$$

⁶³ The VSL is typically estimated as the amount of money someone would need to be compensated at date *t* to raise their mortality risk by some amount $d\pi_t$. On the basis of our model, we can write the discounted utility for someone as of date *t* as $V_t = u(y_t - x) + \tilde{\rho}\pi_t[u(\varepsilon) + \tilde{\rho}\pi u(\varepsilon) + (\tilde{\rho}\pi)^2 u(\varepsilon) + ...]$, where we specifically allow y_t and π_t to differ from all future values since we want to consider changing both while holding lifetime utility constant. Totally differentiating V_t with respect to π_t , we obtain the VSL as the amount income would have to adjust in response to keep lifetime utility constant:

If the relative productivity of investments a_1/a_2 or prices p_1/p_2 depend on educational attainment (e.g., more educated persons may be better at implementing more complicated treatments), then the choice of inputs will naturally depend on education. If education raises the productivity of or lowers the price of x_2 relative to x_1 , then more-educated individuals will choose to use x_2 , while less-educated individuals will favor x_1 . For example, a new treatment that only lowered costs for the most educated should cause only them to switch treatments, increasing the education-health gradient. New treatments that greatly simplify health regimens may reduce costs for the least educated relatively more, having the opposite effect (Goldman and Lakdawalla (2005)).

Knowledge about more health options/inputs should generally improve health, since more knowledgeable individuals may be able to take advantage of more productive or less costly inputs. This is the presumption in much of the health economics literature. However, it is also quite possible that individuals are aware of their health options but are simply misinformed about the productivity of some inputs. In this case, it is not obvious that they should underestimate health productivity. It seems equally plausible that some individuals may overestimate the value of certain health behaviors or treatments (e.g., the value of exercise or not drinking). This should lead to excess investments in health. In this case, better information may actually cause individuals to invest less in their health and become less healthy (albeit better off in utility terms). This suggests that the nature of information deficiencies may be as important as the extent of those deficiencies.

3.2. Evidence on the Effects of Education on Health and Mortality

We first review recent empirical studies that estimate the causal effects of education on health and mortality, focusing attention on studies that most convincingly address concerns about the endogeneity of schooling and unobserved heterogeneity. We then discuss evidence aimed at distinguishing between various mechanisms that might explain the relationship between education and health.

3.2.1 Does Education Affect Health and Mortality?

It is important to recognize that education may be correlated with health and mortality even if schooling has no causal effect on either. For example, healthier individuals are likely to be more efficient producers of human capital through schooling (Grossman (1972a)). These effects may be compounded over time if previous health levels affect present health status. More generally, unobservable individual characteristics (e.g., genetic background, family environment, patience) may jointly affect both education choices and health-related behaviors (Fuchs (1982)). An important goal of recent empirical studies has been to identify the causal effects of education on health or mortality outcomes. These studies generally rely on IV or RD estimation strategies designed to address concerns about the endogeneity of schooling and bias from unobserved heterogeneity. We discuss this more recent literature here, referring the reader to Grossman and Kaestner (1997) or Grossman (2000, 2006) for comprehensive surveys of earlier research.

Many of the most convincing and recent IV and RD studies use changes in compulsory schooling laws as instruments for educational attainment. Estimates from these analyses reflect the effects of education on health for youth, who respond to changes in schooling laws (i.e., less-educated youth constrained from dropping out earlier).⁶⁴ As such, they measure improvements in health associated with lower levels of education (typically secondary school). Few studies employ instruments that would help identify the effects of postsecondary schooling on health (de Walque (2007a) and Grimard and Parent (2007) are notable exceptions).

Most studies estimate health/mortality equations of the following form:

$$H_{idt} = \beta E_{id} + X_{id}\gamma + d_c + d_l + d_t + \varepsilon_{idt}, \qquad (2.11)$$

where H_{iclt} reflects a health outcome for person *i* born in year *c* and location *l* measured as of year *t*. *E* reflects educational attainment, while *X* reflects other characteristics that may affect health. Some studies (especially those analyzing mortality) aggregate across individuals by gender, cohort, and place of birth, but otherwise estimate similar specifications.

Since health and survival rates have generally improved over time and across cohorts, it is important to account for health trends across cohorts. This is especially true for studies using school reforms as instrumental variables, since most reforms increased educational attainment for affected cohorts. Failure to account for secular improvements in health may incorrectly attribute those changes to school reforms, biasing estimates toward finding health benefits of schooling. The extent to which researchers account for general trends in mortality (or other health outcomes) across cohorts is crucial for the credibility of most IV estimates. Some of the best studies, implicitly or explicitly, exploit an RD design by (1) focusing on changes in schooling and health only around the time of a change in school reforms or (2) allowing for general cohort trends.

Although the importance of accounting for general health trends cannot be overstated, RD strategies are likely to underestimate the long-run effects of education on health if there are important cross-cohort spillovers. These spillovers may arise for a number of reasons. For example, firms typically offer a limited set of health insurance options to all employees, reflecting the average demand across all cohorts employed by the firm. Thus, health insurance policies may be relatively homogeneous across

⁶⁴ Angrist and Imbens (1994, 1995) showed that, under certain conditions, IV estimation using discrete instrumental variables identifies the average effect of the endogenous regressor on the outcome for the population that changes its behavior in response to changes in the instrument. They referred to this as the "local average treatment effect."

cohorts despite differences in demand for insurance. Alternatively, social interactions (including between spouses) may exist across cohorts. As a result, demand for smoking, drinking, or exercise may depend on what friends and family members do. RD estimates will fail to incorporate any impact due to these cross-cohort interactions.⁶⁵ Therefore, when cross-cohort spillovers are thought to be important, there is a tension between accounting for secular trends in health via RD estimation and identifying the full long-term effects of education.

Tables 2.4–2.7 summarize IV studies from the past 5 years that estimate the impact of education on health and mortality. All of the studies in these tables employ credible identification strategies and provide estimates with reasonable levels of precision. Tables 2.4–2.6 document the following characteristics of these studies: data sources, important sample characteristics, details on the health outcome of interest, instruments for schooling, and the measure of schooling used in the analysis (e.g., years of schooling). Two sets of estimated effects of schooling on health/mortality are reported: (1) regression-based estimates (including probits) and (2) IV or RD estimates that attempt to address concerns about endogeneity or unobserved heterogeneity. Table 2.4 describes studies that estimate the effect of schooling on mortality, while Tables 2.5 and 2.6 describe studies focused on other health outcomes in the United States (Table 2.5) and Europe (Table 2.6).

Our discussion is organized largely around Table 2.7, which categorizes selected estimates by general health outcome or behavior across all countries. This table further reports the magnitude of estimated effects of schooling (using IV or RD strategies) as a percent of base outcome levels (e.g., percent reductions in mortality or smoking).

Lleras-Muney (2005, 2006) provided the first serious attempt to estimate the causal effect of schooling on mortality. Using US census data from 1960, 1970, and 1980, she analyzed 10-year mortality rates among native-born whites born 1901–1925. She first calculated the mortality rates by gender, year of birth, and state of birth using consecutive censuses. Then, she combined these mortality rates with average years of

cohort t is given by $U(x_{it}) = \alpha_t x_{it} - \frac{\beta}{2} \left[\theta(x_{it}^2 + \gamma [x_{it} - \overline{x}]^2) + (1 - \theta) \sum_{j=1}^J \omega_j (x_{it} - \overline{x}_{t-j})^2 \right]$, where \overline{x}_t reflects the cohort t

average for x and $\sum_{j=1}^{t} \omega_j = 1$. This specification implies that individuals consider their own personal benefits and costs associated with x, but they would also like to choose x to mimic the choices of others in their cohort as well as J previous cohorts. θ reflects the importance of one's own private cost and the desire to conform with others in the same cohort relative to the desire to conform with older cohorts. If we consider a policy that increases the marginal returns to x by Δ at time t=1, so $\alpha_t = \alpha_0 \forall t \le 0$ and $\alpha_t = \alpha_0 + \Delta \forall t > 0$, then in equilibrium, the immediate change in x is given by $\overline{x}_1 - \overline{x}_0 = \frac{\Delta}{\beta}$, while the long-run change is given by $\overline{x}_{\infty} - \overline{x}_0 = \frac{\Delta}{\beta\theta}$. This implies that an RD estimator comparing cohorts 0 and 1 will underestimate the long-run effect of the policy change on x by a factor of θ .

 $^{^{65}}$ Consider a simple example of social interactions, in which the utility associated with health input x for individual *i* from

		Sample	Instrument	Important Features of Estimation Specification	Mortality or Survival Measure	Estimated Effect		
Study	Data					No Instruments	RD/IV Methods	
Mazumder (2008, 2010)	US censuses (1960–2000)	Native whites born 1901–1925	State-specific changes in compulsory schooling ages	2SLS with state- specific linear cohort trends	10-yr mortality rate	- 0.036 (0.003)	0.006 (0.031)	
Lleras-Muney (2005, 2006)	US censuses (1960–1980)	Native whites born 1901–1925	State-specific changes in compulsory schooling ages	2SLS with region-specific linear cohort trends	10-yr mortality rate	- 0.036 (0.004)	- 0.063 (0.024)	
Clark & Royer (2010)	Birth & mortality data, Office for National Statistics, England and Wales (1970–2007)	Residents of England & Wales in 1970–2004, born 1925–1938	National compulsory schooling age increase in 1947	RD with cohorts defined by quarter of birth*	Probability of dying between ages 45 and 69		Women 0.004 (Reduced form r	Men 0.012 not significant)
Albouy and Lequien (2009)	French Censuses and Enchantillon Demographique Permanent (various years 1968–1999),	French individuals born on first 4 days of October in 1920–1925 or	National compulsory schooling age increases for cohorts born in	RD with 1-yr cohort window around schooling age change	Survival from age 15 to 52 Survival from		0.013 (0.011)	
	National Registers on Deaths (1968–2005)	1950–1955 and alive in 1968	1923 & 1953		age 45 to 82			

Table 2.4 Regression Discontinuity and Instrumental Variable Studies of the Effects of Years of Schooling on Mortality and Survival

Notes: Standard errors shown in parentheses. Estimates in bold are statistically significant at 0.05 level.

* Clark and Royer (2010) do not report IV estimates; values reported in table reflect RD estimated effect of 1947 reform on death rates between ages 45 and 69 divided by the effect of the reform on average years of completed schooling.

			Instrument	Schooling Measure	Health	Estimated Effect			
Study	Data	Sample			Outcome	No Instrum	ents	IV Methods	
Mazumder (2008)	Survey of Income and Program	Native Native born 1901–1925 on 36– 0– 6)	State-level compulsory schooling laws	Years of education	Self-reported health fair or poor	- 0.036 (0.001)		- 0.082 (0.034)	
	Participation (1984, 1986– 1988, 1990– 1993, 1996)				Number of nights in hospital last year	- 0.073 (0.0186)		-1.083 (0.767)	
					Trouble with stairs	- 0.025 (0.001)		-0.007 (0.032)	
					Health limitation	- 0.025 (0.0013)		- 0.074 (0.035)	
Oreopoulos (2006)	US censuses (1950–2000)	Natives born 1901–1956, ages 25–84 at time of survey	State-level compulsory schooling laws	Average years of education	Disability that limits personal care	- 0.014 (0.0003)		- 0.025 (0.006)	
					Disability that limits mobility	- 0.020 (0.0004)		- 0.043 (0.007)	
Kenkel, Lillard, & Mathios (2006)	National Longitudinal Survey of Youth, 1979 cohort	Born 1957– 1964. Observed in 1998 (oversamples minorities & poor whites)	High school & GED require- ments & policies	High school graduate indicator	Currently smoke Obese	Women - 0.194 (0.030) 0.005 (0.029)	Men - 0.226 (0.029) 0.013 (0.026)	Women -0.102 (0.124) -0.021 (0.139)	Men - 0.229 (0.088) -0.008 (0.082)

 Table 2.5
 Recent US-Based Instrumental Variable Studies of the Effects of Schooling on Health

Continued

				Schooling Health Estimated Effect		stimated Effect		
Study	Data	Sample	Instrument	Measure	Outcome	No Instruments	IV Methods	
de Walque (2007)	National Health Interview Survey— Smoking Supplements (various years 1983–1995)	Natives born 1937–1956, aged 25+ at the time of the survey	Risk of induction in Vietnam War	Number of years of education above high school	Currently smoke Stopped smoking	- 0.040 (0.004) 0.041 (0.002)	Risk of induction -0.038 (0.020) 0.051 (0.031)	Risk of induction × Risk of being killed -0.040 (0.019) 0.066 (0.029)
				Indicator for college graduate or more	Currently smoke Stopped smoking	- 0.173 (0.015) 0.178 (0.010)	-1.169 (0.642) 2.190 (1.395)	- 0.181 (0.087) 0.297 (0.133)
Grimard and Parent (2007)	Current Population Survey Tobacco	White native-born citizens at least age 25	Vietnam draft cohort indicator (born 1945–	Years of education	Currently smokes everyday	- 0.055 (0.002)	- 0.136 (0.024)	
	Supplements (1995, 1996, 1998, 1999)	Supplements born 1935– 1950) 1995, 1996, 1974 .998, 1999)	Ever smoked	- 0.043 (0.002)	- 0.111 (0.029)			

Table 2.5 Recent US-Based Instrumental Variable Studies of the Effects of Schooling on Health-continued

Notes: Estimates under "No Instruments" column are based on OLS estimation or Probit estimates (marginal effects reported). Estimates under "IV Methods" column report estimates from two-stage least squares or two-stage conditional maximum likelihood estimation (marginal effects reported). Standard errors are shown in parentheses. de Walque (2007) provides *t*-statistics rather than standard errors; standard errors are calculated as the reported estimates divided by reported *t*-statistics. Estimates in bold are statistically significant at 0.05 level.

Table 2.6 Recent European Instrumental Variab	le Studies of the Effect	ts of Years of Schooling on Healtl
---	--------------------------	------------------------------------

					Estimated Effect			
Study	Data	Sample	Instrument	Health Outcome	No Instru	iments	IV/RD Me	thods
Clark & Royer (2010)	Health Survey of England	Residents of England in	National compulsory	Self-reported health fair or Bad	- 0.081 (0.003)		-0.005 (0.014)	
	(1991–2004)	1991–2004 (Bom 1918–1972)	schooling age increases in 1947 and 1973	Long illness	- 0.033 (0.003)		0.023 (0.016)	
				Reduced activity	- 0.021 (0.002)		0.014 (0.012)	
				Obese	- 0.024 (0.003)		0.028 (0.016)	
				Hypertension	- 0.015 (0.003)		-0.025 (0.018)	
				Currently smoke	- 0.072 (0.003)		- 0.035 (0.014)	
				Physically active	0.055 (0.004)		-0.001 (0.019)	
Powdthavee (2010)	Health Survey of England (1991–2007)	Residents of England in 1991–2007 (born	National compulsory schooling age increases in 1947 and 1973	Hypertension (1947 law change)	Women - 0.017	Men - 0.015	Women - 0.084	Men - 0.070
		1929–1939 or 1952–1965)		Hypertension (1973 law change)	-0.009	-0.008	0.014	0.075
Oreopoulos (2006, 2008)	Annual General Household Survey	Native-born British & North	Compulsory schooling age	Self-reported poor health	- 0.037 (0.002)		0.007 (0.008)	
	for England (1983– 1998) & Northern Ireland Continuous Household Surveys (1985–1998)	Ireland residents, (born 1921–1951, ages 25–84 at time of survey)	increases in United Kingdom (1947) & North Ireland (1957)	Self-reported good health	0.065 (0.002)		-0.010 (0.011)	

				Estimated Effect				
Study	Data	Sample	Instrument	Health Outcome	No Instruments		IV/RD Me	ethods
Silles (2009)	General Household Survey for	UK individuals (ages 25–60)	National compulsory schooling age	Self-reported good health	0.026 (0.001)		2SLS 0.045 (0.009)	RD 0.064 (0.003)
	England, Scotland, and Wales (1980–2004)		and 1973	No long-term illness	0.008 (0.001)		0.055 (0.009)	0.075 (0.003)
				No activity- limiting illness	0.008 (0.001)		0.046 (0.008)	0.051 (0.002)
Kemptner, Jürges, & Reinhold (2010)	German Microcensus (1989, 1995, 1999, 2002, 2003)	Germans (born 1930–1960) living and educated in former West German states	Increased state- specific compulsory schooling from 8 to 9 (between 1949 and 1969)	Long-term illness	Women - 0.012 (0.002)	Men - 0.029 (0.001)	Women 0.013 (0.024)	Men - 0.039 (0.019)
				Work-limiting illness	- 0.013 (0.001)	- 0.029 (0.001)	0.036 (0.022)	- 0.037 (0.017)
				Overweight	- 0.047 (0.002)	- 0.031 (0.001)	-0.016 (0.016)	- 0.034 (0.015)
				Obese	- 0.020 (0.001)	- 0.018 (0.001)	0.007 (0.012)	-0.028 (0.015)
				Ever smoked	- 0.005 (0.002)	- 0.026 (0.001)	0.014 (0.011)	-0.012 (0.013)
				Currently smoker	- 0.021 (0.002)	- 0.032 (0.001)	0.001 (0.01)	0.005 (0.012)
Brunello, Fabbri, and	ECHP (1998), SHARE (2004),	European women (from 10	Increased compulsory	Overweight	- 0.022 (0.001)		- 0.044 (0.021)	
Fort (2009)	GSOEP (2002), BHPS (2003), Enquete sur la Sante (2003)	countries) (ages 25–65), born within 7 yr of cohort first affected by school law changes	schooling ages by country	Obesity	- 0.012 (0.001)		- 0.012 (0.013)	

Table 2.6 Recent European Instrumental Variable Studies of the Effects of Years of Schooling on Health-continued

Notes: Estimates under "No Instruments" column are based on OLS estimation or Probit estimates (marginal effects reported). Estimates under "IV/RD Methods" column report estimates from two-stage least squares, IV probit estimation (marginal effects reported), or regression discontinuity (RD) methods. Standard errors are shown in parentheses. Estimates in bold are statistically significant at 0.05 level.

 Table 2.7 Effects of Years of Education on Health and Mortality by Outcome (Best IV/RD Estimates from Selected Studies)

Study		Country	Outcome Details	Sample Average		Estimated Effect	Effect as % of Average	
A.	Mortality							
	Mazumder	United	10-yr	0.213		0.006	2.8%	
	(2008)	States	mortality					
			rates					
	Lleras-	United	10-yr	0.106		-0.063*	-59.4%	
	Muney	States	mortality					
	(2005, 2006)		rates					
	Clark &	United	Probablity	Women:	0.146	0.004	2.7%	
	Royer (2010) ^a	Kingdom	of dying between	Men:	0.221	0.012	5.4%	
			and 69					
	Albouv and	France	37-vr	Ages 15-42:	0.07	-0.013	-18.6%	
	Lequien		mortality	Ages 45–82:	0.40	-0.063	-15.8%	
	$(2009)^{b}$		rate	8				
В.	Self-Reported	Health						
	Mazumder	United	Self-report	0.357		-0.082^{*}	-23.0%	
	(2008)	States	fair or poor					
	Clark &	United	Self-report	0.25		-0.005	-2.0%	
	Royer (2010)	Kingdom	fair or bad					
			Long illness	0.46		0.023	5.0%	
	Oreopoulos	United	Self-report	0.150		0.007	4.7%	
	(2006, 2008)	Kingdom	poor health					
			Self-report good health	0.564		-0.01	-1.8%	
	Silles (2009)	United Kingdom	Self-report good health	0.661		0.064*	9.7%	
		C	No long- term illness	0.685		0.075*	10.9%	
	Kemptner,	Germany	Long-term	Women:	0.15	0.013	8.7%	
	Jürges, & Reinhold	,	illness	Men:	0.20	-0.039*	-19.5%	
	(2010)							
C.	Disability, Lin	nited Mobi	lity, and Activit	.v				
	Mazumder	United	Health	0.423		-0.074*	-17.5%	
	(2008)	States	limitation					
	Oreopoulos	United	Disability	0.092		-0.025*	-27.2%	
	(2006)	States	limiting personal care					

Continued

Study	Country	Outcome Details Sample Average		rage	Estimated Effect	Effect as % of Average	
		Disability limiting mobility	0.128		-0.043*	-33.6%	
Clark & Royer (2010)	United Kingdom	Reduced activity	0.16		0.014	8.8%	
		Physically active	0.44		-0.001	-0.2%	
D. Smoking							
Clark & Royer (2010)	United Kingdom	Currently smoke	0.27		-0.035*	-13.0%	
Kenkel et al.	United	Currently	Women:	0.26	-0.229^{*}	-88.1%	
(2006) ^d	States	smoke	Men:	0.27	-0.102	-37.8%	
de Walque (2007a) ^c	United States	Currently smoke	0.40		-0.040*	-10.0%	
Grimard and Parent (2007)	United States	Currently smoke every day	0.31		-0.136*	-43.9%	
		Ever smoked	0.58		-0.111*	-19.1%	
Kemptner,	Germany	Currently	Women:	0.24	0.001	0.4%	
Jürges &		smoke	Men:	0.36	0.005	1.4%	
Reinhold		Ever smoked	Women:	0.39	0.014	3.6%	
(2010)			Men:	0.63	-0.012	-1.9%	
E. Obesity							
Clark & Royer (2010)	United Kingdom	Obese	0.21		0.028	13.3%	
Kemptner,	Germany	Overweight	Women:	0.44	-0.016	-3.6%	
Jürges, &			Men:	0.66	-0.034*	-5.2%	
Reinhold		Obese	Women:	0.13	0.007	5.4%	
(2010)			Men:	0.16	-0.028	-17.5%	
Brunello, Fabbri, and	10 European	Overweight (women)	0.387		-0.044*	-11.4%	
Fort (2009)	Countries	Obese (women)	0.114		-0.012	-10.5%	
Kenkel et al.	United	Obese	Women:	0.27	-0.021	-7.8%	
(2006) ^d	States		Men:	0.25	-0.008	-3.2%	

 Table 2.7 Effects of Years of Education on Health and Mortality by Outcome (Best IV/RD Estimates from Selected Studies)—continued

Notes:

^aClark and Royer (2010) did not report IV estimates for the effect of education on mortality; values reported in table reflect RD estimated effect of 1947 reform on death rates between ages 45 and 69 divided by the effect of the reform on average years of completed schooling.

^bMortality estimates for Albouy and Lequien (2009) are calculated from survival rates; mortality rates are inferred from their Figures 1 and 2.

^cAverage smoking rates for de Walque (2007a) inferred from his Figure 1 for those with a high school degree. ^dKenkel, Lillard, and Mathios (2006) estimates are for high school graduation rather than years of schooling. *Denotes statistical significance at 0.05 level. completed schooling and compulsory schooling ages (both by gender, year of birth, and state of birth) to estimate the effect of an additional year of education on mortality using standard two sample IV strategies.⁶⁶ As reported in Table 2.4, her IV specification suggests that an additional year of schooling reduces 10-year mortality rates by about 6 percentage points (an effect that is statistically significant and nearly twice the effect estimated by OLS). This estimate is almost certainly greater than the true effect, since it implies a nearly 60% reduction per year of school given baseline 10-year mortality rates of 0.11 in her sample.

Because Lleras-Muney (2005, 2006) controls for region-specific cohort trends, her estimates are identified from state-level deviations from regional trends in schooling and mortality associated with changes in state compulsory schooling ages. However, Mazumder (2008) noted that if states with larger increases in compulsory schooling ages also had greater secular improvements in life expectancy (relative to the regional average), then failing to control for state-specific trends will cause her to overestimate the effect of schooling on mortality. Using an expanded sample to include 1960–2000 US censuses (while continuing to study native whites born 1901–1925), Mazumder (2008, 2010) showed that controlling for state-specific cohort trends (rather than just region-specific trends) produces negligible estimated effects of education on mortality.⁶⁷

Using US data on disease-specific mortality, Glied and Lleras-Muney (2008) estimated that increases in compulsory schooling ages led to greater reductions in death rates for diseases that experienced greater technological improvements in recent years, whether the latter is measured by the number of new active drug ingredients or disease-specific declines in mortality in recent years. Given the absence of state- or region-specific cohort trends in their specification, it seems likely that they overestimate the average effect of schooling on mortality. However, if state-specific cohort trends in mortality do not vary systematically across diseases, then differences in mortality reductions across different diseases should be consistently estimated. In this case, their finding that mortality reductions associated with schooling are greatest for diseases experiencing faster technological improvements is more credible. We return to this issue below.

Two recent international studies, one in the United Kingdom (Clark and Royer (2010)) and the other in France (Albouy and Lequien (2009)), used national increases in compulsory schooling ages and RD designs to estimate the effects of schooling on mortality.

⁶⁶ Her main specification controls for gender, cohort, state of birth, and census year fixed effects. She also controlled for region-specific cohort trends and observed cohort-specific state-of-birth characteristics (e.g., percent urban, percent black, percent employed in manufacturing, average manufacturing wage, and so on). Like Lochner and Moretti (2004), she linked cohorts to compulsory schooling ages in their state of birth when they were age 14.

⁶⁷ Mazumder (2008, 2010) further showed that the key difference in estimated effects arises due to the inclusion of statespecific cohort trends and not the addition of 1990 and 2000 census data. Baseline mortality rates are considerably higher in Mazumder (2008, 2010) than in Lleras-Muney (2005, 2006) because the 1901–1925 birth cohorts are older at the time of the more recent censuses.

They effectively estimated changes in educational attainment and mortality for the first cohorts affected by increases in compulsory schooling ages relative to the last cohorts before the reform. Both studies estimated statistically insignificant effects: estimates by Albouy and Lequien (2009) are modest in size, while those of Clark and Royer (2010) suggest that education actually *increases* mortality.⁶⁸

A growing literature studies the effects of education on self-reported health status, disability, mobility, and activity. Mazumder (2008) used data from the Survey of Income and Program Participation (SIPP) from 1984 to 1996 to estimate the effects of schooling on self-reported health for native-born white Americans born 1901–1925. As with mortality, he instrumented for schooling using state-level compulsory schooling laws, controlling for state-specific cohort trends. He estimated sizeable and statistically significant effects of an additional year of school on the probability someone reports being in fair or poor health (23% reduction) as well as on the probability someone reports a health limitation (17.5% reduction). Estimates focusing on specific limitations suggest that schooling primarily reduces problems with seeing, hearing, and speaking; estimated effects on walking, climbing stairs, or lifting are small and statistically insignificant. Using data from 1950–2000 US censuses and changes in state-level compulsory schooling laws as instruments for educational attainment, Oreopoulos (2006) estimated that an additional year of school reduces disabilities limiting personal care or mobility by about 30% among 25–84-year-old Americans.⁶⁹

A number of recent studies examined the impacts of education on self-reported health and limitations in Europe using changes in compulsory schooling laws. Silles (2009) used the General Household Survey for England, Scotland, and Wales to estimate the effect of education on self-reported health in the United Kingdom. Using an RD design around national increases in compulsory schooling ages in 1947 and 1973, she estimated that an extra year of secondary school increases self-reported good health by about 10%, a modest effect about half as large as Mazumder's estimate for self-reported bad health in the United States. Clark and Royer (2010) exploited the same increases in the compulsory schooling age (also with an RD design) using data from the Health Survey of England (1991–2004). They estimated small and statistically insignificant effects of education on self-reported health, long-term illness, and physical activity.

⁶⁸ Clark and Royer (2010) did not actually estimate the effect of schooling on mortality. Instead, they estimated the reduced form effect of changes in the compulsory schooling age on mortality and (using a slightly different specification) the effect of compulsory schooling age on years of completed schooling. Their estimates suggest that the increase in compulsory schooling increased both education and mortality rates. For comparison purposes, Tables 2.4 and 2.7 report the results from dividing the effect of the 1947 reform on mortality by the reform's effect on educational attainment to arrive at an "IV estimate" (in the spirit of a Wald estimator).

⁶⁹ His sample includes Americans born 1901–1956. His specifications controlled for general national cohort and time trends, but they did not account for differences in health trends across states. Like Lleras-Muney (2005), he did control for time-varying state-level demographic and labor market characteristics.

Using data from the General Household Survey for England and its Northern Ireland counterpart, Oreopoulos (2006, 2008) reached a similar conclusion about the effect of education on self-reported health in the United Kingdom. His analysis focuses on the 1947 increase in the compulsory schooling age in England along with a similar increase in Northern Ireland that took place 10 years later.⁷⁰ Kemptner, Jürges, and Reinhold (2010) estimated the effect of increasing minimum schooling levels in (formerly West) Germany from 8 to 9 years using German Microcensus data. Interestingly, they found that education significantly reduces reported long-term illness (by about 20%) among men but not women. Point estimates for women are positive and sizeable, but less precisely measured.

Based on the studies summarized in Table 2.7, education appears to have a weaker effect on mortality, self-reported health, and physical activity in Europe than in the United States. This is not because education and health are unrelated in Europe—indeed, some of the OLS results in Tables 2.5 and 2.6 imply a stronger relationship in Europe. It is tempting to speculate that the effects of education depend on access to health care, general social welfare and unemployment policies, and the level of over-all inequality.

Although most studies examined self-reported measures of health, Powdthavee (2010) exploited the Health Survey of England (1991–2007) and changes in compulsory schooling in the United Kingdom to estimate the effects of education on hypertension as determined from blood pressure measurements (see Table 2.6). IV-probit estimates exploiting the 1947 increase in the compulsory schooling age (from 14 to 15 years) suggest that an extra year of schooling reduces hypertension by slightly more than 10% for both men and women. Unfortunately, RD estimates using the 1973 increase in the compulsory schooling age (from 15 to 16 years) are much more imprecise, yielding positive but statistically insignificant point estimates. Aside from this study, there is little evidence on the effects of education on objective measures of health.

The final panel of Table 2.7 reports estimates of the effects of education on smoking and obesity, two important contributors to a number of serious and chronic health problems. For these outcomes, studies mostly agree: education significantly reduces smoking but has negligible effects on obesity.⁷¹ Clark and Royer (2010) estimated that an additional year of schooling reduces the probability someone is currently smoking by about 13% in England, very similar to estimates for men in the United States by de Walque (2007a). Unlike all other studies discussed thus far, de Walque (2007a) identifies the effect of postsecondary schooling rather than primary or secondary schooling.

⁷⁰ His specifications allow for general cohort trends; however, he assumed that those trends are the same for England and Northern Ireland.

⁷¹ Many of these studies also find that education increases the probability of quitting smoking conditional on having smoked; although, this is not as easily interpreted since education affects the likelihood of starting smoking.
This is because he uses the risk of being drafted for the Vietnam War to instrument for schooling attainment.⁷² Grimard and Parent (2007) also exploited the effect of the Vietnam draft on male college attendance decisions for cohorts born between 1945 and 1950. Their estimates suggest that additional postsecondary schooling has even greater effects on rates of frequent smoking, reducing the likelihood that someone smokes every day by roughly 40%.73 Kenkel, Lillard, and Mathios (2006) estimated the effects of receiving a GED or finishing high school on smoking for NLSY79 respondents in their late 30s and early 40s using state-specific high school and GED requirements and policies as instruments (e.g., the number of courses required to complete high school, an indicator for whether local districts set graduation requirements, the fraction of youth taking the GED test, an index of GED policies related to score requirements and fees, and the log of per capita education spending averaged over primary and secondary school ages). They estimated very large effects of high school graduation on smoking, with effects for women especially large and statistically significant (88% reduction). Estimated effects of GED receipt are much weaker (not reported here).⁷⁴ In contrast to previous studies, Kemptner, Jürges, and Reinhold (2010) estimated relatively small and insignificant effects of schooling on smoking in Germany. This difference may be due to the cohorts analyzed in their study. In exploiting increases in compulsory schooling (from 8 to 9 years) from 1949-1969, they largely identified the effects of schooling on cohorts who would have made their initial smoking decisions before the harms of smoking were well known. Among these early cohorts, smoking was only weakly correlated with education (e.g., see the relatively small OLS estimates for this study relative to others in Tables 2.5 and 2.6).⁷⁵ By this reasoning, the relatively large estimated effects of schooling on smoking by Kenkel, Lillard, and Mathios (2006) may be explained by the fact that these individuals would all have made their smoking decisions long after the 1964 US Surgeon General widely publicized the dangers of smoking. Consistent with the above model, the effects of education on smoking seem to depend on the perceived impacts of smoking on health.

⁷² In some specifications, he also used the risk of induction multiplied by the probability of being killed to instrument for schooling. When he controlled for the endogeneity of veteran status, de Walque (2007a) estimated even greater effects of education on smoking than reported in Tables 2.5 and 2.7.

⁷³ Grimard and Parent (2007) and de Walque (2007) effectively identify the effects of college for men who chose to avoid the draft. Since these men may be more risk averse (or may value their survival more) than the average male, the effect of college-going on smoking for them may differ from that of a more representative population.

⁷⁴ One concern with this study is that differences in education policies or expenditures across states and over time may affect skill levels conditional on educational attainment. That is, it may mean something different to finish high school in a state that spends a lot on secondary schools and has high graduation requirements relative to a state with low expenditures and weaker standards. This would, in general, violate their exogeneity assumptions; however, Kenkel, Lillard, and Mathios (2006) tested and did not reject their exclusion restrictions in their context, so this may not be a serious problem.

⁷⁵ See de Walque (2010) for a detailed analysis of how smoking differences by education evolved over time in the United States.

The estimated effects of education on weight problems in Europe and the United States are uniformly small. None of the studies we surveyed report statistically significant effects on obesity (body mass index [BMI] of 30 or more); however, a few find small but statistically significant effects on the probability that someone is overweight (BMI of 25 or more). See Table 2.7.

Tables 2.5 and 2.6 do not include a few earlier IV studies of the effects of education on health; however, there are good reasons to question the exogeneity of at least some of their instruments.⁷⁶ Both Auld and Sidhu (2005) and Adams (2002) used parental background as instruments; however, these instruments have generally been dismissed as endogenous in the literature on schooling and earnings (e.g., see Carneiro and Heckman (2002)). Furthermore, Kenkel, Lillard, and Mathios (2006) rejected the validity of parental education levels as instruments in estimating the effects of education on smoking and obesity. Adams (2002) also used quarter of birth as an instrument for schooling; however, a large literature on weak instruments suggests that this may produce IV estimates biased toward OLS estimates (see, e.g., Bound, Jaeger, and Baker (1995)). Arkes (2003) and Auld and Sidhu (2005) used local unemployment rates when individuals are in their teenage years to instrument for educational attainment; however, the strong serial correlation in local labor market conditions over time may lead to bias as mobility rates are often low and contemporaneous local labor market conditions may affect health through a variety of channels. Cameron and Taber (2004) showed that, at least in the estimation of earnings equations, it is important to control for current local labor market conditions when using past conditions as an instrument for schooling. Finally, the Danish study by Arendt (2005) used two national school reforms, which both raised compulsory schooling ages and affected the nature of schooling (e.g., removal of student tracking based on tests, convergence in curricula between village and city schools). The latter raises similar concerns to those with Kenkel, Lillard, and Mathios (2006). Even ignoring any concerns about endogeneity, the IV estimates in Arendt (2005) are extremely imprecise.

Conti, Heckman, and Urzua (2010a, 2010b) offer a departure from the standard IV/RD approach to studying the relationship between education and health. They estimated a multifactor model of schooling, earnings, and health outcomes using data from the British Cohort Study, which has followed all babies born in the United Kingdom during 1 week in 1970 periodically through age 30. Their model assumed that education, postschool earnings, and health behaviors/outcomes depend on family background characteristics, age 10 observable health endowments, and three unobserved latent factors (cognitive ability, noncognitive ability, and an unobserved health endowment). Because their data

⁷⁶ These tables also exclude a few recent IV studies exploiting schooling reforms (e.g., Grabner (2009); Jürges, Kruk, and Reinhold (2009); and Reinhold and Jürges (2010)). While the identification strategies in these analyses are reasonably credible, the reported IV estimates are very imprecise and unable to rule out a large range of effects (including the estimates reported in Tables 2.5 and 2.6).

contain multiple measurements of the three (correlated) unobserved factors, they are able to estimate the joint distribution of these factors as well as their effects on health behaviors and outcomes. They focused attention on the impacts of completing mandatory schooling (through age 16) on smoking, obesity, and self-reported health (fair or bad). Their estimates suggest that the causal effect of education explains 60–70% of the raw differences in smoking by educational attainment, 35–55% of the raw differences in self-reported health, and one-third of the differences in obesity for men (none of the difference for women). Thus, education has important causal effects on smoking and self-reported health for men and women, but it only reduces obesity rates (by a fairly small amount) for men. These findings are largely consistent with the evidence presented in Table 2.7.

The approach of Conti, Heckman, and Urzua (2010a, 2010b) has enabled them to draw a number of other interesting conclusions. First, they showed that the effects of education are greatest for individuals with relatively high cognitive skills and/or low noncognitive skills. Second, they showed that cognitive ability is not a very important determinant of smoking decisions or obesity. There is a modest effect of cognitive ability on self-reported health for women, but little effect for men. Finally, they estimated that noncognitive skills are generally more important for smoking, obesity, and self-reported health. In fact, noncognitive skills are roughly as important as early health endowments in explaining these outcomes at age 30.

Parental Education and Child Health

There is a growing literature examining the effects of parental education on child (often infant) health and mortality. We briefly discuss a few recent studies that exploit IV and RD techniques to address concerns about unobserved heterogeneity. Currie and Moretti (2003) provided the first evidence that maternal education causally affects maternal prenatal behavior and birth outcomes. They exploited openings of 2- and 4-year colleges by county in the United States from 1940 to 1996 as instruments for maternal education. These college openings increased average years spent in college among their sample of 25-45-year-old women considerably (roughly 1 year for four-year colleges and 0.15–0.20 years for 2-year colleges). Relative to a high school graduate, their IV estimates suggest that an extra year of college reduces the likelihood of a low-birthweight child by about 20% and a pre-term birth by about 15%. Examining maternal behaviors, they estimated that an extra year of college reduces smoking during pregnancy by roughly onethird but only increases the incidence of prenatal care by 3%. The impact on maternal smoking is a likely channel for the improvements in birth outcomes. McCrary and Royer (2009) exploited a very different source of exogenous variation in maternal schooling: school entry age laws. State laws often specify that children must be age 5 by a certain date before entering kindergarten. Using data from Texas (1989–2001) and California (1989– 2002), they showed that these laws induce a discontinuity in schooling attainment for women in their sample (age 23 or less) based on date of birth. In contrast to Currie and

Moretti (2003), their RD estimates suggest negligible effects of schooling on low birthweight, prematurity, and infant mortality. In most cases, their estimates are precise enough to rule out impacts as large as those in Currie and Moretti (2003). They also found little evidence to suggest that schooling impacts maternal smoking decisions; however, they found some evidence that it improves prenatal care in Texas.

Lindeboom, Llena-Nozal, and van der Klaauw (2009) used the United Kingdom's 1947 increase in the compulsory schooling age (from 14 to 15) discussed above to estimate the effect of maternal and paternal education on child health outcomes in the United Kingdom. Using National Child Development Study data on all births from March 3–9, 1958, they estimated small and statistically insignificant effects of parental education on birth outcomes (low birthweight, illness at birth) as well as child health outcomes (chronic conditions, overweight) at ages 7–16; however, standard errors for their estimates are sizeable and cannot generally rule out impacts of 20% or more. Finally, Chou, Liu, Grossman, and Joyce (2010) exploited the rapid construction of junior high schools in Taiwan from 1968 to 1973. Using variation in openings across counties as instruments for parental education, they estimated very large effects on birth outcomes. Their estimates suggest that, on average, the expansion of junior high schools increased schooling levels by 0.11–0.16 years and reduced the incidence of low-birthweight children by 5% and infant mortality by 8–19%. The implied effects of an additional year of schooling are enormous—much larger than those of Currie and Moretti (2003).

Given the substantial differences in impacts estimated by these studies, it is difficult to draw any strong conclusions about the effects of parental education on child health. Looking only at the United States and United Kingdom studies, it is tempting to conclude that, in developed countries at least, additional college education improves child outcomes, while additional years of high school do not. Furthermore, the findings of Chou, Liu, Grossman, and Joyce (2010) suggest that educational attainment may be more important in a developing economy like Taiwan, even at lower levels of education. Although these tentative conclusions seem reasonable, there are also important differences in the estimation strategies employed by these studies that complicate any comparison of results. First, the exogenous sources of variation that affect schooling also tend to affect other factors that may influence child health. For example, the variation in school entry ages used by McCrary and Royer (2009) also leads to differences in age-for-grade. Students who start a year later are a year older than their counterparts during every grade at school. This may affect success in school beyond the number of years attended; it may also affect the timing of marriage and fertility or the choice of spouse. In this case, estimated effects of the policy on children will reflect the influence of all of these factors in addition to changes in parental education. Second, additional schooling is likely to affect the timing of fertility decisions. The aforementioned studies typically control for maternal age, but this may be an important channel through which schooling affects infant health and mortality. One final issue that complicates the interpretation of most

of these studies is the role of marriage markets and both parents' education. On the one hand, both parents' education may be important, yet it is often difficult to account for both due to data limitations and the strong correlation between the two (specifications in Lindeboom, Llena-Nozal, and van der Klaauw (2009) are unusual in that they separately identify the effects of both parents' education). Many of the policies used as instruments are likely to affect both parents' education to different degrees, so examining the effects of only one parent's education likely overstates that parent's role. These policies are also likely to affect marriage markets, with effects spilling over to other cohorts/ groups assumed to be unaffected by the policies. All of these issues make estimation of the effects of parental education on child health much more complicated than the analysis of education's impacts on one's own health. Future research on this topic will need to grapple with these in terms of estimation and the interpretation of results.

3.2.2 Why Does Education Affect Health and Mortality?

To the extent that education does have small to modest effects on health and mortality, what are the key operating channels or mediating factors? Education may improve decision-making abilities, which may lead to better health decisions and more efficient use of health inputs (productive efficiency). More educated individuals may also be better at acquiring or interpreting information from doctors, other individuals, or the press/internet. This, too, may lead to better decisions about health, including more appropriate treatments (allocative efficiency). Economists have devoted considerable attention to the distinction between productive efficiency (better health production given a set of inputs) and allocative efficiency (better allocation of health inputs) as it pertains to the relationship between education and health.

There are many other channels through which education may affect health and mortality as shown in Table 2.8, which also lists any personal costs associated with these channels and whether associated benefits are private (accruing to the individual who receives more schooling) or public (accruing to others). In addition to improving the production of health via productive or allocative efficiency, education may (cheaply) improve health by alleviating stress, as emphasized by some social epidemiologists. This may be due to improvements in social standing or socioeconomic well-being that result in fewer stressful situations or less long-term stress (e.g., Marmot (2004) and Marmot and Wilkinson (2006)).

Education may also lead to better health outcomes because it raises income levels. As the above model shows, this increases the demand for better health and may cause individuals to spend more to remain healthy. More educated individuals are more likely to purchase better health insurance or to spend more on costly treatments.⁷⁷ More educated individuals may also choose healthier lifestyles. This may include taking greater safety

⁷⁷ Recent quasi-experimental studies estimate significant health benefits from health insurance across a wide range of outcomes. See McWilliams (2009) for a review of this literature.

Channel	Private Costs	Benefits Private or Public?
1. Reduce stress	None	Private
2. Better decision-making ability/use of inputs	None	Private
3. Better at gathering/interpreting information	None	Private
4. Health insurance	Financial	Private
5. Healthier lifestyle		
a. Safety precautions (e.g., seatbelts, smoke alarms)	Utility, financial	Private
b. Diet	Utility, financial	Private
c. Exercise	Utility, financial	Private
d. Nonsmoking, alcohol moderation, avoiding drugs	Utility, saves money	Private
6. Healthier/safer employment	Lower wages	Private
7. Healthier neighborhoods	Housing prices	Private
8. Healthier peers and friends	None	Public

Table 2.8 Channels through Which Education May Improve Health

precautions (e.g., wearing seat belts, installing smoke alarms), eating a healthier diet, exercising more, drinking alcohol in moderation, or smoking less. In fact, Leigh (1983) estimated that much of the positive effect of schooling on health can be explained by exercise, smoking, and occupational choices; however, more recent studies have reached a different conclusion as discussed below. While it is possible that education directly modifies preferences toward these activities, it seems more likely that causal effects of education on these behaviors are due to costly behavioral changes or to changes in peers, friends, or neighbors.⁷⁸ If, in the absence of health benefits, people would prefer to eat fatty foods, drink, smoke, and drive without a seat belt, then foregoing these activities is costly. Similarly, more educated individuals may choose healthier or safer occupations and employers. As the literature on compensating differentials (Rosen, (1986)) reminds us, these "benefits" come at a cost, usually lower wages. More educated individuals may also choose to pay higher housing prices to live in neighborhoods that offer a healthier environment (e.g., less pollution, traffic, and so on).⁷⁹ In all of these cases, health benefits from increased education are effectively paid for, at least on the margin.

⁷⁸ Education may be correlated with these behaviors due to underlying heterogeneity in preferences as well (e.g. individuals with more patience or self-control may acquire more schooling and prefer a healthier lifestyle without any effect of the former on the latter.). Tables 2.5–2.7 suggest a significant causal effect of schooling on smoking but not obesity.

⁷⁹ See Chay and Greenstone (2005) for evidence on the effects of air quality on housing prices. Chay and Greenstone (2003) and Currie and Neidell (2005) provided evidence that air pollution significantly increases child mortality; how-ever, effects on adult mortality appeared to be weaker (Chay, Dobkin, and Greenstone (2003)). Moretti and Neidell (2009) found evidence of important short-run effects of pollution on health. Lleras-Muney (2010) estimated that ozone pollution significantly increases hospitalizations among military children aged 2–5. Neidell (2004) estimated important adverse effects of pollution on child health, with larger effects on children of lower socioeconomic status.

Finally, education may affect coworkers and the friends one interacts with on a frequent basis. To the extent that more educated persons interact more with other educated persons, effects of education on health may be increased through the desire to conform or the spread of health-related information. These "social multiplier" effects are essentially cheap and provide public benefits not taken into account when schooling decisions are made.

Empirically, economists have largely relied on two broad approaches for identifying the channels through which education may reduce mortality or improve health. One approach examines whether controlling for a number of potential mediating factors (e.g., income, occupation, health knowledge) affects estimates of the education-health gradient. Put another way, these studies explore the extent to which the correlation between education and health can be explained by different factors. A second approach focuses more on testing or measuring the importance of specific factors. For example, Glied and Lleras-Muney (2008) examined which diseases are most affected by schooling to see whether differences in technological innovation are important for the educationmortality gradient. Other studies have examined changes in knowledge or behavior in response to new public information (e.g., information that smoking is unhealthy). We briefly discuss the strengths and limitations of these approaches and review the findings from more recent or influential studies.

Numerous studies from a broad range of social scientists have explored the extent to which the relationship between education and mortality or health can be explained by the well-known differences in income (or economic well-being more generally) associated with schooling. This approach typically begins by regressing some measure of health (including mortality) on education and a few standard demographic characteristics (e.g., race, gender, age). Then, a measure of income or economic well-being is included in the regression, and the extent to which the coefficient on education declines (towards zero) is taken to be the effect of education "explained" or "mediated" by income. We discuss two of these studies, which extend this approach to study the importance of a broad range of mediating factors.⁸⁰

Cutler and Lleras-Muney (2010) explored the importance of the following mediating factors on health behavior: (1) income and economic resources, (2) personality, self-esteem, and sense of control, (3) social integration, (4) health-specific knowledge, (5) cognitive ability/skill, and (6) preferences for risk and time discounting. Their analysis sheds light on the role of several mechanisms in affecting health and mortality through important health and lifestyle choices. Because they looked at such a wide array of mediating factors, they exploited numerous data sources in the United States (the National Health Interview Survey [NHIS], National Survey of Midlife Development

⁸⁰ See Cutler and Lleras-Muney (2006) for a recent review of this literature.

in the United States [MIDUS], the Health and Retirement Survey [HRS], National Longitudinal Survey of Youth [NLSY], and the Survey on Smoking [SOS]) as well as the National Childhood Development Study (NCDS) in the United Kingdom.

While Cutler and Lleras-Muney (2010) estimated the relationship between education and numerous health behaviors (smoking, diet/exercise, alcohol consumption, and preventive care), it is useful to summarize these relationships based on the importance of different behaviors for actual health outcomes. They, therefore, calculated a "mortality-weighted average effect" of the relationship between education and smoking, drinking, and obesity (three of the more important and often-measured behaviors). This summary measures weights education-behavior relationships in proportion to their likely impacts on mortality. We briefly summarize these "mortality-weighted average effects." Cutler and Lleras-Muney (2010) first estimated the extent to which differences in household income (and economic resources more generally) can account for the correlation between education and health behaviors. Their estimates suggest that economic resources can account for 11-32% of this correlation in the United States and United Kingdom. Next, Cutler and Lleras-Muney (2010) controlled for other potential mediators along with household income to explore their additional influence. Of the other types of mediators, only cognitive ability/skills and social integration appear to be important. Using the NLSY, they found that cognitive ability, as measured by a battery of aptitude tests taken during late teenage years, explains about 15% of the correlation between education and health behaviors.⁸¹ Economic resources and cognitive ability/skills combine to explain roughly 30% of this correlation in the NLSY. In the MIDUS data, social integration explains 7% of the education-health behavior correlation after accounting for income; together income and social integration account for nearly 20% of the correlation. The strongest effects of cognitive ability/skills and social integration come from the NCDS in the United Kingdom, where ability accounts for 44% and social integration 15% of the correlation.⁸² Simultaneously controlling for economic resources, cognitive ability/skills, social integration, time preference, and sense of control accounts for 72% of the correlation between education and health behaviors in the NCDS; however, nearly all of this effect can be accounted for by income and cognitive ability/skills alone. In all specifications that explore the role of preferences for risk, time discounting, self-efficacy, or sense of control, these factors are found to explain little if any of the relationship between education and health behaviors. Lastly, Cutler and Lleras-Muney (2010) explored whether individuals who are better informed about the

⁸¹ All respondents in the NLSY were given a battery of 10 tests measuring such skills as math, science, reading comprehension, word knowledge, and electronics. These tests make up the Armed Services Vocational Aptitude Battery, four of which go into the well-known Armed Forces Qualifying Test (AFQT).

⁸² Ability in the NCDS is based on drawing, reading, and math tests administered at ages 7, 11, and 16; although, Cutler and Lleras-Muney (2010) showed that scores at ages 7 are not very important.

health impacts of smoking and drinking are less likely to smoke and drink, and whether this can explain any of the correlation between education and these unhealthy habits. Their findings are mixed. Although knowledge about the harms of smoking appears to explain about 15% of the relationship between education and smoking, knowledge about the harms of drinking do not explain differences in drinking by education.

The main conclusion from Cutler and Lleras-Muney (2010) is that both income and cognitive ability/skills play an important role in explaining the relationship between education and health behaviors. These two factors combine to explain roughly 30–50% of the correlation between education and health behaviors in the United States and nearly 70% of the correlation in the United Kingdom.⁸³ Social integration plays a modest role, and all other potential mediating factors appear to be relatively unimportant.

Although commonly measured health behaviors are important determinants of health and mortality, previous studies have reported that smoking, drinking, exercise, and other health behaviors can only explain about 30% of the relationship between education and health (e.g., Marmot (1994) and Cutler and Lleras-Muney (2006)). Furthermore, examining the effects of mediating factors on health behaviors does not necessarily inform us about the full importance of these mediating factors for health status or mortality, since some of these mediators may affect health through other channels. For example, increases in income may improve health by reducing stress or by allowing individuals to move to safer and healthier neighborhoods.

Ross and Wu (1995) directly examined the importance of the following mediating factors for the relationship between education and health: (1) work and economic resources, (2) social/psychological resources, and (3) health lifestyle. The first two factors are also analyzed in Cutler and Lleras-Muney (2010); however, the third factor is treated very differently in the two studies. Unlike Cutler and Lleras-Muney (2010), who examined the effects of other factors on health behaviors, Ross and Wu (1995) considered the mediating effect of health-related behaviors on health itself. Using data from the 1990 Work, Family, and Well-Being (WFW) Survey and the 1979–1980 National Survey of Personal Health Practices and Consequences (HP), Ross and Wu (1995) examined the extent to which these factors explain the relationship between education and both self-reported health and physical functioning.⁸⁴ In both data sets, education is significantly correlated with better self-reported health and physical functioning.

⁸³ It is not clear whether the difference between the United States and the United Kingdom is real or due to relatively more accurate measures of economic resources and ability in the UK's NCDS. Furthermore, Smith (1999, 2007) showed that correlations between income and health are at least partially due to the adverse effects of disease onset on income and labor supply. The extent of this "reverse causality" may differ across countries.

⁸⁴ Measures for work and economic conditions include household income, employment status, economic hardship, and work fulfillment (WFW only). Social-psychological resource measures include sense of control and social support. Health lifestyle measures include self-reported exercise, smoking, drinking (HP only), and health checkups (HP only). Controls for sex, race, age, and marital status are included in all specifications.

Assuming work and economic resources may affect social/psychological resources, which may in turn, influence health-related habits, Ross and Wu (1995) first added controls for work and economic resources, then controls for social-psychological resources, and finally controls for health behaviors. Consistent with earlier studies examining the relationship between education, income, and mortality (e.g., Elo and Preston (1996) and Kitagawa and Hauser (1973)), they found that economic resources explain 42–46% of the correlation between education and self-reported health and 34–54% of education's association with physical functioning. By comparison, the roles played by both social-psychological factors and health habits are much weaker, each explaining less than 10% of the association between education and either health measure. Ross and Wu (1995) estimated that all three broad classes of factors combine to explain 46–59% (66–71%) of the association between education and self-reported health (physical functioning).

The finding of Ross and Wu (1995) that health behaviors explain little of the overall education-health relationship suggest that the decomposition of the education-health behavior gradient by Cutler and Lleras-Muney (2010) is limited in its implications for health outcomes. Consistent with this finding, Cutler et al. (2010) showed that despite sizeable cross-sectional differences in mortality risk factors (e.g., smoking, obesity, high blood pressure, high cholesterol) by education, trends in these risk factors explain little of the observed widening in mortality gaps (generally and specifically for cancer and cardiovascular diseases) by educational attainment.

Altogether, these "decomposition studies" offer a few lessons but are not without caveats. Taken at face value, these studies suggest that income and ability combined are important for explaining differences in health behaviors or outcomes by education, while differences in preferences, specific health knowledge, and psycho-social factors are relatively unimportant. It should be noted, however, that evidence from Cutler and Lleras-Muney (2010) on the importance of ability relative to income is more difficult to interpret. This is because they control for income when examining the importance of ability, but they do not control for ability when estimating the importance of income. As shown in the Appendix, this ordering should be reversed since postschool income is strongly influenced by adolescent ability conditional on education (e.g., see Cawley, Heckman, Lochner, and Vytlacil (2000)).⁸⁵ Despite the suggestion by Cutler and Lleras-Muney (2010) that income is more important, the reverse is also quite possible. The fact that other studies (e.g., Ross and Wu (1995); Elo and Preston (1996); and Kitagawa and Hauser (1973)) also find that

⁸⁵ In general, if one is interested in the importance of some mediating factor x in explaining the relationship between education and health, then it is important to control for any additional mediating factors that may affect both x and health directly. By contrast, one would not want to control for additional mediating factors affected by x if they do not also affect x themselves. See the Appendix for details.

income plays an important mediating role does not alleviate this concern, since they too neglect measures of ability in their analysis. Conti, Heckman, and Urzua (2010a, 2010b) have argued that noncognitive skills may be more important factors in explaining education-health gradients than cognitive skills, further complicating matters. Altogether, it appears that education has important effects on health behaviors, health outcomes, and mortality via its combined effects on income and ability (both cognitive and noncognitive).⁸⁶ A second lesson from Ross and Wu (1995) and Cutler et al. (2010) is that differences in health behaviors (e.g., smoking, drinking, obesity) by education explain little of the education-health gradient or its evolution in recent decades. An important caveat cannot be overlooked. None of these studies have addressed issues of endogeneity, nor do they simultaneously control for all factors that may directly affect health and which may be correlated with potential mediators. These studies aim to decompose the correlation between education and health, but the relative importance of different factors in terms of correlation may not reflect the relative importance in terms of causal effects.

Rather than decompose education-health gradients into a variety of potential mediating factors or explanations, a number of other recent studies have focused on estimating and testing the implications of specific channels, including the role of productive efficiency, allocative efficiency, and peer effects/social networks. In drawing conclusions from these types of studies, it is important not to focus too much on any one mechanism or disease in trying to understand the link between education and health. Link and Phelan (1995) forcefully and convincingly argued that education and socioeconomic differences are "fundamental causes" of health disparities. That is, these factors affect the financial resources, knowledge, and access to power, which determine the ability and desire of individuals to avoid risks and address health problems. The mechanisms and diseases that lead to health disparities by socioeconomic status continually evolve over time, arguing against a narrow research or policy focus on any one mechanism or disease. Instead, it is important to interpret the body of evidence in light of more basic theories about the role of financial resources, differential costs, and information. We organize our discussion of this evidence along these lines.

If education improves productive efficiency, more educated individuals should (1) be healthier conditional on all health inputs and (2) use fewer inputs conditional

⁸⁶ Smith (1999, 2007) argued that changes in income have relatively small effects on the onset of major health diseases; instead, the onset of disease tends to adversely affect subsequent income. This might suggest that health plays a "mediating" role in determining the effects of education on income and wealth. In decompositions of the relationship between education on health, this would incorrectly be attributed to a mediating role of income. Of course, the correlation between income and health may be largely due to differences in treatment and management of disease by income rather than differences in the likelihood of disease onset as emphasized by Smith (1999, 2007). In this case, education–health decompositions are likely to measure the mediating effect of income.

on input prices and health outcomes (Grossman (1972a, 1972b)). Earlier studies by health economists (e.g., Grossman (1972b); Wagstaff (1986); Erbsland, Ried, and Ulrich (1995); and Gilleskie and Harrison (1998)) largely focused on measuring the role of productive efficiency based on these two implications. Findings from these studies are generally supportive of productive efficiency. Unfortunately, these methods require good measures of all health inputs or input prices. If some inputs or prices are missing, it is extremely difficult to estimate the role of productive efficiency, since unobserved inputs or prices are likely to be correlated with education. In practice, these approaches also typically rely on strong functional form assumptions about the health production function—assumptions that are difficult to verify.

Kenkel (1991) is the first to directly examine the importance of allocative efficiency by estimating whether specific health knowledge affects behavior. Using the same decomposition approach as described earlier, he estimated whether controlling for knowledge about the harms of smoking, drinking, and exercise reduces the correlation between education and those behaviors. His estimates suggest that both knowledge and education affect smoking, heavy drinking, and exercise; however, differences in health knowledge explain less than 20% of the correlation between education and those behaviors (controlling for income, employment, and standard demographic factors). Thus, health-specific knowledge matters, but it does not explain much of the difference in health behavior. This finding is echoed in the more recent study by Cutler and Lleras-Muney (2010).

A few recent studies have taken a more indirect approach to evaluate the role of health-specific knowledge and allocative efficiency. These studies estimated the extent to which new public health information affects the relationship between education and specific health behaviors. De Walque (2010) and Aizer and Stroud (2010) showed that education-smoking gradients increased substantially in the United States after the 1964 Surgeon General's Report announced the dangers of smoking. De Walque (2007b) found that the education-HIV prevalence gradient evolved in a similar way following a major HIV/AIDS prevention campaign in Uganda during the late 1980s and early 1990s. Price and Simon (2009) examined the impact of a medical research study published in 2001, which showed that the risks of complications from attempting a vaginal birth after a caesarean (VBAC) were higher than previously thought. Their analysis shows that VBACs declined sharply after publication of the article, with significantly larger declines among more educated women. In a variety of contexts (and countries), these studies show that new publicly provided information about health procedures or behaviors tends to widen gaps in those procedures/behaviors by education.

As Lange (forthcoming) notes, this type of evidence is consistent with two potential explanations: (1) more educated individuals may gather and process information more easily and/or (2) more educated individuals may respond more to new information

because it is of greater value to them (e.g., see Section 3.1).⁸⁷ Assuming information is exogenously acquired, Lange (forthcoming) showed that examining actual beliefs about the likelihood of a health outcome (or someone's risk level) may help identify whether explanation (1) is important. If more educated individuals are better informed, then their beliefs should be more strongly correlated with objective risks. However, the fact that more educated individuals may have a greater demand for health does not, by itself, make that information any better. If information is exogenously acquired, then there is no reason to expect more educated individuals to have more accurate beliefs due to their greater demand for health. Based on this insight, Lange (forthcoming) estimated the correlation between individuals' beliefs about their risk of cancer and objective risk levels (based on demographics and family history). His estimates suggest that this correlation is stronger for more educated individuals. Consistent with this finding, Aizer and Stroud (2010) showed that knowledge about the harms of smoking increased more quickly among the most educated following the 1964 Surgeon General's warning. Both Lange (forthcoming) and Aizer and Stroud (2010) concluded that those with more education are better informed and that this affects their health decisions (i.e., cancer screenings and smoking decisions). An obvious complication is the likely possibility that information about risks of cancer or smoking-related health problems is not exogenously acquired. Because more educated individuals have a greater demand for health generally, they are also likely to have a greater demand for accurate health information about their objective risks, causing them to invest more in acquiring that information. Thus, differences in endogenously acquired information should reflect differences in demand for health in the same way differences in health behaviors do.

Overall, these studies convincingly show that education affects the acquisition of health-related information. Whether or not improved information comes at a cost is more difficult to determine. Furthermore, does the information translate into important differences in health behaviors and, ultimately, health outcomes? Aizer and Stroud (2010) showed that the widening of the smoking differential corresponds to the emergence of a knowledge gap about the effects of smoking on heart disease. Lange (forthcoming) also showed that more educated individuals are both better informed about their cancer risks and more likely to undergo screening for cancer. Despite these findings, evidence by Kenkel (1991) and Cutler and Lleras-Muney (2010) suggests that differences in health-specific knowledge by education account for only small differences in behavior. Furthermore, differences in behavior explain only a small fraction of differences in health (Ross and Wu (1995)).

⁸⁷ In some cases, it may be useful to distinguish between the potential for education to improve the ability to acquire information and the ability to process that information and make good decisions. In the former case, attempts to better distribute information to lesser educated individuals should be helpful, while the latter suggests efforts to simplify and explain the implications of new information may be needed.

A number of recent studies have explored the extent to which new technologies affect the relationship between education and health. Goldman and Smith (2002) estimated that more educated patients who suffer from diabetes or from HIV better self-manage their diseases by adhering to complex treatments. Most notably, they found that less-educated diabetes patients benefitted more (than more-educated patients) from an experimental treatment regimen with intensive patient monitoring. Because treatments were provided freely as part of the experiment, differences in income are unlikely to explain the differential responses by education. This strongly suggests that factors other than income (e.g., cognitive ability) may play an important role for difficult-to-manage diseases and that differences in these factors by education may help explain the education–health gradient. Unfortunately, it is difficult to know whether these other factors are as important outside the "laboratory" where more intensive treatments also tend to be more costly.⁸⁸

As noted earlier, Glied and Lleras-Muney (2008) found that more educated individuals experience greater survival advantages for diseases that show the fastest technological progress. They argued that more educated individuals may be more likely to adopt new treatments and medicines. However, Goldman and Smith (2005) found no evidence of differential diffusion of hypertension drugs by education over the 1980s and 1990s. Goldman and Lakdawalla (2005) showed that the impacts of new treatments on education-health relationships are likely to depend on whether those treatments are more complicated or simpler than previous treatments. Consistent with Goldman and Smith (2002), they showed that antiretroviral therapies for HIV introduced in the 1990s led to greater health improvements among more educated HIV patients. By contrast, the emergence of β -blockers as a form of treatment for hypertension reduced education disparities in hypertensive cardiovascular disease. More generally, they point out that innovations in health practices, technologies, or pharmaceuticals that raise optimal treatment costs are likely to exacerbate health and mortality differences by education, while innovations that lower treatment costs should do the opposite. Of course, most new technologies tend to be more expensive than their predecessors, at least initially, suggesting that periods of rapid medical innovation are likely to be associated with widening education-health gradients.

It is difficult to know whether differences in disease management and medical treatments are due to differences in the ability to utilize best practices/treatments or to differences in demand for health. Do more educated individuals better self-manage

⁸⁸ Maitra (2010) estimated that the impacts of education on diabetes in the Health and Retirement Survey are largely due to differences in economic resources and health insurance coverage rather than differences in disease management or cognitive ability as emphasized by Goldman and Smith (2002). Maitra (2010) did not dispute the experimental evidence by Goldman and Smith (2001); however, his results may suggest that income and health insurance differences by education are more important for diabetes management once the actual costs of different treatments are factored in.

their diseases because it is easier for them to do so or because they have more to gain from putting forth the effort? Do less educated individuals fail to take the newest and most effective drugs, because they find the regimens difficult or because they are too expensive? Despite serious efforts taken in these studies to narrow the focus on a single mechanism, these open questions reveal the difficulty of the task at hand.

A final issue worth discussing is the spillover effects of education for health. There is some evidence regarding the peer effects associated with smoking and other health outcomes that speaks to the potential for education to have external effects on the health of others. A number of recent studies (e.g., Norton, Lindrooth, and Ennett (1998); Gaviria and Raphael (2001); Powell, Tauras, and Ross (2005); and Fletcher (2010)) have used IV techniques to estimate the influence of peers on smoking and drinking among school-age adolescents. Most of these studies assumed that school peers' family background or neighborhood characteristics (e.g., parental education, family composition, race) have no direct affect on whether someone engages in risky behaviors conditional on the individual's own family background and school characteristics. (They also implicitly assumed that these neighborhood or parental characteristics are uncorrelated with unobserved factors that may affect youth's decisions.) These studies have typically found statistically significant peer effects; however, the magnitudes vary considerably. Using class-level data, Fletcher (2010) showed that incorporating school fixed effects and allowing for some contextual effects of peer background characteristics on smoking reduces the estimated peer effects on smoking by about 30% relative to the approach taken in most previous studies. Krauth (2005) eschewed the strong exogeneity assumptions employed by these studies and instead took a more structural approach to bound the effects of peers on smoking with different assumptions about the correlation in unobserved tastes across individuals. His approach suggests that the effects of peers on smoking are likely to be much smaller than IV estimates suggest. Altogether, these studies suggests that peer effects for risky behaviors like smoking and drinking may amplify the effects of policies on youth; however, Cohen-Cole and Fletcher (2008) estimated insignificant effects of peer obesity levels on one's own obesity among middle and high school students in the United States. Little is known about the extent to which peer effects and social networks are important for adults and their health behaviors. Given the evidence that differences in habits and health behaviors by education only explain a small share of the overall education-health gradient, it seems unlikely that these behavior spillover effects are economically very important. The most important spillover is likely to be from parental education to child health as discussed earlier.

In summary, evidence on the extent of productive efficiency is suggestive but comes with many caveats. Growing evidence more clearly suggests that education affects the allocation of health inputs and behaviors; however, it is still not clear whether these allocations are more efficient for more educated individuals or simply different due to differential demand for health. The fact that differences in income explain one-quarter to one-third of the correlation between education and health/mortality is consistent with a differential demand-based story. The fact that differences in ability (cognitive and/or noncognitive) seem to matter suggests that productive and allocative efficiency may also be important. Peer effects appear to exist among adolescents for some risky behaviors like smoking and drinking, but there is no evidence to suggest that this type of externality is important among adults or for broader measures of health.

3.3. Implications for Education and Health Policy

For public policy purposes, it is constructive to categorize potential channels according to the following: (1) Does the potential benefit to health come at some direct cost to the individual (other than through costs associated with schooling)? (2) Does the health benefit come primarily from one's own schooling or does it result from the schooling of others?

If acquiring more schooling makes someone healthier or live longer, there is every reason to think individuals will take that into account when making their education decisions—if they are aware of the benefit and value it provides. If the benefits are private and known, then there is little case to be made (at least on efficiency grounds) for subsidizing schooling with the aim of improving health outcomes. On the other hand, if the benefits are unknown or are public in nature (e.g., peer effects), then education subsidies may be efficient. Thus, it is clearly important to distinguish between channels implying public versus private benefits. Most mechanisms that link education and health suggest that the benefits are private or at least contained within the family. Empirical evidence on peer effects in health behaviors is limited to a few risky behaviors and adolescents. Although these studies tend to find some role for peers, it is unlikely that the peer effects of education on health more generally are very important. There is some evidence that parental education affects child health outcomes. To the extent that individuals internalize benefits to their future offspring and recognize benefits to their own health when making education decisions, they have a strong incentive to optimally weigh those benefits against the costs of schooling.

If individuals are unaware of the health benefits of education (for themselves or their future children) at the time schooling decisions are made, it is important to distinguish between benefits that must be paid for and those that are freely acquired through schooling. If individuals are unaware of benefits that come without cost (e.g., lower stress, better at gathering/interpreting information, better decision-making), youth may acquire too little schooling. However, failure to fully consider future benefits that come at a personal cost (e.g., better health care or healthier lifestyle choices) is not likely to distort education decisions in the same way it is not important for youth to know how much they will spend on golf during retirement.

Distinguishing between costly and free health benefits of education may also be of interest to researchers interested in quantifying the total net benefits from schooling.

To this end, health benefits (net of their associated costs) should be incorporated. These benefits may be important for accurate cost-benefit analyses of specific educationrelated programs or more generally for evaluating whether individuals appear to appropriately balance the total costs and benefits of schooling. If the total private benefits substantially outweigh the private costs, it may indicate that distortions exist in the education sector (e.g., borrowing constraints, lack of information about the costs or benefits of schooling).

What do we know about the extent to which the health benefits of education are paid for? On the one hand, evidence that health behaviors and habits play little role in explaining the education-health gradient would seem to suggest that differences in these types of costly inputs are not particularly important (even though education does appear to affect smoking behavior). On the other hand, the fact that differences in income by education explain one-fourth to one-third of the education-health gradient is consistent with some role for costly inputs. It is not clear why income should be such an important mediator otherwise. Of course, as noted earlier, none of these studies necessarily identifies the channels through which education causally affects health. Furthermore, much of the education-health relationship is left unexplained. Unfortunately, existing studies are not particularly informative about the extent to which the health benefits of education are effectively paid for through costly changes in behavior or in health inputs.

In calculating the total returns to schooling, it would be useful to put a dollar value on all estimated health benefits. Most benefits are difficult to value (e.g., improvements in self-reported health) even if we could accurately estimate their magnitudes. Fortunately, Eq. (2.10) shows how we can value the estimated reductions in mortality associated with schooling using measures of the value of a statistical life (VSL). Typical VSL estimates of \$3-5 million suggest that a 0.01 reduction in 10-year mortality rates should be valued at roughly \$3-5 thousand in annual gains.⁸⁹ Since the IV estimates in Table 2.4 generally measure the total effect of education on mortality, $\frac{d\pi}{ds}$, rather than the desired partial effect, $\frac{\partial \pi}{\partial s}$, a high-end estimate of the total benefit from mortality reduction is probably around \$9-15 thousand dollars (assuming 10-year mortality rates are reduced by 0.03, roughly half the IV estimates of Lleras-Muney (2005, 2006)). This is quite large relative to estimated increases in annual earnings for an additional year of school; however, it is probably much higher than the actual benefit given the small and statistically insignificant estimates on mortality found in three of the four studies reported in Table 2.4. If education reduces 10-year mortality rates by 0.01-roughly 10% of baseline mortality rates for 45-80 year olds in Lleras-Muney (2005) (a figure more in line with the larger set of estimates on self-reported health, health limitations,

⁸⁹ See Viscusi (1993) and Viscusi and Aldy (2003) for summaries of VSL estimates.

and activity)—and if roughly half of that reduction is "paid for" in the form of costly health investments and behavior changes, then a ballpark figure for the mortality benefits of an extra year of school is probably on the order of \$1500–2500 per year, a nontrivial sum comparable in size to the gains from crime reduction discussed earlier. The value of more general health improvements may also be sizeable.

Of course, even if education does produce health-related externalities or if some of the health benefits of education are not incorporated by youth making their education decisions, it is not clear that education-related policies are the best way to address these problems. Other, more targeted information or public health campaigns may be more efficient policies worth considering. General education policies may also produce different health effects in different environments. Table 2.7 suggests that the health and mortality benefits from education (and education policies) may be weaker in Europe (compared with the United States) where there is greater access to health care, more generous social support, and lower income inequality. We might expect the opposite in developing countries characterized by high inequality and very limited access to health care.

3.4. Suggestions for Future Research on Education and Health

Research on the effects of education on health has accelerated in recent years. Much of that research has justifiably focused on (1) identifying causal effects using IV or RD methods or (2) on identifying the channels by which education influences health. Most studies of the former type exploit changes in compulsory schooling laws in a few developed countries. Looking forward, we are likely to learn the most from exploiting other sources of exogenous variation in schooling at different margins or by looking at a wider range of countries (or by focusing on different subpopulations). Does schooling raise health in developed economies? Do the effects of education differ across countries with very different underlying levels of inequality or access to medical care? Answers to these questions are not only interesting in their own right, but they may also shed additional light on key mechanisms.

Thanks to a number of recent studies, we know much about factors that "explain" the correlation between education and health, but we know little about how these factors actually determine the *causal* relationship. Attempts to identify different channels through which education causally impacts health and mortality should, therefore, be high on the list of priorities. In particular, it is important to better determine the extent to which health and mortality benefits of education are driven by differences in demand, in which case they are likely to be paid for via costly inputs and foregone opportunities. This is of central concern for policy analysis.

In addition to searching for exogenous sources of variation in schooling induced by policies and other natural experiments, efforts to put some structure on the education and health decision problems, combined with additional data on individual abilities and skills as in Conti, Heckman, and Urzua (2010a, 2010b), is also likely to yield new insights about which subpopulations benefit most from education and why.

4. EDUCATION'S EFFECT ON CITIZENSHIP, POLITICAL PARTICIPATION, AND DEMOCRACY

The hypothesis that education encourages and strengthens democracy has a long intellectual history, with Lipset (1959) crediting the basic idea to Aristotle. Early views by Lipset (1959) and Aristotle emphasized the role of education in informing citizens and increasing their capacity to make "good" electoral decisions while resisting demagoguery. Education may also affect both the benefits and costs of voting and other forms of political engagement. For example, education may instill civic and democratic values, either through the explicit design of education systems (especially in democratic countries) or indirectly by improving analytical skills and an awareness of history and a diversity of opinions.⁹⁰ Education may also indirectly affect political participation by altering social networks and peers. By raising wage rates, education may affect the time costs associated with active political participation and voting. As noted by Verba, Scholzman, and Brady (1995), individuals reported that lack of time is the most important reason for political inactivity.

Despite the plethora of hypotheses linking education and democracy, formal economic models of this link are scarce. A notable exception is Glaeser, Ponzetto, and Shleifer (2007), who emphasize the social nature of political action and education's role in facilitating social interaction. They develop a model in which education endogenously affects political participation, assuming more educated individuals are better at persuading others to become politically active (or, alternatively are more persuadable by other active participants). To the extent that these social influences are more valuable to larger political constituencies with less active participation (i.e., broader political regimes), the authors showed that a more educated population is likely to be more inclusive and democratic.

Although democracy and political freedoms are, no doubt, intrinsically valuable, economists have largely been interested in the link between education and democracy based on the potential for democratic institutions to facilitate economic growth.⁹¹ This is evident in two other political economy models that consider the inter-relationship

⁹⁰ Spilimbergo (2009) found that countries which send students abroad for higher education experience greater democratization if the host countries are democratic. This suggests that direct efforts by education systems to teach democratic ideals may be important.

⁹¹ Proponents of the idea that democracy encourages growth are diverse in their reasoning (e.g., see Hayek (1960); Wittman (1989); North (1990); Olson (1993); and Sen (1999)). Others are more skeptical (e.g., see de Tocqueville (1835/1998); Huntington (1968); Buchanan and Tullock (1962); Alesina and Rodrik (1994); Persson and Tabellini (1994); and Besley and Coate (1998)). Evidence on the effects of democracy on growth is mixed (e.g., see Przeworski and Limongi (1993); Barro (1997); Minier (1998); La Porta, Lopez-de Silanes, Shleifer, and Vishny (1999); Tavares and Wacziarg (2001); Doucouliagos and Ulubasoglu (2008); and Papaioannou and Siourounis (2008a)).

between the distribution of education, democracy, and economic growth. These models focus more on the conflict between different political constituencies as defined by education and income rather than potential reasons educated citizens may behave more democratically. Bourguignon and Verdier (2000) assume that power is initially concentrated among a small group of educated elites and that the larger uneducated populace is excluded from the political process. By assumption, only the educated vote in their framework. Their main objective is to understand conditions under which the educated elite will selfishly choose policies to increase access to education and, therefore, broaden the political base. Their results suggest that a more equal initial distribution of education leads to faster democratization and more economic growth. Rajan (2009) instead emphasized the fact that less powerful constituencies have competing education and policy preferences. He argued that in a plutocracy in which political power is determined by wealth, a growth in wealth and power among educated lower-middle class constituents may threaten the economic rents earned by wealthier elites more than the prospect of a fully democratic government in which everyone has equal power regardless of wealth. In this case, the distribution of education may affect decisions to democratize; however, democratization, if it occurs, is likely to be characterized by policy gridlock and weak economic growth.

Beginning with Barro (1999), economists have sought to empirically determine whether increases in a nation's education strengthen democratic institutions. Most studies are macro in nature, using aggregate measures of democracy and education across countries and over time to study the issue. A few microeconometric studies have exploited individual-level data within the United States, United Kingdom, and Germany to estimate the effects of an individual's own education on their likelihood of participating in the political process in one way or another. We review both types of studies below.

4.1. Macro Evidence on Education and Democracy

Empirically, education and democracy are highly correlated across countries. For example, Barro (1999) showed that countries with higher average years of primary schooling also have greater electoral rights and civil liberties. Economists have begun to examine this relationship more closely to determine whether changes in education actually lead to changes in democracy and political freedom, attempting to assign causality. These studies typically estimate an equation of the form:

$$D_{c,t} = \rho D_{c,t-1} + E_{c,t-1}\beta + X_{c,t-1}\gamma + \delta_t + \mu_c + \varepsilon_{c,t}, \qquad (2.12)$$

where $D_{c,t}$ reflects some measure of democracy or political openness in country *c* and period *t*, $E_{c,t-1}$ reflects some measure of a country's education level in period t-1, $X_{c,t-1}$ reflects any number of potential time-varying determinants of democracy, δ_t reflects unrestricted time effects, and μ_c reflects country-fixed effects. Glaeser, La Porta, Lopez-de Silanes, and Shleifer (2004) estimated a version of Eq. (2.12) using standard fixed effects estimation methods (and 5-year periods). Their estimates suggest that lagged education improves a variety of measures of democracy.⁹² However, Acemoglu, Johnson, Robinson, Yared (2005) have pointed out that Glaeser, La Porta, Lopez-de Silanes, and Shleifer (2004) did not account for time effects in their specification, so identification largely comes from widespread international increases in education and democracy over time. Accounting for time effects (using the same specifications and data), Acemoglu, Johnson, Robinson, Yared (2005) found little remaining effect of education.⁹³

More importantly, it is well-known that simple fixed effects estimates of Eq. (2.12) are likely to be biased with short panels. This is because $D_{c,t-1}$ is correlated with all past values of $\varepsilon_{c,t-j}$ for $j \ge 1$. (Note that $E_{c,t-1}$ is also likely to be correlated with past ε 's if there are any feedback effects of democracy on schooling.) Arellano and Bond (1991) showed that if the ε 's are serially uncorrelated, then one can obtain consistent estimates by first-differencing Eq. (2.12) and using lagged values of democracy (and schooling) to instrument for their differences, based on the following moment condition: $E[(\varepsilon_{c,t} - \varepsilon_{c,t-1})D_{c,t-j}] = 0$ for all $j \ge 2$. Using this approach and 5-year differences from 1965 to 1995, Acemoglu, Johnson, Robinson, Yared (2005) estimated small and statistically insignificant effects of average educational attainment on political freedom. They, therefore, concluded that cross-country correlations between education and democracy are due to omitted country-specific factors (e.g., political, social, and economic institutions).⁹⁴

Most recently, Bobba and Coviello (2007) and Castelló-Climent (2008) argued that the Arellano and Bond estimator is likely to suffer from weak instrument problems, since democracy is highly persistent over time (in this case, past values of democracy are only weakly correlated with changes in democracy). A similar problem arises with respect to measures of education if one instruments for its changes using past values. The presence of weak instruments implies that estimates will be imprecise and biased toward the standard first-difference estimator (that does not use any instruments). These authors, therefore, suggest that the "system GMM" estimator of Blundell and Bond (1998) be used instead. This estimator includes the additional moment equation,

$$E[(D_{c,t-1} - D_{c,t-2})(\mu_c + \varepsilon_{c,t})] = 0, \qquad (2.13)$$

⁹² Specifically, they found significant effects on a measure of "executive constraints" and a general measure of democracy (Jaggers and Marshall (2000)) as well as autocracy as measured by Alvarez, Cheibub, Limongi, and Przeworski (2000).

⁹³ Papaioannou and Siourounis (2008b) used a sample of countries that are nondemocratic in 1960 to estimate the effects of education on the probability of making a democratic transition. Their estimates suggest that education significantly increased the probability, intensity, and speed of democratization.

⁹⁴ Also, see Acemoglu, Johnson, and Robinson (2005) on the importance of institutions.

which implies that changes in democracy must be orthogonal to country-fixed effects. Both Bobba and Coviello (2007) and Castelló-Climent (2008) estimated much larger and statistically significant effects of educational attainment on democracy using this estimator. More interestingly, Castelló-Climent (2008) found that the distribution of education in a country is more important than the average level. Her estimates suggest that greater equity in educational outcomes is associated with stronger measures of democracy.⁹⁵

The system GMM estimator has the advantage of using additional information/ instruments; however, the cost is an additional stationarity assumption (Eq. (2.13)) that is quite strong, since it requires that changes in democracy be unrelated to steady-state levels of democracy (i.e., $\mu_c/(1 - \rho)$). Given the widespread increases in democracy over the past few decades, it seems likely that many countries are still converging toward their long-run steady states. In this case, we would expect μ_c to be positively correlated with changes in democracy, violating the additional assumption of the system GMM estimator. In this case, there is no clear preference between the Arellano-Bond estimator and the Blundell-Bond system estimator—the former likely suffers from bias due to weak instruments, while the latter likely suffers from bias due to an invalid moment restriction. Given these concerns, we turn next to microeconometric evidence on the causal effects of education on voting and citizenship.

4.2. Estimating "Causal" Effects of Education on Citizenship

A large literature in political science has demonstrated a strong correlation between educational attainment and political participation, voting, and civic awareness at the individual level. See Nie et al. (1996) for an extensive review. Of course, as noted repeatedly above, correlation need not imply causation. Three recent studies by economists (Dee (2004); Milligan, Moretti, and Oreopoulos (2004); and Siedler (2010)) used individual-level data and instrumental variables strategies to estimate the causal effects of educational attainment on voter registration, voting, support for free speech, and other measures of civic engagement (e.g., newspaper readership, group membership). Table 2.9 summarizes evidence from these studies on voting and voter registration, while Table 2.10 summarizes evidence on other measures of political/social involvement and views.

Dee (2004) examined the effects of education on voting and other forms of civic engagement in the United States using data from the High School and Beyond (HS&B) longitudinal study and the General Social Surveys (GSS). HS&B followed a cohort of high school sophomores in 1980 through 1992, while Dee's analysis of the 1972–2000 GSS considers individuals born 1900–1964. Milligan, Moretti, and Oreopoulos (2004)

⁹⁵ Spilimbergo (2009) used the system GMM estimator to estimate the effects of "sending" postsecondary students abroad to democratic countries on the originating country. His estimates suggest that foreign-educated individuals promote democracy back home if the education is acquired in democratic countries.

Study	Data	Education Measure	Instrumental Variable	Voting/Registration Measure	Sample Mean	IV Estimate	% Effect
A. United States							
Dee (2004)	HS&B	College	Local 2-yr	Registered to vote	0.67	0.218^{*}	32.5
		enroll.	colleges	Voted in last 12 months	0.36	0.176*	48.9
				Voted in 1988 presidential election	0.55	0.178	32.4
Dee (2004)	GSS	Years of school	State school/ Labor laws	Voted in last presidential election	0.73	0.068*	9.3
Milligan et al. (2004)	NES	High school	State school/	Registered to vote	0.82	0.093	11.3
		graduation	Labor laws	Voted in last election	0.68	0.288*	42.3
Milligan et al. (2004)	CPS	High school graduation	State school/ Labor laws	Voted in last election	0.62	0.435*	70.4
B. United Kingdom		0					
Milligan et al. (2004)	BES	Years of school	National minimum school age	Voted in last election	0.85	-0.008	-0.9
Milligan et al. (2004)	EB	Years of school	National minimum school age	Name on electoral list	0.92	-0.014	-1.5
C. Germany							
Siedler (2007, 2010)	ALLBUS	Years of school	National minimum school age	Voted in last election	0.93 ^a	0.005	0.5
Siedler (2007, 2010)	ForsaBus	Years of school	National minimum school age	Voted in last election	0.93 ^a	0.000	0.0

Table 2.9 Effects of Education on Voter Registration and Voting

Notes: HS&B = High School and Beyond; GSS = General Social Surveys; NES = National Election Survey; CPS = Current Population Survey November Voting Supplements; BES = British Election Studies; EB = Eurobarometer surveys for Britain and Northern Ireland.

^aSample means taken from ALLBUS data in Table 3 of Siedler (2007) for those "not affected by the reform."

*Reflects statistically significant at 0.05 level.

Study	Country	Data	Education Measure	Instrumental Variable	Outcome	Sample Mean	IV Estimate	% Effect
A D 1:: 14								
A. Political Interes	st						0 - - 0*	
Milligan et al. (2004)	United States	NES	H1gh school graduation	State school/labor laws	Interested in election	0.30	0.270*	90.0
Siedler (2007, 2010)	Germany	ALLBUS	Years of school	National minimum school age	Interested in politics	0.30^{a}	-0.096*	-32.0
Siedler (2007, 2010)	Germany	ForsaBus	Years of school	National minimum school age	Interested in politics	0.30 ^a	-0.008	-2.7
B. Political Inform	nation							
Dee (2004)	United States	GSS	Years of school	State school/labor laws	Newspaper readership index (0–4)	3.2	0.113*	3.5
Milligan et al.	United States	NES	High school	State school/labor laws	Follow campaign on TV	0.79	0.392*	49.6
(2004)			completion		Follow campaign in newspapers	0.66	0.852*	129.1
					Follow public affairs	0.66	0.544*	82.4
Milligan et al. (2004)	United States	EB	Years of school	National minimum school age	Follow news every day	0.9	-0.007	-0.8
C. Political Views	/Philosophy			C C				
Dee (2004)	United States	GSS	Years of school	State school/labor laws	Allow homosexuals to speak	0.73	0.123*	16.8
					Allow militarist to speak	0.59	0.036*	6.1
Siedler (2007, 2010)	Germany	ALLBUS	Years of school	National minimum school age	Demonstrating is civil right	0.67 ^a	-0.001	-0.1
								Continued

Table 2.10	Effects	of Education	on	Political	Interest,	Information,	Views,	and	Particip	atior
------------	---------	--------------	----	-----------	-----------	--------------	--------	-----	----------	-------

	Table 2.	10 Effects	of Education o	n Political Interest	, Information,	Views,	and Partici	pation—continu
--	----------	------------	----------------	----------------------	----------------	--------	-------------	----------------

Study	Country	Data	Education Measure	Instrumental Variable	Outcome	Sample Mean	IV Estimate	% Effect
D. Political/Social	Participation							
Dee (2004)	United States	HS&B	College enroll.	Local 2-yr colleges	Volunteered in last 12 months	0.37	-0.047	-12.7
Dee (2004)	United States	GSS	Years of school	State school/labor laws	Group memberships	1.8	0.164	9.1
Milligan et al. (2004)	United States	NES	High school grad.	State school/labor laws	Attend political meetings	0.07	0.132	188.6
Milligan et al. (2004)	United Kingdom	EB	Years of school	National minimum school age	Often or occasionally try to persuade others to share views	0.45	0.066*	14.7
					Discuss political matters with friends at least occasionally	0.67	0.095*	14.2
Siedler (2007, 2010)	Germany	ALLBUS	Years of school	National minimum school age	Active in citizen group Participated in demonstration	0.17 ^a 0.09 ^a	-0.050 -0.016	-29.4 -17.8

Notes: HS&B = High School and Beyond; GSS = General Social Surveys; NES = National Election Survey; EB = Eurobarometer surveys for Britain and Northern Ireland. ^aSample means taken from ALLBUS data in Table 3 of Siedler (2007) for those "not affected by the reform." *Reflects statistically significant at 0.05 level. examined political outcome measures in the United States using the biennial National Election Surveys (NES) from 1948 to 2000 and the 1978–2000 waves of the November Voting Supplement to the Current Population Survey (CPS). The former offers more detailed and consistent measures of political activity for a long time period, while the latter offers much larger sample sizes. Milligan, Moretti, and Oreopoulos (2004) also analyzed political behavior in the United Kingdom using the British Election Surveys (BES) for 1964, 1974, 1979, 1983, 1987, 1992, and 1997, as well as 50 surveys from the Eurobarometer (EB) United Kingdom surveys covering 1973–1998. Siedler (2010) studied political outcomes in (formerly Western) Germany using data from even numbered years of the ALLBUS (German Social Survey) from 1980 to 2004 (plus the 1991 survey) as well as the large ForsaBus survey on political attitudes from 1991 to 2006.

All three studies exploit variation in compulsory schooling or minimum work age laws as discussed in previous sections. Cross-state variation in the these laws over time is exploited in the United States and Germany, while increases in minimum schooling ages in United Kingdom occurred nationally in 1947 and 1973. Specifications using variation in these laws identify the effects of secondary schooling on political behaviors and attitudes. By contrast, Dee's (2004) analysis using the HS&B data exploits cross-sectional variation in (1) distance two a 2-year college and (2) number of 2-year colleges in the county of residence in 1983. He, therefore, estimates the causal effect of college enrollment (rather than secondary schooling) on voting and volunteering using these data.⁹⁶

Table 2.9 reveals sizeable effects of schooling on voting behavior in the United States, but much weaker and statistically insignificant effects in the United Kingdom and Germany. In the United States, the effects of high school graduation on voting are comparable to the effects of college enrollment: the latter increases rates of voting by 30–50% (Dee (2004)), while the former increases voting by 40–70% (Milligan, Moretti, and Oreopoulos (2004)).⁹⁷ Education not only affects voting in the United States but also affects voter registration. Milligan, Moretti, and Oreopoulos (2004) argue that differences in registration may drive the impacts of education on voting in the United States. The fact that voter registration is voluntary and left to the individual in the United States, while it is mandatory and the responsibility of local governments in the United States.

⁹⁶ We refer to Dee's (2004) bivariate probit specifications (with instruments excluded from the voting choice equation) as IV-based estimates. We report his estimates of the average marginal effect of enrollment.

⁹⁷ Dee (2004) notes that the estimated effect of a year of completed schooling in the GSS, 0.07, is comparable in magnitude to the effect of college enrolment, 0.18–0.20, after accounting for the fact that college enrollment is associated with roughly 2.5 years of additional schooling, on average.

Table 2.10 summarizes the estimated effects of education on political interest, efforts to acquire information about politics and elections, political views about civil rights and freedoms, and general political and social involvement. Dee (2004) and Milligan, Moretti, and Oreopoulos (2004) estimated sizeable effects of education on a broad range of political domains in the United States. In particular, they found that additional years of high school significantly increase interest in politics, efforts to acquire information about political issues/campaigns, and beliefs in freedom of speech. By contrast, evidence regarding volunteering and group membership, as emphasized in the theory of Glaeser, Ponzetto, and Shleifer (2007), receives little support in US data. There is, however, modest evidence from the United Kingdom that additional secondary schooling increases individuals' efforts to discuss politics with others and to persuade others to share their views (Milligan, Moretti, and Oreopoulos (2004)). Evidence from Germany is generally inconsistent with the view that education encourages democratic ideals (Siedler (2010)).

Overall, the micro-empirical evidence suggests important effects of education on a wide range of political behaviors and views in the United States; however, this is not the case in the United Kingdom or Germany. Evidence from the United States is most consistent with the older views of Aristotle and Lipset that education leads to a more informed and engaged citizenry. The evidence is less supportive of the idea that schooling influences political outcomes by substantially altering the nature of social interaction.

5. CONCLUSIONS

A fast growing literature has established that education and human capital impact a wide range of personal decisions and activities. Education has been shown to reduce crime, improve health, lower mortality, and increase political participation. The social benefits from these impacts can be sizeable. For example, Lochner and Moretti (2004) estimated that high school completion may lower the annual social costs of crime by roughly \$3000 per male graduate. Increasing high school completion rates in the United States by one percentage point would reap a savings of more than \$2 billion. Annual benefits from reductions in mortality are likely to be in the neighborhood of \$1500–2500 per additional graduate; however, there is considerable uncertainty about this value.

Much of the evidence on the causal effects of education on crime, health, and citizenship has come from changes in compulsory schooling laws, which primarily affect secondary schooling choices. School lotteries have also been used to estimate the effects of improvements in middle and high school quality on delinquency and crime. As a result, we know much about the impact of additional years of high school on crime, health, and citizenship, but evidence on the effects of higher education is more scarce. Given the rise in college attendance throughout the world, additional efforts to study this margin are needed. There is good reason to believe that increases in college-going are not likely to yield dramatic benefits from crime reduction (at least in the near future), since studies have shown that education-based interventions and policies appear to reduce crime and delinquency most among the least able, most disadvantaged. A few studies have estimated significant reductions in smoking and improvements in political participation in response to additional years of college, but studies that measure the impacts of higher education on health or citizenship are the exception. There is growing evidence that preschool and school interventions at early ages can reduce delinquency and crime years later. Although some programs (e.g., Perry Preschool, CCPC) have been credited with remarkable reductions in adult crime, other similar programs (e.g., Abecedarian, IHDP) have not. Unfortunately, we do not yet understand these differential impacts.

Much of the evidence is based on the United States. A growing number of very recent studies have begun to analyze the nonproduction benefits of education in the United Kingdom, Germany, and Europe more generally; however, very few studies exploit data from developing countries where education levels are much lower. One might expect substantial differences in the impact of education on crime, health, and citizenship across countries with very different criminal justice, health care, and political systems. Indeed, comparisons across estimates from the United States and Europe seem to suggest that education may improve health and mortality less in Europe, where health care tends to be universal and economic inequality is generally lower. Education also appears to impact voting and political participation less in Europe, where voter registration is required and governments are more active in registering voters. Although it is tempting to speculate about factors that might explain observed differences in estimates across countries, we are far from understanding them. Much can be learned from simply extending successful empirical strategies to a broader range of countries.

In addition to documenting the impacts of education on nonproduction activities, this chapter has also discussed what is known about the underlying mechanisms that produce those impacts. We have argued that education may largely affect crime because it raises human capital and legitimate wage rates. It seems less likely that education reduces crime by altering preferences for risk or impatience. Decompositions of the impacts of education on health suggest that improvements in earnings may explain as much as one-third of the education–health gradient. The fact that education increases knowledge and ability also appears to be important; however, there is little evidence that changes in preferences or psychological traits play important roles. Evidence from the United States on the impacts of education on citizenship and democracy appear to support the views of Aristotle and Lipset that schooling leads to a more informed and engaged citizenry.

What role should government play given the findings of this chapter? Crime reduction is an obvious externality that may justify expenditures on policies that improve the skills of the most disadvantaged (e.g., targeted preschool programs, improvements in school quality in low-income areas, or policies that encourage high school completion). Current evidence suggests that well-targeted education-based programs can be more cost-effective than traditional law enforcement policies once all costs and benefits

are accounted for. Education policies targeted to the most disadvantaged have the added benefit of reducing economic inequality. There is little evidence of important education externalities in the health domain (most gains are private or, at least, contained within the family), so arguments for education interventions based on health gains are likely to be based on equity and social justice or on the argument that individuals are unaware of some health benefits. It is important to acknowledge that equity-based proposals that focus on educating the most disadvantaged are likely to be inefficient if individuals are aware of the full costs and benefits of education. Indeed, some individuals may benefit more (in terms of health and in other ways) from transfers of income than from additional schooling.⁹⁸ However, if individuals are unaware of some health benefits and unknown benefits are achieved without cost, then education-based policies may be justified. Unfortunately, we still do not know the extent to which most health benefits are effectively paid for through higher health care expenditures or costly changes in behavior. These costs must be factored into policy decisions. Finally, it is clear that increases in political participation will affect the democratic process; however, it is difficult to know exactly how and even more difficult to put a value on this.

APPENDIX: ESTIMATING MEDIATING EFFECTS

Researchers are often interested in the effect of education on health (or other outcomes) that operates through or is mediated by some particular variable. This appendix considers this problem in a simple setting. See Alwin and Hauser (1975) for an early treatment of this type of problem. Also, see Hauser and Goldberger (1971) and Goldberger (1972) for discussions of identification and estimation in more general environments.

Consider the following model for health outcome y as a function of education E, and mediating variables x_1 and x_2 :

$$\gamma = E\alpha + x_1\beta_1 + x_2\beta_2 + \varepsilon \tag{2.14}$$

$$x_1 = E\gamma_E + x_2\gamma_2 + \eta_1$$
 (2.15)

$$x_2 = E\delta_E + x_1\delta_1 + \eta_2. (2.16)$$

To focus on the appropriate way to attribute effects of education through a mediating variable, we abstract from concerns about endogeneity, assuming $\varepsilon \perp (E, x_1, x_2)$. The direct, or partial, effect of E on γ holding x_1 and x_2 constant is given by α in Eq. (2.14). We assume $0 \le \gamma_2 \delta_1 < 1$, so the system is well-behaved.

⁹⁸ Of course, income transfers typically create disincentives for work, so education interventions may provide a less inefficient way of improving equity regardless of its health benefits.

By solving for x_1 and x_2 as functions of E, η_1 , and η_2 , then substituting these into Eq. (2.14), we can write γ as a function of E (and error terms) only:

$$\gamma = E\left[\alpha + \frac{(\gamma_E + \gamma_2 \delta_E)\beta_1 + (\delta_E + \delta_1 \gamma_E)\beta_2}{1 - \delta_1 \gamma_2}\right] + \left[\frac{\beta_1 + \beta_2 \delta_1}{1 - \delta_1 \gamma_2}\right]\eta_1 + \left[\frac{\beta_2 + \beta_1 \gamma_2}{1 - \delta_1 \gamma_2}\right]\eta_2 + \varepsilon. \quad (2.17)$$

This defines the "total effect" of *E* on γ : $a = \alpha + \frac{(\gamma_E + \gamma_2 \delta_E)\beta_1 + (\delta_E + \delta_1 \gamma_E)\beta_2}{1 - \delta_1 \gamma_2}$.

Suppose we are interested in the effect of education on health that operates through or is mediated by x_1 . In the model defined by Eqs. (2.14–2.16), this is given by

$$ME_{1} = \gamma_{E}[\beta_{1} + \delta_{1}\beta_{2} + \delta_{1}\gamma_{2}\beta_{1} + \delta_{1}^{2}\gamma_{2}\beta_{2} + \delta_{1}^{2}\gamma_{2}^{2}\beta_{1} + ...]$$

= $\frac{\gamma_{E}(\beta_{1} + \delta_{1}\beta_{2})}{1 - \delta_{1}\gamma_{2}}.$ (2.18)

A common decomposition approach for estimating $M E_1$ is to regress γ on E and x_2 ; then regress γ on E, x_1 , and x_2 (i.e., Eq. (2.14)), subtracting the latter estimated coefficient on E from the former. Estimation of Eq. (2.14) yields $\hat{\alpha} \xrightarrow{P} \alpha$ given the assumption $\varepsilon \perp (E, x_1, x_2)$. Further assuming $E(\eta_1 \mid E, x_2) = 0$ and estimating $\gamma = Ea_2 + x_2b_2 + \nu_2$ yields $\hat{a}_2 \xrightarrow{P} \alpha + \gamma_E \beta_1$. The estimate of the mediating effect of x_1 is then:

$$\widetilde{ME}_1 = \hat{a}_2 - \hat{\alpha} \xrightarrow{p} \gamma_E \beta_1 = ME_1 - \frac{\gamma_E \delta_1(\gamma_2 \beta_1 + \beta_2)}{1 - \delta_1 \gamma_2}$$

This strategy obtains consistent estimates for the effect mediated through x_1 , $M E_1$, if any of three special cases hold: (1) no effect of x_1 on $x_2(\delta_1 = 0)$; (2) no direct effect of E on x_1 ($\gamma_E = 0$); or (3) no effect of either x_1 or x_2 on outcome γ ($\beta_1 = \beta_2 = 0$). In general, this strategy leads to bias, because including x_2 in the specifications eliminates the mediating role x_1 plays through its impact on x_2 . If all parameters are positive, this produces a downward biased estimate of the mediating effect of x_1 .⁹⁹ Also, notice that even if $\beta_2 = 0$ (so x_2 does not directly affect outcome γ), this estimator may be biased if x_1 and x_2 affect each other.

A second common decomposition used to estimate ME_1 regresses γ on E; then regresses γ on E and x_1 and subtracts the latter estimated coefficient on E from the former. That is, estimate $\gamma = Ea + \nu$ and $\gamma = Ea_1 + x_1b_1 + \nu_1$, then subtract the estimate of a_1 from the estimate of a. This strategy is often used (implicitly or explicitly) when x_2 is unavailable in the data, in which case the first estimation strategy is infeasible. As long as $E(\eta_1 | E) = E(\eta_2 | E) = 0$ (and $\varepsilon \perp E$ as assumed throughout), it is clear that $\hat{a} \stackrel{p}{\rightarrow} a$.

⁹⁹ This strategy does, however, identify the effect of increasing E through x_1 if any effects on x_2 could somehow be offset.

If $E(\eta_2 | E, x_1) = 0$, then $\hat{a}_1 \xrightarrow{p} \alpha + \beta_2 \delta_E$. This approach yields the following estimator for the mediating effect of x_1 :

$$\widehat{ME}_1 = \hat{a} - \hat{a}_1 \xrightarrow{p} \frac{(\gamma_E + \gamma_2 \delta_E)(\beta_1 + \delta_1 \beta_2)}{1 - \delta_1 \gamma_2} = ME_1 + \frac{\gamma_2 \delta_E(\beta_1 + \delta_1 \beta_2)}{1 - \delta_1 \gamma_2}$$

This strategy obtains consistent estimates for the effect mediated through x_1 , ME_1 , if any of three special cases hold: (1) no effect of x_2 on x_1 ($\gamma_2 = 0$); (2) no direct effect of E on x_2 ($\delta_E = 0$); or (3) no effect of either x_1 or x_2 on outcome γ ($\beta_1 = \beta_2 = 0$). It is not enough for x_2 to have no direct effect on γ to obtain consistent estimates of ME_1 . In this case, the bias $\frac{\gamma_2 \delta_E \beta_1}{1 - \delta_1 \gamma_2}$ depends on the effect of x_1 on x_2 , x_2 on x_1 , and the direct effect of education on x_2 .

The decision regarding which estimator to use depends on the relationship between the x variables, the direct effects of each x on the outcome y, and the effects of education on each x. Notice that $\widehat{ME}_1 = \widetilde{ME}_1 = ME_1$ if $\beta_1 = \beta_2 = 0$, so both estimation strategies will identify the mediating effect of x_1 in this case. Of course, this is neither surprising nor interesting. If we assume that both x variables affect the outcome of interest, the decision of which estimator to use depends on the relationship between education and both x's and on the interdependence of the two x's on each other. Assuming both x variables affect each other, then the first estimator, \widehat{ME}_1 , yields a consistent estimate of the mediating effect of x_1 if education does not directly affect x_2 while the second estimator, \widehat{ME}_1 , yields a consistent estimate if education does not directly affect x_1 . If we instead assume that education directly affects both x variables, then the choice of estimators depends on whether we can rule out an effect of one x variable on the other. If x_1 does not affect x_2 , then the first estimator will yield a consistent estimate of ME_1 , while the second estimator is consistent if x_2 does not affect x_1 .

These findings imply the following general lessons when estimating the mediating effect of some variable using the standard decomposition approach: On the one hand, failure to control for other variables can bias estimates of the mediating effect when (1) education directly affects those variables and (2) those variables directly impact the "mediating variable" of interest. On the other hand, controlling for other variables can also bias estimates of the mediating effect when (1) the mediating variable directly affects these other variables and (2) education directly affects those variables.

If x_2 is available in the data, it is possible under some conditions to use the two estimators to bound the mediating effect of x_1 . For example, suppose x_1 and x_2 are normalized such that they both have positive effects on γ (i.e., β_1 and β_2 are positive).¹⁰⁰ If x_1 and x_2 have positive cross-effects (i.e., $\gamma_2 \ge 0$ and $\delta_1 \ge 0$) and if education increases both x_1 and x_2 (i.e., $\gamma_E \ge 0$ and $\delta_E \ge 0$), then the two estimation strategies bound the

¹⁰⁰ Note that this can be verified from estimation of Eq. (2.14).

mediating effect of $x_1: \widetilde{ME}_1 \leq ME_1 \leq \widetilde{ME}_1$. Alternatively, under the exogeneity assumptions assumed throughout this Appendix, one could simply construct a consistent estimate of the mediating effect of x_1 using Eq. (2.18) and parameter estimates from Eqs. (2.14–2.16).

ACKNOWLEDGMENTS

For comments on different sections of this chapter, I thank Gary Becker, David Card, Phil Cook, David Deming, Darius Lakdawalla, Fabian Lange, Jens Ludwig, and participants at the Conference on the Handbook of the Economics of Education held in Munich, Germany, and the NBER Economics of Crime Control Conferences in Boston, MA, and Berkeley, CA.

REFERENCES

- Acemoglu, D., Angrist, J., 2001. How Large are Human Capital Externalities? Evidence from Compulsory Schooling Laws. NBER Macroeconomics Annual 2000 15, 9–74.
- Acemoglu, D., Johnson, S., Robinson, J.A., 2005. Institutions as a Fundamental Cause of Long-Run Growth. In: Aghion, P., Durlauf, S.N. (Eds.), Handbook of Economic Growth, vol. 1. Elsevier, Amsterdam, pp. 385–472, Part 1, chapter 6.
- Acemoglu, D., Johnson, S., Robinson, J.A., Yared, P., 2005. From Education to Democracy? Am. Econ. Rev. 95 (2), 44–49.
- Adams, S.J., 2002. Educational Attainment and Health: Evidence from a Sample of Older Adults. Educ. Econ 10 (1), 97–109.
- Aizer, A., Stroud, L., 2010. Education, Knowledge and the Evolution of Disparities in Health. NBER Working Paper 15840.
- Albouy, V., Lequien, L., 2009. Does Compulsory Education Lower Mortality? J. Health Econ. 28 (1), 155–168.
- Alesina, A., Rodrik, D., 1994. Distributive Politics and Economic Growth. Q. J. Econ. 109 (2), 465-490.
- Alvarez, M., Cheibub, J.A., Limongi, F., Przeworski, A. 2000. Democracy and Development: Political Institutions and Material Well-Being in the World, 1950–1990. Cambridge University Press, Cambridge.
- Alwin, D.F., Hauser, R.M., 1975. The Decomposition of Effects in Path Analysis. Am. Sociol. Rev. 40 (1), 37–47.
- Anderson, M., 2009. In: School and Out of Trouble? The Minimum Dropout Age and Juvenile Crime Working Paper.
- Angrist, J.D., Imbens, G.W., 1994. Identification and Estimation of Local Average Treatment Effects. Econometrica 62 (2), 467–475.
- Angrist, J.D., Imbens, G.W., 1995. Two-Stage Least Squares Estimation of Average Causal Effects in Models with Variable Treatment Intensity. J. Am. Stat. Assoc. 90 (430), 431–442.
- Arellano, M., Bond, S., 1991. Some Tests of Specificatino for Panel Data: Monte Carlo Evidence and an Application to Employment Equations. Rev. Econ. Stud. 58 (2), 277–297.
- Arendt, J.N., 2005. Does Education cause Better Health? A Panel Data Analysis Using School Reforms for Identification. Econ. Educ. Rev. 24 (2), 149–160.
- Arkes, J., 2003. Does Schooling Improve Adult Health? RAND DRU-3051.
- Auld, M.C., Sidhu, N., 2005. Schooling, Cognitive Ability and Health. Health Econ. 14 (10), 1019–1034.
- Barnett, W.S., Masse, L.N., 2007. Comparative Benefit-Cost Analysis of the Abecedarian Program and its Policy Implications. Econ. Educ. Rev. 26, 113–125.

Barro, R.J., 1997. Determinants of Economic Growth: A Cross-Country Empirical Study. The MIT Press, Cambridge, Mass.

Barro, R.J., 1999. Determinants of Democracy. J. Polit. Econ. 107 (6), S158-S183.

- Becker, G.S., 1967. Human Capital and the Personal Distribution of Income: An Analitical Approach. Woytinsky Lecture 1, University of Michigan, Institute of Public Administration, Ann Arbor.
- Becker, G., Mulligan, C., 1997. The Endogenous Determination of Time Preference. Q. J. Econ. 112, 729–758.
- Belfield, C.R., Nores, M., Barnett, W.S., Schweinhart, L., 2006. The High/Scope Perry Preschool Program: Cost-Benefit Analysis Using Data from the Age-40 Followup. J. Hum. Resour. 41, 162–190.
- Ben-Porath, Y., 1967. The Production of Human Capital and the Life Cycle of Earnings. J. Polit. Econ. 75 (4), 352–365.
- Besley, T., Coate, S., 1998. Sources of Inefficiency in a Representative Democracy: A Dynamic Analysis. Am. Econ. Rev. 88 (1), 139–156.
- Blau, D., Currie, J., 2006. Pre-School, Day Care, and After-School Care: Who's Minding the Kids. In: Hanushek, E., Welch, F. (Eds.), Handbook of the Economics of Education, vol. 2. Elsevier Science BV, Amsterdam, pp. 1163–1278, chapter 20.
- Blundell, R., Bond, S., 1998. Initial Conditions and Moment Restrictions in Dynamic Panel Data Models. J. Econom. 87 (1), 115–143.
- Bobba, M., Coviello, D., 2007. Weak Instruments and Weak Identification, in Estimating the Effects of Education, on Democracy. Econ. Lett. 96 (3), 301–306.
- Boisjoli, R., Vitaro, F., Lacourse, E., Barker, E.D., Tremblay, R.E., 2007. Impact and Clinical Significance of a Preventive Intervention for Disruptive Boys. Br. J. Psychiatry 191, 415–419.
- Bound, J., Jaeger, D., Baker, R., 1995. Problems with Instrumental Variables Estimation When the Correlation Between the Instruments and the Endogenous Explanatory Variable is Weak. J. Am. Stat. Assoc. 90 (430), 443–450.
- Bourguignon, F., Verdier, T., 2000. Oligarchy, Democracy, Inequality and Growth. J. Dev. Econ. 62 (2), 285–313.
- Braga, A., 2005. Hot Spots Policing and Crime Prevention: A Systematic Review of Randomized Controlled Trials. J. Exp. Criminol. 1, 317–342.
- Brunello, G., Fabbri, D., Fort, M., 2009. Years of Schooling, Human Capital and the Body Mass Index of European Females. IZA Discussion Papers 4667, Institute for the Study of Labor (IZA).
- Buchanan, J.M., Tullock, G., 1962. The Calculus of Consent: Logical Foundations of Constitutional Democracy. University of Michigan Press, Ann Arbor, MI.
- Buonanno, P., Leonida, L., 2006. Education and Crime: Evidence from Italian Regions. Appl. Econ. Lett. 13, 709–713.
- Cameron, S.V., Heckman, J.J., 1993. The Nonequivalence of High School Equivalents. J. Labor Econ. 11 (1, part 1), 1–47.
- Cameron, S.V., Taber, C., 2004. Estimation of Educational Borrowing Constraints Using Returns to Schooling. J. Polit. Econ. 112 (1), 132–182.
- Campbell, F.A., Ramey, C.T., Pungello, E., Sparling, J., Miller-Johnson, S., 2002. Early Childhood Education: Young Adult Outcomes from the Abecedarian Project. Appl. Dev. Sci. 6, 42–57.
- Card, D., 1999. The Causal Effect of Education on Earnings. In: Ashenfelter, O., Card, D. (Eds.), Handbook of Labor Economics, vol. 3. Elsevier Science BV, Amsterdam, pp. 1801–1863, (Chapter 30).
- Carneiro, P., Ginja, R., 2008. Preventing Behavior Problems in Childhood and Adolescence: Evidence from Head Start. Working Paper.
- Carneiro, P., Heckman, J.J., 2002. The Evidence on Credit Constraints in Post-secondary Schooling. Econ. J. 112 (482), 705–734.
- Castelló-Climent, A., 2008. On the Distribution of Education and Democracy. J. Dev. Econ. 87 (2), 179–190.
- Cave, G., Bos, H., Doolittle, F., Toussaint, C., 1993. Jobstart: Final Report on a Program for School Dropouts. Manpower Development Research Corporation, New York.
- Cawley, J., Heckman, J., Lochner, L., Vytlacil, E., 2000. Understanding the Role of Cognitive Ability in Accounting for the Recent Rise in the Economic Return to Education. In: Arrow, K., Bowles, S.,

Durlauf, S. (Eds.), Meritocracy and Economic Inequality. Princeton University Press, New Jersey, pp. 230–265.

- Chay, K., Dobkin, C., Greenstone, M., 2003. The Clean Air Act of 1970 and Adult Mortality. J. Risk Uncertain. 27 (3), 279–300.
- Chay, K.Y., Greenstone, M., 2003. The Impact of Air Pollution on Infant Mortality: Evidence from Geographic Variation in Pollution Shocks Induced by a Recession. Q. J. Econom. 118 (3), 1121–1167.
- Chay, K.Y., Greenstone, M., 2005. Does Air Quality Matter? Evidence from the Housing Market. J. Polit. Econ. 113 (2), 376–424.
- Chou, S.-Y., Liu, J.-T., Grossman, M., Joyce, T., 2010. Parental Education and Child Health: Evidence from a Natural Experiment in Taiwan. Am. Econ. J.: Appl. Econom. 2 (1), 33–61.
- Ciccone, A., Peri, G., 2006. Identifying Human-Capital Externalities: Theory with Applications. Rev. Econ. Stud. 73 (2), 381–412.
- Clark, D., Royer, H., 2010. The Effect of Education on Adult Health and Mortality: Evidence from Britain. Working Paper.
- Clarke, S.H., Campbell, F.A., 1998. Can Intervention Early Prevent Crime Later? The Abecedarian Project Compared with Other Programs. Early Child. Res. Q. 13, 319–343.
- Cohen-Cole, E., Fletcher, J.M., 2008. Is Obesity Contagious? Social Networks vs. Environmental Factors in the Obesity Epidemic. J. Health Econ. 27 (5), 1382–1387.
- Conduct Problems Prevention Research Group, 2007. Fast Track Randomized Controlled Trial to Prevent Externalizing Psychiatric Disorders: Findings From Grades 3 to 9. J. Am. Acad. Child Adolesc. Psychiatry 46, 1250–1262.
- Conduct Problems Prevention Research Group, 2010. Fast Track Intervention Effects on Youth Arrests and Delinquency. J. Exp. Criminol. 6 (2), 131–157.
- Conduct Problems Prevention Research Group, 2011. The Effects of the Fast Track Preventive Intervention on the Development of Conduct Disorder across Childhood. Child Dev. 82 (1), 331–345.
- Conti, G., Heckman, J., Urzua, S., 2010a. The Education-Health Gradient. Am. Econ. Rev. 100 (2), 234–238.
- Conti, G., Heckman, J.J., Urzúa, S., 2010b. Understanding the Early Origins of the Education-Health Gradient: A Framework that can also be Applied to Analyze Gene-Environment Interactions. Working Paper June.
- Cullen, J.B., Jacob, B.A., Levitt, S., 2006. The Effect of School Choice on Participants: Evidence from Randomized Lotteries. Econometrica 74, 1191–1230.
- Cunha, F., Heckman, J.J., 2008. Formulating, Identifying and Estimating the Technology of Cognitive and Noncognitive Skill Formation. J. Hum. Resour. 43, 738–782.
- Currie, J., Moretti, E., 2003. Mother's Education and the Intergenerational Transmission of Human Capital: Evidence from College Openings. Q. J. Econ. 118 (4), 1495–1532.
- Currie, J., Neidell, M., 2005. Air Pollution and Infant Health: What can we Learn from California's Recent Experience? Q. J. Econ. 120 (3), 1003–1030.
- Cutler, D., Lange, F., Meara, E., Richards, S., Ruhm, C., 2010. Explaining the Rise in Educational Gradients in Mortality. Working Paper 15678, National Bureau of Economic Research, Cambridge, MA.
- Cutler, D.M., Lleras-Muney, A., 2006. Education and Health: Evaluating Theories and Evidence. Working Paper 12352, National Bureau of Economic Research.
- Cutler, D.M., Lleras-Muney, A., 2010. Understanding Differences in Health Behaviors by Education. J. Health Econ. 29 (1), 1–28.
- de Tocqueville, A., 1835/1998. Democracy in America. Wordsworth, Hertfordshire (originally published in 1835)
- de Walque, D., 2007a. Does Education Affect Smoking Behaviors? Evidence using the Vietnam draft as an Instrument for College Education. J. Health Econ. 26, 877–895.
- de Walque, D., 2007b. How does the Impact of an HIV/AIDS Information Campaign Vary with Educational Attainment? Evidence from Rural Uganda. J. Dev. Econ. 84 (2), 686–714.
- de Walque, D., 2010. Education, Information and Smoking Decisions: Evidence from Smoking Histories in the United States, 1940–2000. J. Hum. Resour. 45 (3), 682–716.

Dee, T.S., 2004. Are there Civic Returns to Education? J. Public Econ. 88 (9-10), 1697-1720.

Deming, D., 2009a. Better Schools, Less Crime? Working Paper.

- Deming, D., 2009b. Early Childhood Intervention and Life-Cycle Skill Development: Evidence from Head Start. Am. Econ. J.: Appl. Econ. 1 (3), 111–134.
- Donohue III, J.J., Siegelman, P., 1998. Allocating Resources Among Prisons and Social Programs in the Battle Against Crime. J. Leg. Stud. 27, 1–44.
- Doucouliagos, H., Ulubasoglu, M.A., 2008. Democracy and Economic Growth: A Meta-Analysis. Am. J. Polit. Sci. 52 (1), 61–83.
- Durose, M.R., Langan, P.A., 2003. Felony Sentences in State Courts, 2000. United States Department of Justice, Bulletin, NCJ 198821, Bureau of Justice Statistics, Washington, DC.
- Ehrlich, I., 1975. On the Relation between Education and Crime. In: Thomas Juster, F. (Ed.), Education, Income, and Human Behavior. McGraw-Hill Book Co., New York, pp. 313–338.
- Ehrlich, I., Chuma, H., 1990. A Model of the Demand for Longevity and the Value of Life Extension. J. Polit. Econ. 98 (4), 761–782.
- Elo, I.T., Preston, S.H., 1996. Educational Differentials in Mortality: United States, 1979-85. Social Sci. Med. 42 (1), 47–57.
- Erbsland, M., Ried, W., Ulrich, V., 1995. Health, Health Care, and the Environment: Econometric Evidence from German Micro Data. Health Econ. 4 (3), 169–182.
- Farrington, D.P., Gallagher, B., Morley, L., St. Ledger, R.J., West, D.J., 1986. Unemployment, School Leaving and Crime. Br. J. Criminol. 26, 335–356.
- Fella, G., Gallipoli, G., 2009. Education and Crime over the Life Cycle. Working Paper, University of British Columbia.
- Fletcher, J.M., 2010. Social Interactions and Smoking: Evidence using Multiple Student Cohorts, Instrumental Variables, and School Fixed Effects. Health Econ. 19 (4), 466–484.
- Fuchs, V.R., 1982. Time Preference and Health: An Exploratory Study. In: Fuchs, V.R. (Ed.), Economic Aspects of Health. University of Chicago Press, Chicago, pp. 93–120 chapter 3.
- Garces, E., Thomas, D., Currie, J., 2002. Longer-Term Effects of Head Start. Am. Econ. Rev. 92, 999-1012.
- Gaviria, A., Raphael, S., 2001. School-Based Peer Effects and Juvenile Behavior. Rev. Econ. Stat. 83 (2), 257–268.
- Gilleskie, D.B., Harrison, A.L., 1998. The Effect of Endogenous Health Inputs on the Relationship Between Health and Education. Econ. Educ. Rev. 17 (3), 279–295.
- Glaeser, E., Ponzetto, G., Shleifer, A., 2007. Why does Democracy Need Education? J. Econ. Growth 12, 77–99.
- Glaeser, E., Sacerdote, B., Scheinkman, J., 1996. Crime and Social Interactions. Q. J. Econ. 111 (2), 507–548.
- Glaeser, E.L., La Porta, R., Lopez-de Silanes, F., Shleifer, A., 2004. Do Institutions Cause Growth? J. Econ. Growth 9, 271–303.
- Glied, S., Lleras-Muney, A., 2008. Technological Innovation and Inequality in Health. Demography 45 (3), 741–761.
- Goesling, B., 2007. The Rising Significance of Education for Health? Soc. Forces 85 (4), 1621–1644.
- Goldberger, A.S., 1972. Structural Equation Methods in the Social Sciences. Econometrica 40 (6), 979–1001.
- Goldman, D.P., Lakdawalla, D.N., 2005. A Theory of Health Disparities and Medical Technology. Contrib. Econ. Analysis Policy 4 (1), 1–32.
- Goldman, D.P., Smith, J.P., 2002. Can Patient Self-Management Help Explain the SES Health Gradient? Proc. Natl. Acad. Sci. USA 99 (16), 10929–10934.
- Goldman, D.P., Smith, J.P., 2005. Socioeconomic Differences in the Adoption of New Medical Technologies. Am. Econ. Rev. 95 (2), 234–237.
- Gottfredson, D., 1985. Youth Employment, Crime, and Schooling. Dev. Psychol. 21, 419-432.
- Gould, E., Mustard, D., Weinberg, B., 2002. Crime Rates and Local Labor Market Opportunities in the United States: 1977-1997. Rev. Econ. Stat. 84, 45–61.
- Grabner, M.J., 2009. The Impact of Education on Health Using Compulsory Schooling Laws. Working Paper.

- Grenet, J., 2010. Is it Enough to Increase Compulsory Education to Raise Earnings? Evidence from French and Brisish Compulsory Schooling Laws. Working Paper.
- Grimard, F., Parent, D., 2007. Education and Smoking: Were Vietnam War Draft Avoiders also More Likely to Avoid Smoking? J. Health Econ. 26, 896–926.
- Grogger, J., 1995. The Effect of Arrests on the Employment and Earnings of Young Men. Q. J. Econ. 110, 51–71.
- Grogger, J., 1998. Market Wages and Youth Crime. J. Labor Econ. 16, 756-791.
- Grossman, M., 1972a. On the Concept of Health Capital and the Demand for Health. J. Polit. Econ. 80 (2), 223–255.
- Grossman, M., 1972b. The Demand for Health: A Theoretical and Empirical Investigation. Columbia University Press, New York.
- Grossman, M., 1975. The Correlation Between Health and Schooling. In: Terleckyj, N.E. (Ed.), Household Production and Consumption. Columbia University Press for the National Bureau of Economic Research, New York, pp. 147–211.
- Grossman, M., 2000. The Human Capital Model. In: Culyer, A.J., Newhouse, J.P. (Eds.), Handbook of Health Economics, vol. 1. Elsevier BV, Amsterdam, pp. 347–408, Part 1, chapter 7.
- Grossman, M., 2006. Education and Nonmarket Outcomes. In: Hanushek, E., Welch, F., (Eds.), Handbook of the Economics of Education. , vol. 1. Elsevier BV, Amsterdam, pp. 577–633, chapter 10.
- Grossman, M., Kaestner, R., 1997. Effects of Education on Health. In: Behrman, J.R., Stacey, N. (Eds.), The Social Benefits of Education. University of Michigan Press, Ann Arbor, pp. 69–123 chapter 4.
- Guryan, J., 2004. Desegregation and Black Dropout Rates. Am. Econ. Rev. 94 (4), 919-943.
- Hanushek, E.A., Lindseth, A.A., 2009. Schoolhouses, Courthouses, and Statehouses: Solving the Funding-Achievement Puzzle in America's Public Schools. Princeton University Press, Princeton.
- Harlow, C.W., 2003. Education and Correctional Populations. U.S. Dept. of Justice, Bureau of Justice Statistics.
- Hauser, R.M., Goldberger, A.S., 1971. The Treatment of Unobservable Variables in Path Analysis. Sociol. Methodol. 3, 81–117.
- Hawkins, J.D., Catalano, R.F., Kosterman, R., Abbott, R., Hill, K.G., 2005. Promoting Positive Adult Functioning Through Social Development Intervention in Childhood: Long Term Effects from the Seattle Social Development Project. Arch. Pediatr. Adolesc. Med. 159, 25–31.
- Hayek, F.A., 1960. The Constitution of Liberty. University of Chicago Press, Chicago.
- Heckman, J.J., Klenow, P., 1997. Human Capital Policy. Working Paper.
- Heckman, J.J., LaFontaine, P., 2006. Bias-Corrected Estimates of GED Returns. J. Labor Econ. 24 (3), 661–700.
- Heckman, J.J., Lochner, L., Todd, P., 2006. Earnings Functions, Rates of Return, and Treatment Effects: The Mincer Equation and Beyond. In: Hanushek, E., Welch, F. (Eds.), Handbook of the Economics of Education, vol. 1. Elsevier BV, Amsterdam, pp. 307–458, chapter 7.
- Heckman, J.J., Lochner, L., Todd, P., 2008. Earnings Functions and Rates of Return. J. Hum. Capital 2 (1), 1–31.
- Heckman, J.J., Moon, S.H., Pinto, R., Savelyev, P., Yavitz, A., 2009. The Rate of Return to the High/ Scope Perry Preschool Program. NBER Working Paper No. 15471.
- Heckman, J.J., Rubinstein, Y., 2001. The Importance of Noncognitive Skills: Lessons from the GED Testing Program. Am. Econ. Rev. 91 (2), 145–159.
- Heckman, J.J., Stixrud, J., Urzua, S., 2006. The Effects of Cognitive and Noncognitive Abilities on Labor Market Outcomes and Social Behavior. J. Labor Econ. 24 (3), 411–482.
- Hindelang, M.J., Hirschi, T., Weis, J.G., 1981. Measuring Delinquency. Sage Publications, Inc, Beverly Hills.
- Hjalmarsson, R., 2008. Criminal Justice Involvement and High School Completion. J. Urban Econ. 63 (2), 613–630.
- Huntington, S.P., 1968. Political Order in Changing Societies. Yale University Press, New Haven.
- Jacob, B., Lefgren, L., 2003. Are Idle Hands the Devil's Workshop? Incapacitation, Concentration, and Juvenile Crime. Am. Econ. Rev. 93, 1560–1577.
- Jacob, B., Ludwig, J., 2008. Improving educational outcomes for poor children. Working Paper 14550. National Bureau of Economic Research, Cambridge, MA.
- Jaggers, K., Marshall, M.G., 2000. Polity iv Project. Technical report, Center for International Development and conflict Management, University of Maryland.
- Jürges, H., Kruk, E., Reinhold, S., 2009. The Effect of Compulsory Schooling on Health-Evidence from Biomarkers. Working paper, Mannheim Research Institute for the Economics of Aging (MEA), University of Mannheim.
- Kandel, E., et al., 1988. IQ as a Protective Factor for Subjects at High Risk for Antisocial Behavior. J. Consult. Clin. Psychol. 56, 224–226.
- Karoly, L.A., et al. 1998. Investing in Our Children: What We Know and Don't Know about the Costs and Benefits of Early Childhood Interventions. RAND.
- Kemptner, D., Jürges, H., Reinhold, S., 2010. Changes in Compulsory Schooling and the Causal Effect of Education on Health: Evidence from Germany. Working paper 10200, Mannheim Research Institute for the Economics of Aging (MEA), University of Mannheim.
- Kenkel, D., Lillard, D., Mathios, A., 2006. The Roles of High School Completion and GED Receipt in Smoking and Obesity. J. Labor Econ. 24 (3), 635–660.
- Kenkel, D.S., 1991. Health Behavior, Health Knowledge, and Schooling. J. Polit. Econ. 99 (2), 287-305.
- Kitagawa, E.M., Hauser, P.M., 1973. Differential Mortality in the United States. Harvard University Press, Cambridge.
- Kling, J., 2006. Incarceration Length, Employment and Earnings. Am. Econ. Rev. 96 (3), 863-876.
- Kling, J., Ludwig, J., Katz, L., 2005. Neighborhood Effects on Crime for Female and Male Youth: Evidence from a Randomized Housing Voucher Experiment. Q. J. Econ. 120 (1), 87–130.
- Krauth, B., 2005. Structural Estimation of Peer Effects in Youth Smoking. Adv. Health Econ. Health Serv. Res. 16, 201–211.
- La Porta, R., Lopez-de Silanes, F., Shleifer, A., Vishny, R., 1999. The Quality of Government. J. Law Econ. Organiz. 15 (1), 222–279.
- Lally, J.R., Mangione, P., Honig, A., 1988. The Syracuse University Family Development Research Program: Long-Range Impact of an Early Intervention with Low-Income Children and their Families. In: Powell, D., (Ed.), Parent Education as Early Childhood Intervention: Emerging Directions in Theory, Research, and Practice. Ablex, Norwood, pp. 79–104.
- Lange, F., (forthcoming). The Role of Education in Complex Health Decisions: Evidence from Cancer Screening. J. Health Econ.
- Lange, F., Topel, R., 2006. The Social Value of Education and Human Capital. In: Hanushek, E., Welch, F., (Eds.), Handbook of the Economics of Education, vol. 1. Elsevier, Amsterdam, pp. 459–509, chapter 8.
- Leigh, J.P., 1983. Direct and Indirect Effects of Education on Health. Soc. Sci. Med. 17 (4), 227-234.
- Lindeboom, M., Llena-Nozal, A., van der Klaauw, B., 2009. Parental Education and Child Health: Evidence from a Schooling Reform. J. Health Econ. 28 (1), 109–131.
- Link, B., Phelan, J., 1995. Social conditions as Fundamental Causes of Disease. J. Health Soc. Behav. 35 (Extra Issue), 80–94.
- Lipset, S., 1959. Some Social Requisites for Democracy: Economic Development and Political Legitimacy. Am. Polit. Sci. Rev. 53 (1), 69–105.
- Lleras-Muney, A., 2002. Were Compulsory Attendance and Child Labor Laws Effective? An Analysis from 1915 to 1939. J. Law Econ. 45 (2), 2002.
- Lleras-Muney, A., 2005. The Relationship between Education and Adult Mortality in the United States. Rev. Econ. Stud. 72 (1), 189–221.
- Lleras-Muney, A., 2006. Erratum: The Relationship Between Education and Adult Mortality in the United States. Rev. Econ. Stud. 73 (3), 847.
- Lleras-Muney, A., 2010. The Needs of the Army: Using Compulsory Relocation in the Military to Estimate the Effect of Air Pollutants on Children's Health. J. Hum. Resour. 45, 549–590.
- Lochner, L., 2004, Education, Work, and Crime: A Human Capital Approach. Int. Econ. Rev. 45 (3), 811-843.
- Lochner, L., Moretti, E., 2004. The Effect of Education on Crime: Evidence from Prison Inmates, Arrests, and Self-Reports. Am. Econ. Rev. 94 (1), 155–189.
- Long, D., Mallar, C., Thornton, C., 1981. Evaluating the Benefits and Costs of the Job Corps. J. Policy Anal. Manage. 1, 55–76.

- Luallen, J., 2006. School's Out...Forever: A Study of Juvenile Crime, At-Risk Youths and Teacher Strikes. J. Urban Econ. 59, 75–103.
- Ludwig, J., Miller, D.L., 2007. Does Head Start Improve Children's Life Chances? Evidence from a Regression Discontinuity Design. Q. J. Econ. 122, 159–208.
- Machin, S., Marie, O., Vujic, S., (2011). The Crime Reducing Effect of Education. Econ. J. 121 (552), 463–484.
- Machin, S., Meghir, C., 2004. Crime and Economic Incentives. J. Hum. Resour. 39, 958-979.
- Maitra, S., 2010. Can Patient Self-Management Explain the Health Gradient? Goldman and Smith's "Patient Self-Management Help Explain the SES Health Gradient?" (2002) revisited. Soc. Sci. Med. 70 (6), 802–812.
- Marmot, M., 2004. Status Syndrome: How your Social Standing Directly Affects your Health and Life Expectancy. Bloomsbury, London.
- Marmot, M., Wilkinson, R.G., 2006. Social Determinants of Health, second ed. Oxford University Press, Oxford.
- Marmot, M.G., 1994. Social Differentials in Health within and Between Populations. Daedalus 123 (4), 197–216.
- Mazumder, B., 2008. Does Education Improve Health? A Reexamination of the Evidence from Compulsory Schooling Laws. Econ. Perspect. 2–16.
- Mazumder, B., 2010. Erratum: Does Education Improve Health? A Reexamination of the Evidence from Compulsory Schooling Laws. November 16.
- McCormick, M.C., et al. 2006. Early Intervention in Low Birth Weight Premature Infants: Results at 18 Years of Age for the Infant Health and Development Program. Pediatrics 117, 771–780.
- McCrary, J., Royer, H., 2009. The Effect of Female Education on Fertility and Infant Health: Evidence from School Entry Policies using Exact Date of Birth. NBER Working Paper No. 12329.
- McMahon, W.W., 2009. Higher Learning, Greater Good: the Private and Social Benefits of Higher Education. Johns Hopkins University Press, Baltimore.
- McWilliams, M., 2009. Health Consequences of Uninsurance among Adults in the United States: Recent Evidence and Implications. Milbank Q. 87 (2), 443–494.
- Meara, E.R., Richards, S., Cutler, D.M., 2008. The Gap Gets Bigger: Changes In Mortality and Life Expectancy, By Education, 19812000. Health Aff. 27 (2), 350–360.
- Merlo, A., Wolpin, K., 2009. The Transition from School to Jail: Youth Crime and High School Completion among Black Males. Penn Institute for Economic Research Working Paper 09–002.
- Miller, T., Cohen, M., Wiersema, B., 1996. Victim Costs and Consequences: A New Look. Final Summary Report to the National Institute of Justice, Feb.
- Milligan, K., Moretti, E., Oreopoulos, P., 2004. Does Education Improve Citizenship? Evidence from the United States and the United Kingdom. J. Public Econ. 88 (9–10), 1667–1695.
- Minier, J.A., 1998. Democracy and Growth: Alternative Approaches. J. Econ. Growth 3, 241–266.
- Moretti, E., 2004a. Estimating the Social Return to Higher Education: Evidence from Longitudinal and Repeated Cross-Sectional Data. J. Econom. 121 (1–2), 175–212 Higher education (Annals issue).
- Moretti, E., 2004b. Workers' Education, Spillovers, and Productivity: Evidence from Plant-Level Production Functions. Am. Econ. Rev. 94 (3), 656–690.
- Moretti, E., Neidell, M., 2009. Pollution, Health, and Avoidance Behavior: Evidence from the Ports of Los Angeles. Working Paper 14939, NBER.
- Murnane, R.J., 2008. Educating Urban Children. Working Paper 13791, National Bureau of Economic Research, Cambridge, MA.
- Mustard, D., 2001. Racial, Ethnic and Gender Disparities in Sentencing: Evidence from the US Federal Courts. J. Law Econ. 44 (1), 285–314.
- Nagin, D., Waldfogel, J., 1995. The Effects of Criminality and Conviction on the Labor Market Status of Young British Offenders. Int. Rev. Law Econ. 15, 109–126.
- Neidell, M., 2004. Air Pollution, Health, and Socio-Economic Status: The Effect of Outdoor Air Quality on Childhood Asthma. J. Health Econ. 23 (6), 1209–1236.
- Nie, N.H., Junn, J., Stehlik-Barry, K., 1996. Education and Democratic Citizenship in America. University of Chicago Press, Chicago.

- North, D.C., 1990. Institutions, Institutional Change, and Economic Performance. Cambridge University Press, Cambridge and London.
- Norton, E.C., Lindrooth, R.C., Ennett, S.T., 1998. Controlling for the Endogeneity of Peer Substance Use on Adolescent Alcohol and Tobacco Use. Health Econ. 7, 439–453.
- Olds, D.L., et al., 1998. Long-Term Effects of Nurse Home Visitation on Children's Criminal and Antisocial Behavior: 15-Year Follow-up of a Randomized Controlled Trial. J. Am. Med. Assoc. 280, 1238–1244.
- Olson, M., 1993. Dictatorship, Democracy, and Development. Am. Polit. Sci. Rev. 87 (3), 567-576.
- Oreopoulos, P., 2006. Estimating Average and Local Average Treatment Effects of Education when Compulsory Schooling Laws Really Matter. Am. Econ. Rev. 96 (1), 152–175.
- Oreopoulos, P., 2008. Estimating Average and Local Average Treatment Effects of Education when Compulsory Schooling Laws Really Matter: Corrigendum. University of British Columbia Working Paper.
- Oreopoulos, P., Salvanes, K.G., 2009. How Large are Returns to Schooling? Hint: Money Isn't Everything. NBER Working Paper No. 15339.
- Papaioannou, E., Siourounis, G., 2008a. Democratisation and Growth. Econ. J. 118 (532), 1520-1551.
- Papaioannou, E., Siourounis, G., 2008b. Economic and Social Factors Driving the Third Wave of Democratization. J. Comp. Econ. 36 (3), 365–387.
- Pappas, G., Queen, S., Hadden, W., Fisher, G., 1993. The Increasing Disparity in Mortality between Socioeconomic Groups in the United States, 1960 and 1986. New England J. Med. 329 (2), 103–109.
- Persson, T., Tabellini, G., 1994. Is Inequality Harmful for Growth? Am. Econ. Rev. 84 (3), 600-621.
- Powdthavee, N., 2010. Does Education Reduce the Risk of Hypertension? Estimating the Biomarker Effect of Compulsory Schooling in England. J. Hum. Capital 4 (2), 173–202.
- Powell, L.M., Tauras, J.A., Ross, H., 2005. The Importance of Peer Effects, Cigarette Prices and Tobacco Control Policies for Youth Smoking Behavior. J. Health Econ. 24 (5), 950–968.
- Price, J., Simon, K., 2009. Patient Education and the Impact of New Medical Research. J. Health Econ. 28 (6), 1166–1174.
- Przeworski, A., Limongi, F., 1993. Political Regimes and Economic Growth. J. Econ. Perspec. 7 (3), 51-69.
- Rajan, R.G., 2009. Rent Preservation and the Persistence of Underdevelopment. Am. Econ. J.: Macroeconom. 1 (1), 178–218.
- Reber, S., 2007. School Desegregation and Educational Attainment for Blacks. NBER Working Paper No. 13193.
- Reinhold, S., Jürges, H., 2010. Secondary School Fees and the Causal Effect of Schooling on Health Behavior. Health Econ. 19 (8), 994–1001.
- Reynolds, A.J., Temple, J.A., Robertson, D.L., Mann, E.A., 2001. Long-term Effects of an Early Childhood Intervention on Education Achievement and Juvenile Arrest: A 15-Year Follow-up of Low-Income Children in Public Schools. J. Am. Med. Assoc. 285, 2339–2346.
- Reynolds, A.J., Temple, J.A., Robertson, D.L., Mann, E.A., 2002. Age 21 Cost-Benefit Analysis of the Title I Chicago Child-Parent Centers. Educ. Eval. Policy Anal. 24 (4), 267–303.
- Rosen, S., 1986. The Theory of Equalizing Differences. In: Ashenfelter, O.C., Layard, R. (Eds.), Handbook of Labor Economics. , vol. 1. Elsevier BV, Amsterdam, pp. 641–692, chapter 12.
- Rosenzweig, M.R., Schultz, T.P., 1982. The Behavior of Mothers as Inputs to Child Health: The Determinants of Birth Weight, Gestation, and Rate of Fetal Growth. In: Fuchs, V. (Ed.), Economic Aspects of Health. University of Chicago Press, Chicago, pp. 53–92, chapter 2.
- Ross, C.E., Wu, C.-L., 1995. The Links Between Education and Health. Am. Sociol. Rev. 60 (5), 719-745.
- Sabates, R., Feinstein, L., 2008. Effects of Government Initiatives on Youth Crime. Oxf. Econ. Pap. 60 (3), 462–483.
- Sanbonmatsu, L., Kling, J.R., Duncan, G.J., Brooks-Gunn, J., 2006. Neighborhoods and Academic Achievement: Results from the Moving to Opportunity Experiment. J. Hum. Resour. 41 (4), 649–691.
- Schirm, A., Stuart, E., McKie, A., 2006. The Quantum Opportunity Program Demonstration: Final Impacts. Mathematica Policy Research, Inc, Washington, DC.
- Schochet, P., Burghardt, J., Glazerman, S., 2001. National Job Corps Study: The Impacts of Job Corps on Participants' Employment and Related Outcomes. Mathematica Policy Research.

- Schweinhart, L.J., et al., 2005. Lifetime Effects: The High/Scope Perry Preschool Study through Age 40. High/Scope Press, Ypsilanti, MI.
- Sen, A., 1999. Development as Freedom. Anchor Brooks, New York.
- Siedler, T., 2007. Schooling and Citizenship: Evidence from Compulsory Schooling Reforms. Discussion Papers of DIW Berlin 665, DIW Berlin, German Institute for Economic Research.
- Siedler, T., 2010. Schooling and Citizenship in a Young Democracy: Evidence from Postwar Germany. Scand. J. Econ. 112 (2), 315–338.
- Silles, M.A., 2009. The Causal Effect of Education on Health: Evidence from the United Kingdom. Econ. Educ. Rev. 28, 122–128.
- Smith, J.P., 1999. Healthy Bodies and Thick Wallets: The Dual Relation between Health and Economic Status. J. Econ. Perspec. 13 (2), 145–166.
- Smith, J.P., 2007. The Impact of Socioeconomic Status on Health Over the Life-Course. J. Hum. Resour. XLII (4), 739–764.
- Spilimbergo, A., 2009. Democracy and Foreign Education. Am. Econ. Rev. 99 (1), 528-543.
- Tavares, J., Wacziarg, R., 2001. How Democracy Affects Growth. Eur. Econ. Rev. 45 (8), 1341–1378.
- Thornberry, T., Krohn, M., 2000. The Self-Report Method for Measuring Delinquency and Crime. In: Duffee, D., et al. (Eds.), Criminal Justice 2000: Innovations in Measurement and Analysis. National Institute of Justice, Washington, DC, pp. 33–83.
- US Department of Justice, 1994. Profile of Inmates in the United States and in England and Wales, 1991. Washington, DC.
- US Department of Justice, 1999. State Prison Expenditures, 1996. Washington, DC.
- Verba, S., Scholzman, K.L., Brady, H.E., 1995. Voice and Equality. Harvard University Press, Cambridge.
- Viscusi, W.K., 1993. The Value of Risks to Life and Health. Journal of Economic Literature 31 (4), 1912–1946.
- Viscusi, W.K., Aldy, J.E., 2003. The Value of a Statistical Life: A Critical Review of Market Estimates throughout the World. J. Risk Uncertain. 27 (1), 5–76.
- Wagstaff, A., 1986. The Demand for Health: Some New Empirical Evidence. J. Health Econ. 5 (3), 195–233.
- Waldfogel, J., 1994. The Effect of Criminal Conviction on Income and the Trust 'Reposed in the Workmen'. J. Hum. Resour. 29, 62–81.
- Weiner, D., Lutz, B., Ludwig, J., 2009. The Effects of School Desegregation on Crime. NBER Working Paper No. 15380.
- Western, B., 2006. Punishment and Inequality in America. Russell Sage Foundation, New York.
- White, J., Moffitt, T., Silva, P., 1989. A Prospective Replication of the Protective Effects of IQ in Subjects at High Risk for Juvenile Delinquency. J. Consult. Clin. Psychol. 57, 719–724.
- Witte, A.D., 1997. Crime. In: Behrman, J., Stacey, N., (Eds.), The Social Benefits of Education. University of Michigan Press, Ann Arbor, pp. 219–246, chapter 7.
- Witte, A.D., Tauchen, H., 1994. Work and Crime: An Exploration Using Panel Data. NBER Working Paper 4794.
- Wittman, D., 1989. Why Democracies Produce Efficient Results. J. Polit. Econ. 97 (6), 1395–1424.



Overeducation and Mismatch in the Labor Market¹

Edwin Leuven^{*} and Hessel Oosterbeek^{**}

*CREST (ENSAE). Also affiliated with CEPR, CESifo, and IZA **University of Amsterdam, Tinbergen Institute, and TIER

Contents

1.	Introduction	284
2.	Questions	287
3.	Measurement Issues	290
	3.1. Job Requirements	290
	3.1.1 Self-Assessment	291
	3.1.2 Job Analysis	292
	3.1.3 Realized Matches	292
	3.2. Overskilled	294
4.	Incidence	296
	4.1. Aggregate Incidence	296
	4.2. Individual Level: Determinants	298
	4.2.1 Gender	298
	4.2.2 Age	298
	4.2.3 Ethnicity	298
	4.2.4 Ability	299
	4.3. Individual Level: Dynamics	300
5.	Impact on Earnings	302
	5.1. Omitted Variable Bias	304
	5.2. Measurement Error	306
	5.3. Findings	309
6.	Relation to Labor Market Theories	312
	6.1. Human Capital	312
	6.2. Career Mobility	313
	6.3. Job Competition	314
	6.4. Signaling/Screening	314
	6.5. Preferences	315
	6.6. Search and Frictions	315
7.	Conclusion	316
Re	ferences	317

¹ We gratefully acknowledge valuable comments from participants of a Handbook conference in Munich.

Abstract

This chapter surveys the economics literature on overeducation. The original motivation to study this topic reports that the strong increase in the number of college graduates in the early 1970s in the United States led to a decrease in the returns to college education. We argue that Duncan and Hoffman's augmented wage equation—the workhorse model in the overeducation literature—in which wages are regressed on years of overschooling, years of required schooling, and years of underschooling is at best loosely related to this original motivation. Next, we discuss how overschooling and underschooling at the level of individual workers have been measured, and what the incidence of overschooling and underschooling is. We then analyze in more detail Duncan and Hoffman's wage equation. We discuss the potential problems with it due to endogeneity and measurement error, and we review the results from earlier studies using this specification. We conclude that because of the issues concerning endogeneity and measurement error, the estimated returns to required/under/overschooling cannot be interpreted as causal.

Keywords

Mismatch Overschooling Underschooling Wage Equation JEL-Classification: 12

1. INTRODUCTION

Pointing to an analogy between countries' agricultural methods and their perspectives towards education, Gladwell (2008) describes the views of early educational reformers in the United States on overeducation. He refers to the historian Gould who points out that these reformers were very concerned that children got too much education. To illustrate this, Gould cites the US commissioner of education, Jarvis, who in 1872 published a report under the title "Relation of Education to Insanity" in which he claims that of the 1741 cases of insanity he studied, "over-study" was responsible for 205 of them. "Education lays the foundation of a large portion of the causes of mental disorder." Gladwell also quotes Mann—a pioneer of public education in Massachusetts—who believed that "not infrequently is health itself destroyed by over-stimulating the mind."²

More recently, concerns with overeducation were expressed in the 1970s when the supply of educated workers seemed to outpace its demand in the labor market (Berg (1970) and Freeman (1975, 1976)), apparently resulting in a substantial reduction in the returns to schooling. This is illustrated in Fig. 3.1, which is based on data presented

² Gladwell argues that such concerns were totally absent in East-Asian countries like Japan, Singapore, and Korea. He points to the analogy between this difference in perspectives on schooling and the difference in agricultural methods in both parts of the world. Growing rice is much more labor intensive than growing wheat, and moreover, the quality of rice paddies improves after every crop (even with multiple crops per year). In contrast, wheat fields need to be idle once every so many seasons.



Figure 3.1 Ratio of Income/Earnings of College to High-School Graduates. (a) Relative to High-School Graduates of the Same Age. (b) Relative to High-School Graduates Entering the Labor Market at the Same Time.

Sources: Smith and Welch (1978), Table 1 p. 6; Smith and Welch (1978), Table 5 p. 32.

by Smith and Welch (1978). The numbers from 1969 to 1974 (between the vertical lines) in panel A reproduce the key evidence provided by Freeman (1976). The line "New entrants" shows a sharp decline in the ratio of the average income of new entrants with a college degree to the average income of high-school graduates of the same age (between 25 and 34). In a period of just six years, the income premium decreased from 40 to 16%. This decline suggests that supply outpaced demand and caused concerns about overinvestment in college education in the United States. In his book "The Overeducated American," which was published in 1976, Freeman predicts that a situation with substantial oversupply of college graduates is likely to remain for many years to come. Because of this prediction, the book attracted much attention.

Soon after its publication, Freeman's book was critically reviewed by Smith and Welch (1978).³ They first added data for two years prior to the period covered in Freeman's analysis and for two years following his period (the dashed line segments in Fig. 3.1). Adding this information clearly tunes down the spectacular gist of the graph: between 1967 and 1976 the decrease in the income ratio among new entrants is only 8 percentage points instead of 24.

Smith and Welch argued that the evidence is more consistent with an overcrowded labor market among new entrants due to larger cohort sizes than with a situation of overeducation. The line for "New entrants" in panel A compares earnings of college graduates in the age bracket of 25-34 with the earnings of high-school graduates in the same age bracket. But, while college graduates enter the labor market around age 25, high-school graduates enter the labor market at an earlier age. Hence, the graph compares earnings of new entrants of college with earnings of high-school graduates with some years of labor market experience. This means that although the line suggests that the college premium for new entrants has declined, it may as well be explained by new entrants from large cohorts (college graduates of 25-34 in 1973 and 1974) being in a disadvantaged position relative to those with some experience from large cohorts (high-school graduates of 25-34 in 1973 and 1974). To eliminate the effect of labor market experience, Smith and Welch compared the earnings of college graduates in the age bracket of 25-34 with the earnings of high-school graduates of 20-29. This gives the "New entrants" line of panel B. Although the ratio still reaches its lowest value in 1974, the declining trend observed here is much less pronounced.⁴

Although Smith and Welch's review together with the increase in relative earnings of college graduates shortly after 1974 seem sufficient reasons to temper concern about overeducation, the overeducation literature was revitalized by the publication of Duncan and Hoffman's article "The Incidence and Wage Effects of Overeducation" in 1981 in the first issue of the *Economics of Education Review*.⁵ Duncan and Hoffman followed the lead of Eckaus (1964) and Berg (1970) to analyze overeducation by confronting levels of education supplied by workers and education levels "demanded" by jobs. But, where earlier authors confronted these two variables at an aggregate level, Duncan and Hoffman analyze overeducation at an individual level by comparing workers who

³ There were many other critical reviews, including Hammack (1978) and Levin (1977).

⁴ Later studies about wage inequality across skill groups confirm that the drop in the college premium observed by Freeman was a short-lived phenomenon (see Autor and Katz (1999)).

⁵ The top journals (American Economic Review, Journal of Political Economy, Quarterly Journal of Economics, Econometrica, and the Review of Economic Studies) published hardly anything on the topic. To a large extent, this is also true for the top-field journals (Journal of Labor Economics, Journal of Human Resources, Review of Economics and Statistics, and the Industrial and Labor Relations Review). In contrast, five papers in the top 10 of most cited papers published in the Economics of Education Review have overeducation as their topic. These are: Hartog (2000), rank 1; 95 citations, Duncan and Hoffman (1981), 4; 78, Groot and Maassen van den Brink (2000), 5; 68, Dolton and Vignoles (2000), 6; 62, Hartog and Oosterbeek (1988), 9; 55.

end up in a job that matches their level of acquired education with workers who end up in a job that requires more education or less schooling than they acquired. They introduced a specification of the wage equation that allows for the separate estimation of returns for years of education required for the job, years of overeducation, and years of undereducation:

$$\ln w_i = \delta_r S_i^r + \delta_o S_i^o + \delta_u S_i^u + x_i^\prime \beta + \varepsilon_i$$
(3.1)

where w_i denotes the wage, S' is years of education required for the job, S' years of overeducation, and S'' years of undereducation. x_i is a vector of control variables including experience and experience squared. δ_r , δ_o , and δ_u are the returns to required years of education, years of overeducation, and years of undereducation, respectively.

Duncan and Hoffman's seminal article triggered the start of a separate subfield: the economics of overeducation.⁶ This chapter reviews this literature.⁷ The next section analyzes how the approach of Duncan and Hoffman relates to the original overeducation literature, and to which questions their approach potentially addresses.

In Section 3, we describe how different authors have measured required education, and thereby, overeducation and undereducation. We then turn in Section 4 to the incidence of overeducation and undereducation and to the factors that correlate with overeducation and undereducation. In Section 5, we turn to a more detailed discussion of Duncan and Hoffman's wage equation. This section discusses the specification of the wage equation, it gives an overview of the results that have been obtained using this specification, and it discusses problems with this approach due to omitted variable bias and measurement error issues. In Section 6, we discuss how the findings relate to different theories about the functioning of the labor market. These theories include human capital theory, career mobility, job competition theory, signaling/screening, preferences, and search and matching. The final section summarizes and concludes.

2. QUESTIONS

The question whether there is overinvestment in education is interesting and clearly relevant from a policy perspective. Especially in the context of continental European countries where education is heavily subsidized, it is important to know whether education investments pay off. In Freeman's framework, the key indicator for overinvestment in education is the rate of return to education or the college premium.

⁶ We asked Professor Duncan how he perceives the continued success of their specification. He responded that "It took awhile for the simplicity of the required, surplus and deficit categorization to dawn on us, but I think that its transparency led people to remember the results and the analysis. Also, the topic sits in the middle of the human capital, credentialing and labor market institutions literatures, a popular place to be."

⁷ Other reviews of this literature include Groot and Maassen van den Brink (2000); Hartog (2000); McGuinness (2006); Sloane (2003).

Freeman's analysis fits the neoclassical framework. The college wage premium decreases in response to an increase in the supply of highly educated workers. This can happen because firms adjust their production technology to take advantage of the now relatively cheaper and abundant input factor of highly educated workers. It can also happen through a process in which highly educated workers compete for a limited number of skilled jobs by underbidding the wages they demand. Whether firms adjusted their technology can be inferred from changes in the required levels of skills.

In theory, both private and social returns to education investments can be calculated, where the latter takes government expenditures and externalities into account. Estimation and valuation of externalities is, however, difficult. If the return (preferably the social) is low in comparison with returns on investments that are equally risky, this can be regarded as evidence of overinvestment. It is important to realize that this is indicative for overinvestment at an aggregate level. From a policy perspective, this is also the relevant level; overinvestment at the aggregate level can be addressed by reducing public subsidies to education or (in extreme cases) limiting access to education institutes.

The overeducation literature that emerged in the footsteps of Duncan and Hoffman's seminal contribution has been mainly occupied by the following:

- with the incidence of overeducation and its determinants at the individual level, and
- with estimating Duncan and Hoffman's extended wage equation in order to obtain separate estimates of returns to required education, overeducation, and undereducation.

Although Duncan and Hoffman's approach is more detailed and precise than Freeman's macro approach, and can therefore not be less informative, it turns out that the additional information it provides has sometimes led to—in our view—misleading inferences.

Duncan and Hoffman's analysis of overeducation at the individual level reflects a different view on the functioning of labor markets. They argue that overeducation can only be a serious, long-run problem, if changes in the relative supplies of different types of education have little or no effect on the skill composition of labor demand, implying production technologies with fixed skill requirements. Production is not redesigned, jobs are not upgraded, and some workers end up in a job below their skill level. In this view, Duncan and Hoffman argue that jobs are characterized by fixed productivity levels and fixed wages, and individuals working below their skill level produce and earn the same as workers with less schooling in the same job. When they introduce their specification of the wage equation, Duncan and Hoffman claim that it offers "[a] straightforward way to examine the economic effects of over- (and under-) education."

Obviously, from the shares of overeducated, undereducated, and adequately educated workers together with the returns to required education, overeducation, and undereducation, it is possible to recover the return to actually attained schooling which is the statistic on which Freeman would base his diagnosis of overschooling.

To see how the additional information is interpreted, it is instructive to quote the conclusions from some articles that estimated Duncan and Hoffman's wage equation.

[W]ages do not appear to be determined solely on the basis of educational requirements. ... surplus education has a positive and significant effect on wage rates ... [I]t clearly does matter, in analyzing the effects of education on earnings, whether any additional year of education is required or surplus ... [T]he estimated coefficient on years of surplus education is approximately half as large as the coefficient on required years of education ... and this suggests some potential misallocation of educational resources.

Duncan and Hoffman (1981)

This means that workers who are working in occupations that demand less schooling than they actually have (overeducated) get higher wages than their coworkers (holding other characteristics constant) but lower wages than workers with similar levels of schooling who work in jobs in which their schooling equals what is required.

Sicherman (1991)

This suggests that additional schooling is not completely unproductive, but simply that jobs constrain the ability of workers to fully utilize the skills and capabilities they acquire in school. Rumberger (1987), p. 46

Being overeducated lowers an individual's expected earnings, but the schooling that causes an individual to become overeducated generally results in a wage increase.

Rubb (2003b), p. 70

[O]vereducation incurs significant wage costs on the individual and productivity costs on the economy that may well rise if higher education participation continues to expand without a corresponding increase in the number of graduate jobs.

McGuinness (2007), p. 147

More importantly, the returns to surplus and deficit schooling are very low in absolute value and represent only around 45% of the returns to required schooling. Thus, both the human capital model—which would predict equal returns to adequate, over- and under-education—and the job competition model—that would predict zero returns to surplus and deficit years—can be rejected.

Slonimczyk (2008)

In all cases, those with surplus education received a wage premium and those with deficit education suffered a wage penalty. These findings support the idea that productivity on any job is affected by the level of education a worker brings to employment.

Daly, Büchel, and Duncan (2000)

And, finally,

In regard to the wage effects of over- and underschooling, we find that the rate of return to overschooling is positive but lower than the rate of return to adequate schooling, and that the rate of return to underschooling is negative. Also, in general, overschooled (underschooled) workers have wages that are substantially lower (higher) than the wages they would have earned in a job for which they are adequately schooled. In addition, we found that the rates of return to adequate schooling and overschooling (underschooling) decline (rise) as labor market experience rises.

Cohn and Ng (2000), p. 166

All quotes focus on the fact that different components of schooling earn different rewards, and almost all quotes phrase this finding as a causal effect. According to the same quotes, this is considered to be interesting from the perspective of different theories about wage setting. The quotes also suggest that these separate estimates are interesting for their own sake. It is unclear, however, why this is the case. Some quotes suggest that overschooling could (or should) be avoided by reducing the amount of schooling. Given a sufficiently high-average return to attained schooling, this is only useful if it is known in advance who will end up being overeducated and who will end up in a job that matches the actual schooling level. Other quotes suggest that overschooled workers should be assigned to more demanding jobs. It is left unspecified who will create these more-demanding jobs and who gets them.

Alternatively, separate estimates for the three education components are of interest if we want to estimate the efficiency gain from overeducated workers and undereducated workers swapping jobs. This approach expresses that the current allocation of workers to jobs contains mismatches. How many job swaps are possible depends on the levels of attained and required education of the mismatched workers. Although it is unlikely that such job swaps will come about on a voluntary basis (because undereducated workers are likely to loose from it), estimates of returns to undereducation, overeducation, and required education will allow us to compute the potential aggregate wage gain from such a reallocation. This requires, however, that these estimates can be interpreted as the causal effects of moving a person with a given level of education from one job level to another job level. The review in the next sections argues that this is typically not the interpretation that can be given to the estimates from the overeducation literature.

3. MEASUREMENT ISSUES

Before being able to address the questions of the earlier section, one first needs to solve a measurement issue, namely how to define overeducation and undereducation. In the literature, this is typically done by comparing individuals' education with educational requirement of job or occupations. In the next subsection, we describe the various ways in which required schooling has been measured. Subsection 3.2 introduces the related concept of "overskilled" and how people have measured this.

3.1. Job Requirements

Overschooling is usually defined as the difference between a worker's attained or completed level of schooling and the level of schooling required for the job the worker holds. To measure overschooling, most studies therefore start by measuring the required level of schooling. Required schooling has been measured in three different ways. The first method asks workers about the schooling requirements for their job. The second method is based on information that is included in job descriptions. The third method uses information from realized matches. Below we discuss these methods in some more detail. This description is important for later uses (especially Section 5) when we discuss measurement error issues.

3.1.1 Self-Assessment

Measures based on workers' self-assessment rely on questions that ask workers about the schooling requirements of their job. Although this may sound straightforward, it turns out that the exact phrasing varies substantially across studies. Some examples help to illustrate this point.

- 1. How much formal education is required to get a job like yours? (Duncan and Hoffman (1981))
- 2. If someone is applying nowadays for the job you do now, would they need any education or vocational schooling beyond compulsory education? And if so, about how many years of education or vocational schooling beyond compulsory education would they need? (Galasi (2008))⁸
- **3.** Which education, according to you, is the best preparation for the work you are doing? (Hartog and Oosterbeek (1988))
- 4. What kind of education does a person need in order to perform in your job? (Alba-Ramirez (1993))

The first two questions refer to recruitment standards, whereas the last two refer to requirements to perform in the job. But the first and the second question also differ, with the first only referring to formal education, and the second explicitly referring to informal schooling. And also the third and fourth questions differ in an important dimension, with the third asking for the best preparation, and the fourth asking what is needed to perform. It is quite conceivable that the same person gives four different answers to all four questions. Evidence of this is reported by Green, McIntosh, and Vignoles (1999). They asked alumni from the University of Newcastle how much schooling is required to get their current job and how much schooling is required to do their job. A quarter of their respondents gave different answers to these two questions. Moreover, moving from one question to the other may also affect the ordering of answers. The more fundamental problem is that questions differ in the factors respondents are supposed to condition their answers on. It is not clear whether and to what extent differences in framing and phrasing causes differences in the measured levels of required schooling.

⁸ Using these questions, Galasi (2008) reports fractions of overschooling and underschooling that are far off the fractions normally reported in the literature. In a sample of workers from 25 countries, he finds 33% of overschooled workers (with a low 15% in the Netherlands and a high 67% in Ireland) and 59% of underschooled workers (with a low 13% in Estonia and a high 82% in the Netherlands). Only 8% is properly educated (with a low 1% in Turkey and a high 19% in Austria).

Workers' self-assessment of the required amount of schooling of their job is by definition subjective. According to Hartog (2000), respondents may have a tendency to overstate the requirements of their jobs and to upgrade the status of their position. This tendency may differ between persons doing the same job, and may even systematically differ with attributes that have an impact on wages, such as gender or completed years of schooling.

A potential advantage of self-assessment is that it is in principle based on all the relevant information. The downside is, however, that workers may be very poorly informed about the relevant counterfactuals; how can they be informed about the performance in the same job of people with different levels of completed schooling?

A variation of workers' self-assessment of the schooling requirements of their jobs is to ask them directly whether they are overschooled, underschooled or rightly educated for their job (cf. Chevalier (2003) and Verhaest and Omey (2006)).

3.1.2 Job Analysis

A second approach to measure required schooling levels is based on information contained in occupational classifications. A well-known example is the Dictionary of Occupational Titles which contains an indicator for educational requirements in the form of the General Educational Development (GED) scale. This scale runs from 1 to 7. These GED categories are then translated into school-year equivalents (0–18) (cf. Eckaus (1964), p. 184).

Measures based on job analyses are attractive because they are based on the technology of the job. As disadvantages, Hartog (2000) mentions that updates are infrequent and sometimes not so accurate because they are costly. Other disadvantages have been identified by Halaby (1994) specifically pertaining to the use of the GED. He argues that there is no consensus on the conversion of the GED scale to years of schooling. Second, the procedure assumes that there is a fixed required schooling level within an occupation rather than a distribution of required schooling across jobs. For a large number of occupations (47%), it was found that 1960 median schooling was less than the GED requirement, and finally—according to some—the GED scores merely measure social standing of an occupation instead of schooling requirements. All in all, the reliability and the validity of the measure based on job analysis have been called into question (cf. Hartog (2000), p. 132).

3.1.3 Realized Matches

The third method uses information from realized matches (Verdugo and Verdugo (1989)).⁹ In this method, the required amount of schooling for a worker is inferred

⁹ This method goes back to Sullivan (1978) and Clogg (1979).

from the mean of completed schooling of all workers holding the same occupation. Verdugo and Verdugo defined occupations at the three-digit code, and most others applying this method followed that example. People are then defined to be over-schooled or underschooled if their completed level of schooling deviates at least one standard deviation from the mean in their occupation.

Kiker, Santos, and de Oliveira (1997) proposed a method that is closely related; they define the required schooling level in a job as the mode of the completed schooling levels of the people working in that job. People who have more or less than this amount are over/undereducated. Note that this measure does not use the two standard deviations interval around the centrality measure.

The use of realized matches is often regarded as inferior to the other two methods and is only used if the data do not contain one of the other measures. The main reason for this is that the realized match is the result of demand and supply forces and does not only reflect requirements.¹⁰ Moreover, like the method based on job analysis, it ignores variation in required schooling across jobs within an occupation. Furthermore, the cut-off at one standard deviation from the mean is arbitrary.

Another variation of this method was proposed by Quinn and Rubb (2006), who argue that required education may be dynamic due to changes in technology and educational quality. They therefore allow required education to vary with year of birth and survey year. In practice, required education for a given occupation is then equal to the coefficient on the relevant occupation dummy from a regression of actual education on occupation dummies, birth year, and survey year (omitting a constant term).

Groeneveld and Hartog (2004) use the indication used in hiring by the personnel department of the organization as measure of required schooling. They argue that it seems fair to assume that the personnel department has good information on technically required qualifications for a particular job. This measure combines elements of each of the other three measures. It is close to the worker and subjective (as in workers' assessments); it uses insights from personnel experts (as in the job analysis method) and also reflects current market forces (captured by the reference to hiring standards).

Gottschalk and Hansen (2003) argue that the lack of attention in the economics literature to the possibility that college workers were displaced into jobs formerly held by noncollege workers is because of the subjective nature of classifications of occupations as college jobs and noncollege jobs. They give workers' self-assessment and job analysis as examples of this type of classification. As an alternative they propose to define a noncollege job as a job that offers a low college premium. An occupation with a large college premium signals that college workers have skills that are valued by employers in that occupation (p. 455). This approach to measure skill requirement is conceptually

¹⁰ Notice, however, that the same argument applies to methods using information based on hiring standards (such as the question used in Duncan and Hoffman, 1981).

close to Verdugo and Verdugo's approach as their classification is also based on market signals (i.e., employers' willingness to pay a premium for college-educated workers) rather than on subjective judgments. At the same time, these two approaches may produce very different results.

Gottschalk and Hansen (2003) estimate occupation-specific log wage regressions to obtain occupation-specific college premiums. Occupations with a college premium below a certain threshold are classified as noncollege jobs. Using a threshold of 10%, they find that in 1983, 10.7% of all college graduates worked in a noncollege job.¹¹ They also show that their approach gives a rather different picture of the development of the share of college graduates in a noncollege job than would be obtained if workers' self-assessment is used. While their approach shows a 6.6 percentage point decline in this share between 1983 and 1994, based on workers' self-assessment there is a 2.2 percentage point increase.

3.2. Overskilled

At one point, various authors became concerned with the fact that overschooled workers need not be identical to adequately matched workers. Overschooled workers may have lower skill levels, and overschooling does therefore not necessarily imply that someone is overskilled (e.g., Allen and Van der Velden (2001)).

One approach to correct for such unobserved heterogeneity or omitted variable bias is an instrumental variable approach that exploits exogenous variation in the levels of completed schooling and required schooling.¹² So far, finding credible instrumental variables has proven to be very difficult (if not impossible), and researchers have instead constructed measures that somehow attempt to capture workers' excess skill. An example is Chevalier (2003) who criticized the (then) existing measures of overschooling because they implicitly assume that all workers with a given education level are perfect substitutes.

For the case of university graduates, Chevalier distinguishes two ability levels (clever and underachiever) and three job levels (graduate jobs, nongraduate jobs, and upgraded nongraduate jobs). The following allocations are possible in this setup: clever workers can end up in a graduate job or in an upgraded job. In a graduate job, their match is perfect; in an upgraded job, they are genuinely overeducated. Underachievers can end up in a upgraded job in a nongraduate job. In an upgraded job, they are apparently overeducated; in a nongraduate job, they are genuinely overeducated.¹³

¹¹ At the same time, the correlation between the occupation-specific college premium and the wages offered to college graduates within an occupation is positive but weak (0.33).

¹² In Section 5, we show that exogenous variation in just one of these is not enough to identify the return to overschooling.

¹³ By assumption, clever graduates cannot end up in a normal nongraduate job, and underachievers cannot end up in a graduate job.

	Clever	Underachiever
Graduate job	Perfect match	Х
Upgraded nongraduate job	Genuine overeducation	Apparent overeducation
Normal nongraduate job	Х	Genuine overeducation

Neither the worker's skill level nor whether a nongraduate job has been upgraded is observed in his data. Instead Chevalier uses information about workers' satisfaction with the match between their work and their qualifications. Graduates who work in a nongraduate job are considered genuinely overschooled if they are dissatisfied with this match, otherwise they are considered to be apparently overschooled. Using a sample of almost 5000 individuals who graduated from a higher education institution in the United Kingdom in 1985 or 1990, he finds that 10% of these individuals is apparently overschooled, and 6% is genuinely overschooled. From the fact that apparently overschooled workers believe much more than the genuinely overschooled that their degree contributed in getting an interesting job, Chevalier infers that most of the genuinely overschooled are underachievers in nongraduate jobs (the bottom right cell) instead of clever graduates in upgraded jobs (the middle left cell). Further support for this is that within the overschooled group, the apparently overschooled have better credentials than the genuinely overschooled. In a follow-up article, Chevalier and Lindley (2009) find evidence that the genuinely overschooled lack graduate skills such as management and leadership skills.

Chevalier argues that genuinely overschooled people may move to a higher-level job over time, whereas apparently overschooled workers are unlikely to do so. This is consistent with the observation that a substantial fraction of overschooled workers is permanently in this status.

Green and McIntosh (2007), motivated by similar concerns as Chevalier, make the distinction between overqualified and overskilled employees. Someone is defined to be overskilled if she/he disagrees with the statement that the current job offers the opportunity to use the knowledge and skills she/he has, or when the respondent indicated to make little use of past experience, skill, and abilities in the present job. The authors report that the correlation between being overqualified and being overskilled is of moderate strength (0.2). Less than half of the employees who are overqualified report to have skills and abilities they are not using in their job. At the same time, 28% of those who are not overqualified report that they underutilize their skills.

The information that Chevalier, and Green and McIntosh used is admittedly based on subjective self-assessments that probably contain a substantial amount of measurement error. Yet, their results give a strong indication that the estimated return to overschooling in Duncan and Hoffman's wage equation is biased.¹⁴

¹⁴ Overskilling is not the only new concept introduced in the overschooling literature. In fact, there is a whole range of closely related terms that researchers working on overschooling issues use. In a report on skill mismatch in Europe, CEDEFOP (2010) even introduces a glossary of—mostly self-explanatory—terms.

We have seen that the existing approaches to measure overschooling are very much data driven. Ideally one would like to have a theoretical basis for measurement and be explicit about these conditioning factors and whether they should be contemporeneous or relative to the lifecycle, be made on an ex ante basis or ex post. In practice this is not the case. Each method and each survey question either implicitly or explicitly conditions on different factors, whether it is hiring standards, performance requirements, relative prices, etc. It is unclear which method is to be preferred, and importantly the different measures only correlate imperfectly suggesting that measurement error is a concern once these measures are used for analysis.

4. INCIDENCE

In this section, we describe findings regarding the incidence of overschooling and underschooling. We start with the incidence of overschooling and underschooling at an aggregate level and the patterns of these incidences across countries and over time. We then discuss the systematic differences that exist in the measured incidence of overschooling and underschooling due to the use of different measures of required schooling. Finally, we turn to overschooling and underschooling at the level of individual workers. Which individual characteristics correlate significantly with overschooling and underschooling, and how persistent is the individual overschooled or underschooled status over time?

4.1. Aggregate Incidence

Most studies dealing with overschooling report the incidence of overschooling and underschooling observed in the dataset that is used for the analysis. We collected many of these studies, and Table 3.1 reports the (unweighted) means and medians of the reported shares of overschooled and underschooled workers.¹⁵ The first row reports these statistics for all studies together; in the next rows, we report the means and medians for different continents, for time periods, by measure of required schooling and by gender.

The overall means reported by existing studies is 30% for the share of overschooled workers and 26% for the share of underschooled workers. For the share of overschooled workers, the median is equal to the mean; for the share of underschooled workers, the median of the shares reported by existing studies is somewhat below the mean. This indicates that outliers do not play a major role.

The share of overschooling is on average larger in studies that report results for the United States/Canada than elsewhere, whereas the share of underschooled is on average somewhat larger in studies dealing with European countries than in studies dealing

¹⁵ We report the median to check to what extent the mean is sensitive to outliers.

	Fraction Underschooled		Fraction Overschooled	
	Mean	Median	Mean	Median
All studies	0.26	0.21	0.30	0.30
By continent				
Asia	0.21	0.24	0.26	0.29
Australia	n/a	n/a	0.08	0.08
Europe	0.31	0.26	0.30	0.28
Latin America	0.21	0.17	0.24	0.18
United States/Canada	0.16	0.16	0.37	0.35
By decade				
1970s	0.12	0.12	0.40	0.40
1980s	0.23	0.21	0.30	0.32
1990s	0.23	0.20	0.24	0.25
2000s	0.49	0.50	0.39	0.40
By measure of S^r				
Direct self-assessment	0.10	0.10	0.33	0.33
Firm's recruitment	0.15	0.15	0.19	0.19
Job analysis	0.30	0.32	0.34	0.29
Mean	0.15	0.15	0.16	0.13
Mode	0.27	0.27	0.31	0.30
Quinn and Rubb	0.14	0.14	0.14	0.14
Self-assessment	0.32	0.22	0.37	0.33
By gender				
Female	0.21	0.19	0.28	0.31
Both	0.35	0.36	0.34	0.31
Male	0.21	0.19	0.28	0.30

Table 3.1 Incidence of Overschooling and Underschooling

Source: Database constructed from existing studies.

with other countries. This goes against economic priors because education is more subsidized, and labor markets are arguably more imperfect in Europe. Both factors would lead to higher levels of overeducation in Europe compared with America. There appears to be a decline in reported overschooling from the 1970s to the 1990s. This, however, is consistent with the literature on skill-biased technical change. In the 2000s, there is suddenly a sharp increase. This increase is, however, solely attributable to the shares reported in the study by Galasi (2008). As we discussed in Section 3, this study employs a somewhat different method to measure required schooling.

Studies reveal no systematic difference between the reported shares of over/underschooled men and women. Strangely enough, studies that report results for men and women together tend to find larger shares of over/underschooled workers.

Finally, the fractions of overschooled and underschooled workers reported in different studies vary greatly with the method that was used to measure required schooling. The methods using workers' self-assessment or the information from job analysts find shares of overschooling of more than 30%; the same is true for the mode method. The mean method, which uses the two standard deviations bandwidth, finds much lower shares of overschooled (and underschooled) workers. These results support concerns over measurement error.

4.2. Individual Level: Determinants

Many studies estimate probit or similar binary outcome models of the determinants of overschooling and underschooling. The specification of these models varies widely among studies, complicating the comparison of the resulting estimates. There is typically little motivation of why some variables are included as controls and not others. An additional complication is that some of the included regressors such as tenure are potentially endogenous. More or less consistent findings across studies are that young people, women, migrants, and people who are unmarried are more likely to be overschooled.

4.2.1 Gender

The higher incidence of overschooling among women has been addressed by Frank (1978). His explanation is that when the men are the prime income earners in a house-hold and when the choice of location is determined by the man's labor market prospects, women are necessarily more restricted, and this may translate in an increased probability to be overschooled.

4.2.2 Age

Older workers are less likely to be overschooled than their younger colleagues. This fact is consistent with search theory which predicts that workers are increasingly in better matches, but also with the theory of career mobility where workers who are overschooled in their first job have a higher probability to be promoted. It is also consistent with the view that the labor market rewards workers' entire bundle of human capital in which extra schooling can compensate for lack of experience. The observed age effect is also consistent with a gradual upgrading of schooling requirements.

4.2.3 Ethnicity

The number of studies inquiring into the relationship between ethnicity and overschooling is limited (Green, Kler, and Leeves (2007) and Battu, Sloane, Building, Street, and Park (2004) are exceptions). This is probably partially due to problems of comparibility of schooling in the country of origin to schooling in the country of residence. When minorities have not been educated in other countries, a rationale for higher overschooling incidence among minority groups runs similar to the rationales for women and young people. They may be more geographically restricted, and/or they may need the extra schooling to compensate for other shortages in their human capital bundle, such as proficiency in the native language. Of course, higher overschooling incidence among minorities may also reflect plain labor market discrimination.

4.2.4 Ability

Another determinant of being overschooled may be a lower level of ability. The studies that have access to ability measures find that ability and overschooling (given the level of actual schooling) are indeed negatively correlated.

Chevalier and Lindley (2009) construct a measure of unobserved ability as the residual from a first-job earnings equation capturing all observed characteristics of an individual including job characteristics that affect wages more than the control variables included. Including these residuals should then proxy for time-invariant fixed unobservable characteristics. Estimates from multinomial logit models show that a one standard deviation increase in the unobservable component reduces the probability to be apparently overschooled by 2.9 percentage points, whereas it reduces the probability to be genuinely overschooled by 3.7 percentage points. These changes should be compared with base probabilities of 19.6 and 14.9%.

Green and McIntosh (2007) use data from the 2001 Skill Survey in the United Kingdom to examine the relation between being over/underqualified and the extent to which respondents report that they use their skills. Respondents were asked how much they agreed with the statement: "In my current job I have enough opportunities to use the knowledge and skills that I have" and "How much of your past experience, skills and abilities can you make use of in your present job?" If the respondent (strongly) disagreed with the statement or answered (very) little to the second question, she/he was recorded as overskilled. Likewise respondents who (strongly) agreed with the statement "I would perform better in my current job if I possessed additional knowledge and skills," were recorded as underskilled. Green and McIntosh report that the correlation coefficient of being overqualified and overskilled is only 0.20. Less than half of the overqualified employees report that they have skills and abilities that they are not using in their job. Among those who are not overqualified, this percentage equals 28%. Even stronger is the finding for underqualified workers: they are not more or less likely than the nonunderqualified workers to report being underskilled. Green and McIntosh interpret this as an evidence that these employees have gained skills through other routes after the end of their schooling (p. 432).

Very similar results are reported by Allen and Van der Velden (2001) using data from the Dutch wave of the Higher Education and Graduate Employment in Europe project. The respondents graduated in 1990/1991 and were interviewed at the end of 1998. For example, of the respondents with a higher vocational degree and who found a job matching their education level, 13% report that their skills are underutilized.

Among the respondents with the same education but with a job below their education level, this percentage equals 21%.

Green, McIntosh, and Vignoles (1999) find that individuals who had higher scores in a math test when aged 16 were significantly less likely to be overeducated in their working lives; a 10% higher score in the mathematics test is associated with a 2% lower probability of being overeducated. Using the English data from the International Adult Literacy Suvey, they also find that individuals with low quantitative skills are more likely to be overeducated. Using German data, Buchel and Pollmann-Schult (2001) find that a poor grade on the school leaving certificate has a strong effect on the probability of overeducation for those graduating with a vocational degree (after controlling for selection effects).¹⁶

4.3. Individual Level: Dynamics

Overschooling is not necessarily a permanent state. For a better understanding of the overschooling phenomenon, it is helpful to know more about its persistence at an individual level. Is 30% of the workforce permanently overschooled or is a larger fraction overschooled for a shorter while and is being overschooled just a stage in workers' careers?

Sicherman and Galor (1990) developed a theory of career mobility that can partially explain overschooling. In their model, individuals may accept a lower-level job in which the direct return to schooling is lower if, in those jobs for a given level of schooling, the probability of promotion is higher.

To test their model, Sicherman and Galor estimate for 24 different occupations the effect of education on wages and the effect of education on the probability of upward career mobility. They find that the estimated effects are negatively correlated: occupations with a higher wage return to education have on average a smaller effect of education on upward mobility.

Sicherman and Galor (1990) claim that their model implies that "it will be rational for some individuals to spend a portion of their working careers in occupations that require a lower level of schooling than they have acquired" (p. 177). This is explored by Sicherman (1991). Using PSID data for 1976 and 1978, he finds that being overschooled indeed has a positive impact on upward mobility. The size of this impact is, however, quite small. Some of the overschooled workers move upwards, but a large fraction remains overschooled. Various studies corroborate that for many workers overschooling persists. Although this is consistent with substantial mobility constraints on the part of workers, it is difficult to reconcile this with a structural population of underschooled workers.

¹⁶ See also Battu, Belfield, and Sloane (1999) (who control for the degree class and institution type in the United Kingdom), Dolton and Vignoles (2000) and Green and Zhu (2010) (who look at degree classifications and universities' reputation).

A more likely explanation seems that workers are and stay overschooled because they lack the necessary skills to perform more-demanding jobs.

Bauer (2002) uses the German GSOEP data from 1984 to 1998 to analyze the returns to overschooling and underschooling in a panel data context. He does not report in detail the evolution of the mismatch status over time. From the information that he provides, it is evident, however, that rather few employees in Germany change their mismatch status. Using the mean method of Verdugo and Verdugo, 5.3% of the observations change their mismatch status, whereas this percentage is 16.2% when using the mode index. This implies relatively few changes, especially if one realizes that some of the variation in the realized match method is artificial. A worker who does not change his own schooling level and keeps the same job can still change his overschooling status if the mean (mode) schooling level in his occupation changes. Bauer is aware of the fact that the small numbers of a change in the mismatch status cast doubts on whether the wage effects of educational mismatch could be identified in the fixed effects model.¹⁷

Rubb (2003a) uses data from 1992/1993 to 1995/1999 waves of the Current Population Survey to document the dynamics of overschooling in the United States. He measures over/underschooling from realized matches (using Verdugo and Verdugo's mean method) and reports incidence rates of overschooling around 0.14. Rubb reports an annual ouflow from overschooling to adequately schooled of a bit below 0.2. Outflow rates out of overschooling are slightly higher for women than for men, decrease with age, and are not different for recession and expansion years. Rubb concludes that overschooling is not solely a short-run phenomenon that only exists when individuals gain experience or search for a new job.

Battu, Belfield, and Sloane (1999) examine two cohorts of higher education graduates in the United Kingdom. One cohort graduated in 1985, the other in 1990. The 1985 cohort was interviewed 1, 6, and 11 years after graduation; the 1990 cohort was interviewed 1 and 6 years after graduation. Mismatch is measured by responses to the question: "Was the degree gained in 1985 or 1990 a requirement in the job specification for your main employment (including self-employment)?" For men, the fractions of correctly matched workers all circle around 60%, whereas for women, the fraction increases from less than 55% one year after graduation to just above 60% five years later. These fractions hide some movement into and out of correct matches; between 65 and 70% of the workers ever had a job that required their degree. One potential caveat of this study is that some people indicate that they move from a job that requires their degree to a managerial job that does not require their degree. This indicates that part of the "mismatched" workers is not really mismatched.

¹⁷ We discuss these results in Section 5.

Mavromaras, McGuinness, and Fok (2009) analyze the Australian HILDA Survey 2001–2006. The sample consists of 15- to 64-year-old employees. They consider employees who report that they do not use many of their skills in their current job. This is different from the standard overschooling measure in the sense that it considers skill use and not educational attainment. Workers are classified as moderately or severely overskilled (and well matched). In the bottom half of the education distribution, about one-third of the workers is moderately overskilled and 19% severely overskilled. In the top half of the education distribution, these numbers are 27 and 10%. They drop, however, the moderately overskilled from the sample, thereby ignoring transitions from severely overskilled and well matched to moderately overskilled. Average transition rates per year are not reported (but are probably low). They estimate dynamic random effects probit models and find substantial state dependence. When stratifying by education they find that persistence is especially high for individuals with a higher education degree.¹⁸

Careful and systematic study of the incidence and dynamics of overschooling have received relatively little attention compared with the return estimates we review in the next section. This is somewhat unfortunate as this could have increased our understanding of what is actually being measured. Our reading of the cross-sectional evidence on the incidence and persistence of overschooling is that although part of it is frictional, it mostly appears to correlate with unobserved skills.

5. IMPACT ON EARNINGS

The workhorse model in the overeducation literature is the extended version of the Mincerian wage equation introduced by Duncan and Hoffman (1981). Denote the standard Mincerian wage equation as:

$$\ln w_i = \delta_a S_i^a + x_i' \beta + \epsilon_i \tag{3.2}$$

where w_i is individual *i*'s wage, S_i^a attained years of schooling, and x_i a vector of control variables including experience and experience squared. δ_a is the return to (attained) education. Duncan and Hoffman divided S^a into three components: the years of schooling required for the job (*S'*), and the difference between attained and required years, which is labeled years of overeducation (or surplus years) if attained schooling exceeds required schooling (*S*^o) and which is labeled years of undereducation (or deficit years) if attained schooling falls short of required schooling (*S'*). The following identity therefore holds

$$S_i^a \equiv S_i^r + \underbrace{\max(0, S_i^a - S_i^r)}_{S_i^o} - \underbrace{\max(0, S_i^r - S_i^a)}_{S_i^{u'}}$$

¹⁸ See also Lindley and McIntosh (2008), Dolton and Vignoles (2000), and Frenette (2004).

Duncan and Hoffman replace S^a in the Mincerian wage equation by these three components and allows them to have different effects on wages. The wage equation then reads as follows:

$$\ln w_i = \delta_r S_i^r + \delta_o S_i^o + \delta_u S_i^u + x_i^\prime \beta + \varepsilon_i$$
(3.3)

where δ_r , δ_o , and δ_u are the returns to required schooling, overschooling, and underschooling, respectively.

An attractive feature of Eq. (3.3) is that it nests the standard human capital Eq. (3.2) as a special case. By imposing the restriction $\delta_r = \delta_o = -\delta_u$, it is straightforward to statistically test whether the Duncan and Hoffman specification fits the data better than the standard Mincer equation.

Another specification that is nested in the Duncan and Hoffman specification imposes the restriction $\delta_o = \delta_u = 0$ which implies that only the job requirements matter for the wage. This is consistent with the Thurow's job competition model in which marginal productivity is a fixed characteristic of the job and independent of the worker (see Section 6.3 below). Wages are also assumed to be related to the job rather than to the worker. There is no claim of equality between wages and marginal productivity, and it is not clear how wages are determined.

From the perspective of an individual who has to decide on the optimum amount of schooling, the ex ante return δ_a is the more interesting parameter. Separate information on the parameters δ_n , δ_o , and δ_u is only interesting for an individual who knows whether she/he will end up in a job below/above her/his attained schooling level. Alternatively, the separate parameters may also convey relevant information about the risk involved in the schooling decision.

An often cited (and heavily criticized) contribution to the overeducation literature is Verdugo and Verdugo (1989). They estimated the returns to overschooling and underschooling using a specification that differed in three important ways from the specification used by Duncan and Hoffman (1981).¹⁹ First, they measured required schooling through the mean method (see Section 3). Second, instead of including years of over/underschooling, they included dummies for being over/underschooled. And third, instead of required years of schooling, they control for completed years of schooling in the regression. Confusingly, Verdugo and Verdugo discuss their findings as if they controlled for required instead of completed education. In their critiques, neither Cohn (1992) nor Gill and Solberg (1992) clearly point out that this is why Verdugo and Verdugo find effects of opposite sign.²⁰ As Hartog (2000) argues, "[b]uilding up a body of comparable research, over time and place, would be facilitated by deleting this specification." (p. 139).

¹⁹ Surprisingly, Verdugo and Verdugo do not refer to Duncan and Hoffman's article.

²⁰ It is straightforward to show that when estimating $\ln w_i = \alpha_a S^a + \alpha_o S^o + \alpha_u S^u + x'_\beta + \varepsilon_i$ one obtains estimates such that $\alpha_a = \delta_r, \ \alpha_o = \delta_o - \delta_r$ and $\alpha_u = \delta_u + \delta_r$.

A special case of the Verdugo and Verdugo specification naturally arises when considering the impact of overeducation on earnings of university graduates. This is an intensively studied topic in the United Kingdom (examples include Battu, Belfield, and Sloane (1999); Chevalier (2003); Chevalier and Lindley (2009); Dolton and Vignoles (2000); Dolton and Silles (2008)). Also, studies from the United Kingdom that do not exclusively deal with university graduates are often not comparable with studies that apply the Duncan and Hoffman specification as these studies measure education in terms of qualification levels instead of years of education.

Before we review the findings that have been obtained in studies reporting estimates from Duncan and Hoffman's specification in Subsection 5.3, we first discuss problems concerning omitted variable bias and measurement error. In our view, neither problem has been addressed satisfactorily in the overeducation literature. This makes interpretation of the findings less straightforward.

5.1. Omitted Variable Bias

Since the early 1990s, identification issues have been a central theme in labor economics and in the economics of education (see Meghir and Rivkin (2011) for a review in the context of the economics of education). Developments in this area have almost entirely been ignored in the overeducation literature. This is not because identification issues are trivial in this subfield. Clearly if ϵ_i and S_i^a in Eq. (3.2) are correlated—as many returns to schooling studies suggest (see Card (1999) and Ashenfelter, Harmon, and Oosterbeek (1999) for surveys)—then there is no reason to assume that this would not be the case for the error term and the schooling components in Duncan and Hoffman's equation. Hence, also ε_i and S'_i , S^o_i and S^u_i will be correlated. The endogeneity problem is likely to be even more severe in the context of Eq. (3.3) than in the context of Eq. (3.2). Not only are estimates biased because of the nonrandom assignment of individuals to completed schooling levels, we now also have to deal with the nonrandom assignment of individuals to required schooling levels. Even if completed schooling is randomly assigned—or if we have credible instrumental variables for completed schooling—estimates of δ_{i} , δ_{o} , and δ_{u} will still be biased when workers are not randomly assigned to job requirements. And vice versa, when workers are randomly assigned to job requirements—or if we have credible instruments for required schooling-these estimates will still be biased when workers are not randomly assigned to completed schooling levels.

Addressing these endogeneity problems is far from trivial. This is illustrated by Korpi and Tahlin (2009), one of the few studies using instrumental variable methods to estimate returns to over/underschooling. Korpi and Tahlin instrumented the three schooling components in Duncan and Hoffman's specification using number of sibings, place of residence during childhood, economic problems in the family of origin, and disruption in family of origin. Not only are there indications that these instruments are weak, but one may question the validity of these exclusion restrictions. Because there is more than one endogenous variable one must also be careful to verify the rank conditions. If instruments are motivated as instruments for completed years of schooling, they cannot serve as instruments for required schooling, overschooling, and underschooling. This is easily seen by considering the following simplified model (dropping individual subscripts, intercepts, and covariates for convenience):

$$\ln w = \delta_r S^r + \delta_o (S^a - S^r) + \varepsilon \tag{3.4}$$

$$S^a = \alpha_1 z_1 + \alpha_2 z_2 + u_a \tag{3.5}$$

$$S^r = \beta S^a + u_r \tag{3.6}$$

where z_1 and z_2 are (valid) instruments for S^a , and u_a and u_r are error terms. Equation (3.4) is a simplified version of Duncan and Hoffman's wage equation in which years of overschooling and years of underschooling are restricted to have symmetric effects on wages. Equation (3.6) reflects that S^r may vary with S^a ; higher-educated people are more likely to end up in higher-qualified jobs.

Substituting Eq. (3.5) into Eq. (3.6) and the resulting equation along with Eq. (3.5) into Eq. (3.4) gives us the reduced form wage equation:

$$\ln w = (\delta_r \beta + \delta_o - \beta) \alpha_1 z_1 + (\delta_r \beta + \delta_o - \beta) \alpha_2 z_2 + \tilde{\varepsilon}$$
(3.7)

where $\tilde{\epsilon}$ collects the remaining terms. Although we can obtain estimates for α_1 and α_2 from Eq. (3.5) and for β from Eq. (3.6), the two terms in parentheses in Eq. (3.7) are identical and will not suffice to identify δ_r and δ_o separately. This changes when one of the instruments (or a third instrument z_3) has an independent effect on required schooling.

Hence, for instrumental variable techniques to be applicable in this context (where required schooling depends on actual schooling) we do not only need an instrument for completed schooling, but also for required schooling. This severely limits the potential of instrumental variable methods to estimate the returns to overschooling. Although researchers have been reasonably successful in finding and applying instruments for completed years of schooling, finding instruments for required schooling is even more challenging.

As an alternative method to address the endogeneity problem, several articles have applied fixed effects' techniques (Bauer (2002); Dolton and Vignoles (2000); Dolton and Silles (2008); Korpi and Tahlin (2009); Lindley and McIntosh (2008); Tsai (2010)).²¹ Fixed effects' estimates of the returns to over/underschooling are identified from persons who have changed educational level, job level, or both. In both cases, it

²¹ Many of these articles employ data for the United Kingdom and as a consequence do not estimate a specification with years of overschooling but with overschooling dummies.

needs to be the case that relevant unobservables are time invariant. Consider an example of a person who increases her level of completed schooling between waves. She may have been devoting time to study before the change occurred. If this unobserved change in study time also has an impact on wages, fixed effects estimates will be biased.

In practice, fixed effects' estimates will mostly rely on job changes for identification since people rarely change schooling after having entered the labor market. A person experiencing a change in the level of schooling required for the job will therefore probably have changed jobs. Job changes can, however, be preceded, accompanied, or followed by many other changes that are unobserved and affect wages. In such cases, the strict exogeneity assumption that is necessary for fixed effects' estimates to be consistent fails.

As another alternative approach to address endogeneity problems, McGuinness (2007) applies propensity score matching. As the identifying assumption is essentially the same as under ordinary least squares (no endogeneity conditional on observables), it is not surprising that he finds very similar estimates. His claim that this approach addresses omitted variable problems is in our view not realistic. At best, his results show that earlier findings obtained from OLS are not attributable to common support problems.

The above discussion does not imply that overeducation researchers are not aware of the fact that estimated returns to overeducation may reflect differences across people in terms of other unobserved components of their human capital stock or of their motivation. The insight that people may have acquired more schooling than would normally be required for their job to compensate for a shortage in some other human capital components such as experience or ability, motivated the work of, for instance, Chevalier (2003), Allen and Van der Velden (2001), and Green and McIntosh (2007). It does imply, however, that it is extremely difficult to obtain credible estimates of causal effects of being over/underschooled. And, as we have argued in Section 2, obtaining estimates of causal impacts is crucial for the results from Duncan and Hoffman's wage equation to be useful.

5.2. Measurement Error

Another concern in the literature on the returns to schooling is attenuation bias because of measurement error in the years of schooling variable. Given the discussion in Section 3 regarding the problems surrounding the measurement of required schooling, this issue is likely to be even more important when it comes to the estimation of the returns to required schooling. Moreover, both overschooling and underschooling are measured as a difference between two possibly mismeasured schooling levels. It is well known that differencing leads to exacerbation of measurement error problems and is akin to the measurement error problem in returns to schooling studies that use data from twins. As discussed in Section 4, different methods to measure required schooling present rather different results for the incidence of adequately over- and undereducated workers. At the same time, it has been noted that the method used to measure required schooling does not seem to have a large influence on the estimated returns despite these differences in incidence. This apparent consistency has led some commentators to conclude that measurement error bias is not an issue in the estimation of (3). This is of course a fallacy. The difference in incidence indicates that at least one required schooling variable is measured with error. As a consequence, at least one set of estimates will be biased (and different from others if they are measured correctly). But, since estimates based on the various measures give similar results, these estimates all need to be biased.

Only a few studies have made attempts to address the issue of measurement error explicitly in the context of returns to required, overschooling and underschooling. Robst (1994) and Dolton and Silles (2008) use instrumental variable methods, and Tsai (2010) uses a minimum distance approach.

To better understand the role of measurement error in (3), let D_i equals 1 if individual *i* is overeducated and zero otherwise. We can rewrite Eq. (3.3) as follows (dropping control variables)

$$\ln w_i = \delta_r S_i^r + \delta_o (S_i^a - S_i^r) D_i + \delta_u (S_i^r - S_i^a) (1 - D_i) + \epsilon_i$$
(3.8)

We assume that only S^r is measured with error

$$\tilde{S}_i^r = S_i^r + u_i \tag{3.9}$$

where classical measurement error corresponds to the case in which S_i^r and u_i are independent. The classical measurement error assumption typically fails when variables are bounded. To see this note, in the case of (required) schooling, when \tilde{S}_i^r is the highest possible schooling level then u_i must be negative. Similarly, when \tilde{S}_i^r is the lowest possible level then u_i can only be positive. As a consequence the measurement error, u_i will be negatively correlated with the true value S_i^r . Where classical measurement error leads to attenuation bias, the negative covariance leads to bias in the opposite direction.

It has been shown, in the case of single mismeasured regressors that with nonclassical measurement error, OLS estimates tend to be biased towards zero as long as the covariance between the measurement error and the variable is not stronger than the variance of the measurement error (Black, Berger, and Scott (2000)). In this case, the OLS estimate would serve as a lower bounds on the true effect. When there is a very strong negative correlation between the measurement error and the true value, the OLS estimate is also biased and can even change sign (Aigner (1973)).

An additional complication of measurement error in S_i^r in the context of Eq. (3.8) is that it will also lead to measurement error in S_i^o and S_i^u . This happens both directly and

through classification error with respect to the overeducation status D_i . The resulting measurement error will, again, be nonclassical.

Substituting (9) in (8) we get:

$$\ln w_i = \delta_r \tilde{S}^r + \alpha_o \underbrace{(S_i - \tilde{S}^r_i) \tilde{D}_i}_{\tilde{S}^o_i} + \delta_u \underbrace{(\tilde{S}^r_i - S_i)(1 - \tilde{D}_i)}_{\tilde{S}^u_i} + \tilde{\epsilon}$$
(3.10)

where

$$\tilde{\epsilon} = \epsilon - u(\delta_r - \delta_o D_i + \delta_u(1 - D_i)) + (\delta_o + \delta_u)(S_i - \tilde{S}_i^r)(D_i - \tilde{D}_i)$$
(3.11)

and \tilde{D}_i is the misclassified overschooling indicator. There are now two channels which affect our estimates in the presence of measurement error. First, there is the bias channelled by u in \tilde{e} . This not only attenuates the estimate but also leads to a bias in the opposite direction in case of a negative correlation between the measurement error and S^r . The second source of bias comes from classification bias $(D_i - \tilde{D}_i)$.

The standard method to correct for bias resulting from classical measurement error is instrumental variables. This is the approach of Robst (1994) and Dolton and Silles (2008). When measurement error is nonclassical, instrumental variables are, however, also biased. To see why, note that any instrument that is correlated with required/under/overschooling will also be correlated with the residual in (10) because $cov(S^r, u) \neq 0$. It can be shown that with a single mismeasured regressor the IV estimate will be biased upwards and represents an upper bound (Black, Berger, and Scott (2000)). In the context of (10), there are three mismeasured variables, where the measurement error is not mutually independent, and it is unclear whether the results of Black, Berger, and Scott (2000) carry over.

As an illustration of these issues, consider Robst (1994) who recognized early the potential importance of measurement error. He used one measure of required schooling as an instrument for the other (and vice versa). Of the two measures, one is based on workers' assessment and the other on the GED estimate of required education based on three-digit Census codes for occupations. The dataset that Robst used is the 1985 wave of the PSID. The first thing to note is that only 53% of the observations is classified in the same category (adequately educated, overeducated, and undereducated) on both measures. This implies that at least one of the required schooling measures is plagued by substantial measurement error.

Robst instrumented self-assessed required schooling, overschooling, and underschooling by required schooling, overschooling, and underschooling based on the GED measure of required schooling (and vice versa). The return to required schooling as measured by the worker's assessment increases from 0.07 to 0.09 (using required schooling according to the GED measure as an instrument). For the opposite instrumentation the return hardly changes (from 0.08 to 0.09). The return to overschooling, however, changes drastically. Using OLS, the estimated return to a year of overschooling is about 0.05 using both measures of required schooling. When instrumenting one for the other, these returns drop to about -0.02 and are not statistically significant.

Where Robst (1994) estimates the full Duncan and Hoffman specification with three potentially mismeasured variables, Dolton and Silles (2008) analyze a sample of graduates from a UK university where, by definition, all hold the same level of education and can only be overschooled. As a consequence, they estimate a wage equation with only a single overeducation dummy *Di*. They find that university graduates who report to be overschooled in their first job earn less than their nonoverschooled counterparts with an estimated coefficient of -0.16 (s.e. 0.05). The IV estimate increases (in absolute value) as expected to -0.41 (s.e. 0.10). As graduates progress in their career, the wage penalty associated with overschooling increases. The corresponding OLS estimate of the overeducation coefficient for the current job is -0.35 (s.e. 0.05) with the IV estimate at -0.66 (s.e. 0.10). In both cases, the IV estimate is about twice the size of the OLS estimate. Dolton and Silles (2008) also control for individual fixed effects and obtain very similar results.

In a recent article, Tsai (2010) uses the PSID data to estimate fixed effects' versions of (3). The author reports negative point estimates on overschooling (and positive point estimates on underschooling) because the used specification deviates from (3) in that it controls for actual rather than required schooling. It is therefore necessary to take the appropriate transformation of the coefficients to compare the estimates to those of the literature. Doing this, we see a return to required schooling that decreases from about 0.11 in the pooled OLS to 0.02 in the fixed effects' models. Similar changes are observed for the coefficients on overschooling and underschooling. The pooled OLS estimate on overschooling is 0.076, and the fixed effect estimate is 0.015. For underschooling, these numbers are -0.057 and -0.010.

One interpretation of Tsai's results is that the decrease in the estimates when using fixed effects' estimation is caused by measurement error. To address this question, Tsai implements a numerical procedure that assumes classical measurement error on required schooling but allows for classification bias that might result of this. Unfortunately, it is unclear whether the implemented estimation procedure produces consistent estimates. Moreover, the estimated variance of the measurement error (about 16) is six times larger than the variance of (mis)measured required schooling itself (about 2.6) which should be a logical impossibility.

5.3. Findings

This final subsection summarizes the estimates of the returns to required, over-, and underschooling reported in the overeducation literature. Given the unsolved issues concerning omitted variable bias and measurement error, we do not claim that these estimates represent causal effects. It is probably not only the simplicity of Duncan and Hoffman's specification that has been an important factor for its popularity, but also the results originally reported by Duncan and Hoffman (1981) are likely to have contributed. Using the 1976 wave of the PSID, they report a return to completed years of education equal to 0.058 for white men. When they break completed education down into its three components, they find a return of 0.063 to a year of education required for the job, a return of 0.029 for a year of surplus education and a negative return of -0.042 for each year of deficit education. The return to an overeducated year is significantly *and substantially* lower than the return to a required year of education, suggesting that it is indeed worthwhile to break completed education down.

Many studies have replicated Duncan and Hoffman's study using different data and (not always) different measures for required schooling. In Table A, in the Appendix, we tabulate the findings of many of these studies. Table 3.2 summarizes the findings. To average the estimates of returns to required schooling and over/underschooling from various studies, we have used a weighting method which has been developed in the meta-analysis literature (see Raudenbush (1994) or Van Ewijk and Sleegers (2010) for details).

The first row in the table shows that the return to a year of required schooling is around 0.09, to a year of overschooling more or less half of that, and a year of underschooling results in a wage penalty of again around half of the return to a required year of schooling.

The other rows of Table 3.2 break the sample of results down by continent, decade, method to measure required schooling, estimation method, and gender. The results by continent suggest that the absolute values of the returns are a bit larger in the Asian countries in our sample than in Europe and in the United States/Canada whereas differences between European countries and North America are small. Over time we do not observe a consistent pattern. Only the 1990s are characterized by somewhat higher absolute values of the returns estimates. The differences in estimated returns between studies that use self-assessed measures of required schooling, and studies that use the mean or mode method find larger returns on required schooling, and those based on the mode method also find larger absolute returns on overschooling and underschooling.

Differentiating estimated returns to the three components by different estimation methods indicates that studies that use IV tend to find a much lower return to overschooling than studies using OLS. Likewise studies using fixed effects are characterized by a lower return to required years of schooling. It should be noted though that the number of studies applying IV and fixed effects is rather small.

The final rows in Table 3.2 show that returns to different schooling components are not systematically different between men and women. Results based on data that

	δ_r	δ_o	δ_u	N Results
All studies	0.089 (0.003)	0.043 (0.002)	-0.036 (0.002)	151/151/143
By continent				
United States/	0.083 (0.006)	0.046 (0.004)	-0.027 (0.004)	32/32/26
Canada				
Latin America	0.075	0.041	-0.034	
Europe	0.076 (0.003)	0.038 (0.003)	-0.035 (0.004)	94/94/94
Asia	0.135 (0.008)	0.052 (0.003)	-0.042 (0.004)	18/18/18
Australia	0.105	0.05		
By decade				
1970s	0.079 (0.007)	0.043 (0.004)	-0.037 (0.008)	13/13/9
1980s	0.084 (0.005)	0.048 (0.003)	-0.036 (0.003)	37/37/37
1990s	0.113 (0.008)	0.038 (0.004)	-0.026 (0.007)	41/41/38
2000s	0.095 (0.004)	0.046 (0.003)	-0.040 (0.003)	27/27/26
By measure of S ^r				
Self-assessment	0.079 (0.003)	0.041 (0.003)	-0.034 (0.003)	68/68/66
Job analysis	0.075 (0.004)	0.043 (0.002)	-0.032 (0.002)	27/27/23
Mean method	0.108 (0.009)	0.041 (0.007)	-0.025 (0.011)	29/29/27
Modal method	0.101 (0.009)	0.054 (0.004)	-0.056(0.005)	24/24/24
By estimation method				
OLS	0.092 (0.003)	0.043 (0.002)	-0.033 (0.002)	128/128/120
IV	0.095 (0.016)	-0.031 (0.028)	-0.031 (0.032)	4/4/4
FE	0.043 (0.019)	0.031 (0.023)	-0.040 (0.022)	5/5/5
By gender				
Male	0.090 (0.005)	0.047 (0.003)	-0.037 (0.004)	51/51/46
Female	0.101 (0.009)	0.046 (0.005)	-0.042 (0.006)	35/35/32
Both	0.077 (0.004)	0.035 (0.004)	-0.029 (0.004)	63/63/63

Table 3.2 Returns to Required Schooling (δ_r) , Overschooling (δ_o) , and Underschooling (δ_u)

Source: Database constructed from existing studies. Means and their standard errors are obtained by weighting estimates with the inverse of their variance. See Harbord and Higgins (2004) for details.

combine information of men and women tend to report lower returns that are smaller (in absolute value). Table 3.3 presents the results from meta-regressions in which the estimated returns to the three schooling components are regressed on study characteristics. These results reiterate the patterns observed in Table 3.2 in a multivariate framework.

The quotes that we cited in Section 2 demonstrate that the types of results summarized in this subsection have been interpreted as causal effects of overeducation and undereducation on wages. Wordings like: "a positive effect," "holding other characteristics constant," "not completely unproductive," "results in," "returns are very low," and "productivity is affected" are very explicit in that respect. Because of problems with omitted variable bias and measurment error, we believe that such interpretations are not warranted.

Table 3	.3 Met	a-Regr	essions
---------	--------	--------	---------

	δ_r	δ_o	δ_u
Dummy Asia = 1	0.033*** (0.009)	0.010 (0.008)	-0.029** (0.010)
Dummy Europe = 1	-0.011 (0.008)	-0.006 (0.007)	-0.013 (0.009)
United States/Canada,	reference	reference	reference
Mexico, Australia			
Decade 1970s	-0.015 (0.012)	-0.010 (0.010)	0.002 (0.014)
Decade 1980s	-0.014* (0.007)	-0.004 (0.006)	0.014* (0.007)
Decade 1990s	0.018** (0.007)	-0.015^{*} (0.006)	0.025*** (0.007)
Decade 2000s	reference	reference	reference
Measure mean	0.003 (0.007)	-0.004 (0.007)	0.014 (0.008)
Measure job analysis	-0.014* (0.007)	0.007 (0.006)	-0.004 (0.007)
Measure mode	0.000 (0.009)	0.008 (0.008)	-0.010 (0.008)
Measure workers' assessment	reference	reference	reference
Fixed effects	-0.041*** (0.011)	-0.014 (0.010)	0.003 (0.011)
Instrumental variables	0.033 (0.018)	-0.084** (0.031)	-0.006 (0.035)
OLS	reference	reference	reference
Dummy male $= 1$	0.001 (0.007)	0.010 (0.007)	-0.013 (0.007)
Dummy female = 1	0.012 (0.008)	0.010 (0.007)	-0.017* (0.007)
Dummy both genders $= 1$	reference	reference	reference
Constant	0.089*** (0.008)	0.043*** (0.008)	-0.024*** (0.009)
Ν	151	151	143

Source: Database constructed from existing studies. Each column gives the results from a different meta-regression in which studies' estimates are regressed on study characteristics. Each estimate is weighted by the inverse of its variance. See Harbord and Higgins (2004) for details.

*/**/*** = statistically significant at the 10/5/1 percent level.

6. RELATION TO LABOR MARKET THEORIES

The incidence of over- and underschooling and the pattern of estimated returns to required, over-, and underschooling have been interpreted in terms of various labor market theories. In this section, we give a brief summary of these interpretations and the (sometimes) implicit assumptions made for these interpretations.

6.1. Human Capital

According to Mincer's wage equation, only the amount of attained schooling matters for wages. Because Duncan and Hoffman's wage equation nests Mincer's equation as a special case, it is straightforward to test whether the restrictions implied by Mincer's equation can be rejected. This is a test of the joint equality: $\delta_r = \delta_o = -\delta_u$. Without exception, this restriction is rejected by the data. This has been interpreted as evidence against a simple version of the human capital model in which wages are solely determined by attained level of schooling independent of the job to which a worker has been assigned.

The results are, however, consistent with slightly adapted versions of the human capital model. A first possibility is that being in a job that requires less schooling than available, provides an investment opportunity. This idea has been formalized by Sicherman and Galor (1990). We discuss this contribution in more detail in the next subsection.

Also consistent with the human capital model is that overschooling substitutes for other components in a person's stock of human capital, such as training, experience, and innate ability. The finding that overschooling is more prominent among younger workers is consistent with this. It is also consistent with the reported negative correlation between overschooling and measures of ability; see Section 4.

6.2. Career Mobility

Sicherman and Galor (1990) formalized the notion of career mobility as an explanation for the phenomenon of overeducation. They constructed a model of which the assumptions immediately imply that a worker with given innate ability may prefer to start in a job below his ability level if this is compensated by a higher probability to be promoted. According to Sicherman and Galor, this implies that "individuals may consider a lower-level firm in which the direct return to schooling is lower if in those firms, for a given level of schooling, the probability of promotion is higher." To test this implication of their model, the authors estimate for 24 different occupations, the effect of education on wages and the effect of education on the probability of upward career mobility. They find that the estimated effects are negatively correlated: occupations with a higher wage return to education have on average a smaller effect of education on upward mobility. This supports their model.

Sicherman and Galor (1990) also claim that their model implies that "it will be rational for some individuals to spend a portion of their working careers in occupations that require a lower level of schooling than they have acquired" (p. 177). This implication is further explored in Sicherman (1991). He regresses an indicator for upward mobility on indicators for overschooling and underschooling, controlling for *attained* schooling, experience (squared), and some other observables. The results show positive coefficients for the overschooling and underschooling indicators and a negative coefficient for attained schooling. Although Sicherman summarizes his finding with respect to the overschooling coefficient as indicating that "overeducated workers are more likely to move to a higher-level occupation than workers with the *required* level of schooling," the results from the specification that controls for experience (squared) suggest the opposite. Overeducated workers are more likely to move to a higher-level of *attained* schooling who are not over- or undereducated.²² This calls into question why

²² The coefficient for years of attained schooling equals -0.0676, and the coefficient of the overschooling dummy equals 0.2181. Hence, if someone is more than 3 ($\approx 0.2181/0.0676$) years overeducated, his probability to be promoted is smaller than that of the correctly educated workers in the same occupation.

the overeducated invested in schooling; they earn more than the not adequately educated in the same occupation, but at the same time also have a lower promotion probability.

Sicherman also finds that underschooled workers are more likely to be promoted than someone with the same level of schooling working in an occupation requiring this schooling level (with the effect being even bigger than that of the overschooling dummy). This puzzling finding can be explained if the underschooling is highly correlated with unobserved ability.

6.3. Job Competition

Thurow (1975) proposed a theory of job competition which stipulates that wages are solely determined by requirements of the job. Based on their attained schooling, workers are ranked in order of trainability, and the highest-ranked workers are assigned to the highest-ranked job. Attained schooling does, however, not affect earnings directly. This model implies the following testable restriction for Duncan and Hoffman's wage equation: $\delta o = \delta u = 0$. These equalities are almost always rejected by the data.

Job competition is closely related to the notion of crowding out of lower-educated workers by higher-educated workers during recessions. Gautier, J. van den Berg, C. van Ours, and Ridder (2002) find no support for this process. Only for one out of six job complexity levels they find that firms upgraded their workforce in low employment years. As far as crowding out takes place, it is more outflow driven than inflow driven. While this result seems to contradict Thurow's model, the study finds at the same time that at a given level of job complexity workers with relatively many years of schooling are not more productive than their direct colleagues. The authors explain the difference between this result and the results in the overeducation literature by the fact that their analysis takes account of match specific effects. This means that workers with relatively many years of schooling (compared with other workers at the same job level) select themselves into high-wage firms. This result is consistent with Thurow's model and points to the fact that estimation of Duncan and Hoffman's wage equation is not the correct approach to test Thurow's model.

6.4. Signaling/Screening

Since the early 1970s, the human capital model has been contested by the signaling hypothesis of Spence (1973). According to Spence's original signaling model, investments in schooling are efficient from an individual's point of view but do not affect a worker's productivity. Schooling does, however, still have a role in allocating the most productive workers to jobs in which they can be more productive. To the extent that this allocative role of schooling can be replaced by a less-expensive selection mechanim, there is overinvestment in schooling. This source of overinvestment in schooling has not been addressed in the overeducation literature.
6.5. Preferences

The consumptive value of education may explain why people acquire more schooling than is optimal from a more narrow perspective of maximizing lifetime earnings. Since people may differ in their taste for schooling, some will overinvest more than others. If some people extract disutility from attending school, this may explain why some people underinvest in schooling. To our knowledge, this issue has been largely ignored in the economics of education literature.²³

People may also differ in the degree to which they value leisure on the job. Jobs with lower skill requirements are likely to come with more leisure on the job for someone with a given skill level. The opposite would then hold for people working in a job above their acquired skill level. The significantly negative/positive impact of over/underschooling on job satisfaction reported by Hersch (1991), Korpi and Tahlin (2009), Verhaest and Omey (2006) contradict this explanation.

Gottschalk and Hansen (2003) develop a simple model with two sectors (the college and the noncollege sector) and two types of workers (college and noncollege graduates). Both sectors produce according to a production function in which the two types of workers are perfect substitutes. The productivity of college graduates is more similar to that of noncollege graduates in the noncollege sector than in the college sector. Workers are allowed to have heterogeneous preferences with regard to being employed in the two sectors. Some college graduates require a higher premium to work in the college sector than others. Gottschalk and Hansen show that in this framework, in equilibrium, some college workers voluntarily choose to work in the noncollege sector. Their overschooling (Gottschalk and Hansen avoid to use this term) does, in this model, not signal a misallocation of resources or an involuntary assignment across sectors.

6.6. Search and Frictions

The labor market does not operate frictionless, and over/underschooling can be the manifestations of frictions. Albrecht and Vroman (2002) proposed a matching model in which on-the-job search is not possible. They show that in that case, two types of equilibria can emerge: highly educated workers match with skilled and with unskilled jobs, or highly educated workers refuse to take unskilled jobs. Which of these equilibria emerges depends on the gap in productivity between skilled and unskilled jobs and on the share of high-skilled workers in the workforce.

Gautier (2002) allows high-skilled workers to search on the job. In that case, the degree to which high-skilled workers accept simple jobs depends on their relative productivity in such jobs and on their quit rate. Dolado, Jansen, and Jimeno (2009)

²³ Exceptions are Lazear (1977) and Oosterbeek and Van Ophem (2000). Using a Cobb-Douglas utility function $U = \ln N + \alpha \ln s$, the latter study finds that α is on average 1.1 and ranges between 0.24 and 1.68, implying that even those with the lowest taste for education extract positive utility from it.

analyze a model with on-the-job search very similar to that of Gautier, but allow wages to depend on aggregate labor market outcomes. Highly educated workers may end up in unskilled jobs for which they are overqualified but are allowed to engage in on-thejob search on pursuit of a better job. Skill mismatch has in that case a transitory nature. This job-acceptance rule is very similar in spirit to the one proposed by Sicherman and Galor (1990). Here too, we may argue that the empirical regularity that overschooling is for many persistent and is at odds with these models.

7. CONCLUSION

This chapter surveyed the economics literature on overschooling. The initial motivation to study this topic was concerns about a strong increase in the number of college graduates in the early 1970s in the United States and a concurrent decrease in the returns to college education. These concerns were fueled by Freeman's book "The Overeducated American." We reconstruct that fresh evidence that was published shortly after Freeman's book should have been sufficient to temper the concerns. Yet, not too long after, Duncan and Hoffman published their seminal article that started the subfield of the economics of overeducation.

The main vehicle in this literature is an extended version of Mincer's wage equation where actual years of schooling is broken down in years of overschooling, years of required schooling, and years of underschooling. The additional estimates that this specification produces are not very useful from the perspective of the decision to invest in schooling, neither from an individual nor from an aggregate perspective. From both perspectives, it is only the expected (private or social) return on completed schooling that counts. At best are the separate estimates informative about the riskiness of investments in schooling.

The efficiency implications of mismatch are interesting and to our knowledge nobody has pursued them to date. Potentially, estimates of returns to required schooling, overschooling, and underschooling are informative about the costs of mismatch. Such estimates, together with information about the assignment of workers (by completed schooling) to jobs (by required schooling), could allow us to calculate by how much productivity (approximated by the wage sum) could increase by reallocating workers to jobs that require their schooling. Unfortunately, the estimates that have been produced are not suitable for such an exercise because the literature has not been able to separate the impact of mismatch from unobserved ability. Measurement error, which is already an issue for the estimation of the return to completed schooling, is probably much more problematic in the context of required schooling and over/underschooling and poses an additional hurdle to the consistent estimation of the effects of interest.

The micro overschooling literature was born out of an additional variable—required schooling—in the PSID. From that, Duncan and Hoffman could construct measures of

overschooling and underschooling at an individual level, which they could then loosely connect to concerns about overschooling at an aggregate level. From this start onwards, the overschooling literature has lacked substantive hypotheses. Instead, results from Duncan and Hoffman's wage equation have been loosely interpreted in terms of existing theories. But we are reluctant to attach too much weight to these interpretations given that it is unclear to what extent the estimates are reliable.

Our overall conclusion regarding the contributions of the overschooling literature to our understanding of the functioning of labor markets and investment in education is rather pessimistic. The literature has produced dozens of estimates of the returns to required schooling, overschooling, and underschooling. But although some studies express awareness that omitted variable (ability) bias and measurement error may be a concern, these issues have not been addressed in a satisfying way. As a result, the estimates that have been reported cannot be interpreted as causal. Consequently, the estimates are consistent with the plethora of interpretations reviewed in Section 6. Our own reading of the evidence is that omitted variable bias is substantial and possibly explains the entire difference between returns to required schooling and overschooling and underschooling.

Ideally, a review like this one solves existing problems or at least provides a road map as how to proceed. Unfortunately, this chapter does not live up to that ideal. The overeducation/mismatch literature has for too long led a separate life of modern labor economics and the economics of education. We conclude that the conceptional measurement of overeducation has not been resolved, omitted variable bias and measurement error are too serious to be ignored, and that substantive economic questions have not been rigorously addressed. Hence, new contributions seem only worthwhile if they include a serious attempt to tackle (at least one of) these issues.

REFERENCES

- Aigner, D., 1973. Regression with a Binary Independent Variable Subject to Errors of Observation. J. Econom. 1 (1), 49–50.
- Alba-Ramirez, A., 1993. Mismatch in the Spanish Labor Market: Overeducation? J. Hum. Resour. 28 (2), 259–278.
- Albrecht, J., Vroman, S., 2002. A Matching Model with Endogenous Skill Requirements. Int. Econ. Rev. 43, 283–305.
- Allen, J., Van der Velden, R., 2001. Educational Mismatches versus Skill Mismatches: Effects on Wages, Job Satisfaction, and On-the-Job Search. Oxford Econ. Pap. 53 (3), 434–452.
- Ashenfelter, O., Harmon, C., Oosterbeek, H., 1999. A Review of Estimates of the Schooling/Earnings Relationship, with Tests for Publication Bias. Labour Econ. 6, 453–470.
- Autor, D., Katz, L.F., 1999. Changes in the Wage Structure and Earnings Inequality. In: Ashenfelter, O., Card, D. (Eds.), Handbook of Labor Economics, Volume 3A. Elsevier-North Holland, Amsterdam, pp. 1463–1555.
- Battu, H., Belfield, C., Sloane, P., 1999. Overeducation among Graduates: A Cohort View. Educ. Econ. 7 (1), 21–38.
- Battu, H., Sloane, P., Building, E., Street, D., Park, S., 2004. Over-Education and Ethnic Minorities in Britain. Manchester Sch. 72 (4), 535–559.

 Bauer, T., 2002. Educational Mismatch and Wages: A Panel Analysis. Econ. Educ. Rev. 21 (3), 221–229.
 Beneito, P., Ferri, J., Molto, M.L., Uriel, E., 1996. Over/undereducation and Specific Training in Spain: Complementary or Substitute Components of Human Capital? Palgrave.

Berg, I., 1970. Education and Jobs: The Great Training Robbery. Praeger.

- Black, D., Berger, M., Scott, F., 2000. Bounding Parameter Estimates with Nonclassical Measurement Error. J. Am. Stat. Assoc. 95 (451), 739–747.
- Büchel, F., Pollmann-Schult, M. 2001. Overeducation and Skill Endowments. The Role of School Achievement and Vocational Training Quality. IZA Discussion Paper No. 337, Bonn.
- Büchel, F., van Ham, M., 2003. Overeducation, Regional Labor Markets, and Spatial Flexibility. J. Urban Econ. 53 (3), 482–493.
- Card, D., 1999. The Causal Effect of Education and Earnings. In: Ashenfelter, O., Card, D. (Eds.), Handbook of Labor Economics. chap. 30, vol. 3, part 1, pp. 1801–1863. North Holland.
- CEDEFOP, 2010. The Skill Matching Challenge; Analysing Skill Mismatch and Policy Implications. Technical Report, CEDEFOP.
- Chevalier, A., 2003. Measuring Over-Education. Economica 70 (279), 509-531.
- Chevalier, A., Lindley, J., 2009. Overeducation and the Skills of UK Graduates. J. R. Stat. Soc. Ser. A 172, 307–337.
- Clogg, C., 1979. Measuring Underemployment: Demographic Indicators for the United States. Academic Press, New York.
- Cohn, E., 1992. The Impact of Surplus Schooling on Earnings: Comment. J. Hum. Resour. 27, 679.
- Cohn, E., Khan, S., 1995. The Wage Effects of Overschooling Revisited. Labour Econ. 2 (1), 67-76.
- Cohn, E., Ng, Y., 2000. Incidence and Wage Effects of Overschooling and Underschooling in Hong Kong. Econ. Educ. Rev. 19 (2), 159–168.
- Daly, M., Büchel, F., Duncan, G., 2000. Premiums and Penalties for Surplus and Deficit Education Evidence from the United States and Germany. Econ. Educ. Rev. 19 (2), 169–178.
- Dolado, J.J., Jansen, M., Jimeno, J.F., 2009. On the Job-Search in a Matching Model with Heterogeneous Jobs and Workers. Econ. J. 119, 200–228.
- Dolton, P., Vignoles, A., 2000. The Incidence and Effects of Overeducation in the U.K. Graduate Labor Market. Econ. Educ. Rev. 19, 179–198.
- Dolton, P.J., Silles, M.A., 2008. The Effects of Over-Education on Earnings in the Graduate Labour Market. Econ. Educ. Rev. 27 (2), 125–139.
- Duncan, G., Hoffman, S., 1981. The Incidence and Wage Effects of Overeducation. Econ. Educ. Rev. 1 (1), 75–86.
- Eckaus, R., 1964. Economic Criteria for Education and Training. Rev. Econ. Stat. 46, 181-190.
- Frank, R., 1978. Why Women Earn Less: the Theory and Estimation of Differential Overqualification. Am. Econ. Rev. 68 (3), 360–373.
- Freeman, R., 1975. Overinvestment in College Training? J. Hum. Resour. 10, 287-311.
- Freeman, R., 1976. The Overeducated American. Academic Press, New York.
- Frenette, M., 2004. The Overqualified Canadian Graduate: the Role of the Academic Program in the Incidence, Persistence, and Economic Returns to Overqualification. Econ. Educ. Rev. 23 (1), 29–45.
- Galasi, P., 2008. The Effect of Educational Mismatch on Wages for 25 Countries. Budapest Working Papers on the Labour Market BWP—2008/8.
- Gautier, P., 2002. Unemployment and Search Externalities in a Model with Heterogeneous Jobs and Workers. Economica 69, 21–40.
- Gautier, P.J., van den Berg, G.C., van Ours, J., Ridder, G., 2002. Worker Turnover at the Firm Level and Crowding Out of Lower Educated Workers. Eur. Econ. Rev. 46 (3), 523–538.
- Gill, A.M., Solberg, E.J., 1992. Surplus Schooling and Earning: a Critique. J. Hum. Resour. 27, 683–689. Gladwell, M., 2008. Outliers; The Story of Success. Little, Brown and Company, New York.
- Gottschalk, P., Hansen, M., 2003. Is the Proportion of College Workers in Noncollege Jobs Increasing? J. Labor Econ. 21, 449–471.
- Green, C., Kler, P., Leeves, G., 2007. Immigrant Overeducation: Evidence from Recent Arrivals to Australia. Econ. Educ. Rev. 26 (4), 420–432.

- Green, F., McIntosh, S., 2007. Is There a Genuine Under-Utilization of Skills among the Over-Qualified? Appl. Econ. 39, 427–439.
- Green, F., McIntosh, S., Vignoles, A., 1999. Overeducation' and Skills-Clarifying the Concepts. Technical Report.
- Green, F., Zhu, Y., 2010. Overqualification, Job Dissatisfaction, and Increasing Dispersion in the Returns to Graduate Education. Oxf. Econ. Pap. 62 (4), 740–763.
- Groeneveld, S., Hartog, J., 2004. Overeducation, Wages and Promotions within the Firm. Labour Econ. 11 (6), 701–714.
- Groot, W., 1993. Overeducation and the Returns to Enterprise Related Schooling. Econ. Educ. Rev. 12, 299–309.
- Groot, W., 1996. The Incidence of, and Returns to Overeducation in the UK. Appl. Econ. 28 (10), 1345-1350.
- Groot, W., Maassen van den Brink, H., 1997. Allocation and the Returns to Over-Education in the UK. Educ. Econ. 5 (2), 169–183.
- Groot, W., Maassen van den Brink, H., 2000. Overeducation in the Labor Market: A Meta-Analysis. Econ. Educ. Rev. 19 (2), 149–158.
- Halaby, C., 1994. Overeducation and Skill Mismatch. Sociology Educ. 67 (1), 47-59.
- Hammack, F.M., 1978. Review: The "Value" of a College Education. Sch. Rev. 86, 669-679.
- Harbord, R., Higgins, J., 2004. Metareg: Stata Module to Perform Meta-Analysis Regression. Statistical Software Components, Boston College Department of Economics, Boston.
- Hartog, J., 2000. Over-Education and Earnings: Where Are We, Where Should We Go? Econ. Educ. Rev. 19 (2), 131–147.
- Hartog, J., Oosterbeek, H., 1988. Education, Allocation and Earnings in the Netherlands: Overschooling? Econ. Educ. Rev. 7 (2), 185–194.
- Hersch, J., 1991. Education Match and Job Match. Rev. Econ. Stat. 73 (1), 140-144.
- Hung, C.-Y., 2008. Overeducation and Undereducation in Taiwan. J. Asian Econ. 19, 125-137.
- Kiker, B., Santos, M., de Oliveira, M., 1997. Overeducation and Undereducation: Evidence for Portugal. Econ. Educ. Rev. 16 (2), 111–125.
- Korpi, T., Tahlin, M., 2009. Educational Mismatch, Wages, and Wage Growth: Overeducation in Sweden, 1974-2000. Labour Econ. 16, 183–193.
- Lazear, E.P., 1977. Education: Consumption or Production? J. Polit. Economy 85, 569-598.
- Levin, H.M., 1977. Book Review: The Overeducated American by Richard B. Freeman. Harv. Educ. Rev. 47, 226–231.
- Lindley, J., McIntosh, S., 2008. A Panel Data Analysis of the Incidence and Impact of Over-Education. Technical Report, University of Sheffield.
- Mavromaras, K., McGuinness, S., Fok, Y.K., 2009. Overskilling Dynamics and Education Pathways. Discussion Paper 4321, IZA.
- McGoldrick, K., Robst, J., 1996. Gender Differences in Overeducation: A Test of the Theory of Differential Overqualification. Am. Econ. Rev. 86, 280–284.
- McGuinness, S., 2006. Overeducation in the Labour Market. J. Econ. Surv. 20 (3), 387–418.
- McGuinness, S., 2007. How Biased Are the Estimated Wage Impacts of Overeducation? A Propensity Score Matching Approach. Appl. Econ. Lett. 15, 145–149.
- Meghir, C., Rivkin, S.G., 2011. Econometric Methods for Research in Education. In: Hanushek, E., Machin, S., Woessmann, L. (Eds.), Handbook of the Economics of Education. Elsevier-North Holland, Amsterdam, pp. 1–87.
- Ng, Y., 2001. Overeducation and Undereducation and Their Effect on Earnings: Evidence from Hong Kong, 1986-1996. Pacific Econ. Rev. 6 (3), 401–418.
- Oosterbeek, H., Van Ophem, H., 2000. Schooling Choices: Preferences, Discount Rates, and Rates of Return. Empirical Econ. 25, 15–34.
- Oosterbeek, H., Webbink, D., 1996. Over scholing, overscholing en inkomen. Econ. Stat. Ber. 81, 240-241.
- Quinn, M., Rubb, S., 2006. Mexico's Labor Market: The Importance of Education-Occupation Matching on Wages and Productivity in Developing Countries. Econ. Educ. Rev. 25, 147–156.

- Raudenbush, S., 1994. Random Effects Models. In: Cooper, H., Hedges, L. (Eds.), The Handbook of Research Synthesis. Russell Sage Foundation, New York, pp. 301–321.
- Robst, J., 1994. Measurement Error and the Returns to Excess Schooling. Appl. Econ. Lett. 1, 142-144.
- Rubb, S., 2003a. Overeducation: A Short or Long Run Phenomenon for Individuals? Econ. Educ. Rev. 22 (4), 389–394.
- Rubb, S., 2003b. Post-college Schooling, Overeducation, and Hourly Earnings in the United States. Educ. Econ. 11 (1), 53–72.
- Rumberger, R., 1987. The Impact of Surplus Schooling on Productivity and Earnings. J. Hum. Resour. 22 (1), 24–50.
- Santos, M., 1995. Education and Earnings Differentials in Portugal. Ph.D. Thesis, Porto University, Faculty of Economics, Porto.
- Schøne, P., Hardoy, I., 2009. Returns to Pre-Immigration Education for Nonwestern Immigrants: Why so low? Working Paper.

Sicherman, N., 1991. Overeducation" in the Labor Market. J. Labor Econ. 9 (2), 101.

- Sicherman, N., Galor, O., 1990. A Theory of Career Mobility. J. Polit. Economy 98 (1), 169–192.
- Sloane, P., 2003. Much Ado About Nothing? What Does the Overeducation Literature Really Tell Us. In: Büchel, F., De Grip, A., Mertens, A. (Eds.), Overeducation in Europe; Current Issues in Theory and Policy, Edward Elgar, Cheltenham, pp. 11–45.
- Slonimczyk, F., 2008. Skill Mismatch and Earnings: A Panel Analysis of the U.S. Labor Market. Working Paper.
- Smith, J., Welch, F., 1978. The Overeducated American? A Review Article. Technical Report, Rand Corporation.
- Spence, M., 1973. Job Market Signalling. Q. J. Econ. 87, 355-374.
- Sullivan, T., 1978. Marginal Workers, Marginal Jobs: the Underutilization of American Workers. University of Texas Press, Austin.
- Thurow, L., 1975. Generating Inequality. Basic Books, New York.
- Tsai, Y., 2010. Returns to Overeducation: A Longitudinal Analysis of the U.S. Labor Market. Econ. Educ. Rev. 29 (4), 606–617.
- Vahey, S., 2000. The Great Canadian Training Robbery: Evidence on the Returns to Educational Mismatch. Econ. Educ. Rev. 19 (2), 219–227.
- Van der Meer, P., 2006. The Validity of Two Education Requirement Measures. Econ. Educ. Rev. 25, 211–219.
- Van der Meer, P., 2009. Investments in Education: Too Much or Not Enough? Econ. Lett. 102, 195–197.
- Van der Velden, R., Van Smoorenburg, M., 2000. The Measurement of Overeducation and Undereducation: Self-Report vs Job-Analyst Method. Les Cahiers de Lamas.
- Van Ewijk, R., Sleegers, P., 2010. The Effect of Peer Socioeconomic Status on Student Achievement: A Meta-Analysis. Educ. Res. Rev. 5, 134–150.
- Van Smoorenburg, M., Van der Velden, R., 2000. Overeducation and the Effects on Training and Wages for School-Leavers. Research Centre for Education and the Labour Market.
- Verdugo, R., Verdugo, N., 1989. The Impact of Surplus Schooling on Earnings: Some Additional Findings. J. Hum. Resour. 24 (4), 629–643.
- Verhaest, D., Omey, E., 2006. The Impact of Overeducation and Its Measurement. Soc. Indic. Res. 77 (3), 419–448.
- Vieira, J., 1999. Evolution of Wage Structures in Portugal 1986-1992. Ph.D. Thesis, Tinbergen Institute, University of Amsterdam, Amsterdam.

Authors	Pubyr	Country	Gender	N	Measure S ^r	Share $S^a > S^r$	Share $S^a < S^r$	δ_r	s.e. (δ_r)	δ_o	s.e. (δ_o)	δ_u	s.e. (δ_u)
Duncan and Hoffman	1981	US	m	2034	wa	0.14	0.42	0.063	0.004	0.029	0.006	-0.042	0.011
Duncan and Hoffman	1981	US	f	1187	wa	0.09	0.41	0.091	0.005	0.052	0.008	-0.014	0.020
Duncan and Hoffman	1981	US	m	798	wa	0.12	0.49	0.076	0.007	0.040	0.008	-0.048	0.013
Duncan and Hoffman	1981	US	f	662	wa	0.11	0.43	0.105	0.007	0.047	0.010	-0.038	0.024
Rumberger	1987	US	m	669	wa			0.052	0.027	0.031	0.016		
Rumberger	1987	US	m	669	ja			0.061	0.031	0.028	0.014		
Rumberger	1987	US	f	368	wa			0.100	0.051	0.057	0.029		
Rumberger	1987	US	f	368	ja			0.115	0.059	0.061	0.031		
Rumberger	1987	US	fm		wa		0.35						
Rumberger	1987	US	fm	1037	wa		0.27						
Rumberger	1987	US	fm		wa		0.32						
Rumberger	1987	US	fm	1037	ja		0.57						
Hartog and Oosterbeek	1988	NL	fm	540	wa	0.22	0.16	0.071	0.005	0.057	0.007	-0.025	0.008
Hartog and Oosterbeek	1988	NL	m	394	wa			0.076	0.005	0.065	0.009	-0.019	0.010
Hartog and Oosterbeek	1988	NL	f	140	wa			0.052	0.009	0.037	0.012	-0.040	0.016
Hartog and Tsang	1989	US	fm		wa			0.095		0.062		-0.039	
Hartog and Tsang	1989	US	fm		wa			0.078		0.044		-0.021	
Hartog and Tsang	1989	US	fm		wa			0.089		0.051		-0.013	
Hersch	1991	US	m	414	wa	0.14	0.54	0.061	0.008	0.023	0.008	-0.006	0.021
Hersch	1991	US	f	213	wa	0.21	0.45	0.064	0.012	0.022	0.014	-0.035	0.027
Hersch	1991	US	m	414	wa	0.21	0.45	0.066	0.063	-0.315	0.061	0.051	0.165
Hersch	1991	US	f	213	wa	0.21	0.45	0.206	0.107	-0.369	0.119	-0.466	0.227
Sicherman	1991	US	m	3133	wa	0.16	0.41	0.048	0.003	0.039	0.004	-0.017	0.005
Alba-Ramirez	1993	ES	fm	11597	wa	0.23	0.17	0.092	0.001	0.040	0.003	-0.060	0.002
Alba-Ramirez	1993	ES	fm	11597	wa	0.23	0.17	0.058	0.002	0.027	0.002	-0.047	0.002
Groot	1993	NL	fm	1057	mean	0.16	0.16	0.055		-0.019		-0.029	
Robst	1994	US	m	2648	ja	0.09	0.54	0.07	0.004	0.047	0.005	-0.026	0.006
Robst	1994	US	m	2648	ja	0.09	0.54	0.09	0.009	-0.027	0.021	-0.016	0.020
Robst	1994	US	m	2648	wa	0.20	0.32	0.084	0.006	0.049	0.005	-0.042	0.006
Robst	1994	US	m	2648	wa	0.20	0.32	0.088	0.007	-0.018	0.018	-0.032	0.022
Cohn and Khan	1995	US	m	3588	mean	0.12	0.13	0.084	0.005	0.059	0.020	-0.044	0.015
Cohn and Khan	1995	US	m	3133	wa	0.20	0.33	0.048	0.003	0.039	0.004	-0.017	0.005
Cohn and Khan	1995	US	m	3898	wa	0.20	0.33	0.082	0.004	0.059	0.004	-0.039	0.005
Cohn and Khan	1995	US	m	3380	wa	0.20	0.33	0.077	0.004	0.049	0.005	-0.038	0.006

Authors	Pubyr	Country	Gender	Ν	Measure S ^r	Share $S^a > S^r$	Share $S^a < S^r$	δ_r	s.e. (δ_r)	δ_o	s.e. (δ_o)	δ_u	s.e. (δ_u)
Santos	1995	PT	m		mode			0.060		0.053		-0.038	
Santos	1995	PT	f		mode			0.060		0.044		-0.034	
Santos	1995	PT	m		ja			0.083		0.023		-0.037	
Santos	1995	PT	f		ja			0.081		0.019		-0.051	
Allen and Van der	1995				-								
Velden													
Beneito et al.	1996	ES	fm		wa			0.033		0.019		-0.033	
Groot	1996	UK	m	2027	mean	0.10	0.13	0.079	0.004	-0.026	0.007	0.122	0.035
Groot	1996	UK	f	2099	mean	0.08	0.10	0.094	0.004	-0.030	0.007	0.109	0.031
Groot	1996	UK	fm	4126	mean	0.09	0.11	0.088	0.003	-0.028	0.005	0.119	0.024
Oosterbeek and Webbink	1996	NL	m	986	ja			0.092	0.000	0.052	0.009	-0.033	-0.008
Oosterbeek and Webbink	1996	NL	f	740	ja			0.079	0.008	0.063	0.012	-0.011	-0.014
Halaby	1994												
McGoldrick and Robst	1996	US	f		wa		0.32						
McGoldrick and Robst	1996	US	m		wa		0.30						
McGoldrick and Robst	1996	US	f		ia		0.50						
McGoldrick and Robst	1996	US	m		ja		0.50						
Kiker et al	1007	DT	fm	30366	ja mode	0.17	0.36	0.083	0.032	0.062	0.024	-0.053	0.021
Kiker et al	1997	PT	fm	30366	wa	0.38	0.33	0.005	0.039	0.035	0.024	-0.051	0.021
Kiker et al	1007	DT	f	11130	mode	0.19	0.35	0.100	0.030	0.054	0.021	-0.052	0.020
Kiker et al	1997	PT	f	11130	mode	0.15	0.23	0.078	0.039	0.029	0.021	-0.052	0.020
Kiker et al	1007	DT	1	10206	wa	0.16	0.46	0.086	0.033	0.027	0.024	-0.054	0.022
Kiker et al	1997	PT	m	19206	wa	0.44	0.20	0.000	0.038	0.005	0.024	-0.034	0.021
Groot and Massen van	1997	NI	111	17200	mean	0.11	0.27	0.061	0.050	0.028	0.015	-0.026	0.017
den Brink	1777	IL			mean			0.001		0.020		0.020	
Vieira	1999	РТ	fm		ia	0.48	0.24	0.069	0.027	0.039	0.015	-0.030	0.012
Vieira	1999	PT	m		ja	0.53	0.22	0.072	0.028	0.039	0.015	-0.032	0.012
Vieira	1999	PT	f		ja	0.32	0.28	0.054	0.021	0.031	0.012	-0.025	0.010
Vieira	1999	PT	fm		ja	0.44	0.26	0.070	0.027	0.041	0.012	-0.031	0.012
Vieira	1999	PT	m		ja	0.50	0.24	0.073	0.028	0.041	0.016	-0.032	0.012
Vieira	1999	PT	f		ja ia	0.31	0.32	0.058	0.023	0.033	0.013	-0.025	0.010
Vieira	1999	PT	fm		j" ia	0.38	0.33	0.084	0.033	0.045	0.018	-0.037	0.014
Vieira	1999	PT	m		j ia	0.45	0.29	0.085	0.033	0.041	0.016	-0.037	0.015
					5			0.000	5.000	0.0.1	5.010	0.007	0.010

Authors	Pubyr	Country	Gender	Ν	Measure S ^r	Share $S^a > S^r$	Share $S^a < S^r$	δ_r	s.e. (δ_r)	δ_o	s.e. (δ_o)	δ_u	s.e. (δ_u)
Vieira	1999	РТ	f		ja	0.26	0.41	0.077	0.030	0.047	0.018	-0.035	0.014
Green et al.	1999	UK	fm	4025	wa	0.20	0.29						
Green et al.	1999	UK	fm	2482	wa	0.20	0.32						
Cohn and Ng	2000	HK	m	166649	mean	0.28	0.38	0.110	0.000	0.040	0.001	-0.040	0.001
Cohn and Ng	2000	HK	f	102694	mean	0.24	0.32	0.090	0.001	0.050	0.001	-0.060	0.001
Cohn and Ng	2000	HK	m	179889	mean	0.28	0.37	0.130	0.000	0.050	0.001	-0.040	0.001
Cohn and Ng	2000	HK	f	120777	mean	0.25	0.31	0.150	0.001	0.040	0.001	-0.050	0.001
Cohn, Johnson and Ng	2000	US	m		mode			0.140		0.110		-0.095	
Cohn, Johnson and Ng	2000	US	f		mode			0.160		0.110		-0.095	
Cohn, Johnson and Ng	2000	US	m		mode			0.140		0.090		-0.090	
Cohn, Johnson and Ng	2000	US	f		mode			0.180		0.080		-0.090	
Daly et al.	2000	US	m	1784	wa	0.16	0.39	0.061	0.004	0.045	0.005	-0.034	0.009
Daly et al.	2000	US	f	1119	wa	0.11	0.37	0.090	0.005	0.061	0.007	-0.036	0.016
Daly et al.	2000	US	m	2138	wa	0.21	0.32	0.078	0.004	0.054	0.006	-0.016	0.008
Daly et al.	2000	US	f	1726	wa	0.17	0.34	0.109	0.005	0.086	0.007	-0.025	0.011
Daly et al.	2000	DE	m	2035	wa	0.07	0.14	0.090	0.003	0.049	0.008	-0.078	0.014
Daly et al.	2000	DE	f	1031	wa	0.07	0.21	0.090	0.005	0.066	0.008	-0.038	0.022
Van der Velden and	2000	NL	fm	56300	wa	0.07	0.23	0.067	0.002	0.047	0.002	-0.070	0.006
Van Smoorenburg													
Van der Velden and	2000	NL	fm	57900	ja	0.10	0.65	0.067	0.002	0.054	0.002	-0.091	0.005
Van Smoorenburg					-								
Van Smoorenburg and	2000	NL	fm	2917	wa	0.02	0.40	0.076	0.010	0.055	0.010	-0.055	0.020
Van der Velden													
Van Smoorenburg and	2000	NL	fm	2755	wa	0.02	0.40	0.066	0.010	0.048	0.010	-0.071	0.020
Van der Velden													
Van Smoorenburg and	2000	NL	fm	2782	wa	0.02	0.40	0.061	0.010	0.044	0.010	-0.043	0.020
Van der Velden													
Vahey	2000	CA	m	569	wa	0.24	0.3						
Vahey	2000	CA	f	424	wa	0.17	0.32						
Ng	2001	HK	m	0	mean	0.12	0.14	0.140	0.001	0.080	0.002	-0.020	0.001
Ng	2001	HK	f	0	mean	0.18	0.09	0.130	0.001	0.080	0.003	-0.040	0.001
Ng	2001	HK	m	0	mode	0.28	0.38	0.110	0.000	0.040	0.001	-0.040	0.001
Ng	2001	HK	f	0	mode	0.24	0.32	0.090	0.001	0.050	0.001	-0.060	0.001
Ng	2001	HK	m	569	mean	0.12	0.14	0.140	0.000	0.060	0.002	-0.010	0.001
Ng	2001	HK	f	424	mean	0.13	0.14	0.170	0.001	0.040	0.002	-0.020	0.001
Ng	2001	HK	m	569	mode	0.28	0.37	0.130	0.000	0.050	0.001	-0.040	-0.001

Authors	Pubyr	Country	Gender	Ν	Measure S ^r	Share $S^a > S^r$	Share $S^a < S^r$	δ_r	s.e. (δ_r)	δ_o	s.e. (δ_o)	δ_u	s.e. (δ_u)
Ng	2001	HK	f	424	mode	0.25	0.31	0.150	0.001	0.040	0.001	-0.050	0.001
Ng	2001	HK	m	211712	mean	0.14	0.14	0.150	0.000	0.060	0.002	-0.020	0.002
Ng	2001	HK	f	154864	mean	0.13	0.14	0.210	0.001	0.050	0.002	-0.040	0.001
Ng	2001	HK	m	211712	mode	0.35	0.28	0.150	0.000	0.050	0.001	-0.050	0.001
Ng	2001	HK	f	154864	mode	0.34	0.23	0.180	0.001	0.040	0.001	-0.070	0.001
Bauer	2002	DE	m	13364	mean	0.10	0.12	0.107	0.003	0.090	0.009	-0.100	0.016
Bauer	2002	DE	f	5273	mean	0.16	0.11	0.125	0.007	0.052	0.023	-0.115	0.018
Bauer	2002	DE	m	13364	mode	0.21	0.31	0.084	0.003	0.066	0.004	-0.050	0.005
Bauer	2002	DE	f	5273	mode	0.37	0.30	0.087	0.006	0.045	0.007	-0.070	0.006
Bauer	2002	DE	m	13364	mean	0.10	0.12	0.069	0.002	0.062	0.005	-0.087	0.010
Bauer	2002	DE	f	5273	mean	0.16	0.11	0.060	0.004	0.045	0.008	-0.063	0.012
Bauer	2002	DE	m	13364	mode	0.21	0.31	0.071	0.002	0.063	0.002	-0.065	0.003
Bauer	2002	DE	f	5273	mode	0.37	0.30	0.064	0.004	0.057	0.004	-0.062	0.004
Bauer	2002	DE	m	1824	mean	0.10	0.12	0.012	0.004	-0.010	0.007	-0.013	0.011
Bauer	2002	DE	f	922	mean	0.16	0.11	-0.005	0.006	-0.019	0.011	0.022	0.016
Bauer	2002	DE	m	1824	mode	0.21	0.31	0.092	0.005	0.086	0.005	-0.092	0.005
Bauer	2002	DE	f	922	mode	0.37	0.30	0.085	0.007	0.084	0.007	-0.088	0.007
Rubb and Quinn	2002	MX	m		mean			0.065		0.031		-0.055	
Rubb and Quinn	2002	MX	m		mode			0.058		0.032		-0.011	
Rubb	2003	US	m	417090	mean	0.10	0.13	0.133	0.001	0.054	0.000		
Rubb	2003	US	f	303071	mean	0.11	0.09	0.150	0.001	0.054	0.004		
Büchel and van Ham	2003	DE	fm	5143	wa		0.19						
Groeneveld and Hartog	2004	NL	fm	2782	firm	0.15	0.19	0.087	0.001	0.025	0.002	-0.003	0.002
Verhaest and Omey	2006	BE	fm	2436	direct wa	0.10	0.33	0.041	0.004	0.028	0.004	-0.020	0.003
Verhaest and Omey	2006	BE	fm	2436	wa	0.03	0.42	0.033	0.003	0.021	0.004	-0.028	0.009
Verhaest and Omey	2006	BE	fm	2436	ja	0.05	0.52	0.042	0.003	0.022	0.004	-0.016	0.007
Verhaest and Omey	2006	BE	fm	2436	mode	0.18	0.14	0.064	0.005	0.030	0.007	-0.016	0.004
Verhaest and Omey	2006	BE	fm	2282	direct wa	0.10	0.33	0.032	0.016	-0.039	0.018	-0.033	0.024
Verhaest and Omey	2006	BE	fm	2282	wa	0.03	0.42	0.023	0.014	-0.030	0.016	0.028	0.041
Verhaest and Omey	2006	BE	fm	2282	ja	0.05	0.52	0.032	0.015	-0.058	0.018	-0.073	0.032
Verhaest and Omey	2006	BE	fm	2282	mode	0.18	0.14	0.067	0.016	-0.042	0.028	-0.009	0.018
Van der Meer	2006	NL	fm	2248	ja	0.38	0.28	0.062	0.002	0.042	0.003	-0.036	0.003
Van der Meer	2006	NL	fm	2248	ja	0.26	0.21	0.071	0.002	0.041	0.004	-0.028	0.003
Van der Meer	2006	NL	fm	2431	ja	0.37	0.30	0.059	0.002	0.038	0.004	-0.037	0.003
Van der Meer	2006	NL	fm	2431	ja	0.25	0.24	0.068	0.002	0.035	0.004	-0.029	0.003

Authors	Pubyr	Country	Gender	Ν	Measure S ^r	Share $S^a > S^r$	Share $S^a < S^r$	δ_r	s.e. (δ_r)	δ_o	s.e. (δ_o)	δ_u	s.e. (δ_u)
Van der Meer	2006	NL	fm	2473	ja	0.33	0.33	0.068	0.002	0.047	0.004	-0.044	0.004
Van der Meer	2006	NL	fm	2473	ja	0.14	0.28	0.075	0.003	0.047	0.004	-0.036	0.005
Quinn and Rubb	2006	MX	m	1442	mean	0.26	0.11						
Quinn and Rubb	2006	MX	m	1442	mode	0.36	0.30						
Quinn and Rubb	2006	MX	m	1442	quinnRubb	0.14	0.11						
Quinn and Rubb	2006	MX	m	1092	mean	0.17	0.18						
Quinn and Rubb	2006	MX	m	1092	mode	0.34	0.43						
Quinn and Rubb	2006	MX	m	1092	quinnRubb	0.12	0.16						
Quinn and Rubb	2006	MX	m	1288	mean	0.16	0.20						
Quinn and Rubb	2006	MX	m	1288	mode	0.28	0.42						
Quinn and Rubb	2006	MX	m	1288	quinnRubb	0.14	0.16						
Quinn and Rubb	2006	MX	m	1123	mean	0.17	0.21						
Quinn and Rubb	2006	MX	m	1123	mode	0.25	0.48						
Quinn and Rubb	2006	MX	m	1123	quinnRubb	0.15	0.12						
Quinn and Rubb	2006	MX	m	4945	mean	0.19	0.17	0.085	0.005	0.043	0.008	-0.030	0.010
Quinn and Rubb	2006	MX	m	4945	mode	0.31	0.40	0.076	0.005	0.048	0.007	-0.036	0.008
Quinn and Rubb	2006	MX	m	4945	quinnRubb	0.14	0.14	0.090	0.006	0.049	0.008	-0.039	0.010
Green et al.	2007	AU	m	8000000	ja		0.07	0.100	0.000	0.050	0.000		
Green et al.	2007	AU	m	8000000	ja		0.09	0.110	0.000	0.050	0.000		
Hung	2008	TW	fm	1606	wa	0.12	0.46	0.086	0.004	0.063	0.006	-0.069	0.011
Hung	2008	TW	fm	1606	mean	0.14	0.17	0.111	0.006	0.065	0.009	-0.058	0.008
Galasi	2008	AT	fm	415	wa	0.35	0.47	0.077	0.011	0.031	0.015	-0.029	0.028
Galasi	2008	BE	fm	519	wa	0.62	0.25	0.078	0.016	0.019	0.017	-0.008	0.022
Galasi	2008	CZ	fm	544	wa	0.44	0.50	0.076	0.010	0.072	0.019	-0.007	0.029
Galasi	2008	DK	fm	672	wa	0.39	0.53	0.063	0.010	0.013	0.012	-0.032	0.017
Galasi	2008	EE	fm	722	wa	0.13	0.79	0.136	0.010	0.049	0.011	-0.030	0.030
Galasi	2008	FI	fm	787	wa	0.39	0.53	0.087	0.005	0.011	0.009	-0.031	0.011
Galasi	2008	FR	fm	641	wa	0.65	0.27	0.129	0.012	0.017	0.025	-0.056	0.011
Galasi	2008	DE	fm	688	wa	0.71	0.20	0.116	0.010	0.049	0.024	-0.039	0.013
Galasi	2008	UK	fm	601	wa	0.63	0.28	0.147	0.037	0.054	0.028	-0.058	0.031
Galasi	2008	GR	fm	355	wa	0.18	0.77	0.057	0.015	0.026	0.012	-0.011	0.037
Galasi	2008	HU	fm	440	wa	0.64	0.31	0.129	0.016	0.053	0.030	-0.071	0.022
Galasi	2008	IS	fm	239	wa	0.48	0.48	0.076	0.009	0.029	0.018	-0.003	0.012
Galasi	2008	IE	fm	609	wa	0.22	0.67	0.100	0.026	0.065	0.019	-0.023	0.031
Galasi	2008	LU	fm	438	wa	0.49	0.45	0.095	0.010	0.086	0.016	-0.004	0.012
Galasi	2008	NL	fm	607	wa	0.82	0.15	0.124	0.010	0.009	0.026	-0.037	0.011

Authors	Pubyr	Country	Gender	N	Measure S ^r	Share $S^a > S^r$	Share $S^a < S^r$	δ_r	s.e. (δ_r)	δ_o	s.e. (δ_o)	δ_u	s.e. (δ_u)
Galasi	2008	NO	fm	910	wa	0.51	0.41	0.077	0.006	0.040	0.011	-0.021	0.011
Galasi	2008	PO	fm	481	wa	0.37	0.59	0.099	0.010	0.068	0.013	-0.017	0.020
Galasi	2008	PT	fm	265	wa	0.49	0.33	0.133	0.009	0.048	0.016	-0.008	0.020
Galasi	2008	SK	fm	332	wa	0.38	0.47	0.088	0.008	0.051	0.021	-0.010	0.015
Galasi	2008	SI	fm	433	wa	0.78	0.18	0.095	0.021	-0.019	0.056	-0.097	0.037
Galasi	2008	ES	fm	418	wa	0.44	0.50	0.084	0.011	0.040	0.010	-0.036	0.013
Galasi	2008	SE	fm	923	wa	0.51	0.40	0.076	0.004	0.023	0.009	-0.028	0.009
Galasi	2008	CH	fm	689	wa	0.68	0.22	0.086	0.008	0.029	0.017	-0.029	0.010
Galasi	2008	TR	fm	252	wa	0.71	0.28	0.109	0.017	0.047	0.029	-0.062	0.018
Galasi	2008	UA	fm	508	wa	0.64	0.24	0.069	0.023	-0.015	0.045	-0.110	0.034
Korpi and Tahlin	2009	SE	fm	12124	wa			0.068	0.001	0.027	0.002	-0.026	0.003
Korpi and Tahlin	2009	SE	fm	12124	wa			0.061	0.002	0.023	0.002	-0.023	0.002
Korpi and Tahlin	2009	SE	fm	6233	wa			0.034	0.003	0.009	0.003	-0.018	0.003
Korpi and Tahlin	2009	SE	fm	6233	wa			0.207	0.053	-0.176	0.099	-0.371	0.206
Korpi and Tahlin	2009	SE	fm	6233	wa			0.222	0.068	-0.203	0.124	-0.431	0.265
Korpi and Tahlin	2009	SE	fm	3285	wa			0.048	0.006	-0.041	0.010	-0.038	0.012
Schøne and Hardoy	2009	NO		926154	mean	0.15	0.11	0.094	0.000	0.044	0.000	-0.042	0.000
Van der Meer	2009	NL	fm	6906	mean			0.089	0.005	0.057	0.004	-0.014	0.003

Table A Overview of Over/Underschooling Studies-continued

Note: Pubyr refers to year of publication. Country codes: AT, Austria; AU, Australia; BE, Belgium; CA, Canada; CH, Switzerland; CZ, Czech Republic; DE, Germany; DK, Denmark; EE, Estonia; ES, Spain; FI, Finland; FR, France; GR, Greece; HK, Hong Kong; HU, Hungary; IE, Ireland; IS, Iceland; LU, Luxemburg; MX, Mexico; NL, Netherlands; NO, Norway; PT, Portugal; SK, Slovakia; SI, Slovenia; SE, Sweden; TR, Turkey; TW, Taiwan; UA, Ukraine; UK, United Kingdom; US, United States. Gender: f, female; m, male; fin, female and male. Measure: direct wa, direct self-assessment; firm, firm's recruitment; ja, job analysis; mean, mean method; mode, method; quinnRubb, Quinn and Rubb's method; wa, worker self-assessment. S', required years of schooling; S', actual years of schooling; δ_{μ} , return on years of overschooling; δ_{μ} , return on years of underschooling. S.e. (x), standard error of x.



Migration and Education

Christian Dustmann^{*} and Albrecht Glitz^{**1}

*Economics Department, University College London, London WC1E6BT, UK and Centre for Research and Analysis of Migration (CReAM)

**Department of Economics and Business, Universitat Pompeu Fabra, 08005 Barcelona, Spain and Centre for Research and Analysis of Migration (CReAM)

Contents

1.	Introduction	329
2.	Empirical Evidence	331
	2.1. Educational Attainment of Migrants	331
	2.2. Migration and Acquisition of Education	339
	2.3. Return Migration	343
	2.4. The Next Generation	346
3.	The Migrant	349
	3.1. The Migration Decision and Human Capital Investment	349
	3.2. A Simple Model of Return Migration and Investment in Human Capital	351
	3.2.1 Skill Enhancement and Return Decisions of Migrants	353
	3.2.2 Skill Investment in the Home Country	354
	3.2.3 The Migration Decision	355
	3.3. Implications of the Model	355
	3.3.1 The Optimal Investment in Human Capital and the Optimal Migration Duration	355
	3.3.2 Permanent Migration	356
	3.3.3 Temporary Migration with Exogenous Return Date	357
	3.3.4 Temporary Migration with Endogenous Return Date	358
	3.3.5 The Optimal Investment in Learning in the First Period	360
	3.4. Empirical Studies	362
	3.4.1 Assimilation and Adaptation	362
	3.4.2 Extensions of the Basic Approach	365
	3.4.3 Return Migration	368
	3.4.4 Language	370
	3.4.5 Downgrading and Ethnic Networks	374
	3.4.6 Observed Postmigration Schooling Investment and Learning Centers	375
4.	The Effect of Migration on the Skill Base and Educational Attainment of Nonmigrants	376
	4.1. The Selection of Migrants	377
	4.1.1 A Multiple-Skill Model of Migrant Selection	378

¹ We would like to thank Jérôme Adda, Andrew Mountford, Ian Preston, and Yoram Weiss for helpful comments, and Andrew Aitken, Gianandrea Lanzara, Anna Raute, and Anna Rosso for excellent research assistance. Albrecht Glitz acknowledges the support of the Barcelona GSE Research Network, the Government of Catalonia, and the Spanish Ministry of Science (Project No. ECO2008-06395-C05-01). Christian Dustmann acknowledges the support of the NORFACE programme on migration.

	4.1.2 Skill Prices, Productive Capacity, and Selection	382
	4.1.3 Explaining Different Selection Patterns	384
	4.2. Selection and Return Migration	387
	4.3. Learning, Multidimensional Skills, and Return Migration	388
	4.4. Empirical Studies	391
	4.4.1 Emigration and Human Capital Investment in the Sending Country	392
	4.4.2 Immigration and Human Capital Investment in the Receiving Country	395
5.	The Children of Immigrants	397
	5.1. Immigrants and Intergenerational Mobility	399
	5.2. A Model of Intergenerational Mobility of Immigrants	401
	5.3. Empirical Evidence	404
	5.3.1 Schooling Outcomes of Immigrant Children	404
	5.3.2 Intergenerational Transmission of Human Capital	407
	5.3.3 Intergenerational Transmission and Language	410
	5.3.4 Intergenerational Transmission and Ethnic Networks	411
6.	Conclusion	412
Re	ferences	429

Abstract

Sjaastad (1962) viewed migration in the same way as education: as an investment in the human agent. Migration and education are decisions that are indeed intertwined in many dimensions. Education and skill acquisition play an important role at many stages of an individual's migration. Differential returns to skills in origin and destination country are a main driver of migration. The economic success of the immigrant in the destination country is to a large extent determined by his or her educational background, how transferable these skills are to the host country labor market and how much he or she invests in further skills after arrival. The desire to acquire skills in the host country that have a high return in the country of origin may be another important reason for a migration. From an intertemporal point of view, the possibility of a later migration may affect educational decisions in the home country long before a migration is realized. In addition, the decisions of migrants regarding their own educational investment and their expectations about future migration plans may affect the educational attainment of their children. But migration and education are not only related for those who migrate or their descendants. Migrations of some individuals may have consequences for educational decisions of those who do not migrate, both in the home and in the host country. By easing credit constraints through remittances, migration of some may help others to go to school. By changing the skill base of the receiving country, migration may change incentives to invest in certain types of human capital. In addition, migrants and their children may create externalities that influence educational outcomes of nonmigrants in the destination country. This chapter will discuss some of the key areas that connect migration and education.

Keywords

Migration Education Human Capital Return Migration Immigrant Selection Second-generation Immigrants

1. INTRODUCTION

Starting with Homo erectus about 2 million years ago, the first humans were organized in groups of hunters and gatherers. These groups were nomadic, and the economic foundation of their existence was making migrations a necessity, forcing them to move continuously according to food supplies. Constant migration was the normality. Today, populations that pursue a nomadic lifestyle are the exception. The abandonment of the hunting and gathering lifestyle by humans about 10,000 years ago, when this lifestyle was replaced by one based on agriculture, changed the technology of subsistence production. It led to social and political structures that built on nonnomadic forms of economizing. If migrations took place nevertheless, then for two main reasons: either because people were forced to move by natural disasters or man-made circumstances (such as persecution due to distinct political or religious views) or because economic prospects seemed more favorable in other regions. Immigrant-receiving countries today draw distinctions between these two different motives for movement. The Geneva Convention of 1951 defines a refugee as any person "who owing to a well-founded fear of being persecuted for reasons of race, religion, nationality, membership of a particular social group or political opinion, is outside the country of his nationality and is unable or, owing to such fear, is unwilling to avail himself of the protection of that countryor return to it." Its signatories committed to not sending an individual back into a situation of possible persecution. According to United Nations figures, 7.1% of all international migrants in 2005 are refugees from their countries of origin.²

In this chapter, we will deal mainly with migrations due to the second motive: movements that are due to individual decisions based on some optimizing considerations. We will focus our discussion on international migrations, although much of what we discuss also holds for migrations within national borders.³ Deliberate migrations are driven by economic motives and considerations and can therefore be subjected to economic analysis and investigation. Hicks noted in his *Theory of Wages* in 1932 (p. 76) that it is *differences in net economic advantages, chiefly differences in wages*, [that] *are the main causes of migration*.

Education, in turn, is a main determinant of wages, both in the country of origin and the potential destination country.⁴ Although the decisions about how much education to obtain and whether to migrate are often sequential, individuals may in many

² Source: United Nations, Department of Economic and Social Affairs, Population Division. International Migrant Stock: The 2008 Revision, 2009.

³ We follow the convention in most of the literature in economics and define an "immigrant" as an individual who resides and works in a country other than the country where he or she was born. This is the standard definition of immigrants in the Anglo-Saxon countries; however, some countries (like Germany) define immigrants based on nationality rather than country of birth.

⁴ We will use the terms country of destination and host country and the terms country of origin and home country interchangeably throughout this chapter.

cases make these choices simultaneously, choosing education at home with a view to migrating later. Once migrated, choices about skill acquisition in the host country are crucial for the economic payoff of the migration decision. These investment decisions are, in turn, related to future migration plans and therefore the forms that migrations take over the individuals' life cycles. Furthermore, acquisition of education may be the sole reason for a migration—student migrations are an example. Some countries, such as the United States, the United Kingdom, and Australia, are established "learning centers," offering educational products to an international market. The acquisition of skills abroad that are more productive at home may also take place on the job, for instance, through the acquisition of language skills or learning of particular production technologies. Thus, individual migration decisions and decisions about educational attainment are strongly intertwined, and we will shed light on this relationship in Section 3 of this chapter.

But migrations may also affect educational attainments of those who do not choose to move, both in the destination and the origin countries. Migrations may lead to a change in the skill base of both sets of countries, affecting average levels of education and possibly generating educational externalities and new incentives for human capital investments. For instance, if immigration is selective in the sense that only better-able individuals move, then this may enhance the skill base in the destination country, while it may deplete the skill base in the country of origin, with consequences for the populations of those who have chosen not to move. However, there may be situations where migration enhances the skill base of both countries. Immigration may lead to a specialization of nonmigrant workers in the destination country in areas where they have a competitive advantage. It may also lead to an improvement of the economic situation of migrants' families, for instance, through remittances, thus enabling children to attend school instead of working. Remittances may also help to create educational infrastructures that foster educational attainment of those who do not have family members who migrate. We will discuss immigrant selection and the interrelation between migration and education of those who do not move in Section 4 of this chapter.

Finally, migration is a long-term process, with many immigrants bringing their children with them or founding families in their host countries. In this context, it is important to understand how the children of immigrants are accommodated by the host country's education system and how they perform in it relative to native children. What is the role intergenerational transmission of human capital plays for the longer run integration process and how does it affect the immigrants' long-term convergence to the native population? These are key questions in the current debate about immigration, in particular in those countries that only recently received large immigrant populations. We will discuss the educational achievement of the children of immigrants and their intergenerational mobility in Section 5 of this chapter. There are therefore three related cornerstones to this chapter. These cover the key economic aspects of the individual migration decision, their connection to education and the skill selection of immigrants, and the nature of intergenerational spillovers. The analysis of these three cornerstones provides a comprehensive overview of the economic connections between migration and education.

Before considering these three interrelated areas in detail, we begin in Section 2 by setting the scene with the presentation of some empirical evidence about immigrant populations in OECD countries, their educational attainment, their labor market performance, and how they compare to natives and to those in their home countries who decided not to move. We will demonstrate the large diversity of immigrant populations in different countries, which is often due to historical reasons, such as colonial pasts, historical trade links, or particular recruitment policies at earlier points in time. At the same time, the descriptive evidence we present shows many common features across populations of immigrants in different destination countries. We will return to some of these features later in the individual sections that deal with the migrant, the nonmigrants, and the children of immigrants where we will embed them into a more structural framework. In order to ensure comparability, we focus on OECD countries for which detailed and standardized data are available.

2. EMPIRICAL EVIDENCE

2.1. Educational Attainment of Migrants

When thinking about the educational attainment of migrants, two points of reference naturally come to mind: the educational attainment of the migrants relative to the native population in the destination country and the educational attainment of the migrants relative to their compatriots who remained in their country of origin. To provide an overview, Table 4.1 shows the educational attainment of the foreign-born (FB) population in the 10 most important immigrant-receiving OECD countries around the year 2000.⁵ These 10 countries together host 86% of the around 76 million foreign-born individuals aged 15 years and older who live in one of the 28 OECD countries for which data are available (data for Chile and Iceland are not available). Not surprisingly, given the heterogeneity in countries of origin and migration policies in place, there is substantial variation in the educational composition of the foreign-born population across

⁵ We report OECD data for the year 2000, as these are the most recent ones that include comparable information on educational attainment. With the exception of Germany and France, the overall stocks of migrants have further increased across the OECD countries listed in Table 4.1 between the year 2000 and 2008, with an overall growth rate of approximately 37% (based on data from the International Migration Database). The most noticeable change over this period took place in Spain where the foreign-born population tripled to around 6.4 million in 2008. However, with the exception of Japan, whose stock of foreign workers increased particularly fast in this period, the countries listed in Table 4.1 remain the main OECD destination countries also in 2008.

Destination Country	Number of FB	Share of FB	Main Countries of Origin (Share of FB Population)	Share of FB with Low Education	Share of FB with Medium Education	Share of FB with High Education	Share of NB with Low Education	Share of NB with Medium Education	Share of NB with High Education
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Australia	3,860,215	26.0	United Kingdom (26.1), New Zealand (8.2), Italy (5.6)	36.8	32.6	30.6	43.8	31.2	25.1
Canada	5,355,210	22.4	United Kingdom (11.4), China (5.9), Italy (5.9)	22.1	31.8	46.1	22.9	38.3	38.8
France	5,600,198	11.7	Algeria (21.6), Morocco (12.3), Portugal (10.1)	48.4	29.9	21.7	33.8	44.1	22.1
Germany	7,831,959	11.5	Turkey (15.2), Poland (13.1), Russia (11.9)	40.1	42.8	17.2	14.6	61.4	24.0
Italy	2,020,934	4.1	Switzerland (8.9), Germany (8.3), Morocco (6.8)	50.2	35.4	14.4	57.3	31.9	10.8
Netherlands	1,419,940	11.2	South and Central America and Caribbean (20.4), Indonesia (12.5), Turkey (11.2)	45.6	32.8	21.6	32.4	43.8	23.8
Spain	1,914,920	5.5	Morocco (14.5), Ecuador (9.9), France (7.8)	52.8	22.7	24.5	60.9	15.9	23.2
Switzerland	1,454,185	24.1	Italy (15.9), Germany (12.1), Serbia and Montenegro (9.1)	38.8	34.6	26.6	16.3	60.8	22.9
United Kingdom	4,503,466	9.4	Ireland (11.7), India (10.1), Pakistan (6.7)	39.1	21.6	39.3	51.3	25.5	23.2
United States	31,389,926	14.5	Mexico (26.3), Philippines (4.3), Puerto Rico (4.1)	35.3	33.9	30.9	12.5	53.2	34.3

Table 4.1 Educational Attainment of Immigrants in 10 Biggest Immigrant-Receiving OECD Host Countries, around 2000

Note: Data taken from Database on Immigrants in OECD Countries (DIOC) provided by the OECD. Baseline population aged 15 years and older. Ten countries with highest number of foreignborn (FB) and main countries of origin are determined using all education categories and age groups. Columns (7) to (9) refer to the native-born (NB). For each destination country, the three biggest countries of origin are reported. Low educational attainment means up to lower secondary education, medium educational attainment means completed upper secondary education, and high educational attainment means tertiary education. Education shares are calculated for the population aged 25–64 years and are reported in percentages. Observations with unknown education level and unknown place of birth were excluded from the calculation of the education shares.

destination countries. In Canada, for example, only 22% of the foreign-born population aged 25-64 years report lower secondary education ("low education") as their highest educational attainment, whereas in France, Italy, and Spain, about 50% of the foreignborn population have at most completed lower secondary education. At the other end of the educational spectrum, the share of foreign-born individuals with tertiary education ("high education") exceeds 30% in Australia, Canada, the United Kingdom, and the United States, whereas in Germany and Italy, this share is below 20%. The composition of the foreign-born population in terms of their educational attainment becomes particularly relevant when seen in relation with the educational attainment of the native-born population (NB). Based on such comparisons, one can broadly divide the OECD destination countries into two groups: one group characterized by a high-skilled foreign-born population consisting of Australia, Canada, and the United Kingdom, and one group characterized by a low-skilled foreign-born population consisting of France, Germany, the Netherlands, Switzerland, and the United States, with Spain's and Italy's foreignand native-born populations showing relatively comparable educational structures. For example, in Australia, Canada, and the United Kingdom, the share of the foreign-born population with high education exceeds the share of the native-born population with high education by 5.5, 7.3, and 16.1 percentage points, respectively. On the other hand, in France, Germany, and the United States, the share of the foreign-born population with low education exceeds the share of the native-born population with low education by 14.6, 25.5, and 22.8 percentage points, respectively.⁶

But the relative educational attainment of the foreign-born in their destination countries is only one side of the coin. The other natural comparison group consists of the migrants' compatriots who have remained in their country of origin. Table 4.2 provides some evidence for the 11 biggest immigrant-sending countries within the group of OECD countries.⁷ Though again subject to substantial variation, a prominent feature in these figures is that for the majority of origin countries, the share of movers with high education is substantially larger than the share of stayers with high education. For example, while about 23, 22, and 34% of the native-born British, French, and Americans still living in their country of birth have high education (Column (7)), the corresponding shares among the group of those who left these countries are 40, 40, and 61% (Column (4)), respectively.

⁶ For a detailed analysis of worldwide international mobility by educational attainment, see Docquier and Marfouk (2006).

⁷ The figures are constructed from OECD data by aggregating across all OECD countries all foreign-born by country of origin, restricting the sample of origin countries to OECD countries, and then selecting the 11 countries from which the largest number of foreign-born individuals originated. Since not all potential countries of origin are separately recorded for each OECD country, this is not entirely accurate, but, given that for each OECD country, the most important countries of origin are separately reported (and often many more), this procedure should correctly pick up the 11 biggest OECD immigrant-sending countries. Of course, there are additional important non-OECD countries of origin such as (in descending order) China, India, the Philippines, Russia, and Vietnam for which, however, we do not observe the educational attainment of the population in the home country in the OECD data.

Country of Origin	Share of People Living in Other OECD Countries	Share of Movers with Low Education	Share of Movers with Medium Education	Share of Movers with High Education	Share of Stayers with Low Education	Share of Stayers with Medium Education	Share of Stayers with High Education
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Mexico	13.3	68.7	24.5	6.8	70.7	14.6	14.7
United Kingdom	6.8	22.0	37.7	40.3	51.3	25.5	23.2
Germany	4.6	20.8	44.2	34.9	14.6	61.4	24.0
Italy	4.8	51.0	32.9	16.1	57.3	31.9	10.8
Poland	6.8	19.0	53.8	27.3	19.0	67.2	13.8
Turkey	4.4	69.9	22.6	7.5	77.2	14.1	8.8
Portugal	14.5	68.4	24.8	6.8	77.7	11.9	10.4
France	2.4	27.4	32.3	40.3	33.8	44.1	22.1
Canada	4.5	11.2	38.5	50.3	22.9	38.3	38.8
South Korea	n.a.	10.3	37.7	52.0	n.a.	n.a.	n.a.
United States	0.4	10.7	28.6	60.7	12.5	53.2	34.3

Table 4.2 Educational Attainment of Movers and Stayers, around 2000

Note: Data taken from Database on Immigrants in OECD Countries (DIOC) provided by the OECD. Baseline population aged 15 years and older. Eleven OECD countries with highest number of native-born residing in a foreign OECD country (listed in descending order) are determined using all education categories and age groups (for South Korea, data on native-born residing in country of birth are not available). Share of people living in other OECD countries is relative to overall population currently living in the country of origin. Low educational attainment means up to lower secondary education, medium education attainment means completed upper secondary education, and high educational attainment means tertiary education. Education shares are calculated for the population aged 25–64 years and are reported in percentages. Observations with unknown education level and unknown place of birth were excluded from the calculation.

It appears that for these countries of origin, the better-educated individuals are the more mobile ones. However, for other countries, the picture looks different. For example, for Mexico, the main source country of US immigration, the share of movers with only low education is relatively similar to that of the stayers (around 70%), whereas the share of movers with high education (around 7%) is significantly lower than in the group of stayers (15%). Turkey, which is the main country of origin for Germany, and Portugal exhibit similar patterns. For these countries, those who decide to emigrate appear to come predominantly from the middle of the educational spectrum.

The aggregate figures in Tables 4.1 and 4.2 conceal the substantial variation in immigrants' educational attainment that exists across origin countries for any given destination country (Table 4.1) and across destination countries for any given origin country (Table 4.2). Focusing on the former, Table 4.3 shows for each of the 10 main

		L	ow Education				Hi	gh Education		
Destination Country	Minimum Share of Low Education across Origin Countries	Country of Origin with Minimum Share	Maximum Share of Low Education across Origin Countries	Country of Origin with Maximum Share	Standard Deviation of Low- Education Shares within Destination Country across Origin Countries	Minimum Share of High Education across Origin Countries	Country of Origin with Minimum Share	Maximum Share of High Education across Origin Countries	Country of Origin with Maximum Share	Standard Deviation of High- Education Shares within Destination Country across Origin Countries
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Australia	7.9	United States (1.2)	71.6	Malta (1.2)	19.4	6.5	Malta (1.2)	63.3	United States (1.2)	19.3
Canada	6.9	South Korea (1.4)	61.7	Portugal (2.9)	16.6	11.7	Portugal (2.9)	66.5	South Korea (1.4)	16.2
France	15.6	United Kingdom (1.4)	75.1	Turkey (3.0)	18.4	4.3	Portugal (10.1)	57.0	United Kingdom (1.4)	13.7
Germany	9.8	France (0.7)	73.0	Turkey (15.2)	19.1	4.2	Turkey (15.2)	49.7	France (0.7)	12.0
Italy	21.7	United States (2.2)	84.0	Senegal (1.4)	17.5	3.9	Senegal (1.4)	37.7	United States (2.2)	8.1
Netherlands	14.9	France (0.8)	75.0	Turkey (11.2)	17.9	6.0	Turkey (11.2)	76.4	United States (0.6)	16.9
Spain	28.1	Cuba (2.4)	82.4	Portugal (2.8)	15.6	8.5	Portugal (2.8)	40.3	Cuba (2.4)	10.6
Switzerland	7.4	United States (1.2)	81.1	Portugal (6.5)	25.7	2.9	Portugal (6.5)	71.9	United States (1.2)	20.3
United Kingdom	10.8	United States (2.8)	74.1	Bangladesh (3.2)	20.4	15.4	Bangladesh (3.2)	70.5	United States (2.8)	16.9
United States	4.9	Japan (1.4)	69.1	Mexico (26.3)	20.4	6.5	Mexico (26.3)	75.1	India (3.1)	19.2

Table 4.3 Variation in Educational Attainment across Foreign-Born Populations from Different Origin Countries in OECD Host Countries, around 2000

Note: Data taken from Database on Immigrants in OECD Countries (DIOC) provided by the OECD. Ten countries with highest number of foreign-born and the twenty main countries of origin for each destination country are determined using all education categories and age groups. The minimum and maximum education shares are obtained from the sample of the 20 biggest origin countries for each destination country. Figures in parentheses after origin country names show the percentage share of the given origin country in the overall foreign-born population in the corresponding destination country. Low educational attainment means up to lower secondary education, medium educational attainment means completed upper secondary education and high educational attainment means tertiary education. Education shares are calculated for the population aged 25-64 years and are reported in percentages. Observations with unknown education level and unknown place of birth were excluded from the calculation. The standard deviation of educational shares within each host country is unweighted with respect to the origin countries.

immigrant-receiving OECD countries the foreign-born populations with the highest and lowest educational attainment. For instance, looking at the last row of the Table 4.3, of all the main countries of origin of immigration to the United States, the foreign-born population originating from Mexico are the least well educated with 69.1% having only low education (Column (3)). On the other hand, the most highly educated group in the United States is the group of Indians of whom 75.1% have high education (Column (8)). For Germany, the country of origin with the highest share of immigrants with only low education is Turkey with 73.0% (Column (3)), compared with French immigrants of whom only 9.8% have low education (Column (1)). The (unweighted) standard deviations reported in Columns (5) and (10) provide a summary measure of the extent of educational heterogeneity of each country's immigrant population across different countries of origin, showing that in many cases, a given destination country attracts immigrants with high educational background from some countries, but with poor educational background from other countries.

Table 4.4 shows for each of the 11 main immigrant-sending countries in the OECD the destination countries that receive the highest and the lowest educated group of its emigrating population. For instance, looking at the first row of the table, only 5.0% of all Mexicans living in Sweden have low education (Column (1)), whereas this is the case for 69.0% of all Mexicans living in the United States (Column (3)). Similarly, only 7.0% of Poles living in the Czech Republic have high education (Column (6)), whereas the corresponding share in the United Kingdom is 48.7% (Column (8)). Clearly, highly heterogeneous subgroups of individuals from a given country of origin decide to move to specific host countries, as again summarized by the standard deviations reported in Columns (5) and (10). We will discuss possible reasons for these differences below.

The educational attainment of the foreign-born population serves as a key indicator of their performance in the host country's labor market. However, even if the foreignborn population in a given host country is as well educated as the native-born population in terms of the level of formal qualification or completed years of schooling, they are unlikely to perform equally well in the labor market. The educational skills immigrants bring with them may not be easily transferable to the host country's labor market (e.g., due to language deficiencies), and a highly skilled immigrant is unlikely to command the same wage as a native-born worker with the same educational background, at least in the first few years after arrival.

We illustrate this in Column (1) of Table 4.5, which shows the median wage of the foreign-born relative to the median wage of the native-born in a selected set of OECD destination countries. Columns (2) to (4) show the corresponding wage ratios separately by education group. With the exception of Australia, the foreign-born earn overall less than the native-born, in particular in the United States where the median wage gap amounts to 21%. While in the lowest education group, immigrants tend to earn slightly more than comparable natives (with the exception of France and

Table 4.4 Variation in Educational Attainment of Emigrant Population from OECD Origin Countries across Different OECD Host Countries, around 2000

			Low Educatio	n			н	ligh Educatior	ı	
Country of Origin	Minimum Share of Low Education across Destination Countries	Destination Country with Minimum Share	Maximum Share of Low Education across Destination Countries	Destination Country with Maximum Share	Standard Deviation of Low- Education Shares within Origin Country across Destination Countries	Minimum Share of High Education across Destination Countries	Destination Country with Minimum Share	Maximum Share of High Education across Destination Countries	Destination Country with Maximum Share	Standard Deviation of High- Education Shares within Origin Country across Destination Countries
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Mexico	5.0	Sweden (0.0)	69.0	United States (99.1)	20.4	6.5	United States (99.1)	77.8	United Kingdom (0.1)	20.9
United Kingdom	5.6	United States (23.4)	46.1	Spain (3.0)	13.8	26.1	Italy (1.5)	57.0	France (2.4)	11.2
Germany	8.1	United States (34.2)	49.1	Italy (5.4)	13.7	12.9	Italy (5.4)	43.3	Canada (5.8)	9.9
Italy	31.2	United States (22.0)	72.3	Belgium (5.5)	13.4	8.2	Germany (14.2)	36.5	Spain (1.0)	10.4
Poland	11.6	Sweden (1.7)	31.1	Czech Republic (1.2)	6.1	7.0	Czech Republic (1.2)	48.7	United Kingdom (2.8)	14.0
Turkey	14.2	United States (4.0)	81.8	Austria (5.4)	20.8	2.5	Austria (5.4)	52.7	United States (4.0)	15.2
Portugal	51.5	United States (16.4)	82.4	Spain (4.2)	9.8	2.2	Luxembourg (3.0)	20.5	United Kingdom (2.6)	5.8
France	6.9	United States (16.6)	51.2	Italy (11.1)	19.4	12.1	Poland (2.9)	65.9	United Kingdom (7.5)	19.5
Canada	0.7	Japan (0.6)	35.7	Italy (2.1)	11.4	16.7	Italy (2.1)	91.8	Japan (0.6)	20.2
South Korea	3.2	New Zealand (1.4)	13.7	Denmark (0.7)	3.8	27.5	New Zealand (1.4)	78.3	United Kingdom (1.0)	15.9
United States	1.8	Japan (3.9)	32.7	Mexico (12.8)	9.4	37.2	Mexico (12.8)	82.9	Japan (3.9)	16.3

Note: Data taken from Database on Immigrants in OECD Countries (DIOC) provided by the OECD. Eleven OECD countries with highest number of natives residing in a foreign OECD country (in descending order) are determined using all education categories and age groups. The minimum and maximum education shares are obtained from the sample of the 10 biggest OECD destination countries for each origin country. Figures in parentheses after destination country names show the percentage share of the corresponding origin country's emigrant population to other OECD countries, who live in the given destination country. Low educational attainment means up to lower secondary education, medium educational attainment means completed upper secondary education, and high educational attainment means tertiary education. Education shares are calculated for the population aged 25–64 years and are reported in percentages. Observations with unknown education level and unknown place of birth were excluded from the calculation. The standard deviation of educational shares within each home country is unweighted with respect to the OECD host countries.

	Overall		Medium Education	High Education	High Education (Men)		High Education (Women)	
Destination Country		Low Education			Obtained in Home Country	Obtained in Host Country	Obtained in Home Country	Obtained in Host Country
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Australia	1.07	1.11	1.02	0.98	0.99	0.93	0.94	1.02
Canada	0.95	1.07	0.94	0.89	0.86	0.95	0.79	0.99
France	0.90	0.94	0.96	0.90	0.88	0.86	0.77	1.10
Germany	0.93	1.14	0.93	0.91	0.86	1.00	0.83	0.95
Netherlands	0.85	1.11	1.02	0.98	n.a.	n.a.	n.a.	n.a.
Portugal	0.94	1.00	0.76	0.80	0.49	0.88	0.52	1.00
Sweden	0.93	0.91	0.94	0.92	0.81	0.88	0.89	0.95
Switzerland	0.89	1.05	0.93	0.96	n.a.	n.a.	n.a.	n.a.
United States	0.79	1.07	0.83	0.93	0.80	1.04	0.79	1.13

Table 4.5 Median Wages of Foreign-Born Relative to Native-Born by Broad Educational Categories, 2005/2006

Note: Data taken from International Migration Outlook 2008 (OECD (2008)), Chart I.13, Chart I.15, Table I.14. Median hourly wages of the foreign-born are expressed relative to median hourly wages of the native-born in the same group. Sample restricted to those aged 15–64 years who are in dependent employment. Low educational attainment means up to lower secondary education, medium educational attainment means completed upper secondary education, and high educational attainment means tertiary education.

Sweden), they earn substantially less than natives in the medium- and high-education group. In Canada, France, and Portugal, for example, immigrants with high education earn at least 10% less than natives with high education. Overall, wages of immigrants seem to rise more slowly with educational attainment than wages of natives, suggesting that immigrants face a lower return to schooling in the host country than natives do.

Columns (5) to (8) focus on the high-education group and show relative wages of foreign-born men and women with high education, now distinguishing between individuals who received their tertiary education in their home country and those who received it in the host country. The difference between the relative returns reported thus addresses, at least in part, the issue of transferability of human capital from the origin to the destination country.⁸ The less transferable the skills acquired in the home country, the lower their return relative to the return obtained from domestically acquired tertiary education. With the exception of men in France and Australia, the figures show that wages are always higher for foreign-born graduates with host country–specific tertiary education.

2.2. Migration and Acquisition of Education

As pointed out in the previous section, it is important to distinguish between education acquired in the home country before migrating and education acquired in the host country. In particular, minors who arrive together with their parents and young adults who immigrate to attend one of the host country's universities add to their existing stock of human capital by acquiring further formal host country–specific education. With education being a tradable good, some countries have specialized in its production. For instance, Australia's third largest export article (after coal and iron ore) is international education.⁹

This specialization in the provision of educational services is apparent from Table 4.6, which shows for the 10 biggest immigrant-receiving OECD countries the share of foreign students enrolled in tertiary education. Across all destination countries, foreign

⁸ Although a lack of transferability of human capital will be necessarily reflected in lower relative wage ratios, it cannot be easily distinguished from differences in the quality of education obtained because of different educational systems and institutions in the origin and destination countries. Although the OECD data reported in Tables 4.1 to 4.5 are based on the International Standard Classification of Education (ISCED 1997) that, in principle, is based on the educational content of the programs assessed rather than institutional idiosyncrasies, an accurate and consistent measurement of actual educational achievements across different countries remains difficult (see also Hanushek and Zhang (2009)).

⁹ In 2007/2008, international education contributed 13.7 billion AUD to the Australian economy, measured through export earnings, which is the sum of international student expenditure on tuition fees, goods, and services related to living in Australia, and tourism associated with visits from relatives (see http://www.idp.com/research/statistics/education_export_statistics.aspx).

Destination Country	Share of Foreign Students in Tertiary Education	Share of Foreign Students in Advanced Research Programmes	Number of Foreign Students in Tertiary Education	Index of Change (1998–2008)	Retention Rates (2007)
	(1)	(2)	(3)	(4)	(5)
Australia	20.6	23.3	230,635	211 ^a	30.0 ^b
Canada	13.1	38.6	185,781	565	14.7-18.8
France	11.2	39.8	243,436	164	27.4
Germany	10.9	n.a.	245,522	143	29.5
Italy	3.0	7.0	60,448	260	n.a.
Netherlands	6.8	n.a.	40,795	300 ^c	15.0
Spain	3.6	24.0	64,906	224	n.a.
Switzerland	20.3	45.9	45,583	187	n.a.
United	19.9	47.7	462,609	221	27.0^{d}
Kingdom					
United States	3.4	n.a.	624,474	145 ^a	n.a.

Table 4.6 Number and Share of Foreign Students in Tertiary Education in OECD Countries, 2008

Note: Data taken from OECD iLibrary. Foreign students defined as noncitizen students except for Australia and the United States where they are defined as nonresident students due to missing information on noncitizen status. Index of Change 1998–2008 in the number of foreign students is given for total tertiary education and relative to 1998 (1998 = 100). Data on retention rates taken from Table I.8 of the International Migration Outlook 2010 (OECD (2010)). Retention rates are calculated as the number of international students who change their legal migration status, for example, from "student" to "work" or "family formation," divided by the number of international students who do not renew their student permit.

^aBase year figure in 1998 covers noncitizen students, whereas figure in 2008 covers nonresident students.

^bFigure for Australia estimated by Australian Department of Immigration and Citizenship.

^cIndex of change calculated relative to 1999.

^dFigure for the United Kingdom refers to 2005/2006 and is taken from ICMPD (2006).

students constitute a significant fraction of the student population, with their share often exceeding 10%. In Switzerland and the United Kingdom, about one in five students is a foreign student, a fraction that increases further to around 45% when restricting attention to advanced research programs such as doctorates. In terms of absolute numbers, the United Kingdom and the United States are the two main destination countries for foreign students, hosting around 460 and 620 thousand students in 2008, respectively. Column (4) shows by how much the number of foreign students in tertiary education has changed over the decade between 1998 and 2008. Across the board, foreign student numbers have increased substantially. This is not only true in countries that started from a relatively low base such as Italy and the Netherlands (where student numbers increased by 160 and 200%, respectively) but also in countries that were already popular destinations in 1998 such as Australia, the United Kingdom, and the United States (where numbers increased by 111, 121, and 45%, respectively). Overall, in the 10 OECD countries listed in Table 4.6, the number of foreign students increased by 80.8% between 1999 and 2008

which, given an increase in the corresponding overall stock of immigrants over the same period of around 42.5%,¹⁰ indicates that the acquisition of formal education in foreign countries is not only a widespread phenomenon, but also one that has been rapidly gaining importance in recent years.

The last column in Table 4.6 shows estimated retention rates of graduates in their host countries following the completion of their studies. Although the calculation of these figures is not unproblematic (for details, see OECD (2010)), the estimates show that between 15 and 30% of all foreign graduates remain in their host countries after graduation, evidently to a large extent for work purposes.¹¹ These numbers suggest that the acquisition of education in countries that have established themselves as "learning centers" is a main reason for migration and that individuals choose to return to their countries of origin in order to apply the skills acquired (see Dustmann, Fadlon, and Weiss (2010) for modeling of such migrations and our discussion in Sections 3.2 and 4.3). Of course, these numbers may also partly be driven by regulations that do not allow individuals to remain after the completion of their studies. Indeed, in many countries, particularly in Europe, existing policies make it difficult for foreign students to stay and obtain a work permit. In recent years, the transition from study to work has been facilitated in many student destination countries, for example, by enabling students to work while studying, or by extending the period granted to search for work following the completion of study (see ICMPD (2006), for a comparative study on retention policies in a large number of industrialized countries). For instance, several countries have recently started to issue a special residence permit to foreign graduates for the purpose of seeking a job, including the United Kingdom (for 1-2 years, introduced in 2004/ 2005), Germany (for 1 year, introduced in 2005), and France (for 6 months, introduced in 2006).¹² Other special provisions introduced to facilitate the transition from study to work for foreign graduates include the allocation of extra points for a degree from a national institution of higher education (Canada, Australia, and New Zealand), a waiver of an obligatory work experience record (Australia and the Czech Republic), the exemption from the regular quota for "key workers" (Austria), and a specific category

¹⁰ The calculation of this figure is based on the data from the International Migration Database and refers to the change of the sum of the foreign-born population in Australia, Canada, France, the Netherlands, Spain, and the United States and the population with foreign citizenship in Germany, Italy, Switzerland, and the United Kingdom between 1999 and 2008. The missing immigrant stocks for Canada and France were linearly extrapolated from available figures in 2001 and 2006 (Canada) and in 1999 and 2006 (France).

¹¹ Retention rates are calculated as the number of international students who change their legal migration status between 2006 and 2007, for example, from "student" to "work" or to "family formation," divided by the number of international students who do not renew their student permit. On average 61% of international students change their status for work-related reasons (OECD (2010)).

¹² While in France and Germany these jobs have to correspond to the graduate's qualification and are subject to labor market testing, permit holders in the United Kingdom are free to take up any employment they like. Contrary to the United States and Canada, the applicants in these countries do not already need to have a job offer at hand.

with a special quota for foreign graduates (Australia, Italy, United States). These policy changes are a reflection of the increasing global competition in attracting and retaining highly skilled workers. Kato and Sparber (2010) show that students are indeed responsive to such immigration policies. Studying the effect of the sharp reduction in the number of available H-1B visas—the primary means of legal employment for college-educated foreign nationals in the United States—in 2003, they find that this restrictive immigration policy has had a negative impact on the quality of prospective international applicants to US colleges. The intuition is that only the high-ability students are affected by the new policy, since low-ability foreign students are unlikely to find US employment even in the absence of visa quotas. Faced with a reduced prospect of gaining access to the US labor market after graduation, the "best and brightest" of the global talent pool appear to choose other host countries to provide them with both valuable education and the possibility of applying this education in the local labor market.

For the United States, which is host to the largest number of foreign students in the world, there are no overall retention rates of graduate students available. However, Finn (2007) shows that for the group of foreigners who have earned a doctorate in an American university, the retention rate 5 years after they received their degree is around 65–70%. In a cross-country comparison, this is likely to be at the upper end of the spectrum of retention rates.

In terms of the students' origin, there is once again substantial heterogeneity across destination countries. Looking at Table 4.7, three factors appear important for the choice of students where to obtain education: geographical distance, language, and former colonial ties. For instance, the majority of students in Australia (79.3%) originate from Asia, and in particular from China, which supplies a quarter of all foreign students. In Europe, Italy and Switzerland are countries that attract mostly foreign students from other European countries, whereas France has a large share of students from Algeria), and Spain has a relatively large share of students from Spanish-speaking South America (40.7%, of which 21.4% are from Colombia and 16.7% from Peru). In the two biggest student-receiving countries, the United Kingdom and the United States, most foreign students originate from Asia (39.5 and 67.2%, respectively, of which 28.1 and 26.3% are from China).

Figure 4.1 shows the development of the overall number of foreign students studying in the 10 main OECD destination countries listed in Table 4.6 between 1999 and 2008. There is a clear upward trend, in particularly from 2001 onwards, with student numbers increasing from a little more than 1.2 million in 1999 to more than 2.2 million in 2008. Most of this increase is driven by increasing numbers of students from Asia and, to a lesser extent, from Europe and Africa. Although we cannot tell from these data whether these students came to their destination countries for the sole purpose of studying or whether they already arrived as young children with their parents and are

Destination Country	Sending Countries (Share of All Foreign Students in Percentage)	Share Europe	Share North America	Share South America	Share Asia	Share Africa
	(1)	(2)	(3)	(4)	(5)	(6)
Australia	China (25.0), India (11.5), Malaysia (8.1)	4.4	3.5	0.9	79.3	3.2
Canada	China (19.5), India (5.6), US (5.4)	12.0	10.1	3.1	50.2	11.7
France	Morocco (11.1), China (8.6), Algeria (7.7)	21.3	3.5	3.7	21.0	43.5
Germany	China (10.4), Turkey (9.7), Poland (5.7)	47.5	2.5	2.7	37.5	9.0
Italy	Albania (19.5), Greece (7.5), Romania (5.2)	56.5	1.9	7.6	15.0	10.8
Netherlands	Germany (40.6), China (8.4), Belgium (5.4)	69.3	2.3	3.3	19.9	5.0
Spain	Colombia (8.7), Morocco (8.4), Peru (6.8)	31.4	11.3	40.7	3.7	11.7
Switzerland	Germany (24.0), Italy (10.8), France (10.3)	75.7	2.5	2.9	9.0	5.4
United Kingdom	China (11.1), India (7.1), Ireland (5.8)	34.7	6.6	1.3	39.5	14.9
United States	China (17.7), India (15.2), Korea (11.1)	11.2	10.1	5.0	67.2	5.7

 Table 4.7 Origin of Foreign Students in Tertiary Education in OECD Countries, 2008

Note: Data taken from OECD iLibrary. Foreign students are defined as noncitizen students except for Australia and the United States where they are defined as nonresident students due to missing information on noncitizen status. Shares refer to students enrolled in tertiary education, both full time and part time.

expecting a more permanent stay, the important role of many destination countries in providing education to noncitizens is clearly discernible.

2.3. Return Migration

Figure 4.1 suggests that many migrations today are undertaken for the purpose of acquiring education. According to the International Passenger Survey, for example, about 40% of all migrants arriving in the United Kingdom in the year 2009 cited



Figure 4.1 Origin of Foreign Students in Tertiary Education between 1999 and 2008.

Note: Data taken from OECD iLibrary. Graph shows total number of students in the 10 main OECD countries listed in Table 4.6 by continent of origin. Foreign students are defined as noncitizen students except for Australia and the United States where they are defined as noncitizen students from 1998 to 2003 and as nonresident students from 2004 to 2008. Numbers refer to students enrolled in tertiary education, both full time and part time. Numbers for Australia in 1999 and Canada in 2001–2003 and 2005 are missing and were linearly interpolated using the numbers in adjacent years.

as their main reason for migration the desire to pursue formal studies, up from around 23% in the year 2000. Migrations of this type are likely to be temporary. But temporary migrations are a more general phenomenon and widespread also among classical labor migrations. There are about 2.5 million temporary workers arriving in the OECD countries per year, mostly seasonal workers and young working holidaymakers with work permits for a duration of less than 1 year (OECD (2008)). The temporary character of these migrations has important implications for the type of immigrants' educational attainments and their investments in human capital as we will see in Section 3.3. Although until recently the analysis of immigrants' earnings and human capital investments has largely assumed migrations to be permanent, modern migrations seem to be characterized by different patterns. Indeed, numbers suggest that a large fraction of the foreign-born population will at some point return to their home country.

One way to study the time dimension of migration is to look at the duration of stay in the host country. Columns (1) to (3) of Table 4.8 show the share of the foreign-born population in the main OECD destination countries that have been in their host country for less than 5 years, more than 10 years, and more than 20 years. Clearly these

Destination Country	Share Duration <5 years	Share Duration >10 years	Share Duration >20 years	Ratio Outflow/ Inflow × 100 1998–2008	5-Year Re-emigration Rate
	(1)	(2)	(3)	(4)	(5)
Australia	13.6	77.5	55.0	9.8	n.a.
Canada	14.8	70.0	50.2	n.a. ^a	23.7 ^b
Switzerland	23.1	62.4	31.1	51.3	n.a.
Germany	3.6	79.7	n.a.	86.0	n.a.
Spain	40.4	49.0	34.6	14.3 ^c	n.a.
France	8.1	82.7	65.9	n.a.	n.a.
United	17.0	70.2	n.a.	42.6	39.9
Kingdom					
Italy	37.8	34.4	9.9	n.a. ^a	n.a.
Netherlands	9.7	71.6	n/a	29.3	28.1
United States	20.1	63.7	35.7	n/a ^a	19.1

Table 4.8 Share of Immigrants by Duration of Stay, Outflow/Inflow Ratio, and 5-Year Re-emigrationRate (around 2000)

Note: Data taken from Database on Immigrants in OECD Countries (DIOC) provided by the OECD. Duration shares refer to the foreign-born population. Observations with unknown length of duration were excluded from the calculation. The ratio of the aggregate outflow to the aggregate inflow of the foreign population between 1998 and 2008 is calculated using data from the OECD's International Migration Database. Ratios are based on data of foreign nationals from population registers for all countries except Australia and the United Kingdom, for which data based on residence permits or other sources were used. Data on re-emigration rates after 5 years are taken from Table III.1 (p. 171) of the International Migration Outlook 2008 (OECD (2008)). Relevant entry period for the United Kingdom was 1992–1998, for the Netherlands 1994–1998, and for the United States 1999.

^aData on outflows were missing for these countries.

^bFigure taken from Aydemir and Robinson (2008).

^cRatio for Spain refers to period 2002–2008.

cross-sectional figures can only be indicative of the temporary nature of migrations since, for example, a high share of short durations could be either due to actual short migration durations or due to a large number of very recent arrivals (as, e.g., in the case of Italy and Spain). However, with the exception of France, the share of migrants who have already been living in their host countries for more than 20 years (and could therefore be considered as permanent) does hardly ever significantly exceed 50%. In the United States, for example, only 35.7% of the foreign-born population have already lived in the country for more than 20 years. As these figures may be driven by changing cohort sizes of the inflows of immigrants, it is useful to look at the outflows of the foreign populations from their host countries. Column (4) shows the outflow/inflow ratio of foreign nationals over the period 1998 and 2008. This ratio ranges from 9.8% in Australia, over 51.3% in Switzerland, to 86.0% in Germany. This suggests that there are indeed significant flows of foreign individuals out of their host countries. These out-migrations may be back to the countries of origin but could also be to an alternative host country. Outflow/inflow ratios are still only a crude measure of return migration, as they do not necessarily relate to the same

individuals. For a more precise measure, one requires data that follow immigrant entry cohorts over time. Column (5) in Table 4.8 provides some estimates of the share of foreign-born individuals that re-emigrate from their host country within the first 5 years of arrival. The estimates show the generally substantial extent of re-emigration, ranging from 19.1% in the United States to 39.9% in the United Kingdom.¹³ Evidently, relatively short migration spells are a widespread phenomenon and, although the destination of the re-emigrating population is generally not observable, one can assume that a large fraction constitutes return migration to the country of origin.¹⁴ As we will see later, this particular migration pattern has important implications for an immigrant's behavior both in the host and in the home country, and therefore requires particular attention. For an informative overview of the return migration issue including a detailed description of methodological approaches to measure it, see OECD (2008, Part III).

2.4. The Next Generation

The focus of the descriptive evidence presented so far has been on the educational attainment and investment, as well as the labor market performance of the working-age immigrant population. Given that a substantial fraction of immigrants will remain in the host country for a considerable amount of time, their children, whether born in the home country before arrival or thereafter, will spend a large part or even their entire childhood in the host country, passing through its educational system and making educational investment decisions along the way. These decisions have wider consequences not only for the performance of this next generation of immigrants in both the host country more generally, for example, through the immigrants' impact on the fiscal balance (see, e.g., Storesletten (2000) and Dustmann, Frattini, and Halls (2010)) or their integration prospects (see, e.g., Constant and Zimmermann (2008)). Given the often substantial differences in family backgrounds and language proficiencies, it is not surprising that in many destination countries immigrant children do significantly

¹³ The OECD also provides corresponding re-emigration rates for Ireland (60.4%), Belgium (50.4%), and Norway (39.6%). Additional studies that estimate comparable 5-year re-emigration rates are Borjas and Bratsberg (1996) for the United States (17.5%), Bijwaard (2004) for the Netherlands (35%), Shortland (2006) for New Zealand (23%), Dustmann and Weiss (2007) for the United Kingdom (40% males and 55% females), Bratsberg, Raaum, and Sorlie (2007) for Norway (50%), Jensen and Pedersen (2007) for Denmark (55%), and Aydemir and Robinson (2008) for Canada (23.7% males).

¹⁴ Nekby (2004) is one of the few who distinguishes between return migration and secondary migration to a third country, using data for Sweden for the period 1991–2000. According to her results, the share that constitutes return migration is around 90% for Nordic immigrants, 70% for Western Europeans and North Americans, 50% for Eastern Europeans, 40% for Asians, and around 30% for Africans. Bratsberg, Raaum, and Sorlie (2007) estimate the return migration share for Norway and find similar magnitudes. Over the period 1967–2003, the share of those who left Norway to return to their home country is 93% for Danes and Swedes, 86% for US Americans, 87% for UK immigrants, 78% for Turks, 81% for Iraqis, 70% for Somalis, and 33% for Vietnamese immigrants.

 Table 4.9 Raw difference in PISA Test Scores between Students with Immigrant Background and Native Students

Difference of Performance between Students with an Immigrant Background Who Speak a Language at Home that Is Different from the Language of Instruction and Native Students

Destination			Students		
Country	Reading	Mathematics	Reading	Mathematics	
	(1)	(2)	(3)	(4)	
Australia	8.8*	15.7**	-4.4	-4.2	
Canada	-2.4	-2.7	-16.1**	-1.1	
France	-25.4**	-38.9**	-31.7**	-66.7**	
Germany	-57.1**	-58.1**	-81.6**	-92.8**	
Italy	-60.7**	-38.6**	-79.4**	-22.2	
Netherlands	-52.1**	-53.8**	-61.4**	-86.9**	
Spain	-41.1**	-47.8**	-46.0**	-26.1*	
Switzerland	-57.3**	-69.5**	-78.3**	-81.7**	
United Kingdom	-11.5	-14.6*	-36.5**	-26.6**	

Source: PISA 2006; reading scores for the US: PISA 2003. Native students are defined as those born in the country of assessment with both parents also born in the country of assessment. Immigrant students are either those born abroad with both parents also born abroad (first generation) or those born in the country of assessment but both parents born abroad (second generation). Students with a mixed background are excluded. Values are computed using the final weights provided by PISA. Stars indicate that the difference between the immigrant and the native average score is statistically significant at the 1% level (**) and at the 5% level (*).

worse at school than their native counterparts. To illustrate this, we use data from the PISA (Programme for International Student Assessment) study that provides internationally comparable measures of proficiency in reading, mathematics, and science of students aged 15 years in 18 OECD countries, as well as information on a large set of student and school characteristics.¹⁵ Table 4.9 shows the raw differences in test scores in reading and mathematics between native children and children with an immigrant background, which include both children born abroad and children born in the host country. Both in reading and in mathematics, students with an immigrant background score significantly lower than native students. With an average native score in these tests of about 500, the test score gaps are substantial with a relative magnitude of around 10% in a number of countries (Germany, Italy, the Netherlands, and Switzerland). The only exceptions to these patterns are Australia, Canada, and to a lesser extent the United Kingdom, where students with an immigrant background do as well as, or better (in the case of Australia), than their native student counterparts. An obvious explanation for this heterogeneity across destination countries is the difference in socio-economic

¹⁵ For detailed information on the PISA study, see OECD (2007).

characteristics between the corresponding immigrant populations, largely driven by the selectiveness of each host countries' migration policies. Australia, Canada, and the United Kingdom are all countries that have been applying a point system to select the immigrants they admit, which strongly favors individuals with characteristics conducive to their performance in the labor market, such as education and language skills. Through the intergenerational transmission of human capital, these beneficial characteristics are likely to be reflected in their children's performance at school. For illustration, we report in Columns (3) and (4) the test score gaps (relative to natives) of students with an immigrant background who do not speak the language of instruction at home. In all but three cases (Canada, Italy, and Spain in mathematics), these gaps are substantially larger than the gaps for the entire immigrant student population. For example, in the United Kingdom, those who do not speak English at home score 36.5 points below the native average in reading and 26.6 points below the native average in mathematics while overall the test score gaps only amount to 11.5 and 14.6 points, respectively. We will analyze the role language and parents' education play in explaining the achievement gaps between immigrant and native students more systematically in Section 5.3.

Do the lower test scores of immigrant children at age 15 carry over into their adulthood? One measure to assess this is to compare the average school-leaving age of firstand second-generation immigrant adults relative to the native population. Figure 4.2



Figure 4.2 Gaps in Educational Attainment and Log Wages of First- and Second-Generation Immigrant Men Relative to Native Men.

Source: Algan, Dustmann, Glitz, and Manning (2010). Data sources are the French Labour Force Survey 2005–2007, the German Microcensus 2005–2006, and the UK Labour Force Survey 1993–2007. Data points reflect estimated gaps in age left education and log wages of different origin groups of firstand second-generation immigrant men relative to native men. Additional controls in the regressions from which these estimates were obtained are a quadratic in year of birth, region dummies, and time dummies in the age left education censored regression, and a quadratic in potential experience, region dummies, and time dummies in the linear wage regression. reproduces estimates for men obtained by Algan, Dustmann, Glitz, and Manning (2010) for France, Germany, and the United Kingdom using the latest available data sources. The corresponding results for women show broadly similar patterns. The differences in the average age left education shown in the left-hand side scatterplot in Fig. 4.2 are estimated coefficients on dummy variables for the main immigrant groups in each country, obtained from a censored linear regression. Additional controls in these regressions are a quadratic in year of birth, region dummies, and time dummies. Each point represents the educational gap relative to natives of firstgeneration immigrants (x-axis) and second-generation immigrants (y-axis) for a particular country (or country group) of origin. For reference, we include a 45° line and a fitted linear regression line. The scatterplot shows that for those immigrant groups in Germany and France, who started with the biggest disadvantage relative to natives, there is some improvement in the educational attainment from one generation to the next (in the sense of a later age when individuals leave education). However, there is still a significant difference in schooling remaining for these groups in the second generation. This persistence in educational differences between natives and immigrants translates into differences in labor market outcomes, as illustrated in the right-hand side scatterplot in Fig. 4.2, which shows estimated relative earnings gaps, again taken from the Algan, Dustmann, Glitz, and Manning (2010) study in which the only included control variables in the estimated linear earnings equations are a quadratic in potential experience, region dummies, and time dummies. Most importantly, the regressions do not control for the differences in educational attainment illustrated in the left-hand side scatterplot. The figures show that for most immigrant groups there is some improvement in earnings from one generation to the next, in part due to the improved educational attainment. However, the majority of adult second-generation immigrants in these countries still experience a substantial wage disadvantage (of the order of 10% on average) relative to their native counterparts. Both scatterplots also reveal a significant correlation between first- and second-generation immigrants' education levels and earnings. Despite some convergence, those immigrant groups who started with the biggest disadvantage relative to natives in the first generation continue to be the most disadvantaged in the second generation. We will get back to this issue in our discussion of intergenerational mobility of immigrants in Section 5.

3. THE MIGRANT

3.1. The Migration Decision and Human Capital Investment

In this section, we investigate the key drivers of individuals' decisions of whether to emigrate, whether and when to return, and how these decisions interact with decisions about education and skill acquisition. In its simplest possible form, the migration decision is based on a comparison of expected lifetime earnings in the current region of residence and in an alternative region, to which the migrant has the possibility to emigrate. In most cases, there is more than one possible destination region in the choice set of potential migrants. For simplicity, we will abstract from that and consider only one potential destination country.¹⁶

What are the factors that determine the emigration decision? Abstracting for the moment from amenities that arise from living in the home or potential host country, the decision problem of the potential migrant is based on the comparison of the net monetary returns of that decision. In the simplest possible model, where migrations are permanent and the acquisition of human capital is completed before the migration, these will depend on the skill prices in the origin and destination country, as well as the degree to which skills acquired in the origin country are transferable to the economy of the destination country. In a more dynamic setting, the migrant compares lifetime net discounted earnings in the two countries, allowing for the possibility that additional human capital investment is undertaken.

After migration, individuals will potentially acquire further skills in the host country. As skills obtained in the home country are not always fully transferable to the host country's labor market, new immigrants should have lower earnings than natives, even when they belong to the same skill group as measured, for instance, by the years of schooling obtained. However, the subsequent transfer of existing skills, facilitated, for example, through the acquisition of complementary skills like language, and the acquisition of new skills lead to an increase in earnings, possibly at a faster rate than that of comparable native workers. A large literature has developed around measuring this process (starting with Chiswick's, (1978) seminal paper), and we will review some of this literature in Section 3.4.

As we have shown in Section 2, many migrations are temporary, with immigrants remaining for a limited amount of time in the host country, and then returning back home. This behavior can be optimal despite consistently more favorable economic conditions in the destination country (see Dustmann (2003) and Dustmann (1994a, 1995), for an early analysis of different return motives). But if immigrants plan to return to their home countries (or to move on to a third country), then this may affect many aspects of their behavior, including their human capital investment. In particular, any investment decisions in further skills will now depend not only on the return to these skills in the host country but also on the return to these skills back in the home country. Thus, understanding the distinct forms of migration is key to understanding immigrants' human capital investment behavior.

¹⁶ Papers by Dahl (2002), Grogger and Hanson (2008), Bishop (2008), Ortega and Peri (2009), Kennan (2010), and Kennan and Walker (2010) consider the choice problem of individuals when deciding between more than one potential destination region.
Return decisions may be driven by preferences (if, for instance, the home country provides the migrant with amenities that are valuable itself or complementary to consumption) or purchasing power considerations (if, for instance, the host country currency has a high value in the home country). However, they may also be the outcome of an optimizing strategy that consists of obtaining human capital in the destination country in order to apply it in the origin country. Examples of this type of human capital are university education or foreign language skills. In that case, return migration is the outcome of an optimal human capital investment plan over the individual's life cycle. In addition, the possibility of a migration later in the individual's life may already induce human capital investment in the home country. One reason may be that skills acquired at home have a high return in the host country. Another reason may be that skills acquired at home are a prerequisite for the possibility to acquire further valuable skills abroad, either directly because of minimum education requirements due to immigration policies or because of the sequential nature of skill accumulation.

3.2. A Simple Model of Return Migration and Investment in Human Capital

In what follows, we set up a simple model that serves to clarify some key ideas of migrants' decision processes. The model shows how individuals reach a decision of whether to emigrate and how this decision is intricately linked to the human capital they accumulate over their life cycle. It illustrates how the possibility of a temporary migration affects optimal human capital investment profiles, what these profiles imply for individuals' earnings and their growth over time, and how these patterns depend on initial observable skills and ability. The model will also help us to structure the vast empirical literature that exists on migrants' education and skill investment decisions in their home and host countries. This literature has focused particularly on the analysis of earnings profiles of immigrants as a reflection of their human capital investments, on how these profiles depend on the time horizon of the migration as well as the language proficiency of the migrant, and on the issues of skill transferability, skill downgrading, and the role of ethnic networks.

Our model has essentially three periods. Life is finite and time flows continuously. The duration of life is T+2. We have illustrated the timing of the model in Fig. 4.3. The first two periods are "learning periods," and they are of unit length. In the first



Figure 4.3 Timing of Model.

period, individuals live in the home country. In that period, individuals do not work but can acquire education and choose how much to acquire. Acquisition of education is costly, and individuals differ in their efficiency to acquire education. After the first period, individuals decide whether to emigrate.¹⁷ In the second period, individuals live either in the home or in the host country, depending on whether they have chosen to emigrate after the first period. During this second period, individuals have another opportunity to acquire education. They divide their time between learning and working in the labor market. Thus, abstracting from direct costs of education such as fees, the cost of acquiring further education is equal to the opportunity cost of forgone earnings. We think about this period as a period where postsecondary education is obtained. This may take the form of vocational training or college education. The third period has length T. If individuals decide not to emigrate after the first period, they will spend both the second and the entire third period in the home country. If individuals decide to emigrate, then they have the possibility to return to the home country either right after the second period, or after a duration t in the host country, with $t \leq T$. Therefore, the length of the migration is given by t + 1, and the remaining time in the home country after remigration is T - t. A permanent migration corresponds to the case where t = T.

In case no migration takes place, individuals stay at home until death, which occurs at T. Although there is no explicit learning in the third period, we allow the return back in the home country of human capital acquired in the host country in the second period to increase with the length of stay in the host country.

In our model, individuals make a number of choices. These choices are made at the start of the first period and at the start of the second period. To solve the model, we first consider the decision problem at the start of the second period. In case no migration takes place, the individual decides about the optimal investment in learning in the home country in the second period. In case a migration does take place, the individual decides about the optimal investment in learning in the home country in the second period. In case a migration does take place, the individual decides about the optimal investment in learning in the host country and the optimal time to spend in the host country. Given these decisions, we then go back to the start of the first period, where individuals decide about the optimal investment in their education while they are still in the home country. This decision depends on the returns to any such investment in the future, given the optimal decisions about human capital investment and the duration of migration at the start of the second period. Finally, the migration decision is based on a comparison of the net present value of their lifetime earnings if migration does or does not take place.

We will first consider the decisions at the start of the second period. We will then consider the educational investment decision before a potential migration has taken place and the migration decision itself.

¹⁷ For simplicity, we assume that immigrants have only the opportunity to emigrate after the first period.

3.2.1 Skill Enhancement and Return Decisions of Migrants

In case migration takes place, the individual maximizes period 2 and period 3 earnings by choosing the optimal investment in period 2, s_D^* , and the optimal time of return, t^* :

$$\max_{s,t} F^{D}(s,t) = \omega_{D}X(1-s) + t[\omega_{D}(X+f(s,A,X))] + (T-t)[\omega_{O}(X+\gamma(t)f(s,A,X))],$$
(4.1)

where ω_i , j = O, D, is the rate of return to human capital X (acquired in the home country in period 1) in either origin country O or destination country D, s is the time investment in human capital acquisition in the second learning period (the first period in the host country in case of a migration), and t and T denote the time in the host country and the total length of the last period of life, respectively. We assume throughout this section the typical case in which $\omega_D > \omega_O$. The function f(s, A, X) translates human capital investment s in the host country in the second period into skills, where A is the ability of the individual. We make the following standard assumptions: $f_s > 0$, $f_{ss} < 0$, $f_{sA} > 0, f_{sX} > 0, f(0, A, X) = 0$. This means that skills are produced with decreasing returns and that ability and the existing stock of human capital are complementary to the production of new human capital (see Ben-Porath (1967), and, for empirical evidence of such complementarity, Chiswick and Miller (1994) or Friedberg (2000)). The parameter $\omega_{0}\gamma(t)$ is the rate of return to human capital acquired in the host country back in the home country. Notice that $\gamma(t)$ may increase with the time the migrant stays abroad after the second learning period, which reflects the possibility that staying abroad increases immigrants' rate of return on human capital in the home country through on-the-job learning. Further note that if $\gamma(0) < \frac{\omega_D}{\omega_O}$, human capital acquired in the host country in the second period is less valuable back home at the beginning of the third period than in the host country. Conversely, if $\gamma(0) > \frac{\omega_D}{\omega_O}$, human capital acquired in the host country has a higher return back home.

In case migration does not take place, the individual only decides about human capital investment in the second period, s_O^* :

$$\max_{s} F^{O}(s) = \omega_{O} X(1-s) + T[\omega_{O}(X+g(s,A,X))],$$
(4.2)

where g(s, A, X) translates human capital investment in the home country in the second period into skills and is subject to the same standard assumptions as f(s, A, X).

In this model, return migration is induced by the possibility to acquire human capital in the host country that is yielding a high rate of return in the home country (see Dustmann (1994a, 1995); Borjas and Bratsberg (1996); Domingues Dos Santos and Postel-Vinay (2003); and Dustmann, Fadlon, and Weiss (2010), for a similar formulation).¹⁸ A return

¹⁸ We will only discuss human capital accumulation as a return motive in this chapter. There are other motives for why immigrants may want to return, such as higher purchasing power of the host country currency in the home country or consumption amenities in the home country. See Dustmann (1994a, 1995) for a detailed discussion.

to the home country may happen at the beginning of the third period if $\gamma(0) > \frac{\omega_D}{\omega_O}$. An example is "student migrations," where a migration takes place to acquire skills abroad that have a higher return in the home country.¹⁹ Return migration can also be induced by a high return in the home country to human capital acquired in the host country "on the job".²⁰ In our model, this is reflected by $\gamma'(t) > 0$. Even if $\gamma(0) < \frac{\omega_D}{\omega_O}$ so that an immediate return after the second period is not optimal, returning before *T* may be an optimal strategy.

3.2.2 Skill Investment in the Home Country

After having chosen the optimal duration abroad t^* and the optimal investment in learning s^* , where we denote the payoff of these choices as $F^D(t^*, s_D^*)$ in case of emigrating and $F^O(s_O^*)$ in case of not emigrating, we will now consider the first-period problem. Before making the migration decision, individuals have the possibility to acquire education in the home country in the first period. To allow for this, we assume that X, the human capital stock after the first period, is a function of first-period investment i: X = X(i). We further assume that the amount of skills acquired in the first period is concave with respect to investment: $X_i > 0$, $X_{ii} \leq 0$. The choice of i will depend on the possibility of a future migration. In case of migrating, the value function is then given by

$$V^{D} = \max F^{D}(s_{D}^{*}(i), t^{*}(i), i) - C(i, A),$$
(4.3)

where C(i, A) is the cost of investing in education in the home country, which has the properties $C_i > 0$, $C_{ii} > 0$, $C_A < 0$, $C_{iA} < 0$: costs are increasing in investment, convex, and lower as well as increasing at a slower rate for high-ability individuals. In case of remaining in the home country, the value function is given by

$$V^{O} = \max_{i} F^{O}(s^{*}_{O}(i), i) - C(i, A).$$
(4.4)

¹⁹ There is relatively little direct empirical evidence on the returns of foreign education in the home country of an immigrant after he or she returned. One obvious problem for empirical analysis is the selection of both those who decide to study abroad and those who return to their home country, which makes it difficult to identify a causal effect. Oosterbeek and Webbink (2006) exploit a discontinuity in awarding a specific grant to Dutch students for studying abroad. Their OLS estimates show a wage gain of around 4–7% for graduates who studied abroad relative to graduates who did not. However, their RD estimates, though of broadly similar magnitude, are inconclusive due to large standard errors. Wiers-Jenssen and Try (2005) find a wage premium of around 3.5% for Norwegian workers who graduated abroad, whereas Palifka (2003), using survey data covering a complete cohort of graduates from a single Mexican university, finds a premium of around 20% 6 months after graduation for graduates who spent at least some time studying abroad.

²⁰ There is evidence that, for migrants who returned to their home country, the work experience acquired abroad enhances earnings by more than the work experience acquired in the home country. Reinhold and Thom (2009) analyzed earnings of Mexican emigrants who returned from the United States. They find that, for these immigrants, the labor market experience accumulated in the United States increases earnings by twice as much as the experience accumulated in Mexico. Papers by Barrett and O'Connell (2001) and Iara (2006) report similar findings for Ireland and migrants who returned to Eastern Europe from Western European countries. Co, Gang, and Yun (2000) report a wage premium for having been abroad for female return migrants in Hungary.

3.2.3 The Migration Decision

The choice whether to migrate will depend on the comparison of the optimal value functions:

$$V = \max[V^{D} - k + \eta_{D}, V^{O} + \eta_{O}].$$
(4.5)

Assume that η_j , j = O, D, are iid extreme value distributed error terms capturing heterogeneity in the choice, and k is the cost of migration. It follows that for a given individual, the probability of a migration is given by

$$\operatorname{Prob}_{\operatorname{migrate}} = \frac{1}{1 + e^{V^O + k - V^D}}.$$

Thus, if we assume that abilities follow a distribution G, with support $[0, \overline{A}]$, then the proportion of individuals who migrate from country O to country D is given by

Proportion_{migrate} =
$$\int_{0}^{\overline{A}} \operatorname{Prob}_{migrate}(A) dG(A)$$
.

This model is simple, but instructive, as it allows for a variety of cases that have been studied in the empirical literature. For a permanent migration, the second-period problem in Eq. (4.1) corresponds to the simple human capital model that underlies the early empirical papers on immigrant assimilation (see, e.g., Chiswick (1978)). Adding the possibility of return migration with a predetermined migration period leads to more complex empirical specifications, as we will illustrate below. The problem becomes even more difficult if the migrant chooses the time of return optimally. The model also allows consideration of the relationship between migration and return migration, and human capital accumulation in the host country. It includes the special case of student migrations, in which some countries are "learning centers," as documented in Table 4.6, and provide education that has a high return in the home country.

The choices made in the first period add additional insight into learning incentives induced by migration possibilities. For instance, acquisition of education in the home country in the first period may be a prerequisite for acquiring further and higher education in the host country in the second period. Furthermore, the model allows for the possibility that human capital in the home country is acquired because it has a high return in the host country. In Section 3.3, we will explore some of the implications of this model for empirical work. We will then discuss the empirical literature, using the model as a framework of reference.

3.3. Implications of the Model

3.3.1 The Optimal Investment in Human Capital and the Optimal Migration Duration

We first consider the problem of the individual at the beginning of the second period. For those who do not migrate, the optimal investment in the second period s_O^* is

simply obtained by differentiating Eq. (4.2) with respect to *s* and equating the additional forgone earnings in the learning period from an extra unit of time investment, $\omega_O X$, to the benefit arising from a higher earnings potential in the subsequent period, $\omega_O g_s T$.

For those who migrate, the problem at the beginning of the second period corresponds to the decision about how much to invest in host country human capital and when and whether to return to the home country. The FOCs are given by

$$\frac{\mathrm{d}F^{D}(s,t)}{\mathrm{d}s}: -\omega_{D}X + f_{s}[t\omega_{D} + (T-t)\gamma(t)\omega_{O}] = 0$$
(4.6a)

$$\frac{\mathrm{d}F^{D}(s,t)}{\mathrm{d}t}:\omega_{D}(X+f(.))-\omega_{O}(X+\gamma(t)f(.))+\omega_{O}(T-t)\gamma'(t)f(.)=0.$$
(4.6b)

The equilibrium condition in Eq. (4.6a) for the optimal investment in human capital while being in the host country indicates that the cost in terms of forgone earnings in the learning period from an additional unit of time investment (first term) must be equal to the benefit arising from a higher earnings potential in the subsequent period (as in the case of no migration). This, in turn, depends on the time spent in the host country t, on the increase in productivity in the second period through investments in human capital in the first period f_s , and on the transferability of human capital acquired abroad to the home country's labor market $\gamma(t)$, which may depend on the time spent in the host country.

The optimal migration duration derived from Eq. (4.6b) depends, for the optimally chosen human capital investment, on a comparison between spending a marginal unit of time in the host country and spending the same unit of time back in the home country. The individual chooses the optimal s and t simultaneously. The optimal human capital investment s_D^* will change in response to changes in exogenous parameter (for instance, the rate of return to human capital ω_D) directly, and indirectly, because any parameter change affects the optimal migration duration, t^* , which in turn changes investment. Given our assumptions about f(.) and $\gamma(t)$, it is straightforward to show that, in case of an interior solution, a unique optimum for s and t exists. To derive the comparative statics is likewise straightforward. We will now investigate some special cases, which relate to the empirical literature in the area.

3.3.2 Permanent Migration

Assume first that the migration is permanent, that is, t = T. One way to generate a permanent migration in our model is to assume that $\gamma = 1$, $\gamma' = 0$, and $\omega_D > \omega_O$: human capital acquired abroad has the same value at home, the value is not increasing with experience abroad, and the rental rate on human capital is higher in the host country.

In that case, the optimal investment in human capital during the second period is chosen so that $\omega_D X = \omega_D f_s T$: forgone earnings in the second period are equalized to the gain from human capital investment in the third period. Given our assumptions of the function f(.), we immediately obtain a number of results: First, human capital investment s in the host country increases in T. A direct implication of this is that immigrants who arrive early in life will invest more into their skills, as their payoff period is longer (Friedberg (1992), Schaafsma and Sweetman (2001), and Wilkins (2003), explore this issue empirically). Second, investment may increase or decrease with the stock of human capital X upon arrival. Well-educated immigrants may lose considerably when spending time in learning activities. On the other hand, as human capital is productive in its own production, well-educated immigrants acquire additional skills more effectively. The larger the complementarity between skills upon arrival and the acquisition of additional skills (f_{sX}), the more likely it is that human capital investment after immigration is higher for well-educated immigrants (see Borjas (2000), for a detailed discussion). Finally, the higher the ability A of immigrants, the higher is their human capital investment.²¹

How is wage growth for permanent migrants in the host country related to these parameters? Wage growth from period 2 to period 3 is given by $\Delta w_D = \omega_D(f(.) + sX)$. Thus, for a given investment *s*, wages grow faster the higher the skills the individual has upon arrival. However, the effect of an increase in the initial skill level at arrival on skill investment itself is ambiguous so that the overall effect of higher education at arrival on wage growth is likewise ambiguous.²² Wage growth is unambiguously positively related to the level of the immigrant's ability, as this raises third-period wages per unit of investment and human capital investment itself. Finally, wage growth is higher for immigrants who arrive at a younger age, as their investment in the second period is higher. We will contrast these results with the empirical literature in Section 3.4.

3.3.3 Temporary Migration with Exogenous Return Date

The next case we will consider is that of a temporary migration, in which the length of the migration period is exogenously determined. This could, for instance, be the result of a contract migration or migrations that are restricted to a limited time period for other reasons. Let \overline{i} denote the exogenously set migration duration We assume that the optimally chosen migration duration would be longer than the predetermined one, so that the constraint is binding, and (for simplicity) that $\gamma(\overline{i}) = \gamma < \frac{\omega_D}{\omega_O}$ so that

²² d
$$\Delta w_D = \omega_D \left[(f_X + s) + (f_s + X) \frac{\mathrm{d}s}{\mathrm{d}X} \right] \mathrm{d}X.$$

²¹ The results follow from totally differentiating Eq. (4.6a) after setting t = T, where the total differential is given by $dsf_{ss}T = dx(1 - f_{sx}T) - dTf_s - dAf_{sa}T$.

human capital acquired in the host country is less valuable back home than in the host country.²³ In that case, the optimal investment in skills *s* is chosen so that $-\omega_D X + f_s[t\omega_D + (T-t)\gamma\omega_O] = 0.$

As before, investment in human capital increases with ability A, and the effect of a higher level of skills upon arrival is ambiguous. Furthermore, an increase in the migration duration \overline{t} increases human capital investments: the longer the immigrants are allowed to stay in the host country, the higher is their human capital investment in host country–specific skills. Finally, notice that investments in human capital depend on the degree to which these skills are transferable to the home country's labor market, that is, the magnitude of γ . If γ is small, then, for any given migration duration \overline{t} , investments will be low. It is immediately obvious that temporary migrations pose a serious problem for empirical analysis: as we will discuss below, both \overline{t} and γ are usually not observed. Both introduce heterogeneity in earnings profiles that is likely to be correlated with many of the typical regressors in an earnings equation.

We can now again investigate wage growth from period 2 to period 3 in the host country. As before, individuals with higher ability A will have faster wage growth, and the effect of an increase in the skill level X upon arrival is ambiguous. However, wage growth will now be the larger the higher the transferability of human capital from the host country labor market to the home country labor market, γ . Furthermore, the longer the contract migration period \overline{t} , the faster the wage growth. Thus, if migrations are nonpermanent, there will be heterogeneity in the slope of immigrant's wage profiles that is determined by the transferability of human capital, as well as the length of the migration period.

3.3.4 Temporary Migration with Endogenous Return Date

So far we have assumed that the migration is permanent, or that t is exogenously given. We will now relax that assumption. The optimal migration period is determined by condition (4.6b) in conjunction with the choice of the optimal second-period investment s:

$$\omega_D(X+f(.)) - \omega_O(X+\gamma(t)f(.)) + \omega_O(T-t)\gamma'(t)f(.) = 0.$$
(4.6b)'

The first term in Eq. (4.6b)' is the return to each unit of time spent in the country of destination. It is constant for the optimally chosen *s*. The second term is the opportunity cost of staying abroad: it is the forgone earnings in the home country by staying abroad. If $\gamma'(t) > 0$, then this term increases with *t*. Finally, the third term is the additional gain from staying one more unit of time abroad through accumulation of additional skills that increase the value of human capital in the home country. If $\gamma''(t) \le 0$, this additional gain will unambiguously decrease with time *t* spent in the host country. The reason is

²³ We do not consider here the case $\gamma > \frac{\omega_D}{\omega_O}$, as this may lead to a return before \bar{t} .



Figure 4.4 Immediate and Postponed Return Migration.

that although on-the-job experience in the host country is valued back home, there is less and less time remaining to reap the returns from applying human capital acquired in the host country back in the home country.

In our simple model, and assuming that $\omega_O < \omega_D$ (i.e., the rental rate on home country–specific human capital is higher in the host country), a return migration will occur for two reasons. First, the return on human capital acquired in the second period in the host country is higher at home, $\gamma(0) > \frac{\omega_D}{\omega_Q}$, and the accumulation of home country relevant skills by staying in the host country is sufficiently slow so that directly after the second period $\omega_D(X + f(.)) < \omega_O(X + \gamma(0)f(.)) - \omega_O T\gamma'(0)f(.)$. In this case, the (constant) marginal gain from delaying return by one period is lower than the marginal cost of staying. Although emigration itself is optimal after the first period, the individual decides to return immediately after the second period. This situation is depicted in the left panel of Fig. 4.4. Migrations that are characterized by this pattern are student migrations, or migrations that take place predominantly for the purpose of acquiring particular skills or experience in the country of destination. In Column (5) of Table 4.6 in Section 2, we demonstrate that these student migrations are frequent, and that many—typically around 70%—are terminated after education has been acquired.

Second, a return migration may occur even if $\omega_D(X + f(.)) > \omega_O(X + \gamma(0)f(.)) - \omega_O T\gamma'(0)f(.)$ —which means that it is optimal for the migrant to initially remain in the host country after the second period—as long as human capital acquired while working abroad increases the earnings potential of the immigrant in her home country ($\gamma'(t) > 0$). In this case, the marginal costs of staying in the host country increase with time spent there due to the increasing forgone earnings in the home country, and the immigrant may choose an optimal t^* so that $0 < t^* < T$.²⁴ The right panel of Fig. 4.4 illustrates this situation.

²⁴ A sufficient condition for the marginal costs of staying to be increasing in *t* is that $\gamma''(t) \le 0$: the gain from remaining an additional unit of time abroad (in terms of enhancing the home country skill stock) decreases with time in the host country.

It is apparent that, if the return time is optimally chosen, the analysis of immigrants' earnings paths becomes more complex than before, as the optimal migration duration may affect human capital investment and is in turn affected by the optimal skill accumulation. The earnings paths of immigrants who choose their migration duration optimally will therefore depend on parameters that determine their return choice as well. This further complicates the analysis of immigrants' earnings profiles, as we will discuss below.²⁵

3.3.5 The Optimal Investment in Learning in the First Period

So far we have not considered the decisions in the first period. Having solved the second-period problem by choosing the optimal migration period t^* and the optimal investment s^* , the individual will now choose the optimal investment in learning in the first period. Assuming that individuals are endowed with a base level of productivity (which could include compulsory schooling: $X(0) = X^0$), the optimal investment in the case of migrating is given by (invoking the envelope theorem)

$$\omega_D(1 - s_D^*(i))X_i + t^*(i)\omega_D X_i(1 + f_X) + (T - t^*(i))\omega_O X_i(1 + \gamma(t^*(i))f_X) = B^D(i) = C_i(i, A),$$
(4.7a)

and in case of nonmigrating by

$$\omega_{\rm O}(1 - s_{\rm O}^*(i))X_i + T\omega_{\rm O}X_i(1 + g_X) = B^{\rm O}(i) = C_i(i, A), \tag{4.7b}$$

where $X_i = \frac{\partial X}{\partial i}$.

Thus, in the migration case, the individual will compare the marginal cost of investing in the first period (which are costs invoked by effort, and possibly monetary cost) with the marginal benefit, which is the impact of an extra unit of investment in the first period on future lifetime earnings.²⁶ Given our assumptions about the cost function and the learning technology in the first period, the individual will invest in learning in period 1 if the expressions on the left-hand side of Eqs (4.7a) and (4.7b) are larger than the marginal cost for the first unit of investment. Note that—as the marginal cost schedule decreases in ability—higher-able individuals will always invest more in learning. Once the optimal investment in the first period i^* has been obtained for both the migration case (together with s_D^* and t^*) and the nonmigration case (together with s_O^*), the migration decision of the individual is based on a comparison of $V^D - k$ and V^O .

²⁵ Derivation of the partial effects is straightforward, although tedious. For the assumptions made and for $\omega_D - \omega_O > 0$, an increase in ω_D decreases investments in human capital *s*, whereas an increase in ω_O leads to an increase in investments. Those with higher ability *A* will invest more, whereas the effect of an increase in human capital upon arrival *X* on *s* is ambiguous.

²⁶ To simplify the analysis, we assume here that the preference shocks η_D and η_O are drawn *after* the investment decision is made.

This framework allows us to explore a number of interesting cases. Let us first consider the simple case in which migration is permanent and no investment in human capital after period 1 is allowed (s = 0). In this case, conditions (4.7a) and (4.7b) simplify to $\omega_D X_i(1 + T) = C_i(i, A)$ and $\omega_O X_i(1 + T) = C_i(i, A)$. If (as we assume throughout) $\omega_D > \omega_O$, the return to the period 1 investment is clearly larger in the case of a migration and investment in learning in the first period will be higher for the case of a permanent migration than for the case of nonmigration. This is the core of the argument by Mountford (1997). In his model, individuals have an (exogenous) probability of migrating π so that the optimal investment is given by

$$(\pi\omega_D + (1-\pi)\omega_O)X_i(1+T) = C_i(i, A)$$

There are two insights from this relationship. First, there is (for a given π) a particular ability level A, only above which it will be worthwhile to invest in learning. Second, even if the probability of emigration is small, individuals will invest more in learning, as long as the return in the country of destination is sufficiently high. Thus, an increase in π may lead to more accumulation of human capital than in the nonmigration case. Furthermore, although emigrants take with them the human capital they acquire in the home country (which is usually associated with a brain drain), some of those who acquired more skills remain in the home country and may therefore increase the overall per capita level of skills in that country, compared with the case where no migration is possible. This may then lead to a brain gain rather than a brain drain. Thus, the country of origin could overall benefit from a migration of skilled workers—see Mountford (1997) for an insightful discussion. In Section 4.4, we discuss papers that investigate the empirical relevance of this hypothesis.

Another situation that is encompassed by this model is the acquisition of human capital in the home country as a *prerequisite* to enter the destination country. Suppose the potential host country has particular entry requirements such as a specific educational degree. The recently introduced point-based immigration system in the United Kingdom and similar existing systems in Australia and Canada reflect this scenario. Thus, if (for optimally chosen s^* and t^*) the value of migrating is sufficiently higher than the value of nonmigrating, then individuals will invest in education in the home country to obtain the critical level of X^{\min} that then allows an emigration in the next period, given that

$$V^D - k |_{X(i) \ge X^{\min}} > V^O.$$

Again, such a policy will lead to a selection of high-ability immigrants to invest in the minimum necessary level of education, as for them the cost of acquiring education is lower. Another (but similar) situation occurs if learning in the second period in the country of destination requires a certain level of education to be obtained at home. For instance, PhD studies in the United States may require a Bachelor's degree in the country of origin. In that case, optimal investment in the home country will take this requirement into account.

3.4. Empirical Studies

3.4.1 Assimilation and Adaptation

The first generation of papers that studies the performance of immigrants in their countries of destination, starting with Chiswick (1978), concentrates on the earnings profiles of immigrants after arrival in their destination country, viewing these as a reflection of the human capital investments undertaken by the migrants and the skill transferability between origin and destination country. These studies do not distinguish between permanent and temporary migrations, and there is no consideration of immigrants having undertaken investments in the home country with a view of obtaining returns in the host country as illustrated in the previous section. The key question these studies address is whether immigrants perform similarly, worse, or better than natives with the same set of characteristics. This depends on two factors: (1) their quality and (2) their effort to invest in further knowledge.

Why is this important and why have so many papers been published that address this issue? Mainly because the relative position of immigrants in the distribution of earnings determines the contribution they make to the host country economy. Higher earners contribute more to tax and benefit systems and may increase per capita GDP. For many years, the study of immigrant assimilation was perhaps the largest empirical literature on immigration in the economic discipline (see Table 4.L1 at the end of this chapter for a comprehensive overview of studies).

Chiswick's (1978) work suggests that immigrants—although starting with a lower level of earnings than comparative natives—experience a higher earnings growth and eventually outperform natives after about 10–15 years.²⁷ He obtains these results by "augmenting" a simple Mincer wage equation and allowing immigrants to have conditional on education and potential experience at entry—different entry wages than natives, as well as different earnings growth. Earnings of immigrants grow because of two types of work experience: (1) experience accumulated in the home country and (2) experience accumulated in the host country. Experience accumulated in the host country has two components: (1) new, host country–specific human capital and (2) human capital that allows already existing knowledge to be used in the destination country. An example for the latter is language proficiency. Chiswick concludes that the

²⁷ Holding other characteristics constant, Chiswick's results show that the earnings of the foreign-born are 9.5% lower than those of the native-born after 5 years, equal after 13 years, and 6.4% higher after 20 years in the country.

foreign-born seem to be able to compensate any potential earnings disadvantage arising from initially lower host country-specific human capital by greater investments in training, higher work motivation, and greater ability, due to being positively selected. Some subsequent papers supported these findings (Carliner (1980) and De Freitas (1980)); thus, at least for the United States, the early literature on immigrant assimilation draws a picture of immigrants as being high achievers, who—after initial disadvantages—outperform natives through ability, hard work, and investment in their human capital and productivity.

However, this positive picture of immigration to the United States was soon challenged by a series of papers starting with Borjas (1985). Borjas argues that estimation of earnings equations based on simple cross-sectional data—as in Chiswick (1978) does not allow a distinction between cohort and years since migration effects. An immigrant who has been in the United States for 10 years in 1970 arrived in 1960, while an immigrant who has been in the United States for 20 years in 1970 arrived in 1950. Thus, if the composition of immigrants changed over time (as it had since the abolition of country quotas - originally established by the US Immigration Act of 1924 - through the Immigration and Nationality Act of 1965 significantly increased the share of immigrants from South and Central America), estimated earnings profiles based on cross-sectional data may over- or understate the earnings growth of immigrants. In the case of the United States, Borjas argues that entry wages of subsequent cohorts have gone down so that a cross-sectional analysis overestimates the earnings paths of immigrants. He shows that distinction between cohort and years since migration effects is possible by simply adding an additional census year to the data. More specifically, Borjas (1985) proposes the so-called synthetic panel methodology in which earnings of migrants and natives are given by the following two equations:²⁸

$$\gamma_{it}^{I} = \alpha^{I} + \beta^{I} E D_{i} + \gamma^{I} E X_{it} + \delta^{I} Y S M_{it} + \sum_{m} \lambda_{m}^{I} C_{im} + \sum_{k} \pi_{k}^{I} T_{ik} + \varepsilon_{it}^{I}$$
(4.8a)

$$\gamma_{it}^{N} = \alpha^{N} + \beta^{N} E D_{i} + \gamma^{N} E X_{it} + \sum_{k} \pi_{k}^{N} \mathbf{T}_{ik} + \varepsilon_{it}^{N}, \qquad (4.8b)$$

where γ_{it}^{I} and γ_{it}^{N} are log earnings of individual *i* in year *t*, T_{ik} is an indicator variable for the year in which individual *i* is observed that is set equal to unity if k = t, and π_{k}^{I} and π_{k}^{N} are time effects on log earnings for immigrants and natives, respectively. The variable C_{im} is an indicator variable for the year *m* in which individual *i* arrived in the host country, and *ED*, *EX*, and *YSM* measure educational attainment, potential (overall) labor market experience, and potential labor market experience in the

²⁸ Originally, many studies did not allow the effect of education and experience to vary between immigrants and natives, assuming $\beta^I = \beta^N$ and $\gamma^I = \gamma^N$.

United States (years since migration), respectively.²⁹ The parameter of interest is given by $\theta = \delta^I + \gamma^I - \gamma^N$, with immigrants' earnings converging to those of natives when $\theta > 0.^{30}$ The coefficient β^I shows the return to education obtained in the home country on the host country labor market (assuming the migrant worker arrived in the host country as an adult). It thus reveals, in conjunction with β^N , the transferability of human capital between home and host country. This setup is called the synthetic panel methodology since it typically uses repeated cross-sectional data, for instance from US Censuses, to construct a pseudo-panel of cohorts that can be followed over time without actually observing any worker more than once.

A fundamental problem with estimating Eq. (4.8a) is that years since migration equal the difference between calendar year of observation and the cohort entry year so that these variables are perfectly collinear. This means that the coefficients δ^I , λ^I_m , and π^I_k cannot be separately identified without imposing additional identification restrictions.

To identify the model, Borjas (1985, 1995a) assumes equal time effects for immigrants and natives, that is, $\pi_t^N = \pi_t^I$, while allowing cohort quality to vary freely over time. In this case, time effects are effectively estimated from the native earnings equation which in turn ensures identification of the cohort effects in the immigrant equation. In contrast, Chiswick (1978) who only had one cross section of data at his disposal assumed in addition to constant time effects that cohort effects did not change over time so that, after normalization, $\lambda_m^I = 0$, for all m.³¹

Using data from the 1970 and 1980 US Census, Borjas (1985) shows that the quality of immigrants admitted to the United States declined over time. As a consequence, the positive impact of the years since migration variable in cross-sectional earnings equations is picking up not only the intended effect of US-specific human capital accumulation but also the effect of the higher quality of earlier immigrant cohorts. Separating these effects by looking at within-cohort earnings growth reveals that the assimilation profiles of immigrants' earnings are significantly flatter than previously estimated, with the true growth rate being up to 20 percentage points lower in some immigrant cohorts so that the point of overtaking happens much later in the life cycle, if at all. Borjas (1995a) confirmed these results in a follow-up study that included the 1990 US Census: the decline in cohort quality continued till the 1980s, albeit at a slower rate than in the 1970s, and for the bulk of first-generation immigrants, earnings parity with the typical native-born worker will never be reached over the life cycle.

²⁹ To simplify the notation, we ignore higher-order terms of years since migration and experience.

³⁰ In a similar setting, LaLonde and Topel (1992) define assimilation differently as occurring if $\delta^I > 0$, hence comparing the economic value of spending an additional year in the host country relative to a year spent in the home country.

³¹ Assume for simplicity that cohort effects are linear, so that $\sum \lambda_m^I C_{im} = \lambda^I C_{im}$. As $C_{im} = T_{ik} - YSM_{it}$, the parameter Chiswick estimates on YSM is $\delta^I - \lambda^I$ clearly if $\lambda^I < 0$ (cohort quality deteriorates), the estimate is upward biased.

Motivated by these first studies, a large literature has developed that examines the earnings paths of immigrants for different countries and different time periods. In Table 4.L1, we provide an extensive overview, focusing on the estimated returns to education and experience, both in the home and the host country, and the transferability of home country human capital to the host country labor market.

3.4.2 Extensions of the Basic Approach

Although more flexible than Chiswick's (1978) cross-sectional approach, there are a number of restrictive assumptions underlying the synthetic panel methodology in its standard formulation in Eqs (4.8a) and (4.8b). First, it assumes that—although the entry wage of different immigrant cohorts may differ-their wage growth is the same. Our model in Section 3.2 shows that this assumption may be quite restrictive. We show that immigrants who are more able do not necessarily start off with higher earnings, as they may initially invest more in their human capital, but that their wage growth is likely to be steeper than that of less able immigrants.³² If we distinguish cohort quality by the amount of *measured* human capital, X, then wage growth is—as described in our model likewise affected. Thus, the assumption that wage growth is the same for different entry cohorts if these differ in terms of their average ability or their measured human capital appears quite strong. As in the benchmark study by Borjas (1995a), it should be justified in each individual case, in particular since the common interpretation of the estimated entry wages of different immigrant cohorts as a measure of their "quality" hinges crucially upon the validity of this assumption. Duleep and Regets (1999, 2002) and Green and Worswick (2004) provide a detailed discussion of the issues involved in the estimation of immigrant earnings profiles in the context of a human capital investment model. They make a strong case for not relying on entry earnings as a measure of relative cohort quality, with Green and Worswick (2004) suggesting instead a more comprehensive measure based on the estimated present value of all future earnings in the host country.

Another strong assumption in the standard synthetic panel methodology is that business cycle and time effects for natives and immigrants are the same. This assumption implies that macroeconomic trends and transitory shocks, as well as aggregate labor market conditions, affect immigrants' and natives' earnings in the same way. This is unlikely, as immigrants and natives have—as we illustrate in Section 2—usually different skills and are allocated to different occupations and industries. Dustmann, Glitz, and Vogel (2010) illustrate for Germany and the United Kingdom that the economic cycle has—even conditional on education, potential experience, and industry allocation—a

³² This seems to suggest that it is important to use wages rather than earnings for assimilation studies. However, even wages may reflect an increased human capital investment, if contracts are of the Lazear (1979) type in which employees accept lower wages in return for training.

stronger impact on the employment of immigrants than on the employment of nonimmigrants and that these differences are more pronounced for non-OECD immigrants than for OECD immigrants. Two studies by Barth, Bratsberg, and Raaum (2004, 2006) for Norway and the United States, respectively, argue that failure to consider these differences may severely bias the assessment of the earnings assimilation process of immigrant workers. As a solution, the authors suggest to augment the earnings equations by including measures of local unemployment and allowing their impact on earnings to vary between immigrants and natives. Conditional on unemployment, time effects can then be assumed to be equal for both groups. Using data from the Current Population Survey (CPS) from 1979 to 2003, Barth, Bratsberg, and Raaum (2006) show that wages of immigrants in the United States are indeed more sensitive to changes in local unemployment than wages of natives. As a result, since the nativeimmigrant wage gap reduces during economic expansions, the standard estimation strategy with equal time effects yields upwardly biased estimates of both the cohort quality of recent immigrant arrivals in the United States and of the immigrant wage growth, as the wage effects of the improving labor market conditions in the 1990s are erroneously attributed to immigrant quality and wage assimilation.

In light of our model in Section 3.2, another shortcoming of the standard model stated in Eqs (4.8a) and (4.8b) is that there is no distinction between returns to education obtained in the home country and returns to education obtained in the host country. The proposed strategy yields meaningful estimates under the assumption that all immigrants arrive in the host country after they finished education. In this case, β^I measures the returns to education obtained in the home country. However, if some immigrants arrive at an age when they are still in the process of obtaining formal education, the estimated parameter β^{I} compounds the potentially different returns to education obtained in the home and host country. In a study for Israel, Friedberg (2000) explicitly distinguishes education obtained in the host country from education obtained in the home country. She shows that the return to an additional year of schooling obtained in Israel is 10.0% for natives and 8.0% for immigrants, whereas the return to schooling obtained in the immigrants' home countries is only 7.1%. She also finds very low returns to work experience accumulated before arrival. An additional year of experience in the country of origin yields a return of only 0.1% compared with 1.1% for an additional year of experience in Israel's labor market and a 1.7% return to experience for natives. The finding of low returns to home country education and experience in comparison with host country education and experience has been confirmed in a number of additional studies for a variety of destination countries, for example, Kossoudji (1989), Schoeni (1997), and Bratsberg and Ragan (2002) for the United States; Beggs and Chapman (1988a, 1988b) for Australia; Kee (1995) for the Netherlands; Schaafsma and Sweetman (2001) for Canada; Cohen-Goldner and Eckstein (2008) for Israel; Sanromá, Ramos, and Simón (2009) for Spain;

and Basilio and Bauer (2010) for Germany (for details of these studies, see Table 4.L1). Thus, transferability of human capital from home to host country tends to be quite low in many migration contexts. The only exception appears to be human capital that was acquired in developed countries of origin, which typically yields relatively high returns in developed host countries (see, e.g., Schoeni (1997), Friedberg (2000), or Bratsberg and Ragan (2002)). This could be either because home and host country are more similar in terms of cultural, institutional, and technological aspects of their economies so that skills are easily transferable or because more-developed countries of origin simply have higher quality education systems. Interestingly, immigrants from developed countries also receive higher returns to human capital acquired in the host countries after their arrival compared with migrants from less-developed countries, pointing toward complementarities between education obtained at home and education obtained in the host country (see, e.g., Sanromá, Ramos, and Simón (2009) and Basilio and Bauer (2010)). Such complementarities are also supported by the observation that obtaining education in the host country tends to have a positive effect on the return to home countryspecific education (see Friedberg (2001)). One reason is that host country education enables the migrant to transfer their premigration skills more effectively to the host country's labor market.

A related literature concerned with the transferability of human capital has studied the extent of overeducation of immigrants that is defined as the difference between the formal qualifications held by the immigrants and the typical qualifications required in the occupations they hold (see, e.g., Chiswick and Miller (2007, 2008), Green, Kler, and Leeves (2007), Lindley and Lenton (2006), Nielsen (2007), and Sanromá, Ramos, and Simón (2008)). The main findings from this literature show that immigrants are more likely to be overeducated than natives, but that with time in the host country, this difference in overeducation relative to natives decreases, a pattern reminiscent of the assimilation of immigrants' earnings to those of natives over time.

One important implication arising from the theoretical model set up in Section 3.2 is that the expected time the migrant will spend in the host country has an important effect on the decision to invest in host country–specific human capital, as it determines the time horizon over which the benefits from such investments can be reaped by the immigrant. The longer the horizon, the higher are the investment incentives. Even under the assumption that migrations are permanent, this implies that immigrants who arrive at a younger age should have more incentives to invest in host country–specific human capital and thus experience a larger initial earnings gap and steeper earnings profile. Wilkins (2003) confirms these predictions using Australian survey data for 1997, distinguishing four age-at-migration groups: 0–14, 15–24, 25–34, and 35+ years of age. His results show that, for a given stock of human capital at the time of migration, initial wages of immigrants who arrive as children are significantly lower, at least 15%, than those of any other age-at-migration group, but their wage growth with time in Australia is

significantly higher.³³ More explicitly focusing on the human capital acquisition, Gonzalez (2003) shows that for Mexicans arriving in the United States before the age of 19, each year of delayed entry results in about 0.25–0.30 less years of overall schooling and, because this reduction in schooling is due to less US-specific education, significantly lower future earnings. This negative relationship between the eventual educational attainment of immigrants who arrive in the host country in their youth and their age at arrival is a fairly consistent finding in the literature (see, e.g., Hirschman (2001); Chiswick and DebBurman (2004); Cortes (2006); and Perreira, Harris, and Lee (2006)).

3.4.3 Return Migration

Relaxing the assumption that all migrations are permanent and allowing for nonpermanent migrations, the estimation of immigrant earnings profiles becomes far more complex. Consider first the case of a temporary migration, in which the return time is exogenously given and where this constraint is binding (in the sense that the migrant would otherwise wish to stay longer). As we have shown in Section 3.3.3, in that case, the immigrant's investment in learning in the host country depends on the level of skills upon arrival and on the expected economic opportunities in the home country, which are directly affected by the return to any human capital investment when back home. Estimating equations as stated in Eqs (4.8a) and (4.8b) would therefore omit an important set of conditioning variables. The evolution of earnings of the migrant in the host country (measured by the return to experience and the return to the number of years since migration) should depend on the length of the migration. This in turn should also depend on the labor market characteristics in the home country, introducing additional heterogeneity if immigrants come from different origin countries. Neglecting these variables may lead to biased estimates of earnings profiles.³⁴

The situation becomes more complex when return migrations are chosen by the immigrant. In that case, investment in human capital in the host country and the optimal migration time are chosen simultaneously and should be modeled accordingly. Table 4.8 in Section 2 shows that return migrations are very common, and in most cases, returns and total migration durations are chosen by the migrant. In principle, the

³³ Friedberg (1992) and Borjas (1995a) find that age at migration has an important overall negative effect on immigrant earnings in the United States. According to their results, a worker who arrived at age of 30 earns about 5% less than one who already arrived at age of 20, all else equal. See also Schaafsma and Sweetman (2001) and van Ours and Veenman (2006) for related work for Canada and the Netherlands, respectively.

³⁴ For instance, our model in Section 3.2 suggests lower initial earnings (due to larger human capital investments), but a steeper earnings profiles for immigrants who have a longer expected duration in the host country. Thus, assimilation profiles will depend on the duration of migration. Omission of variables that capture this in the estimation of earnings profiles will lead to sample-specific returns to time in the host country, which depend on the distribution of anticipated migration durations. Our model also suggests that an increase in migration durations will lead to steeper earnings paths for higher-able immigrants, which adds further identification problems.

remigration decision and the human capital investment decision have to be estimated simultaneously. This poses a number of difficulties for the empirical researcher. Although the simple model we describe earlier is deterministic, remigration decisions in the real world are unlikely to remain unrevised over the migrants' migration history. Thus, even if (as is possible now in some register data sets) completed migration histories were observable, the completed migration duration may have been different than the migration duration that was intended when human capital decisions were made. But what matters for economic decisions is the *expected* migration duration at the time a decision is taken, and not the actual migration duration.

The data thus required are information on the *expected* duration of a migration rather than on the *completed* duration of a migration. Unfortunately, these return intentions are usually unobserved. An exception is the German Socio-Economic Panel, which asks a boost sample of immigrants in each wave how long they would like to remain in Germany, and whether they would like to return home at all. In an early paper, Dustmann (1993) uses this information to estimate earnings profiles of immigrants. Only about a third of all male immigrants intend to stay in Germany for 30 more years or forever, whereas slightly more than 60% of immigrants intend to return to their home countries within the next 10 years, most of them before they reach retirement age. Allowing assimilation profiles to vary by the intended years of stay in Germany, he finds that "permanent" immigrants have indeed steeper earnings profiles than "temporary" immigrants. After 5 years of residence, an additional year in the host country improves immigrants' earnings by 0.4% if the total intended duration of stay is 10 years, 1.05% if it is 20 years, and 1.26% if it is 30 years. Dustmann (1997, 1999, and 2000) provides additional evidence of differential labor market behavior of immigrants with different return intentions.

Following this line of argument, Cortes (2004) suggests that one of the main reasons for the steeper earnings profile of refugee migrants compared with economic migrants in the United States is the implicit difference in their expected duration of stay. As refugees are typically unable or unwilling to return to their home countries for fear of persecution or violent conflict, they have a longer time horizon in the host country and therefore more incentives to invest in country-specific human capital. Her empirical findings support this hypothesis, as do those of Khan (1997) who finds a higher propensity of Cuban and Vietnamese refugees in the United States to invest in schooling compared with other foreign-born immigrants.

Thus, although—as we show in Section 2—return migrations nowadays are likely to be the rule rather than the exception, the empirical literature has so far largely ignored the implications for the estimation of immigrants' earnings profiles. Careful estimation of earnings profiles of immigrants with different migration plans, taken in conjunction with their human capital investment decisions, requires modeling of the processes of human capital investments and return plans simultaneously. This needs to be addressed within a well-defined structural setting. An additional problem with return migration, apart from the behavioral reasons stated earlier, is that it is likely to be selective, in the sense that those who return are not randomly chosen. Returning migrants may be either those who do not perform very strongly in the host country's labor market or those who perform above average. In the latter case, for example, the average quality of a given immigrant cohort in the host country will decrease over time, leading to an underestimation of the true earnings profiles of immigrants of that cohort relative to natives. Lubotsky (2007) addresses this problem by using longitudinal earnings data from US Social Security records that allow following individual migrants over time. His results show that in the US case, out-migrants are negatively selected, implying that previous studies have systematically overestimated the wage progress of immigrants who remained in the United States, by a factor of around 2.³⁵ We will discuss some reasons for selective immigration and out-migration in Section 4.1.

3.4.4 Language

One dimension of human capital that deserves particular attention in the context of migration is language capital. Language is, on the one hand, a crucial human capital factor for the productivity of immigrants in the host country. Not only is language important in its own right, but it is complementary to many other skill components. For instance, a qualified physician is unlikely to be able to work as a general practitioner when she does not master the language of the host country. On the other hand, investments in language skills are likely to be of little use in the home country. For instance, a migrant from Bosnia to Sweden is unlikely to benefit much from speaking Swedish after having returned home.³⁶ Thus, although being very important as a complement to existing and future skills, language may at the same time be less transferable to other countries' labor markets in the future. In any case, the improvement in language skills over the time spent in the host country is an important driver of the observed earnings assimilation profiles of immigrants in their host countries.

A key question in this context concerns the return to language capital: what is the percentage increase in earnings if an immigrant speaks the host country language well as compared to speaking it poorly? This parameter has important policy implications, as it helps assessing the benefits of language schemes or of selective migration policies that discriminate according to language proficiency. However, this parameter is difficult to measure for several reasons. First, immigrants who acquire language proficiency may

³⁵ It is, however, not clear that the hypothetical assimilation profile of immigrants had no return migration taken place is the interesting policy parameter. If the interest in wage profiles of immigrants is driven by their potential contributions to the economy and the tax and benefit system, what matters are those immigrants who remain in the host country.

³⁶ English may be an exception, with the acquisition of English being an important reason for a migration in the first place. It is not surprising in this context that the most popular destination countries for tertiary education are English-speaking countries: the United Kingdom and the United States (see Table 4.6).

be positively selected, thus introducing a classical selection bias in estimations that regress economic outcomes on language proficiency measures. Second, most available language measures are self-reported. This introduces two types of measurement error: (1) a classical measurement error, due to the interviewer reporting with error and (2) a systematic measurement error, due to the fact that individuals have different "scales" on which they assess their own language skills: the same proficiency may be evaluated as "poor" by one individual and as "good" by another individual. We will discuss below attempts to address these problems after reviewing the literature and its main findings.

In much of the literature, the return to language proficiency is obtained by estimating a standard earnings equation in which a measure of language skills is added as an additional regressor (see, for instance, early work by Carliner (1981); McManus, Gould, and Welch (1983); Grenier (1984); Kossoudji (1988); Tainer (1988); Rivera-Batiz (1990, 1992); Chiswick (1991); Chiswick and Miller (1992, 1995); and Dustmann (1994b)). These studies rely on self-reported language information in survey questionnaires, typically on either a 4-point or 5-point scale, and ignore the problems pointed out earlier. In all these papers, language proficiency is found to be strongly positively associated with earnings in the host country. For instance, for a sample of illegal immigrants in the United States, Chiswick (1991) estimates that immigrants who can read the English language well or very well have earnings that are about 30% higher than those of immigrants with low English reading skills. He also finds that reading skills dominate speaking skills and that the latter does not have an additional separate effect on earnings. For a more representative sample of adult foreign-born immigrants in the United States, Chiswick and Miller (1992) report that English-language fluency is associated with around 17% higher earnings. Dustmann (1994b) estimates that immigrants in Germany who speak German well or very well earn about 7% more than immigrants who speak German on an intermediate level, badly, or not at all. Similarly, those who have good or very good German writing skills earn between 7.3% (males) and 15.3% (females) more than those with bad or no German writing skills.³⁷

Language proficiency is also found to have a complementary effect on the transferability of preimmigration human capital in the form of education and experience. Chiswick and Miller (2002, 2003) show that language skills enhance the return to human capital obtained before migration so that a migrant's greater proficiency in the languages spoken in the host country enhances the effects on earnings of his or her preimmigration schooling and labor market experience. These results hence support the hypothesis that language is an important complementary skill to other forms of

³⁷ Additional studies show an earnings advantage associated with host country language fluency of 12% in Canada (Chiswick and Miller (1992)), 8% in Australia (Chiswick and Miller (1995)), and 12% in Israel (Chiswick (1998) and Chiswick and Repetto (2001)).

human capital: if immigrants cannot conduct a conversation in the host country language, human capital acquired prior to immigration cannot be translated into higher earnings in the host country.

The importance of language as a factor to enhance the productivity of other forms of human capital is also demonstrated in studies that investigate the capacity of different immigrant communities in acquiring further human capital. Sanromá, Ramos, and Simón (2009) show that returns to schooling obtained in Spain are significantly higher for immigrants from Latin America (4.4%), who speak Spanish, than for immigrants from other less-developed countries such as those situated in Eastern Europe (3.6%). Beggs and Chapman (1988b) show that the return to schooling in the Australian labor market in 1981 was 9.0% for the native-born, 8.4% for immigrants from English-speaking countries, and only 4.9% for immigrants from non-English-speaking countries. These findings are suggestive for language being important for the acquisition of further skills, although estimates may be compromised by selection, and do not isolate the effect of language from other country of origin-specific factors that may be driving the differential returns to human capital.

Language proficiency is also a key factor in explaining the educational outcomes of the children of immigrants. Dustmann, Machin, and Schönberg (2010) show that the single most important factor explaining achievement gaps between children of immigrants and natives in the United Kingdom is language spoken at home. In Section 5.3.1, we will discuss the importance of language for children of foreign-born parents in more detail.

As we discussed earlier, a key difficulty in determining the impact language has on economic outcomes is selection, likely leading to an upward bias in the return to language proficiency in straightforward earnings equations, and measurement error in self-reported language measures. Dustmann and van Soest (2001, 2002) were the first to argue that measurement error may lead to a substantial downward bias in simple OLS regressions, which possibly overcompensates the upward bias through selection. To illustrate the possible magnitude of the attenuation bias, they use repeated information on self-reported language proficiency from a panel of immigrants in Germany. Assuming that from year t to year t+1, deterioration in language proficiency is not possible, Dustmann and van Soest (2001) estimate that 85% of the within-individual variance and at least 24% of the overall variance in language measures are due to unsystematic measurement error, in the sense that it varies unsystematically over time. They discuss as a further difficulty of self-reported language information that individuals may have different scales of evaluation. In a cross section, these individual-specific scales cannot be distinguished from measurement error. However, with panel data, and if differences in scales across individuals are constant over time, such a distinction is possible. Dustmann and van Soest (2001) develop an estimator that separates time-varying from time-persistent misclassification. Further, to address the

endogeneity problem, they use parental education as an instrument for language proficiency *conditional* on individuals' education, noting that this is less problematic than instrumenting individuals' education with parental education. The findings show that overreporting language ability is more frequent than underreporting and that there is substantial time-persistent misclassification. According to their results, the return to a one standard deviation increase in true German language fluency decreases from 2.8 percentage points to 0.9 percentage points once unobserved heterogeneity is taken into account. However, controlling subsequently for both time-varying measurement error and time-persistent misclassification, the return to German language fluency increases to approximately 7.3 percentage points. Thus, measurement error may lead to a large downward bias of the estimated return to language proficiency that overcompensates any upward bias due to unobserved ability.³⁸

If repeated information on language ability is available, an alternative way to address the endogeneity problem is by conditioning on individual-specific effects (or estimating difference equations). However, the downward bias through measurement error in the language variable will be greatly enhanced by such techniques. In most panel data sets that contain repeated information on language ability, immigrant populations have been resident for a large number of years, so that the noise-to-signal ratio is too large to allow estimation.³⁹ Berman, Lang, and Siniver (2003) use repeated information on the language proficiency of male immigrants from the former Soviet Union, who moved to Israel after 1989, focusing on the first few years after arrival in which typically the largest improvements in language skills take place. They find large wage gains of language proficiency for workers in high-skilled, but not low-skilled professions as well as evidence for an upward ability bias in cross-sectional estimates, particularly for workers in lowskilled professions.

Bleakley and Chin (2004) present a further strategy to address the endogeneity problem of language proficiency. Based on census cross sections, they devise an IV strategy that exploits the psychobiological phenomenon that young children tend to learn languages more easily than adolescents and adults. Focusing on childhood immigrants, an immigrant's age at arrival in the host country is therefore a strong predictor of his or her language proficiency later in life. It can be used as a valid instrument once its effect on earnings through other channels than language is controlled for. Bleakley and Chin (2004) use immigrants from English-speaking countries as a control group to net out the effects of age at arrival that are not associated with language. Their findings

³⁸ Dustmann and Fabbri (2003) find a similarly large downward bias due to measurement error for the United Kingdom, whereas Dustmann and van Soest (2004) compare parametric and semiparametric estimators to address measurement error in language variables.

³⁹ See Dustmann and van Soest (2002) for a discussion.

show that OLS estimates of the returns to English language fluency in the United States are severely downward biased, which is unexpected if selection is the only problem. They explain this by the IV estimator possibly revealing a local average treatment effect, and by measurement error. Using data on language test results, they estimate that attenuation bias due to measurement error may lead to a reduction of the estimated coefficient by one-half, which is similar in magnitude to the effects found by Dustmann and van Soest (2001).

One particular feature of language capital is that it is in most cases not transferable to the country of origin. Thus, in the formulation of our model in Section 3.2, language capital should be sensitive to the duration individual immigrants would like to spend in the country of destination. This hypothesis is analyzed by Dustmann (1999) who investigates the impact of immigrants' intended duration of stay on their language skills. As migration durations are endogenous in a language equation, he uses an indicator for whether parents who are residing in the home country are still alive as an instrument for the planned migration duration. The findings show that an increase in the total intended duration in Germany by 10 years is associated with a 5 percentage point higher probability of being fluent in German.

3.4.5 Downgrading and Ethnic Networks

The low wages immigrants often receive upon arrival may be partly explained by initial "downgrading," possibly due to a lack of important complementary skills that allow individuals to fully utilize their human capital in the host country's labor market. The stereotypical cab-driving physician vividly captures this phenomenon. Friedberg (2001) and Eckstein and Weiss (2004) study directly the type of jobs immigrant workers perform after arrival using data for Israel. They find substantial occupational downgrading of Russian immigrants who arrived in Israel in the 1990s. Although these immigrants worked in Russia predominantly as engineers, managers, physicians, and teachers, their most important occupations in Israel turned out to be occupations such as service workers, locksmiths/welders, and housemaids. However, over time, particularly highly educated immigrants climb up the occupational ladder. Eckstein and Weiss (2004) show that the proportion of highly educated immigrants working in high-paid professional occupations increases from about 30% at arrival to about 70% 20 years later, compared with an increase from 60 to 80% for equally educated natives over the same time interval. Overall, around 17% of immigrants' wage growth in the first 10 years after arrival in Israel can be attributed to occupational transitions. Mattoo, Neagu, and Özden (2008) provide similar evidence of "underplacement" of immigrants in the US labor market, where in particular skilled immigrants from countries with lower expenditures on tertiary education and non-English languages of instruction, such as Latin American or Eastern European countries, tend to end up in unskilled jobs.

An important consequence of this occupational downgrading is that an allocation of immigrants to particular skill groups based on observed measurable skills such as their education—for example, in order to assess with which subgroup of the native work-force they are most likely to compete in the labor market—is likely to be highly inaccurate and not reflecting the true section of the labor market in which the immigrants are active. Dustmann, Frattini, and Preston (2008) illustrate how, due to downgrading, an allocation of immigrants in the United Kingdom based on their observed education levels misrepresents their true position in the native wage distribution: although these immigrants are on average significantly better educated than natives, they earn wages at the lower end of the native wage distribution in the United Kingdom.

Not only complementary skills (such as language) may be important for immigrants to being able to fully utilize their human capital, but also the reduction in informational deficiencies with respect to the host country's labor market. Here ethnic networks may play an important role. Bartel (1989) and Jaeger (2007) demonstrate the tendency of immigrants to settle in areas where there are already established communities of their ethnic group. Chiswick and Miller (2005) show that living in a region of the United States with a high linguistic concentration of the immigrant's mother tongue has a negative effect on the immigrant's own English language skills, which in turn tends to reduce his or her earnings potential. This would speak against ethnic networks operating to the advantage of immigrants. However, straightforward correlations of ethnic segregation and economic outcomes may be affected by a sorting problem. In two papers, Edin, Fredriksson, and Aslund (2003) and (Damm (2009)) use random dispersal policies of refugee immigrants in Sweden and Denmark to investigate the effects of living in enclaves on labor market outcomes. By using the ethnic concentration in the initial assignment area (Edin et al. (2003)) and the past inflow of assigned conationals (Damm (2009)) as an instrument, these authors convincingly address the sorting problem. They find that living in an ethnic enclave has positive effects on wages and employment, in particular for workers who have low skill levels. Dustmann, Glitz, and Schönberg (2010) find similar evidence of a positive effect of obtaining a job through an ethnicity-based network on wages and job stability in the German context. This speaks in favor of networks as a mechanism to reduce informational uncertainties.

3.4.6 Observed Postmigration Schooling Investment and Learning Centers

Most of the assimilation literature discussed so far draws conclusions about the human capital investment of immigrants after arrival in the host country indirectly from the observed earnings patterns. A more direct approach, given suitable data, is to look at the actual acquisition of additional education by immigrants and the factors that determine it. Using data from the 1976 Survey of Income and Education (SIE) and the 1980 US Census and focusing on the years of schooling obtained after migration and the

enrolment status as dependent variables, Khan (1997) finds that the acquisition of human capital of foreign-born adult men decreases with age at migration, and is higher for refugee immigrants and those who are naturalized, and in states with low tuition fees and better quality of schooling. In the SIE data, she also finds that preimmigration schooling up to the postbachelor professional level is a substitute for schooling in the United States, a finding in support of an earlier study of male Hispanic immigrants by Borjas (1982). In contrast, Chiswick and Miller (1994), who study adult immigrants in Australia, find a positive effect of preimmigration schooling and occupational status on postimmigration schooling, concluding that these are complementary.

As briefly pointed out in Section 3, one reason for immigration can be the acquisition of human capital in a host country. This was a particular aspect of our model, which encompasses migration situations where the sole purpose of a migration is the acquisition of human capital that has a high value upon return to the home country (see Section 3.3.4 for details). This phenomenon is particularly pronounced in higher education and in countries such as Australia, the United States, and the United Kingdom, which receive large numbers of foreign students to study at their universities (compare Table 4.6).⁴⁰ In the United States, for example, foreign-born students (mostly from India, Taiwan, South Korea, and China) accounted for 31% of all PhD recipients in 2006, with even higher shares in specific fields such as physical science (44%), engineering (59%), and economics (59%).⁴¹ Bound, Turner, and Walsh (2009) provide an excellent overview of the latest developments in the US context. In the United Kingdom, foreign students account for 42% of all PhD recipients and 55% of all recipients of a Master's degree in 2007/2008.⁴² Many students who acquire doctoral degrees stay on after completing their studies. Finn (2007), for example, estimates that about 58 (71)% of foreign citizens who received a PhD in science or engineering from a US university in 1991 (1999) are still living in the United States in 2001.

4. THE EFFECT OF MIGRATION ON THE SKILL BASE AND EDUCATIONAL ATTAINMENT OF NONMIGRANTS

In the previous section, we discuss the relationship between education and migration from the perspective of the migrant. In this section, we address the issue from the perspective of those who have chosen not to migrate both in the origin and in the destination country. Our focus will be on the consequences of migration for the skill base and the acquisition of education in the two countries. Migration can affect the skill base of

⁴⁰ This type of immigration is institutionalized in many host countries by issuing specific visas created explicitly to permit temporary study (e.g., the F-1 visa in the United States or the Student Visas in the United Kingdom).

⁴¹ Source: National Science Foundation: Science and Engineering Doctorate Awards 2006. Own calculations. Foreign students are defined as non-US citizens with temporary visa.

⁴² Source: Higher Education Statistics Agency. Foreign students are identified as those with non-UK domicile.

the origin country directly, by changing the skill composition and the talent base. Here the question of selection—who migrates—becomes important. Migration can also affect the skill base of the origin country indirectly by generating incentives to invest in learning and skill acquisition. In the destination country, besides the direct effect because of the inflow of immigrants, migration may change the skill base through responses of the native population, by creating incentives for additional skill accumulation, or for specializing in particular skills where natives have a comparative advantage. Migration can also create spillover effects, for example, through complementarity of the migrant population with the existing populations. In this section, we will discuss some of these aspects.

We start with reinvestigating one of the key questions in the literature on migration, which has important consequences for the issues we raise here: Who migrates? In an early paper, Borjas (1987) uses the Roy (1951) model to relate the skills and abilities of immigrants to the distribution of wages and earnings in the host and home country. His analysis provides deep insights and has been empirically tested in a number of subsequent papers. However, many of these papers have focused on a particular case of Roy's model, where skills are one-dimensional. Here we will reexamine the original Roy model and explore more closely the implications of multidimensionality in skills. We believe that in the context of migration, this will provide much additional insight.

4.1. The Selection of Migrants

We will start with addressing the question of who migrates. In Section 3, we discuss the incentives to emigrate from the perspective of the potential migrant and show that these depend—among others—on the capacity of the individual to produce knowledge, which we termed "ability." The optimal migration plan—in the simplest setting where the return to human capital is higher in the host country—usually provides higher migration incentives to those who have a lower cost of human capital production: those with higher ability. In those considerations, we only looked at the migration decision of a single individual, the "average" individual. We did not compare this individual to other individuals in the origin or destination country by characterizing a distribution of skills. Further, we thought about "skills" as a one-dimensional concept—an individual who has more skills is more productive in both countries.

In this section, we give up this assumption by introducing multiple skills, which added up and weighted by skill prices that may differ across the two countries—determine the productive capacity, or human capital, of an individual in a particular country. We argue that viewing skills as a multidimensional concept, with different prices in different countries, is particularly sensible in the context of migration. We investigate the selection of individuals along the distribution of these skills and state the conditions for positive and negative selection.

Our considerations are based on the Roy (1951) model that we will formalize as a multiple-skill model (concentrating on the special case of two skills), and in which we

allow for the possibility that one skill has a higher price in one country, whereas the other skill has a higher price in the other country.⁴³ This generates the possibility of "nonhierarchical sorting" (to use the terminology of Willis (1987)), in which those who are most productive in the host country migrate and those who are most productive in the home country do not migrate.⁴⁴ We will develop this aspect of the Roy model which, as we believe, has not received sufficient attention in the migration context. In our view, thinking about migration as a decision that considers the prices for *multiple* skills is appropriate in a world where diversely structured national economies trade their comparative and absolute advantages on globalized markets. We show that some of the observed migration patterns that seem not compatible with the one-dimensional skill version of the Roy model can be accommodated by a multidimensional skill model. Drawing on Dustmann, Fadlon, and Weiss (2010), we then show how the basic static Roy model can be extended to a dynamic Roy model by allowing for learning of skills in the two countries so that each country is characterized not only by prices for skills but also by learning opportunities.

4.1.1 A Multiple-Skill Model of Migrant Selection

Consider two countries, an origin country (*O*) and a destination country (*D*). Further, suppose individuals have two latent skills, S_1 and S_2 (this can be easily generalized to more skills). We will here refer to these skills as "analytical skills" (S_1) and "manual" or "trade" skills (S_2). Suppose the two countries have different technologies and industry structures. Thus, we can think about the two countries rewarding the two skills differently according to the two equations:

$$Y_{Di} = \ln \gamma_{Di} = \mu_D + b_{D1}S_{1i} + b_{D2}S_{2i} = \mu_D + u_{Di}$$
(4.9a)

$$Y_{\rm Oi} = \ln \gamma_{\rm Oi} = \mu_{\rm O} + b_{\rm O1} S_{1i} + b_{\rm O2} S_{2i} = \mu_{\rm O} + u_{\rm Oi}.$$
(4.9b)

In Eqs (4.9a) and (4.9b), b_{j1} and b_{j2} represent the prices for the two skills in country *j*, j = O, *D*. Notice that this setup allows for many interesting combinations. For instance,

⁴³ The Roy model goes back to a paper by Andrew D. Roy published in the Oxford Economic Papers in 1951. In this paper, Roy develops the implications of *multidimensional* abilities for occupational choice, the structure of wages, and the earnings distributions. The model has in later years been formalized and developed further (see, e.g., Heckman and Honoré (1990); Willis and Rosen (1979); and Willis (1987)).

⁴⁴ Borjas does, in principle, consider this case, which he terms "refugee sorting," but he does not develop its implications in much detail. Most of the literature (e.g., Chiquiar and Hanson (2005); Orrenius and Zavodny (2005); McKenzie and Rapoport (2007); Ibarrarán and Lubotsky (2007); Belot and Hatton (2008); Fernández-Huertas Moraga (2010); and Kaestner and Malamud (2010)) considers a special case of the Roy model, where skills are one-dimensional, which leads to hierarchical sorting. A very interesting and insightful extension is provided by Gould and Moav (2010) who distinguish between observable skills (such as education) and unobservable skills. Bertoli (2010) considers the case in which there is uncertainty about the earnings potential in the destination country, showing that such uncertainty leads to negative selection becoming more likely.

if countries specialize in particular industries and exchange goods in global markets, then in one country the price for skill 1 may be high and the price for skill 2 low, whereas in the other country the price for skill 2 may be high and that of skill 1 low. If both countries are equipped with the same distribution of skills, migration in both directions may create a more efficient skill allocation.

To compare this with the notation used in much of the literature (e.g., Borjas (1987)), we combine the weighted skills to two measures $u_D = \ln K_D = b_{D1}S_1 + b_{D2}S_2$ and $u_O = \ln K_O = b_{O1}S_1 + b_{O2}S_2$, where K_j is the productive capacity of a person if he or she works in country *j*. Therefore, we can characterize every worker by either a pair of latent skills (S_1 and S_2) or a pair of productive capacities in the two countries (K_D and K_O).

We can think of μ_j as the log of the rental rate to human capital in country *j* so that $\gamma_j = e^{\gamma_j} = R_j K_j$, with $\mu_j = \ln R_j$. The rental rate of human capital in the country of destination, *D*, for example, could be persistently higher if it had a superior technology and if it regulated the inflow of immigrants so that only some of those who wish to enter are allowed in.

We assume that both countries have identical distributions of the two skills S_1 and S_2 before migration and that these distributions are normal and independent with mean zero and variance 1: $S_k \sim N(0,1)$.⁴⁵ It then follows that the random variables Y_D and Y_O are likewise normally distributed, with means μ_D and μ_O and variances and covariance.⁴⁶

$$Var(Y_D) = Var(u_D) = \sigma_D^2 = b_{D1}^2 + b_{D2}^2; \quad Var(Y_O) = Var(u_O) = \sigma_O^2 = b_{O1}^2 + b_{O2}^2 \quad (4.10a)$$

$$Cov(Y_D, Y_O) = Cov(u_D, u_O) = \sigma_{DO} = b_{D1}b_{O1} + b_{D2}b_{O2}.$$
 (4.10b)

We define $\sigma^2 = Var(u_D - u_O) = \sigma_D^2 + \sigma_O^2 - 2\sigma_{DO} = b_{D1}^2 + b_{D2}^2 + b_{O1}^2 + b_{O2}^2 - 2b_{D1}b_{O1} - 2b_{D2}b_{O2}$, which is the variance of the difference in the log of productive capacity between country *D* and country *O*. Further, let $u = (u_D - u_O)/\sigma$ and $z = (\mu_O + k - \mu_D)/\sigma$, where *k* are migration costs (in time-equivalent units). Also, let

$$\sigma_{DU} = Cov(u_D, u) = (\sigma_D^2 - \sigma_{DO}) / \sigma = [b_{D1}(b_{D1} - b_{O1}) + b_{D2}(b_{D2} - b_{O2})] / \sigma$$
(4.11a)

and

$$\sigma_{\rm OU} = Cov(u_{\rm O}, u) = (\sigma_{\rm DO} - \sigma_{\rm O}^2) / \sigma = [b_{\rm O1}(b_{\rm D1} - b_{\rm O1}) + b_{\rm O2}(b_{\rm D2} - b_{\rm O2})] / \sigma.$$
(4.11b)

These covariances are the weighted sums of the differences in skill prices between host and home country, where the weights are the skill prices for the host and home country, normalized by σ . Notice that $\sigma_{DU} = \sigma_{OU} + \sigma$ so that $\sigma_{DU} = \sigma_{OU} > 0$.

⁴⁵ The latter assumption simplifies notation but can easily be relaxed.

⁴⁶ Notice that productive capacities are correlated, although we assume that the skills S_1 and S_2 are independent.

Finally, define the correlation between the log of productive capacities in the home and host country as

$$\rho = Corr(u_D, u_O) = \sigma_{DO} / (\sigma_D \sigma_O) = (b_{D1} b_{O1} + b_{D2} b_{O2}) / [(b_{D1}^2 + b_{D2}^2)(b_{O1}^2 + b_{O2}^2)]^{1/2}.$$
(4.12)

We are now ready to establish the different migration scenarios and to compare the wages of those who decide to migrate and those who decide not to migrate. It follows from Eqs (4.9a) and (4.9b) that an individual will migrate from country *O* to country *D* if $Y_{Di} - k > Y_{Oi}$, or

$$\mu_{D} - \mu_{O} + (b_{D1} - b_{O1})S_{1i} + (b_{D2} - b_{O2})S_{2i} > k$$

$$\Leftrightarrow (u_{D} - u_{O})/\sigma > (\mu_{O} + k - \mu_{D})/\sigma$$

$$\Leftrightarrow u > z$$

$$(4.13)$$

Denoting the density function and the cumulative distribution function of the standard normal distribution by $\phi(\cdot)$ and $\Phi(\cdot)$, the expected earnings of individuals who decide to emigrate are given by⁴⁷

(I)
$$E(Y_D|Y_D - k > Y_O) = \mu_D + \sigma_{DU}[\phi(z)/(1 - \Phi(z))]$$

Likewise, the expected earnings of those in the home country who decide not to migrate are given by

(II)
$$E(Y_O|Y_O \ge Y_D - k) = \mu_O - \sigma_{OU}[\phi(z)/\Phi(z)]$$

How much would those who decide to migrate earn in the home country, and how much would those who decide not to migrate earn in the host country? These two counterfactuals are given by

(III)
$$E(Y_{\rm O}|Y_{\rm D} - k > Y_{\rm O}) = \mu_{\rm O} + \sigma_{\rm OU}[\phi(z)/(1 - \Phi(z))]$$

(IV)
$$E(Y_D|Y_O \ge Y_D - k) = \mu_D - \sigma_{DU}[\phi(z)/\Phi(z)].$$

It follows from (I)–(IV) that the selection of migrants depends on the size and the relative magnitude of the covariances σ_{DU} and σ_{OU} . We can distinguish three regimes.

Regime 1: $\sigma_{DU} > 0$ and $\sigma_{OU} > 0$. It follows that the mean earnings of those who decide to emigrate are higher than the mean earnings in the host country (I) and higher than the mean earnings in the home country (III). On the other hand, the mean earnings of those who decide not to migrate are lower than the mean earnings in the host country (IV) and lower than the mean earnings in the home country (II). This case is one of positive selection of immigrants: those who migrate have higher than average earnings in both countries, and those who do not migrate have lower than average earnings in both countries.⁴⁸ A necessary condition for regime 1 is a higher variance

⁴⁷ See Johnson and Kotz (1972), Heckman (1979), and Heckman and Honoré (1990) for details.

⁴⁸ Note that mean earnings always refer to the premigration period.

of the earnings distribution in the destination country compared with the origin country and a sufficiently high correlation between the productive capacities in both countries, $\frac{\sigma_D}{\sigma_O} > 1$ and $\rho > \frac{\sigma_O}{\sigma_D}$ (see Borjas (1987, 1999)).

In terms of the underlying skill distribution, regime 1 states that the return to both skills must be sufficiently large in the host country. It follows from Eqs (4.11a) and (4.11b) that the sum of the price differentials for the two skills between host and home countries, weighted by the host country prices, is larger than the sum of the price differentials weighted by home country prices. This is certainly the case if the returns to both skills are higher in the host country.

One special case, which is frequently assumed in the literature on the selection of immigrants (see e.g., Chiquiar and Hanson (2005) and Belot and Hatton (2008)), is that $u_D = cu_O$, where *c* is some constant. This occurs if either the return to one skill equals zero in both countries (for instance, $b_{D2} = b_{O2} = 0$), or the ratios of skill prices are equal in the two countries $(b_{D2}/b_{D1} = b_{O2}/b_{O1} = c)$. In both cases, the correlation between u_D and u_O , ρ , is equal to one. Notice that the two cases have different interpretations. In the first case, the skill distribution reduces to one dimension. In the second case, the skill distribution is still two-dimensional (both "analytical" and "manual/trade" skills are needed in the two countries), and individuals may still have different endowments of the two skills; however, the production technologies in the two countries are such that the skill price proportions are exactly equal. A particular case is the one where c = 1, which implies identical skill prices in both economies.

Regime 2: $\sigma_{DU} < 0$ and $\sigma_{OU} < 0$. This case leads to opposite conclusions to regime 1.

Regime 1 and 2 are the two cases that are usually considered in the migration literature. They correspond to the "positive selection" and "negative selection" scenarios in Borjas (1987).

Regime 3: $\sigma_{DU} > 0$ and $\sigma_{OU} < 0$. In this case, the mean earnings of those who decide to migrate are higher than the mean earnings in the host country (I), but they are lower than the mean earnings in the home country (III). On the other hand, the mean earnings of those who do not migrate are lower than the mean earnings in the host country (IV), but they are higher than the mean earnings in the home country (II). Thus, those who migrate have a below-average productive capacity in their origin country, but an above-average productive capacity in their origin country, but an above-average productivity level in the home country (as individuals with below-average productive capacity leave the country) and in the host country. Thus, if the initial skill distribution is the same in the two countries, this situation may lead to a "brain gain" in both countries.⁴⁹

Borjas (1987) refers to regime 3 as "refugee sorting," the underlying idea being that highly skilled individuals are discriminated against in dictatorial systems, receiving a

⁴⁹ It is important to define brain drain or brain gain. We think about *brain gain* as an event that increases *per capita* productivity in either country, and a *brain drain* as an event that decreases *per capita* productivity in either country.

return for their skills that is below average, while being rewarded according to market prices in countries that accommodate refugees. However, this case has many more interesting implications for the study of modern migrations. As we discuss earlier, global trade has led national economies to focus on particular industries, such as manufacturing or financial services. These industries may have different skill requirements across more than one dimension, and many migrations we observe today may be a response to these processes. The idea that migration is a response to skill demands along more than one skill dimension is compatible with the literature on task usage and polarization, which argues that jobs can be characterized by multiple tasks, such as cognitive, routine, and manual tasks (see, for instance, Autor, Levy, and Murnane (2003); Goos and Manning (2007); and Acemoglu and Autor (2010)).

Regime 3 is "nonhierarchical," in the sense that both those who migrate and those who do not migrate have above-average earnings in the country of their choice: individuals are sorted based on their *comparative* advantage. Note that the case $\sigma_{DU} < 0$ and $\sigma_{OU} > 0$ is not possible if we allow for regime 3, as it would contradict $\sigma_{DU} - \sigma_{OU} > 0$. Notice further that the assumption $u_D = cu_O$, which is made in many papers that study the selection of immigrants, rules out regime 3.

A special case of regime 3 occurs when each skill is only priced in one of the two countries; for example, $b_{D1} = b_{O2} = 0$. Now the correlation between productive capacities in the two countries is zero: the productive capacity of an individual in one country does not give any insight about his or her productive capacity in another country. An individual who possesses skill S_1 will only be able to obtain a return in the home country, whereas an individual with skill S_2 will only obtain a return in the host country.

4.1.2 Skill Prices, Productive Capacity, and Selection

It follows from Eqs (4.11a) and (4.11b) that whether migration is selective in terms of productive capacity depends on the underlying skill prices. Changes in these prices will change the type of migration that occurs and the nature of selection. Which regime characterizes a particular migration situation depends on the two expressions

$$\sigma_{DU} = (\sigma_D^2 - \sigma_{DO}) / \sigma = [b_{D1}(b_{D1} - b_{O1}) + b_{D2}(b_{D2} - b_{O2})] / \sigma$$

and

$$\sigma_{\rm OU} = (\sigma_{\rm DO} - \sigma_{\rm O}^2) / \sigma = [b_{\rm O1}(b_{\rm D1} - b_{\rm O1}) + b_{\rm O2}(b_{\rm D2} - b_{\rm O2})] / \sigma,$$

where $\sigma_{DU} = \sigma_{OU} + \sigma$. To illustrate how the different regimes depend on skill prices, consider Fig. 4.5a, where we have fixed $b_{D1} = 1$, $b_{D2} = 2$, $b_{O2} = 1$, and we allow b_{O1} to vary between 0 and 5. The dashed and dotted lines in the figure are σ_{DU} and σ_{OU} , respectively. For b_{O1} in the range between 0 and 1.62, $\sigma_{DU} > 0$ and $\sigma_{OU} > 0$; thus,



Figure 4.5 Selection Scenarios.

we are in regime 1, with migration being positively selective. In the range where b_{O1} is between 1.62 and 3, $\sigma_{DU} > 0$ and $\sigma_{OU} < 0$, and we are in the nonhierarchical regime 3, where those who would do best in the host country migrate and those who would do best in the home country do not migrate. Finally, above $b_{O1} = 3$, both σ_{DU} and σ_{OU} are negative; we are in regime 2 where migration is negatively selective. This demonstrates that the selection of immigrants in terms of their productive capacity depends on relative skill prices, which may change over time.

In Fig. 4.5, we plot the corresponding variances and the covariance (Fig. 4.5b), as well as the correlation coefficient (Fig. 4.5c). In the range where positive selection occurs, the variance of productive capacity is higher in the destination country, and the correlation between skills is high. In the range where negative selection occurs, the variance of productive capacity is higher in the country of origin, and the correlation between skills is lower. Notice that there is a range where the variance of productive capacity is higher in the destination country; yet, we are in regime 3, where we cannot hierarchically sort immigrants in terms of their average productive capacity.

The migration decision of an individual migrant is based on a comparison of individual earnings in the home and host country. Using Eqs (4.9a), (4.9b), and (4.13), an individual will emigrate if $Y_D - k > Y_O$, or

$$S_2 > \frac{\mu_{\rm O} - \mu_{\rm D} + k}{b_{\rm D2} - b_{\rm O2}} - \frac{b_{\rm D1} - b_{\rm O1}}{b_{\rm D2} - b_{\rm O2}} S_1.$$

4.1.3 Explaining Different Selection Patterns

There is by now a large empirical literature that attempts to assess the direction of migrant selection. Most of the papers in this literature draw on Borjas (1987) as an underlying theoretical framework, but consider the special case where $u_D = cu_O$. The evidence these papers establish is mixed. Some papers (including Borjas' (1987) original analysis) find evidence that is compatible with the predictions of the simple one-dimensional skill model, namely that selection is positive from country O to country D if skill prices are higher in country D, and that selection is negative if skill prices are lower in country D. Examples are Cobb-Clark (1993) or Ramos (1992), who find that, consistent with negative selection, nonmigrants in Puerto Rico are more educated than individuals migrating from Puerto Rico to the United States and that those individuals migrating back from the United States. Others (e.g., Feliciano (2005), Orrenius and Zavodny (2005), McKenzie and Rapoport (2007), and Belot and Hatton (2008)) find limited or no evidence that is compatible with this simple model.

In an influential paper, Chiquiar and Hanson (2005) analyze migration from Mexico to the United States. They argue that as the return to schooling is higher in Mexico, individuals with high levels of schooling are less likely to migrate. They compare actual wage densities for residents of Mexico with counterfactual wage densities that would be obtained were Mexican immigrants paid according to skill prices in Mexico, thus comparing the conditional distributions whose means are given by (I) and (III). The findings suggest that, were Mexican immigrants in the United States paid according to Mexican skill prices, they would fall disproportionately in the middle and upper middle of Mexico's wage distribution. As Chiquiar and Hanson (2005) point out, this does not support negative selection, but rather suggests intermediate or moderate positive selection of Mexican immigrants. This empirical finding is not compatible with the original model they started off with, which is our model discussed earlier, but restricted to the special case where $u_D = cu_O$. To reconcile the empirical evidence with the model, they introduce nonlinear migration costs. They assume that costs are large, but decrease in schooling at a decreasing rate, so that the net advantage of migration is highest for those in the middle of the distribution of skills. We reproduce their explanatory graph in Fig. 4.6, which illustrates the case of constant migration costs (Y_D)



Figure 4.6 Nonlinear Migration Costs.

and varying migration costs that are nonlinear in schooling $(Y_D^*)^{50}$ Note that skills are one-dimensional and migration costs k in the latter case are equal to $k = \exp(\mu_{\pi} - \delta_{\pi}S)$. The way we have drawn the figure for the constant migration cost case $(\delta_{\pi} = 0)$ implies that the rent on human capital is higher in the destination country (United States), $\mu_D > \mu_0$, but the return to skill S is higher in the origin country (Mexico). Thus, in this case, those with levels of skill S below S* will emigrate, but those with levels of skill S above S* will remain in Mexico—which is what Chiquiar and Hanson expected to find in the data. However, if $\delta_{\pi} > 0$ and if the fixed costs of migration are sufficiently high, those at the low end of the skill distribution, below level S^L, may find it too costly to emigrate, whereas for those in the middle of the skill distribution, between level S^L and level S^U, migration is advantageous. Chiquiar and Hanson conclude that such nonlinear cost schedules may provide a possible explanation for the observed migration pattern from Mexico to the United States.⁵¹

Nonlinear migration costs are one reason why the one-dimensional model may not fit the data. Another reason may be that the one-dimensional model is overly restrictive and omits an important aspect of migration decisions. Our discussion in the previous section has illustrated that migration decisions may be taken by considering the prices of multiple skills in the home and potential host country. Clearly, education is a onedimensional measure of skills, which may for instance reflect well the academic skills of individuals but may measure less well manual and trade skills. If the latter are highly valued in the destination country and are more prevalent for individuals in the middle

⁵⁰ The case of migration costs that are linear in schooling is straightforward in that it either leaves the overall selection pattern unchanged or reverses it entirely, depending on the pace at which migration costs decrease with educational attainment.

⁵¹ For a further discussion of how different assumptions regarding the migration costs affect predicted selection patterns for example, whether migration costs are assumed to be fixed in time-equivalent units or in monetary units—see Rosenzweig (2007) and Hanson (2010).

	S ₁	S ₂	b _{D1}	b _{D2}	b ₀₁	b ₀₂	$Y_D(=b_{D1}S_1+b_{D2}S_2)$	$Y_O(=b_{O1}S_1 + b_{O2}S_2)$	$Y_D - Y_O$
Low	1	0.5	1	2	2	1	2	2.5	-0.5
Medium	2	2.5	1	2	2	1	7	6.5	0.5
High	3	2	1	2	2	1	7	8	-1

Table 4.10 Example Intermediate Selection

of the education distribution, then migration patterns like those illustrated by Chiquiar and Hanson (2005) can be explained without making specific assumptions about migration costs, as the following example illustrates.

Consider the case where the skill distribution is two-dimensional. Suppose further that the first skill (which we termed "analytical" skill) is strongly correlated with education, but the second skill (which we termed "manual" or "trade" skill) is more strongly present in individuals with intermediate years of education. A multidimensional education system like the German apprenticeship system could produce such patterns, with individuals with intermediate levels of education having often gone through vocational training in crafts and trades. In countries without well-developed vocational training schemes, those with intermediate years of education may still have specialized in manual- or trade-related skills, as development of many of these skills requires some basic level of schooling. Measured education may reflect mainly skill S_1 , but not necessarily skill S_2 .

For the case of Mexican–US migration, the manual- or trade-related skill dimension, although probably not highly valued in Mexico (as these skills are in plentiful supply), may command a relatively high price in the United States. This seems to be compatible with the occupational distribution of Mexican immigrant workers in the United States, with the three main occupations falling exactly into this category of skills: cooks (6.1% of workers), construction laborers (5.8%), and grounds maintenance workers (4.9%).⁵²

In Table 4.10, we give a numerical example. We distinguish among three education groups, "low," "medium," and "high." We have chosen the skill prices such that skill 1 has a higher return in the origin country, and skill 2 has a higher return in the destination country. Skill 1 increases with education but skill 2 is particularly high for those in the middle of the education distribution, and less developed for those who are either low educated or highly educated. For simplicity, suppose migration costs are zero and the rental rate of human capital is identical in both countries and normalized to zero. The differentials between wages that can be obtained in the destination and the origin country are reported in the last column of Table 4.10. Those who are low educated will not

⁵² Figures calculated from the 2009 IPUMS CPS sample using all Mexican-born individuals in the United States who are currently at work and aged above 16 years.
migrate and are better off in their country of origin, whereas those with intermediate levels of education will gain from migration, and those with high levels of education will again lose. Thus, this scenario creates a migration situation where only those in the middle of the education distribution will want to emigrate.

The scenario corresponds to the empirical evidence of intermediate selection provided by Chiquiar and Hanson (2005) without introducing nonlinear migration cost. This illustrates the capacity of the multidimensional Roy model to accommodate different migration patterns that remain unexplained in a simple one-dimensional skill model.

To conclude, we believe that the full potential of the Roy model has not been explored in the migration context. The situation where individuals are characterized by multiple skills, and where these skills are priced differently in different countries, is, in our view, important in an ever more globalized world economy where individual countries specialize in particular industries.

4.2. Selection and Return Migration

The framework above explains selection of immigrants but does not address return migration. Borjas and Bratsberg (1996) use the one-dimensional Roy model to explore the selection of emigrants and those who return. We will briefly introduce the main features of their model and extend it below to a multidimensional skill model.

Using our notation, log earnings in the origin and destination country in their model are given by $Y_O = \mu_O + u_O$ and $Y_D = \mu_D + u_D$, respectively. They further assume that $u_O = cu_D$. Thus, $Var(u_D) = 1$, $Var(u_O) = c^2$, and $Corr(u_D, u_O) = 1$ so that the variance of earnings is higher in the host country when c < 1. They allow workers to have three options: (1) to stay at home, (2) to migrate temporarily, or (3) to migrate permanently. A return migration may be optimal for the same reason we discuss in Section 3.3.4: staying abroad for a period t increases human capital that is valuable at home by an amount κ . Thus, earnings when emigrating and returning are given by $Y_{DO} = t(\mu_D + u_D) + (1 - t)(\mu_O + u_O + \kappa)$. No migration will take place if $Y_O > Y_D$ and $Y_O > Y_{DO}$; a permanent migration will take place when $Y_D > Y_O$ and $Y_D > Y_D$. We illustrate the ensuing regimes in Fig. 4.7 for the case where c < 1 (which is the case where the variance of earnings—and the price of skills—is higher in the destination country).

The figure shows the distribution of skills, where those with the lowest skills (below the threshold $\mu_O - \mu_D - \kappa(1-\tau)/\tau$) will decide not to emigrate, those with the highest skills (above the threshold $\mu_O - \mu_D + \kappa$) will decide to emigrate and remain permanently, and those between the two thresholds will decide to emigrate but return home after spending time *t* abroad. Thus, those who return have higher skills than those who have not emigrated, but lower skills than those who decide to remain permanently.



Figure 4.7 Selection with Return Migration.

Further, an increase in the rent on human capital in the country of destination μ_D will lead to a shift in the thresholds to the left, whereas an increase in the value of human capital acquired abroad back in the home country, κ , leads to a widening of the distance between the two thresholds and thus to increased return migration. It is easy to show that selection of emigrants and return migrants will be exactly the opposite when the price of skills is higher in the country of origin. We will now extend this model to a multidimensional setting.

4.3. Learning, Multidimensional Skills, and Return Migration

The multidimensional model described in Section 4.1 is a static model, in the sense that it assumes that the skills individuals have cannot be augmented. In Section 3, we illustrate—within a one-dimensional skill framework—that many migrations take place for the purpose of skill accumulation. Student migrations, for example, have increased by more than 80% between 1999 and 2008, constituting an increasingly important component of international migration as illustrated in Fig. 4.1.

In a recent paper, Dustmann, Fadlon, and Weiss (2010) extend the two-dimensional framework and develop a dynamic Roy model with learning, where migration and return migration decisions do depend (for given skill prices) not only on the skills individuals possess, but also on the learning opportunities in the origin and destination countries, that is, where skills can be acquired more efficiently. As in Section 4.1, they allow skills to command different prices in different countries but, in addition, add the possibility that individuals can accumulate skills in different countries at different rates.

We will briefly sketch their model using the same notation as earlier. Human capital is again an aggregate that summarizes individual skills:

$$Y_{j} = \ln \gamma_{j} = \mu_{j} + \ln K_{j}(t) = \mu_{j} + b_{j1}S_{1}(t) + b_{j2}S_{2}(t),$$

where j = O, D, and where the skills S now carry a time index as they can be augmented in a "learning by doing" way. Dustmann, Fadlon, and Weiss (2010) concentrate on the case where $b_{D1} > b_{O1}$, $b_{O2} > b_{D2}$, which allows for nonhierarchical sorting.

Individuals accumulate skills S_1 and S_2 while working. However, the extent of human capital accumulation differs between the two countries, due to different learning rates and different prices for the two skills. Denoting the rate of accumulation of skill *S* in country *j* by γ_{jS} and assuming that skill S_1 can be accumulated at a faster rate in country *D*, whereas skill S_2 can be accumulated at a faster rate in country *O*, one obtains $\gamma_{D1} > \gamma_{D2}, \gamma_{O2} > \gamma_{O1}$. Assuming continuous time, a person who works in country *D* accumulates local productive capacity ($K_{D,D}$) and productive capacity applicable in the country of origin ($K_{O,D}$) at rates

$$\frac{\dot{K}_{D,D}}{K_{D,D}} = b_{D1}\gamma_{D1} + b_{D2}\gamma_{D2} \equiv g_{DD}; \quad \frac{\dot{K}_{O,D}}{K_{O,D}} = b_{O1}\gamma_{D1} + b_{O2}\gamma_{D2} \equiv g_{OD}.$$

Further, human capital is accumulated in country O at rates

$$\frac{\dot{K}_{O,O}}{K_{O,O}} = b_{O1}\gamma_{O1} + b_{O2}\gamma_{O2} \equiv g_{OO}; \quad \frac{\dot{K}_{D,O}}{K_{D,O}} = b_{D1}\gamma_{O1} + b_{D2}\gamma_{O2} \equiv g_{DO}.$$

The parameter g measures the rates at which productive capacity for either country can be augmented in each country. For example, the parameter g_{OD} measures the rate at which productive capacity for the origin country can be acquired in the destination country D. This depends on the rate at which the two skills S_1 and S_2 are acquired in country D (γ_{D1} and γ_{D2}), and the prices these skills command in country O (b_{O1} and b_{O2}).

The model is analyzed under certainty, with infinitely long-lived agents, and a fixed interest rate. There is a fixed learning period: learning can take place only until age T, and remains constant thereafter, so that substitution between learning abroad and at home occurs. The time line is given in Fig. 4.8. Individuals are born at 0, emigrate at τ , and have the possibility to return at ε . The length of the learning period is given by T, and the return time ε may be before or after T.

Three cases are distinguished that relate to the intensity at which staying in the host country affects human capital in the two countries, referred to as *partial transferability*,



Figure 4.8 Timeline Return Migration.

strong transferability, and super transferability. We will here briefly consider the first two cases.

Partial transferability characterizes a situation where $g_{DD} > g_{OO} > g_{OD}$, whereas strong transferability characterizes a situation where $g_{OD} > g_{DD} > g_{OO}$. Thus, with partial transferability, experience in the host country leads to less accumulation of human capital that is applicable to the home country than experience in the home country. Furthermore, experience in the home country leads to less accumulation of human capital that is applicable to the home country than experience in the host country leads to accumulation of human capital applicable in the host country. Thus, those who have a high endowment of K_D will emigrate (and they will do so at the first possible opportunity, $\tau = 0$, as this maximizes the amount of human capital that can be accumulated). Furthermore, as experience in the host country enhances human capital in the host country by more than human capital in the home country, those who have decided to emigrate will never return.

With strong transferability, time in the host country will enhance human capital applicable in the home country by more than human capital in the host country and by more than time in the home country. Country D is a "learning center." Again, those who decide to emigrate will do so at $\tau = 0$, but now some of them will return prior to T.

In Fig. 4.9, we display the marginal gain and marginal cost schedules from delaying the return back home for the case of strong transferability.⁵³ The intersection of these two curves presents the optimal return time. The cost of a delayed return rises with the time in the host country, as the migrant's home country human capital K_O increases at a faster rate than his or her host country human capital K_D ($g_{OD} > g_{DD}$). A return will occur if the two schedules cross at $\varepsilon < T$. If the schedules cross at $\varepsilon > T$, Dustmann, Fadlon, and Weiss (2010) show that return will occur either at T, or the migration will be permanent.

Consider now the question who leaves and who will return. In the case of partial transferability, migrations will either not occur, or they will be permanent, as the gap between home and host country human capital will increase with the migration duration. This situation is not dissimilar to the two-dimensional Roy model we have discussed in Section 4.1. The migration decision is based on $S_2(0) > \frac{\Omega}{b_{D2} - b_{O2}} - \frac{b_{D1} - b_{O1}}{b_{D2} - b_{O2}}S_1(0)$, where now the skills S_1 and S_2 are evaluated at $\tau = 0$, and Ω depends on the present value of lifetime earnings and is endogenously determined. As before, who migrates (and the type of selection) depends on the skill prices.

⁵³ Notice that this situation is similar to the one-dimensional model we discuss in Section 3.3.4, where a return is triggered by an increase in human capital that is valuable in the home country. The reason for the increasing marginal gain schedule is that Dustmann, Fadlon, and Weiss (2010) allow for imperfect transferability of productive capacity. They assume that the rent on human capital in the destination country is initially the same as in the origin country (R_O) but converges to the rent on human capital in the destination country (R_D) . The discontinuity in the marginal cost schedule at *T* results from the assumption that learning can only take place until *T*.



Figure 4.9 Return Migration in Two-Dimensional Skill Framework with Learning.

More interesting is the case of strong transferability. In this case, return migration occurs, given the assumptions on skill prices and their accumulation rates in the host and home country ($b_{D1} > b_{O1}$, $b_{O2} > b_{D2}$, and $\gamma_{D1} > \gamma_{D2}$, $\gamma_{O2} > \gamma_{O1}$). Figure 4.10 illustrates the migration and remigration choices in the $S_1(0)$ and $S_2(0)$ space. Those with relatively more $S_1(0)$ (which commands a higher price in the country of destination) will emigrate and, among them, those with relatively more $S_2(0)$ will return home.

In Borjas and Bratsberg (1996), the motive for a return is the same as the one we discuss here: the time in the host country may increase immigrants' home country human capital by more than their host country human capital. However, they do not consider learning (which in our case affects both emigration and remigration decisions), and they consider the return decision in the one-dimensional case. The model collapses to the one-dimensional case by assuming that $S_1 = S_2$ (or, alternatively, by assuming $b_{D1} = b_{O1} = 0$). In that case, the two-dimensional distribution of S_1 and S_2 collapses to a one-dimensional distribution along the 45° line (or along the $S_2(0)$ axis if $b_{D1} = b_{O1} = 0$). It is obvious that in that case, selection will be either positively or negatively selective, according to the relative skill prices—similar to Fig. 4.7 in Section 4.2. For the way Fig. 4.10 is drawn, emigration and remigration will be negatively selective.

4.4. Empirical Studies

The selection of immigrants and return migrants and the effect of migration on the skill base of the origin and destination country are important to understand the consequences of migration for those who do not migrate. The last sections have developed



Figure 4.10 Selection in Two-Dimensional Skill Space with Learning.

a framework within which these processes can be understood and analyzed. We have also discussed some empirical papers that try to assess the direction of selection, and we have suggested some possibilities of how to think about selection in a multidimensional skill world.

In this section, we will discuss some of the empirical papers that study additional ways in which migrations affect nonmigrants and their skills and skill accumulation. Much of this work can be directly related to the theoretical models we discussed in the previous sections.

4.4.1 Emigration and Human Capital Investment in the Sending Country

Emigration, or the possibility of emigration, may have important consequences for the skill accumulation in the origin country through several channels. First, emigrants may redistribute some of the surplus they capture to family members back home by means of remittances, which may then be used for educational investments that would otherwise not have been possible due to a lack of available funds and credit constraints (for a formal model describing this mechanism, see Rapoport and Docquier (2006)). A number of studies suggest that this is happening. Yang (2008), for instance, studies how sudden shocks in exchange rates due to the Asian financial crisis in 1997 affected child schooling and educational expenditures in the Philippines through their effect on remittances, taking advantage of the diverse set of host countries Filipino emigrants are located in. He finds positive effects on human capital accumulation in the origin households. Similarly, Cox Edwards and Ureta (2003) and Acosta (2006) provide evidence

that remittances increase educational attainment of children in El Salvador, and López-Córdova (2005) shows that municipalities in Mexico that receive more remittances have higher literacy and school attendance rates among their 6- to 14-year-old children.

Second, the possibility of a future migration may increase the incentives to invest in further education. This point has first been made by Mountford (1997), and we have illustrated the basic idea in Section 3.3.5: if the return to education is higher in a potential destination country and if there is a positive probability of a future migration, then this will lead to a higher incentive to invest in human capital. Although higher returns to education in the host country have a negative direct effect on the home country's skill base by inducing skilled emigration, it encourages human capital formation in the longer run. Mountford shows that this can potentially be beneficial for the country of origin ("beneficial brain drain"), as long as the probability of an actual future emigration is smaller than 1 (see also Stark, Helmenstein, and Prskawetz (1997); Vidal (1998); and Beine, Docquier, and Rapoport (2001)). Such uncertainty of migration could result, for example, from the unpredictability and restrictiveness of migration policies in potential destination countries.

In a series of empirical studies, Beine, Docquier, and Rapoport (2001, 2008) and Beine, Defoort, and Docquier (2010) assess the possibility of a beneficial brain drain using both cross-sectional and panel data for a large set of developing countries. Their findings provide some evidence that higher emigration rates may indeed have a positive effect on average human capital levels. For example, their work shows that in those countries that are characterized by low levels of human capital, low income and relatively low emigration rates of skilled workers (not exceeding 20-30%), the net effect on the average human capital level of the remaining population is positive, implying a beneficial brain drain effect. Such a positive incentive effect of emigration on human capital accumulation may be further reinforced in the presence of positive externalities of human capital in production (Stark and Wang (2002)) and perpetuated through the intergenerational transmission of human capital from one generation to the next. Schiff (2005) takes a more sceptical view regarding the likelihood and magnitude of a beneficial brain drain through the incentive effect of skilled emigration, a view supported by empirical evidence provided by Lucas (2005) and Checchi, De Simone, and Faini (2007). These studies suggest that in many countries that experience emigration of their skilled workers, the net effect on the average educational attainment of those remaining in the country is likely to be negative.⁵⁴

⁵⁴ This is more likely for those countries where skilled emigration rates are excessively high, such as many sub-Saharan African and Central American countries where they often exceed 40%. See also Marchiori, Shen, and Docquier (2009) who come to this conclusion based on an overlapping-generations general equilibrium model.

Third, selective emigration, as discussed in the previous section, may affect skill acquisition in the origin country by changing the existing skill base, which in turn affects the return to education. For example, if emigrants are predominantly high skilled, then the reduction in the relative supply of high-skilled workers in the home country should lead to an increase in skilled workers' wage rates. Mishra (2007) analyzes this relationship for the Mexican case. Following an empirical strategy first suggested by Borjas (2003), she finds that Mexican emigration to the United States has had a significant positive effect on Mexican wages, a conclusion supported by evidence provided by Hanson (2007) and Aydemir and Borjas (2007). According to Mishra's results, a 10% decrease of Mexican workers in a skill group (measured by education and experience) due to emigration increases average wages in that group by 4%. Since emigrants are disproportionately drawn from the middle of the Mexican skill distribution (high school graduates and those with some college education), wages of mediumeducated workers in Mexico have increased the most as a result of emigration. The resulting increase in the return to education may induce more individuals to invest in additional schooling.

Lastly, as discussed by Dustmann, Fadlon, and Weiss (2010), emigration and return migration may lead to an increase in the productive human capital stock in the sending country if a sufficiently large fraction of the migrations are temporary and the returning migrants accumulate valuable human capital while being abroad. One channel through which human capital of return migrants may lead to economic growth in their home countries is through facilitating knowledge and technology spillovers from the typically more advanced host countries. Domingues Dos Santos and Postel-Vinay (2004) derive the theoretical conditions required for such an overall positive effect to occur.

Although each of the four channels presented may individually lead to a positive effect of migration on educational attainment in the origin country, there are also counteracting factors, likely to depend on the particularities of the considered migration situation that tend to reduce educational attainment. For example, although the positive income effect through remittances may well alleviate credit constraints and lead to higher investments in education, the absence of a parent, and in particular of a mother, is likely to negatively affect overall parental inputs into the children's development. It may also force children to undertake additional household chores or other work to help maintaining the household. Cortés (2010), for instance, provides evidence that children of migrant mothers in the Philippines are approximately 10 percentage points more likely to be lagging behind in school compared with children of migrant fathers.

Also, the possibility of a future migration may create opposite incentive effects if the return to education in occupations potential emigrants consider as attainable is lower in the destination country than in the origin country (this could be due to a high return to

trade (manual) skills and low returns to academic skills), or if migration is seen as an alternative to the acquisition of education. Due to these counteracting factors, the overall effect of migration on human capital acquisition in the home country is a priori ambiguous. McKenzie and Rapoport (2010) empirically assess this net effect of migration for the Mexican case based on data from the National Survey of Demographic Dynamics using historical migration rates to instrument for current migration. Their findings show a significant negative effect of migration on school attendance and educational attainment of 12- to 18-year-old boys and 16- to 18-year-old girls: living in a migrant household is estimated to lower the probability of completing high school by 13% for males and 14% for females. These effects are somewhat mitigated for children from a poorer background, which is consistent with a more prominent role of credit constraints in these households. Overall, most of the negative effect of migration on educational attainment is due to young males migrating themselves rather than attending school and young females dropping out of school to work at home.⁵⁵ For a detailed discussion of this literature, including earlier work, see Commander, Kangasniemi, and Winters (2004) and Docquier and Rapoport (2009).

4.4.2 Immigration and Human Capital Investment in the Receiving Country

In the last section, we considered the possible channels by which emigration can affect skill accumulation and education of those in the country of origin who do not emigrate. We now turn to the destination country. Again, there are various channels by which immigration may affect the accumulation of skills and education. First—and similar to what we discussed in the previous section—immigration may affect the price of skills by changing the relative factor supplies in the host country's economy. For example, a large inflow of low-skilled immigrants could lead to a decrease in the relative wages of low-skilled workers and an increase in the relative wages of high-skilled workers. Most of the existing papers do not find large wage effects of immigration, although there is still a lot of controversy in the literature investigating this issue, which includes studies by Card (1990); Altonji and Card (1991); Butcher and Card (1991); Borjas, Freeman, and Katz (1997); Card (2001); Friedberg (2001); Borjas (2003); Dustmann, Fabbri, and Preston (2005); Manacorda, Manning, and Wadsworth (2011); Ottaviano and Peri (2011); Dustmann, Frattini, and Preston (2008); and Glitz (2011).⁵⁶

In the simplest model of an economy with one sector and unskilled and skilled labor, the effect of immigration depends on assumptions about the elasticity of capital supply and the share of skilled immigrants relative to the share of skilled workers in

⁵⁵ See also De Brauw and Giles (2006) and Antman (2009) who similarly find a negative overall effect of migration on investments in education in China and Mexico, respectively, and Hanson and Woodruff (2003) who find a positive effect on the schooling of less-educated mothers' teenage daughters in Mexico.

⁵⁶ For a comprehensive overview of this literature, see Okkerse (2008).

the native population. We have seen in Table 4.1 in Section 2 that this share differs substantially across countries so that we cannot expect immigration to have the same effects in different countries. If immigration is unskilled and affects skill prices by, for example, generating higher returns to education, then this may create incentives to invest in education. Findings for the United States by Jackson (2010) show that a 1% increase in relatively unskilled immigrant labor increases the rate of native college enrolment by 0.33%. This crowd-in effect is driven primarily by natives aged 18–24 because of their higher sensitivity to the returns to college education relative to older natives. In a more specific case, if immigrants are complementing labor supply of highly skilled women by providing domestic services that are otherwise not available or considerably more costly, then this is likely to increase the return to higher education for women. Such immigration should then lead to not only higher female labor supply (as shown by Cortés and Tessada (2009); Cortés and Pan (2009); and Farré, González, and Ortega (2009)) but also more human capital investment.

Second, immigration may create incentives for native-born workers to specialize in areas where they have a comparative advantage, for instance through their knowledge of the host country language or of existing networks. Such specialization may well imply the accumulation of additional skills. More generally, if immigrants and natives have different comparative advantages in the labor market, then immigration will lead to shifts in the skill distribution of the native workforce. For example, Peri and Sparber (2009) show that natives reallocate their task supply toward occupations that are more communication and language intensive when faced with an inflow of immigrants that increases the supply of manual and physical labor intensive occupations. Task specialization also extends to the highly educated segment of the labor market, where immigrants with a graduate degree tend to specialize in occupations demanding quantitative and analytical skills, whereas native-born graduates specialize in occupations requiring interactive and communicative skills (Peri and Sparber (2008)).

Third, immigrants may affect the human capital acquisition of natives through their presence in the educational system, both on the tertiary and prior levels. Borjas (2006a) finds that admitting foreign students to doctoral programs has a negative effect on the earnings of native doctoral recipients in the corresponding field. According to his results, a 10% increase in the supply of doctorates in a specific field reduces earnings of competing workers by 3–4%, half of which is due to a shift toward lower-paid post-doctoral appointments. He also finds some evidence of foreign students crowding out white native men (Borjas (2007)) which, although in line with earlier evidence provided by Hoxby (1998), stands somewhat in contrast to other findings showing no significant crowd-out effect of native students (Jackson (2010)). Crowding-out of native students has also been found at the secondary school level. Betts and Fairlie (2003) found evidence that for every four immigrants who arrive in public high schools, one native student switches to a private school and that this "native flight" is particularly

pronounced among white native students and in response to the arrival of non-Englishspeaking immigrant children. Gould, Lavy, and Paserman, (2009) show that the presence of immigrants during elementary school has a negative long-term effect on the probability of passing the high school matriculation exam in Israel, which enables students to attend college. Neymotin (2009), on the other hand, provides evidence that immigration did not negatively affect the SAT scores of native high school students in California and Texas.

Finally, in a wider context, immigration may affect the stock of human capital in the host country by its contribution to new innovations. For instance, Hunt and Gauthier-Loiselle (2010), using the 2003 National Survey of College Graduates, show that the large number of immigrants with science and engineering degrees in the United States add significantly to the number of patents granted: a single percentage point increase in the immigrant college graduates' population accordingly leads to an increase in patents per capita by 9–18%. Hunt (2010) and Kerr and Lincoln (2010) find complementary evidence regarding immigrants' contribution to general productivity-increasing activities such as patenting, publishing, and company start-ups.

5. THE CHILDREN OF IMMIGRANTS

In the previous sections, we have considered the relationship among education, skill accumulation, and migration. We focused on the way education interacts with migration decisions and how migration affects skill accumulation and the skill base of those populations who do not migrate. In this context, we have touched on the educational achievements of the children of emigrants in the home countries and on the incentives to invest in education for young people in the host countries as a result of immigration. In this section, we will focus explicitly on the children of immigrants in the countries of destination. We will study their educational achievements and ensuing labor market outcomes and relate these to their parent generation and their peers born to native parents.

The educational achievement of the children of immigrants is one of the key issues in the immigration debate in many countries. Underachievement of immigrant children is often seen as a major factor for the long-term segregation of immigrant communities, and educational achievements of immigrant children in comparison to their native-born peers are considered an important indicator of successful immigration policy (see, e.g., OECD (2006)). As Table 4.9 in Section 2 shows, the test score results for children with a migration background, based on the PISA data, are indeed on average below those of children of native-born parents, although there is quite some variation across countries. This is in line with the limited set of findings for individual countries that are based on alternative standardized performance measures. For instance, Rampney, Dion, and Donahue (2009) show that the reading (mathematics) score gap in the National Assessment of Educational Progress (NAEP) between White and Hispanic students in the United States in 2008 was 9.2 (6.4)% for the 9-year-olds, 9.7 (7.9)% for the 13-year-olds, and 8.8 (6.7)% for the 17-year-olds.

As Fig. 4.2 shows, educational achievement of immigrant children is related to the educational achievement of their parents. This suggests that the relatively low achievement of immigrant children in some countries is at least partly due to the lower educational background of their parent generation and does not necessarily reflect the failure of the host society to educate second-generation immigrants. However, in the public debate, the comparison between children born to immigrants and native parents is often *unconditional* on parental characteristics—a comparison that seems not very meaningful. If the first generation of immigrants is less educated than the native population (which is the case for a number of countries, see Section 2), then even a strong educational progress of immigrant children may still result in educational outcomes that are lower than those of natives. We will discuss this issue below.

How should we then think about the *dynastic* assimilation of immigrant communities? What is it exactly that determines the educational outcomes of immigrant children? There are at least three factors that matter: First, as just argued, the educational achievements of immigrant children are linked to their parental background. There is a large literature on the intergenerational transmission of human capital, for example Behrman and Taubman (1985); Solon (1992); Zimmermann (1992); Björklund and Jäntti (1997); Corak and Heisz (1999); or Blanden, Goodman, Gregg, and Machin (2004), which shows that there is considerable intergenerational immobility across generations in a large number of outcomes.

Second, immigrant children may be differently affected by the institutional setting and support structures of the host countries' education systems. Another large literature in economics studies the different features of educational production and how it relates to resources, institutions, and parental input (see, e.g., Todd and Wolpin (2003); Cunha, Heckman, Lochner, and Masterov (2006); or Hanushek and Woessmann (2008, 2010)). Even without differential access to these educational institutions across populations in the same country, the same resources may affect immigrant children's educational outcomes in a different way than those of children born to native-born parents—for example, due to complementary forms of knowledge that are not sufficiently available in immigrant communities⁵⁷—and prevent immigrants' children from making full use of existing educational support structures.

Third, the social context in which immigrant children grow up is an important determinant of their educational outcomes. Borjas (1992, 1995a) was among the first to emphasize what he calls the effect of "ethnic capital" on the educational achievements

⁵⁷ One of those is, for instance, language. Dustmann, Machin, and Schönberg (2010) illustrate that language is the key factor that holds children of ethnic minority background back in the first years of full-time education.

of immigrant children. The ethnic context of immigrants shapes their own, and their parents' values and incentives. For instance, immigrant children who grow up in an ethnic environment that is characterized by particular educational and occupational choices are likely to make similar choices. Ethnic capital summarizes things, such as peer effects, role models, and community ambitions, which may all be strongly related to immigrants' educational outcomes. This may reduce the overall educational attainment of second-generation immigrants, but it may also enhance it. A good example is the success of south and south east Asian immigrant children that may be partly due to the high value attached to education within their communities (which in turn may be related to the Confucian tradition, which emphasizes the importance of study). In the United Kingdom, for example, Chinese students, who at the beginning of their compulsory schooling at age 6/7 start off with a 7.6% of a standard deviation lower English test score than White British students, outperform their native counterparts by more than 50% of a standard deviation by the time they reach age 15/16 (Dustmann, Machin, and Schönberg (2010)).

There are other important factors that relate to the educational achievement of immigrant children. Parental considerations that affect their own investment in learning may also influence decisions about their children. For instance, we have seen that in the case of temporary migrations, and if skills acquired in the host country are only partially transferable to the home country labor market, immigrants will invest less in their human capital than they would in case of a permanent migration. In the same way, it is not unlikely that parental decisions about the educational investments of their children are affected by where parents see their children's future. For instance, if a migrant household intends to return to the home country, and if this planned return migration is including the children, then this may induce parents to invest less in their children's education than they would do in the case of a permanent migration. Dustmann (2008) has made this point and provides empirical evidence that supports this hypothesis.

As earlier, we will commence with a formal discussion, laying out some of the key issues. We will then review and interpret the empirical evidence that exists to date.

5.1. Immigrants and Intergenerational Mobility

The empirical approach to study the relationship between outcomes of parents and outcomes of children is to regress a permanent outcome measure of the child on a permanent outcome measure of the parent by applying a Least Squares estimator to the regression equation

$$\log \gamma_{it}^j = \alpha^j + \rho^j \log \gamma_{it-1}^j + \varepsilon_{it}^j, \tag{4.14}$$

where $\log \gamma_{it}^{j}$ and $\log \gamma_{it-1}^{j}$ are some permanent measures for outcomes (such as education, wealth, or earnings) of a child and parent, respectively, belonging to group *j* (which could be immigrants and their children, or natives and their children). For

simplification, in what follows, we will refer to the outcome as "earnings." According to Eq. (4.14), the earnings of family i's child is determined by family i's parental earnings and other influences ε (which we will discuss further below). The parameter α^j can be thought of as the average effect of these other influences on earnings, which can differ between groups. Assuming that the variances of $\log \gamma_{it}^j$ and $\log \gamma_{it-1}^j$ are the same, ρ^j is the population correlation coefficient between $\log \gamma_{it}^j$ and $\log \gamma_{it-1}^j$. Assume that the ε_{it}^j are iid distributed with mean zero and $Var(\log \gamma_{it}^j) = Var(\log \gamma_{it-1}^j) = \sigma_{\gamma}^{2,j}$, so that estimation of Eq. (4.14) gives a consistent estimate of ρ^j , $\hat{\rho}^{j.58}$ The coefficient ρ^j represents the fraction of economic advantage (in terms of earnings, education, or wealth) that is on average transmitted across the generations. It is called the intergenerational correlation coefficient or transmission parameter. A coefficient close to zero suggests high intergenerational mobility, whereas a coefficient close to one indicates low mobility. The coefficient $(1 - \rho^j)$ is often referred to as the degree of regression to the mean.

For immigrants, the study of intergenerational mobility has a particular significance. If first-generation immigrants are disadvantaged, in the sense that they are less well educated or have lower earnings than the native population, then immigrant and native populations may differ for many generations, depending on the magnitude of ρ .

To illustrate this point, consider Eq. (4.14) and index outcomes of immigrants and natives by *I* and *N*, respectively. Further, allow the intergenerational transmission parameter to differ between the two groups so that $\rho^N = \rho^I + \zeta$. Then, the wage differential between the two populations in generation *t* is given by

$$E(\log \gamma_t^N) - E(\log \gamma_t^I) = \alpha^N - \alpha^I + \rho^N(E(\log \gamma_{t-1}^N) - E(\log \gamma_{t-1}^I)) + \zeta E(\log \gamma_{t-1}^I).$$
(4.15)

Consider the case where $\zeta = 0$ (intergenerational transmission ρ is the same in the two populations) and assume for simplicity that $\alpha^N = \alpha^I$. In this case, the nativeimmigrant gap in outcomes disappears from one generation to the next only if $\rho = \rho^N = \rho^I = 0$. If $\rho = 1$, the initial outcome differential will be fully transmitted to the next generation. If ρ is smaller than 1, inequality between the two groups will fall and assimilation across groups will take place, but not within one generation. The magnitude of ρ determines the speed of convergence. For example, for $\rho = 0.4$, a 20% average earnings disadvantage for immigrants in the parent generation translates into an 8% earnings disadvantage in their children's generation. Now, suppose "other influences" determining outcomes as captured by the parameter α differ across the two groups. If $\alpha^N - \alpha^I > 0$, the difference in outcomes in the next generation may even be larger than in the parent generation, despite regression to the mean, as indicated by $\rho < 1$.

⁵⁸ If the variance of log wages differs across the two generations, the OLS estimator $\hat{\rho}$ measures $\rho \sigma_{yt} / \sigma_{yt-1}$.

Now, consider the case where $\zeta \neq 0$: the intergenerational transmission parameter differs between the two groups. It follows from Eq. (4.15) that if $\zeta > 0$. (i.e., intergenerational mobility in the advantaged groups, natives, is smaller than in the disadvantaged group), outcome differentials in the next generation may even be larger than those in the previous generation despite regression to the mean in both groups. The degree of assimilation between the two groups across generations depends on the parameters ζ , ρ^N , and $\alpha^N - \alpha^I$.

In general, earnings of immigrants in the next generation will converge to the earnings of natives (in the sense of $E(\log \gamma_t^N) - E(\log \gamma_t^I) < E(\log \gamma_{t-1}^N) - E(\log \gamma_{t-1}^I))$ if

$$\frac{\alpha^{N} - \alpha^{I}}{(1 - \rho^{N})E(\log \gamma_{l-1}^{I})} + \frac{1 - \rho^{N} + \zeta}{(1 - \rho^{N})} < \frac{E(\log \gamma_{l-1}^{N})}{E(\log \gamma_{l-1}^{I})}$$

Thus, if mean log earnings of natives are larger than those of immigrants in generation t-1, there will always be convergence as long as $\alpha^N - \alpha^I = 0$, $\zeta = 0$ and $\rho^N = \rho^I = \rho < 1$. On the other hand, both a higher ζ (less intergenerational mobility of natives) and a higher $\alpha^N - \alpha^I$ (stronger effect of other influences on native earnings than on immigrant earnings) lead to slower convergence. For sufficiently high values of either of these parameters, the difference in expected earnings between immigrants and natives in the next generation could be larger than in the previous generation.

What is the interpretation of the parameters ρ and α , and how do they relate to an underlying structural model and its parameters? This is what we will explore next. We will show why these parameters are likely to differ between immigrants and natives and generate different intergenerational correlation coefficients and intercepts and hence different intergenerational persistence of outcomes for these groups. We will also demonstrate why the assumption that ε_{it} is iid is unlikely to hold in reality, which may affect the actual estimation of the parameter ρ .

5.2. A Model of Intergenerational Mobility of Immigrants

Becker and Tomes (1979) develop a theoretical model for the intergenerational transmission of wealth and human capital, and Solon (2004) provides a parameterization that derives an intergenerational transmission function of the type illustrated above. In the following, we will draw on Solon's formulation, simplify, and adjust it to emphasize what we believe are some key insights for the study of educational attainments of the children of immigrants in comparison with natives. The model is a permanent income model of intergenerational mobility with parental investments in the child's earnings potential. Consider a one-person household with one child, situated in the host country. There are two periods. In the first period (period t - 1), the parent has earnings equal to y_{t-1} and the child is in full-time education. In the second period (period t), the parent retires while the child participates in the labor market and has earnings γ_t , which depend in magnitude on investments in education in the first period. The parent is altruistic and maximizes an intertemporal utility function by choosing first-period savings S_{t-1} , and investment in the child's human capital in the first period, I_{t-1} :

$$V = u(c_{t-1}) + \delta[u(c_t) + \gamma \log \gamma_t], \qquad (4.16)$$

where u(.) is the parent's utility from consuming c_t and c_{t-1} in periods t and t-1, respectively, and δ is a discount factor. The parameter γ is an altruistic weight. If $\gamma = 0$, the parent does not consider the child's earnings in period t. Assume that parental investments I_{t-1} translate into human capital of the child (H_t) according to the following production technology:

$$\log H_t = h = \theta \log I_{t-1} + e_t. \tag{4.17}$$

The parameter θ is a technology parameter measuring the productivity of investments. This parameter can be viewed as the "talent" or the "ability" of the child but may also be related to institutional settings and school quality. The term e_t is the human capital the child receives without any direct parental investments. This term represents the attributes endowed upon the child, depending on characteristics of the parents, the child's upbringing, genetic factors, environment, and luck. It may also depend on existing networks, as well as the lack of opportunity to move out of social and economic structures from one generation to the next. Becker and Tomes (1979) refer to this term as endowments of capital that "are determined by the reputation and 'connections' of their families, the contribution to the ability, race, and other characteristics of children from the genetic constitutions of their families, and the learning, skills, goals, and other 'family commodities' acquired through belonging to a particular family culture. Obviously, endowments depend on many characteristics of parents, grandparents, and other family members and may also be culturally influenced by other families." The influence of cultural factors and family background may be particularly important for immigrants, and we will discuss the implications in Section 5.2.3. These factors include what Borjas (1992) calls "ethnic capital," the quality of the environment in which parental investments are made.

Human capital translates into earnings of parents and children according to the following relationship:

$$\log \gamma_{t-1} = \mu_{t-1} + r_{t-1}h_{t-1}, \qquad (4.18a)$$

$$\log y_t = \mu_t + r_t h_t. \tag{4.18b}$$

Our formulation allows for different "rental rates" on human capital in the different periods, μ , as well as different returns to human capital, *r*. It follows from

Eqs (4.17) and (4.18.b) that the child's earnings in the second period are related to parental investments by

$$\log \gamma_t = \mu_t + r_t \theta \log I_{t-1} + r_t e_t. \tag{4.19}$$

The parent's consumption in period t-1 equals $c_{t-1} = y_{t-1} - I_{t-1} - S_{t-1}$, where y_{t-1} are earnings in period t-1. For simplicity, we assume that the parent cannot borrow against the child's future earnings and does not bequeath financial assets to the child. As the parent retires in period t, period t consumption is equal to period t-1 savings, $c_t = S_{t-1}$.

Choosing a simple logarithmic utility function for consumption and substituting Eq. (4.19) for the child's earnings in Eq. (4.16), the optimization problem of the parent can be expressed as

$$\max_{S, I} V = \log(\gamma_{t-1} - I_{t-1} - S_{t-1}) + \delta[\log S_{t-1} + \gamma(\mu_t + r_t \theta \log I_{t-1} + r_t e_t)].$$
(4.20)

Maximizing Eq. (4.20) with respect to savings and investment, and solving the first-order conditions for the optimal investment I_{t-1} yields

$$I_{t-1} = \frac{\delta \gamma \theta r_t}{(1+\delta) + \delta \gamma \theta r_t} \gamma_{t-1}.$$
(4.21)

The term in the numerator (which is the same as the second term in the denominator) is the expected discounted utility gain to one log unit of parental investment in the child's human capital. The first term in the denominator is the expected lifetime utility gain from one log unit of additional lifetime consumption. Thus, investments in the child as a fraction of income equal the expected fraction of utility resulting from this investment. Simple comparative statics show that investment in the child's human capital increases with altruism γ , the productivity of investment θ , the return to human capital r_{t} , and the discount rate δ .

Human capital and earnings of the child are related to human capital and earnings of the parent as follows:

$$h_{t} = \theta \log \left[\frac{\delta \gamma \theta r_{t}}{(1+\delta) + \delta \gamma \theta r_{t}} \right] + \theta \mu_{t-1} + \theta r_{t-1} h_{t-1} + e_{t}$$
(4.22a)

$$\log \gamma_t = \mu_t + r_t \theta \log \left[\frac{\delta \gamma \theta r_t}{(1+\delta) + \delta \gamma \theta r_t} \right] + \theta r_t \log \gamma_{t-1} + r_t e_t.$$
(4.22b)

Equations like (4.22a) and (4.22b) are usually estimated in the literature when regressing education (or earnings) of children on those of their parents. Consider first Eq. (4.22a). The level of education obtained by the child depends on all the parameters that affect investment. These may differ between immigrants and natives. If, for

instance, the rate of return to skills r_t , is perceived to be lower for immigrants, the level of human capital acquired by immigrant children will also be lower as will be their earnings. Neal (2006), for example, discusses statistical discrimination as one reason that has frequently been brought forward to explain the low attainment levels of black youth in the United States, in the sense that employers are not likely to view them as skilled no matter what their level of education is. Also, a lower price for human capital, μ_{t-1} , in the parent generation of immigrants relative to natives due to, for instance, a lack of important complementary skills such as language, leads to a lower level of education of their children. Finally, education of children depends on "inherited" traits and institutional features such as access to schooling and school quality, which are all captured in the term e_t . If these differ between the immigrant and native population, both their education and earnings may differ due to this channel, too (see, e.g., Parent (2009)).

Suppose now that we regress log earnings (or education) of immigrant children on log earnings (or education) of their parents, following much of the literature that we will discuss below. What does the OLS coefficient we estimate measure? Following Becker and Tomes (1979) and Solon (2004), assume that e_t follows an AR(1) process, reflecting a serial correlation in the parent's and the child's human capital endowments, so that $e_t = \lambda_0 + \lambda_1 e_{t-1} + v_t$, where v_t is a white noise error term and $\lambda_1 \in (0, 1)$. As discussed earlier, these endowments may include ability but may also be related to networks, ethnic reference groups, or other "ethnic capital." Then, in steady state, the probability limit of the OLS estimate of the coefficients on h_{t-1} and γ_{t-1} are given by $\frac{\theta_{t-1} + \lambda_1}{1 + \theta_{t-1} + \lambda_1}$ and $\frac{\theta_{t} + \lambda_1}{1 + \theta_{t} + \lambda_1}$, respectively. Thus, the coefficient estimate of ρ in the simple model we discussed at the beginning of Section 5.1 is larger the larger the return to human capital and the productivity of human capital production, r_t and θ , and it also increases in the correlation in heritable traits, λ_1 . All these parameters can differ between populations. For instance, if the heritability parameter is larger in the immigrant population because family structures are tighter, the intergenerational correlation coefficient will also be larger for this population, implying less mobility from one generation to the next.

5.3. Empirical Evidence

5.3.1 Schooling Outcomes of Immigrant Children

We start in this section with the first important period in an immigrant's life, his or her childhood. To this end, we return to the data from the PISA study that we already introduced in Section 2. As we have seen in Table 4.9, with the exception of Australia and Canada, the student population with immigrant background tends to score significantly lower than the native population in both mathematics and reading. One of the main explanations for these differences could be the lack of fluency in the host country language. The last two columns of Table 4.9 indeed show that students who speak a different language at home than the language of instruction at school fare particularly badly. In addition to language, the general skill level of the parents is likely to play a major role in the ability of their children to acquire human capital and, given that in most host countries the immigrant population is less educated than the native population, may contribute significantly to any observed differences in the raw test scores.

Table 4.11 shows a number of descriptive statistics for our 10 most important immigrant-receiving OECD countries. For the sake of brevity, the immigrant student populations we consider include all children whose parents are both foreign-born, no matter whether the children themselves were born in the host country.⁵⁹ As Column (4) shows, in many countries the majority of immigrant students speak a language at home that is different from the language of instruction at school.⁶⁰ This pattern is particularly pronounced in Italy and the United States, where the share of those speaking a foreign language at home exceeds 60%. Given the complementarity between language and human capital accumulation, this is bound to affect the performance of these children in the different proficiency tests. Columns (5) and (6) show that relative to the native student population, immigrant students have parents with, on average, significantly lower educational attainment. With the exception of Italy and Spain, the share of native students whose both parents have low educational attainment (measured as not exceeding lower secondary education), is fairly low, ranging between about 2 and 15%, whereas it ranges between 10 and 40% for immigrant students (not considering Australia and Canada). The differences in parental educational attainment are particularly large in France, Switzerland, and the United States. These results also hold when looking at alternative measures of economic status, such as the Highest International Socio-Economic Index of Occupational Status (HISEI) of the parents (not reported), which is designed to capture those features of occupations that convert education into income (Ganzeboom, De Graaf, de Leeuw, and Treiman (1992)).

So how do these differences in language and family background contribute to the measured raw test score gaps between natives and immigrants? Column (7) shows again the raw proficiency gaps in mathematics between immigrant students and native students. As we already discussed in Section 2, immigrant students do substantially worse in all destination countries except Australia and Canada, with the largest gaps arising in Germany, the Netherlands, and Switzerland. However, controlling for language

⁵⁹ The results we present in this section draw on work by Dustmann, Frattini, and Lanzara (2010).

⁶⁰ Whether a student speaks a foreign language at home is obviously only an imperfect measure of language skills as it may very well be that a student is proficient in the language of instruction but still speaks another language at home with his or her parents.

	% of Stu	dent Population	% Foreign	Language at Home	% Low E	ducation Families	PISA Tes	PISA Test Score Gap (Mathematics)		
Destination Country	Natives	Immigrants	Natives	Immigrants	Natives	Immigrants	Unconditional	Conditional on Language	Conditional on Language & Parents' Education	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	
Australia	60.7	21.2	0.3	33.9	11.0	9.3	15.7**	15.3**	11.1**	
Canada	67.9	20.7	0.3	48.3	2.4	6.3	-2.7	-1.8	-3.2	
France	76.8	19.7	0.3	34.8	10.9	38.9	-38.9**	-32.5**	-21.0*	
Germany	83.8	11.2	0.4	55.0	11.7	24.8	-58.1**	-31.7**	-29.3**	
Italy	91.0	6.4	0.1	66.9	27.3	22.1	-38.6**	-28.0*	-33.2	
Netherlands	82.5	10.6	0.1	46.1	8.3	29.7	-53.8**	-52.4**	-46.3**	
Spain	89.1	6.9	0.3	33.0	37.5	25.3	-47.8**	-49.3**	-55.4**	
Switzerland	62.5	3.4	0.1	56.9	14.9	35.8	-69.5**	-47.1**	-44.0**	
United Kingdom	84.3	9.1	0.1	38.0	4.2	10.5	-14.6*	-6.7	-7.2	
United States	80.1	13.7	0.4	62.3	2.4	25.0	-22.9**	-2.2	4.8	

Table 4.11 PISA Outcomes—Summary Statistics and Regression Results

Source: PISA 2006. Columns (1) and (2) report the share of native and immigrant students in the PISA student population. Native students are defined as those born in the country of assessment with both parents as well born in the country of assessment. Immigrant students are those born either in the country of assessment or in another country with both parents foreign-born. The missing category is students with mixed background. Columns (5) and (6) show the percent of families where both parents have low education. Values of summary statistics are computed using the final weights provided by PISA. Columns (7) to (9) report the proficiency gaps in mathematics of immigrant relative to native students. The values are the estimated coefficients of a regression of PISA scores on a dummy for immigrant status (the omitted category is natives). Column (7) reports unconditional regressions; Column (8) adds a dummy for determining whether the language of assessment is spoken at home; and column (9) adds dummies for the higher educational level of either parent ("low": no schooling, primary education, lower secondary educatior; "high": tertiary education, postgraduate education). The regressions are run separately for each country. All coefficients and standard errors are estimated according to the Unbiased Shortcut procedure (OECD (2009)), using the replicate weights provided by PISA. Stars indicate that the difference between the immigrant and the native average score is statistically significant at the 1% level (*).

reduces these gaps significantly as shown in Column (8). The reduction is particularly large in Germany, Italy, Switzerland, the United Kingdom, and the United States. In the latter two countries, including an indicator for language spoken at home actually closes the test score gap entirely, with none of the remaining small differences being statistically significant. Adding control variables for the educational attainment of the parents in Column (9) has a further mitigating effect on the test score gaps between natives and immigrants, in particular in France, the Netherlands, and the United States. The only major exception is Spain where the test score gap actually widens once controls for parental education are included, owing to the, on average, better educational background of immigrant children's parents (compare Columns (5) and (6)). Very similar patterns hold for immigrant and native students' reading proficiency (not reported) where, naturally, the impact of language spoken at home is an even stronger determinant of performance. Both language and the educational attainment of the parents thus go a long way in explaining the large gap in mathematics and reading skills between native and immigrant students (for additional evidence, see e.g., Entorf and Minoiu (2005) and Schnepf (2007)).

The significance of language spoken at home for the, at least initial, achievements of immigrant children at school is also documented in other work. In a recent paper based on the British school census, Dustmann, Machin, and Schönberg (2010) investigate the school curriculum of children from ethnic minority backgrounds and compare it to children from nonminority families, from the age of 5 until the age of 16. Their findings show that just before they start school, ethnic minority children significantly underperform in early cognitive tests compared with white British-born children. However, by the end of compulsory schooling at age 15/16, most ethnic minority groups catch up with (Bangladeshi, Pakistani, and black non-Caribbean pupils) or even overtake (Indian and Chinese pupils) their white British counterparts (in key stage 4 tests). The analysis shows that improvements in their proficiency of the English language is the single most important contributor to the catch-up of ethnic minority pupils relative to white British pupils, accounting for up to two-thirds of their relative progress. The importance of language proficiency, in particular, for school performance and integration more generally has been recognized by many host countries by implementing a variety of policies and practices to support immigrant students' acquisition of the language of instruction. For a detailed overview of these policies across OECD countries, see OECD (2006).

5.3.2 Intergenerational Transmission of Human Capital

The results presented in Table 4.11 demonstrate that parental background and language spoken at home matter importantly for the educational success of the children of immigrants. We now briefly review the empirical evidence regarding the intergenerational

transmission of human capital and the long-term assimilation of second-generation immigrants.⁶¹

Chiswick (1977) and Carliner (1980) were the first to look at the intergenerational aspect of immigrant earnings in the United States. These studies do not relate one generation's earnings to the earnings of its parent generation in the spirit of Eq. (4.14) but compare directly the earnings of different generations of immigrants using crosssectional data from the 1970 US Census. They distinguish male first- and secondgeneration immigrant workers and native workers, with the latter defined as individuals who have two native-born parents. Both studies' key finding is that the earnings of second-generation immigrant workers are higher than those of natives.⁶² In addition, Carliner (1980) finds that second-generation immigrant men also earn more than first-generation immigrant men from the same ethnic group. Pointing again to the substantial changes in cohort quality over the course of the twentieth century and the bias this can induce in cross-sectional analyses, Borjas (1993) employs grouped data from the 1940-1970 Censuses to isolate cohort effects from intergenerational earnings mobility. He compares the relative earnings of foreign-born men in 1940 with the relative earnings of their potential offspring 30 years later, in 1970. His findings show that the earnings of second-generation immigrants are strongly correlated with the earnings of the corresponding first generation, with an estimate of the intergenerational correlation coefficient of around 0.45. Hence, about half of the differences in relative economic status across different ethnic groups in one generation persist into the next generation. Using more recent data reaching up to the year 2000, but based on a similar methodology of grouping immigrants and their potential offspring, Card, DiNardo, and Estes (2000) and Borjas (2006b) show that the intergenerational correlation of earnings has remained relatively unchanged over the last decades: native-born children of immigrants can expect to close 50-60% of the gap in relative earnings experienced by their father's generation. Card, DiNardo, and Estes (2000) also estimate the intergenerational correlation in the years of education obtained and find a very stable estimate over time for both sons and daughters in the range of 0.41-0.47. These estimates are comparable with those we report in Fig. 4.2 for the pooled sample of immigrant groups in France, Germany, and the United Kingdom, where the slopes of the regression lines with respect to education and log wages are 0.53 and 0.36, respectively.⁶³ Overall, the

⁶¹ For an overview of the literature on intergenerational mobility, see Solon (1999, 2002), Corak (2004, 2006), and D'Addio (2007).

⁶² Note that Chiswick only looks at white second-generation immigrant men who, at the time of the study, had predominantly a European background.

⁶³ Dustmann and Theodoropoulos (2010) analyze the educational attainment and economic behavior of ethnic minority immigrants and their children in Britain and compare it with that of their white British-born peers, showing that Britain's ethnic minority immigrants and their children are on average better educated than their white native-born peers, and that groups with better educated parents have higher levels of education.

empirical evidence suggests that most of the strong intergenerational linkages between immigrant fathers and their native-born children work through education.

With increasing data availability and the passing of time, it has become possible in some cases to analyze the correlation between first- and third-generation immigrants. Using data from the 1910 US Census and the General Social Surveys to link a sample of American-born workers to their grandparents who arrived in the United States during the first Great Migration at the beginning of the twentieth century, Borjas (1994, 2006b) estimates an intergenerational correlation in relative wages between the first and the third generation of about 0.22, which implies that 22% of the wage gap between any two groups in the immigrant generation persisted into the third generation. Note that this is compatible with a correlation coefficient across subsequent generations of about 0.5, similar to the one estimated in the studies discussed above. Current differences in economic status among first-generation immigrants are thus likely to shape the labor market experience of their offspring for generations to come.⁶⁴ Table 4.L2 at the end of this chapter summarizes the existing literature on the intergenerational mobility in earnings and educational attainment of immigrants across a variety of different countries. Column (7) reports estimated intergenerational correlation coefficients based on specifications such as the one in Eq. (4.14). Overall, the evidence suggests that intergenerational mobility is lower for immigrants than for natives, higher for immigrant women than for immigrant men, relatively high in Scandinavian countries (with an estimated correlation coefficient for men of around 0.1–0.2), relatively low in the United States (with estimates of around 0.5–0.6), and somewhere in the middle in Canada and Germany (with estimates of around 0.2-0.4).

As we discussed earlier, the investment decisions of immigrants in their own education, and that of their children, may be shaped by their return intentions. Dustmann (2008) extends the standard permanent income model of intergenerational mobility as sketched in Section 5.2 by allowing for the possibility of return migration. As we have already alluded to on various occasions throughout this chapter, the prospect of returning home has important consequences for an immigrant's human capital investment in the host country. In an intergenerational context with altruistic parents, such considerations also affect the parents' investment in the human capital of their children (under the assumption that the child's perceived return probability increases with the return probability of the parent) and thus the measured intergenerational earnings mobility. More specifically, Dustmann (2008) shows that as long as the return to human capital is higher and the preference for consumption lower in the host than in the home country, investments in the children's human capital will increase with the probability of a permanent migration. This is because the latter increases the expected monetary gain from an additional unit of human capital for the child, and because it decreases the expected

⁶⁴ Deutsch, Epstein, and Lecker (2006) and Hammarstedt (2009) provide evidence on the relative outcomes of first-, second-, and third-generation immigrants in Israel and Sweden, respectively.

utility gain from consuming in the home country, leading to a reduction in parents' savings for future consumption and an increase of their investments in their children. Using German panel data on father–son pairs that include information on parental return intentions, Dustmann (2008) shows, first, that parental investment in children's education increases with the permanent migration probability of the parent and, second, that the son's permanent earnings increase with the father's permanent migration probability, conditional on father's permanent earnings and education. Accounting for measurement error in parental earnings by using repeated wage observations reveals substantial downward bias in a standard estimation of the intergenerational earnings correlation, increasing the parameter estimate from about 0.140 to 0.344. The corresponding estimates for native father–son pairs are 0.177 and 0.251, which, although not statistically different, suggest less intergenerational mobility for immigrants than for natives.

5.3.3 Intergenerational Transmission and Language

In Section 3.4.4, we discussed language as one of the key human capital characteristics determining the economic outcomes of immigrants in their host country. Since the language skills of parents are likely to at least partly determine the language skills of their children, they could be an important factor underlying the observed persistence in economic status between first- and second-generation immigrants. A number of studies have investigated the link between the language proficiency of children and their parents. For Australia, Chiswick, Lee, and Miller (2005) find strong links between parents' measured and unmeasured determinants of language proficiency and the language skills of their children, in particular between mothers and their children. Bleakley and Chin (2008) show that parental language skills have a significant positive causal effect on US-born children's ability to speak English. Interestingly, this positive effect is only present while the children are young but fades out by the time they reach middle school. However, the poorer language skills when young turn out to have detrimental long-term consequences for the children's educational outcomes in terms of drop-out rates, attendance of ageappropriate grades and attendance of preschool.⁶⁵ Unlike Bleakley and Chin (2008), who use data from the 2000 US Census, and Chiswick, Lee, and Miller (2005), who use data from the 1996 Australian Census, Casey and Dustmann (2008) use repeated information on both parents and their children from the German Socio-Economic Panel. This allows them, first, to address the problem of measurement error that is widespread in self-reported data on language proficiency (see Dustmann and van Soest (2001)), second, to avoid sample selection due to children leaving the parental household, and third, to analyze the association between parental language proficiency and children's later

⁶⁵ As in Bleakley and Chin (2004), the authors use the parents' age at arrival interacted with a dummy for non-English-speaking country of origin as an instrument for their English language skills, making this the probably most convincing strategy to deal with the endogeneity of parental language skills.

economic outcomes. The results from this study show a significant and sizeable effect of parental language fluency on that of their children. Although for males, language proficiency does not significantly affect any of the labor market outcomes considered (wages, labor market participation, employment, and unemployment), it has a beneficial effect for the labor market outcomes of women, in particular those who were born abroad but arrived in Germany before the age of 10. This differential pattern could be due to women's higher propensity to work in occupations where language fluency is important, such as white-collar jobs in the service sector. Overall, the empirical evidence so far suggests a strong intergenerational transmission of language skills, in particular at younger ages of the second-generation immigrants, which may contribute to the relatively low intergenerational mobility in educational attainment and earnings that characterizes many immigrant groups in the host countries studied.

5.3.4 Intergenerational Transmission and Ethnic Networks

In an important contribution, Borjas (1992) extends the standard framework for analyzing the intergenerational transmission of human capital by assuming that ethnicity acts as an externality in the human capital accumulation process. In the model outlined in Section 5.2, such an externality would be captured in the term e_t . This implies that a correctly specified economic model of intergenerational mobility should not only include parental inputs as a determinant of the children's skills but also the average quality of the ethnic environment in which the child is raised, the so-called "ethnic capital". As long as ethnic capital plays an important role in the intergenerational transmission of skills, ignoring it in a regression based on individual level data may lead to a severe underestimation of the true persistence in earnings across generations. Using data from the General Social Surveys and the National Longitudinal Surveys of Youth that include information on both the respondents' and their parents' educational attainment and occupation, as well as the respondents' wages (NLSY only), Borjas (1992) finds overall intergenerational correlations of educational attainment, occupations, and log wages of around 0.35–0.40, 0.57, and 0.60, respectively, where all of these estimates reflect the sum of the effects due to parental variables on the one hand, and ethnic capital on the other hand. The latter, measured by the mean of the characteristic in the corresponding ethnic group, has a positive and significant effect of roughly similar (for education and wages) or greater (occupations) magnitude as the corresponding parental variable, suggesting an important role in the intergenerational transmission process. Neglecting ethnic capital will thus lead to an underestimation of the intergenerational correlation coefficient and hence to an overestimation of the speed of economic convergence of ethnic groups across generations.⁶⁶ Aydemir, Chen, and Corak

⁶⁶ In later work, Borjas (1995b) shows that segregation into particular neighbourhoods could be one reason for the external effects of ethnicity, a point that has been reemphasized by Nielsen, Rosholm, Smith, and Husted (2003) and Rooth and Ekberg (2003). See also Cutler, Glaeser, and Vigdor (2005).

(2009) work with a somewhat broader interpretation of what the average group characteristics may capture, including social factors such as discrimination or lack of access to good schools and credit markets. Using grouped data from the 2001 Canadian Census, they employ quantile regression methods to separate the influence of social capital from the influence of broader social institutions. Their findings suggest that social institutions limit intergenerational earnings mobility and that parental education is the key ingredient necessary to circumvent the restrictions imposed by such social institutions.

6. CONCLUSION

This chapter addresses the relationship between migration and education. What seems at first view a small and rather specific area in the research on the Economics of Migration turns out to be at its front and center. The chapter attempts to provide a first overview of the issues we consider important when studying migration in relation to education.

Overall, this is a rapidly growing field, illustrated by the rising number of papers over the last decades,⁶⁷ and there are exciting new research avenues at its frontier. It is also an area that reflects the challenge to single economies to develop competitive structures that prevail on increasingly globalized markets, and that are based on a flexible and highly responsive skill base. Both education and migration are key ingredients to achieve this.

Our first reference is to Sjaastad (1962), who viewed migration—as education—as an *investment in the human agent*. As the various sections of this chapter show, migration decisions and decisions about learning and human capital investments are indeed closely related. Migration is not only intertwined with human capital investment decisions of those who move but also has important consequences for education and knowledge acquisition of those who do not move, both in the home and in the host countries. Migrations are dynamic and dynastic processes, forming countries for generations to come, and one of the key determinants of the success of the children of immigrants is their educational attainment. We decided to focus in this chapter on three aspects that we believe are the cornerstones of the connection between migration and education: the economic aspects of the individual migration decisions and how they relate to the acquisition of education, the connection between the acquisition of education and the skill selection of immigrants, and the nature of intergenerational spillovers. Although we attempted to be exhaustive in our coverage, we have almost certainly missed important additional contributions that investigate these subjects.

The chapter commences with a section (Section 2) that provides an overview of the stylized facts that connect immigration and education. The following three sections

⁶⁷ Searching on Google for papers written in Economics, Business, or Finance with migration, immigration, or emigration and education, human capital, or skill in the title gives 36 papers between 1991 and 1999, 40 papers between 2000 and 2004, and 65 papers between 2005 and 2009.

(Sections 3–5) each start off with a discussion of a simple theoretical framework, which helps to structure the large empirical literature that exists in each of the areas considered. In Section 3, where we focus on the migrant, we show that educational choices and the accumulation of skills are inherently connected to migration and remigration decisions. We emphasize that decisions about nonpermanent forms of migration are key to understand educational choices and decisions about skill acquisition, and we demonstrate the challenges for obtaining estimates of immigrant career paths that are generated if migrations are nonpermanent and if migration decisions are taken in conjunction with decisions on human capital investment. In our view, this is an area where many research questions are still unexplored. For instance, as we show in Section 2, in many cases the acquisition of education rather than the pursuit of higher wages may be the main motive of migration: a hypothesis that is supported by the growing fraction of student migrations in the overall migration flows. Also, the forms of migration and implied career paths of immigrants have been changing, with temporary migrations today being the rule rather than the exception. Yet, most papers that study career paths of immigrants are still assuming permanent forms of migration.

In Section 4, we discuss the way in which migration affects educational choices and skill accumulation of individuals who do not migrate, both in the home and in the potential host countries. This area overlaps with many issues in development economics. We argue that while, as in the quote of Sjaastad, the *return* to education has been the main motive for migration, it is the acquisition of education itself that is becoming an important trigger for migration movements, and we explore the consequences for the destination and the origin countries. Another important aspect, from the perspective of both sets of countries, is who migrates. The answer to this question has important implications for the effect of immigration on the economies of both countries, through mechanisms such as the brain gain and the brain drain. We argue that additional insights can be gained when considering an application of the Roy (1951) model to the migration context that takes account of the multidimensionality of skills, in order to be able to explain recent migration patterns. Modern economies have specialized in different industries to gain competitiveness in international markets. As a consequence, the return to different skills may differ across countries, changing the incentives underlying individual migration decisions. Yet, most of the literature that studies the selection of immigrants focuses on a special case of the Roy model where skills are one-dimensional.

In the final section, we take a more dynastic view of immigration. Here, we focus on the children of immigrants, their educational achievements, and their human capital accumulation and ensuing career paths. This long-term aspect of immigrant integration and assimilation is likely to be a particular focus of research over the next decade due to the increasing number of countries that have recently experienced significant increases in their foreign-born populations. The existing evidence we discuss suggests that education is the key factor determining both the degree and the pace of the economic integration of immigrants and their descendants.

Table 4.L1 Summary of the Literature on Immigrants' Earnings Assimilation

Authors	Country	Data (DA) and Sample (SA)	Dependent Variable (DV) and Identification Assumption (IA)	Main Immigrant Groups	Main Results	Rates of Return to Schooling: Immigrants (RRI), Natives (RRN)	Rates of Return to Experience: Immigrants (RREI), Natives (RREN); Rate of Return to Years Since Migration (RRYSM)
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Baker and Benjamin (1994)	Canada	DA: Public Use Microdata Files of the Canadian Census for 1971, 1981, and 1986. SA: Men aged 16–64 years who report 40 or more weeks of work in the previous year.	DV: Log annual earnings (the sum of wage and self-employment earnings). IA: Equal time effects.	1950s and 1960s: Britain, United States, and Europe; 1970s and 1980s strong increase in inflow from Asia, Africa, and Latin America.	Entry earnings are falling across successive immigrant cohorts, whereas their rates of assimilation are uniformly small. The results confirm US evidence of "permanent" differences across arrival cohorts. The authors find small or negative rates of assimilation for most cohorts over the sample period. The results are robust to the choice of different base group. The authors find a relative decline for returns to experience for immigrants educated outside Canada. Using data from three censuses, the authors fail to reject the usual cohort fixed-effect specification.	RRI: 4.8% (1971), 4.4% (1981), 4.9% (1986). RRN: 7.3% (1971), 6.6% (1981), 7.6% (1986).	Cross-sectional estimates: RREI: 3.3% (sq: -0.0006) (1971), 3.7% (sq: -0.0006) (1981), 4.3% (sq: -0.0007) (1986). RREN: 4.6% (sq: -0.0007) (1971), 5.2% (sq0.0009) (1981), 5.9% (sq0.0009) (1986). RRYSM: n.a.
Barth, Bratsberg, and Raaum (2004)	Norway	DA: Register data for 1980, 1990, 1992–1996. SA: Individuals aged 25–64 years.	DV: Log annual earnings. IA: Equal time effects after allowing for differential effect of local unemployment.	Nordic Countries, (non-Nordic) OECD countries, Eastern Europe, Asia, Africa, and Latin America.	Nordic and OECD immigrant men catch up to the earnings of natives after 15–18 years, for all other groups of men earnings do not converge. Non-OECD men earn 30% less than natives after 25 years. Similar patterns for women. Non-OECD women earn around 18% less than natives after 25 years, whereas OECD women earn 10% less. Standard methodology would understate assimilation effects by 10–20%. Early cohorts have higher earnings than recent cohorts. (All conditional on education, for an immigrant arriving at age 25).	n.a.	n.a.
Barth, Bratsberg, and Raaum (2006)	United States	DA: Current Population Survey (CPS) from 1979 to 2003. SA: Individuals aged 21–64 years (and not enrolled in school).	DV: Log hourly wage rate (constructed). IA: Equal time effects after allowing for differential effect of local unemployment.	Mexico, other Central and South American countries, Asia, Africa, United Kingdom, and Commonwealth, Europe.	Wages of immigrants are found to be more sensitive to unemployment than wages of natives. A 10% increase in the unemployment rate reduces wages of immigrant men aged 31–39 years by 1.7% and those of natives by 0.3%. The traditional synthetic panel methodology assuming equal time effects estimates significant assimilation effects in terms of wages. For males, the standard method predicts immigrant wage growth over 20 YSM to exceed the one of natives by 15–17 pp. The proposed methodology reveals much smaller assimilation effects. The positive bias in the standard method arises from a negative trend in unemployment in the data, attributing to wage effects of improving labor market conditions during the 1990s to wage assimilation.	n.a.	RREI: Low education: males 1.5% (sq: -0.0004), females 1.7% (sq: -0.0012). High school: males 3.8% (sq: -0.0020), females 2.6% (sq: -0.0014). College: males 5.2% (sq: -0.0028), females 5.7% (sq: -0.0036). RREN: Low education: males 3.8% (sq: -0.0021), females 2.1% (sq: -0.0012). High school: males 5.0% (sq: -0.0024), females 3.7% (sq: -0.0020). College: males 6.5% (sq: -0.0026), females 6.5% (sq: -0.0034). RRYSM: Low education: males 2.0% (sq: -0.0005), females 0.7% (sq: -0.0007), females 1.4% (sq: -0.0004). College: males 2.5% (sq: -0.0006), females 2.5% (sq: -0.0006).

Continued

Authors	Country	Data (DA) and Sample (SA)	Dependent Variable (DV) and Identification Assumption (IA)	Main Immigrant Groups	Main Results	Rates of Return to Schooling: Immigrants (RRI), Natives (RRN)	Rates of Return to Experience: Immigrants (RREI), Natives (RREN); Rate of Return to Years Since Migration (RRYSM)
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Basilio and Bauer (2010)	Germany	DA: German Socio-Economic Panel (GSOEP) 1984–2007 for West Germany. SA: Individuals aged 16–64 years.	DV: Log hourly real wage (constructed). IA: Equal time effects.	Turkey, Eastern Europe and Former Soviet Union, Ex- Yugoslavia, high- income OECD.	The native-immigrant earnings gap at the time of arrival can be largely explained by different regional sources of human capital. For immigrants as a whole, foreign schooling and labor market experience is valued lower in the German labor market than domestic schooling and experience. The authors find evidence for heterogeneity in the returns to human capital across origin countries. Immigrants from high-income countries earn similar returns as natives and earn the highest return to their foreign human capital among all immigrant groups.	RRI education abroad: for men 4.5%, for women 6.2%. RRI education in Germany: for men 5.5%, for women 6.2%. RRN: for men 7.2%, for women 6.8%.	RREI experience abroad: for men 0.3%, for women 0.1%. RREI experience in Germany: for men 1.3%, for women 0.9%, RREN: for men 0.9%, for women 0.6%, RRYSM from restricted model: for men 0.8%, for women 0.5%. Only foreign labor market experience for immigrants from high-income OECD countries has positive returns, 0.7% for men and 0.9% for women.
Beggs and Chapman (1988b)	Australia	DA: 1973 ANU Social Sciences Mobility Survey and 1981 Australian Census. SA: Wage or salary-earning men aged 30–64 years.	DV: Log hourly income. IA: Equal time effects.	English-speaking immigrants (mostly from the United Kingdom and Ireland), non- English-speaking immigrants (mostly from Italy and Greece).	Authors analyze assimilation profiles of immigrants both using single cross-sectional data and using the time dimension of the data. Migrants from non- English-speaking countries entering Australia in 1965 perform significantly better between 1973 and 1981 than predicted from the 1973 cross section. This finding is consistent with the view that the quality of non- English-speaking immigrants arriving in Australia has increased over the 1960s. Migrants from English- speaking countries perform similarly no matter whether estimates are based on cross-sectional data or based on time-series data.	RRI for schooling abroad: non-English-speaking immigrants 2.5% (1973), 4.9% (1981), English- speaking immigrants: 8.9% (1973), 8.4% (1981). RRI for schooling in Australia: non-English-speaking immigrants 2.4% (1973), 0.8% (1981), English- speaking immigrants 0.9% (1973), -0.9% (1981). RRN: 10.5% (1973), 9.0% (1981).	Cross-sectional estimates: RREI for experience abroad: non-English-speaking immigrants 0.6% (sq: -0.0002) (1973), 1.1% (sq: -0.0002) (1981), English- speaking immigrants 3.4% (sq: -0.0005) (1973), 0.9% (-0.0002) (1981). RREN: 2.6% (sq: -0.0004) (1973), 2.1% (-0.0003) (1981). RRYSM: non- English-speaking immigrants 0.5% (sq: 0.0001) (1973), -0.3% (sq: 0.0001) (1981), English-speaking immigrants 0.1% (sq: -0.0000) (1973), 0.9% (sq: -0.0001) (1981).
Bell (1997)	United Kingdom	DA: General Household Surveys (GHS) 1973–1992. SA: Immigrant men aged 18–64 years who are working more than 30 hours per week.	DV: Log gross weekly wages. IA: Equal time effects.	Caribbean, India, Europe, and Old Commonwealth.	Large changes in the national-origin mix of immigrants in the United Kingdom in the postwar period. Immigrants have on average more years of schooling than natives, and this gap has risen over successive cohorts. Most disadvantaged group are immigrants from the Caribbean. However, that disadvantage diminishes relatively fast with time spent in the United Kingdom. Immigrants who arrive without any labor market experience typically experience only a small wage penalty. White immigrants earn a wage premium upon arrival but quickly assimilate to the earnings of natives.	RRI: Caribbeans 4.4%, Indians 3.7%, Whites 6.5%. RRN: 7.7%.	RREI: Caribbeans 2.2% (sq: -0.0001), Indians 2.9% (sq: -0.0005), Whites 3.5% (-0.0005). RREN: 5.8% (sq: -0.0010). RRYSM: Caribbeans -0.6% (sq: 0.0002), Indians -1.0% (sq: 0.0002), Whites -1.6% (sq: 0.0003).

Table 4.L1 Summary of the Literature on Immigrants' Earnings Assimilation—continued

Table 4.L1 Summary of the Literature on Immigrants' Earnings Assimilation—continued

Authors	Country	Data (DA) and Sample (SA)	Dependent Variable (DV) and Identification Assumption (IA)	Main Immigrant Groups	Main Results	Rates of Return to Schooling: Immigrants (RRI), Natives (RRN)	Rates of Return to Experience: Immigrants (RREI), Natives (RREN); Rate of Return to Years Since Migration (RRYSM)
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Borjas (1985)	United States	DA: 1970, 1980 US Census. SA: Men aged 18–54 years in 1970 and 28–64 years in 1980.	DV: Log hourly wage rate (constructed). IA: Equal time effects.	Mexico, Cuba, Other Hispanic, Asian, Whites, Blacks.	Only white immigrants overtake the earnings of statistically comparable white native workers after 10–15 years. All other groups have slower or even negative rates of convergence to their specific native comparison groups. Quality of immigrant cohorts declined over time.	RRI: between 2.8% (Mexicans) and 7.1% (Asians) (1970), between 2.9% (Mexicans) and 5.9% (Asians) (1980). RRN: between 4.9% (Cubans) and 6.9% (Whites) (1970), between 4.6% (Cubans) and 6.0% (Other Hispanics) (1980). Return to schooling significantly lower for Mexicans, Cubans and other Hispanics relative to native counterparts. For other groups about the same.	Cross-sectional estimates. RREI: between 0.5% (Asians, sq: 0.0000) and 4.0% (Whites, sq: -0.0007) (1970), between -0.2% (Cubans, sq: -0.0000) and 3.9% (Whites, sq: -0.0006) (1980). RREN: between 2.6% (Blacks, sq: -0.0004) and 11.2% (Cubans, sq: -0.0026) (1970), between 0.6% (Blacks, sq: -0.0000) and 3.5% (Asians, sq: -0.0005) (1980). RRYSM: n.a.
Borjas (1995a)	United States	DA: 1970, 1980, 1990 US Census. SA: Men aged 25–64 years.	DV: Log hourly wage rate (constructed). IA: Equal time effects.	Mexican, Other Hispanic, Asian, Whites.	Relative wages of immigrants grow by about 10 pp during the first two decades after arrival, very little thereafter. For 1970 and 1980 cohorts, eventual wage gap is about 5–10 pp. Immigrants who arrived in the late 1980s (1970, 1960) earned about 19.3% (13.4%, 8%) less than natives at the time of entry (all conditional on education, for an immigrant arriving at age 20).	RRI: 4.7%, RRN: 6%	RREI: 8.8% (sq: -0.0016). RREN: 9.4% (sq: -0.0015). RRYSM: 1.9% (sq0.0004)
Bratsberg and Ragan (2002)	United States	DA: 1970, 1980 and 1990 US Census and National Longitudinal Survey of Youth. SA: Foreign-born men aged 25–64 years (worked positive hours and earned at least US \$1000 wage or salary income in 1989 and not enrolled at school at time of census)	DV: Log of weekly earnings. IA: Equal cohort effects.	Mexico and other Central American Countries, South America, United Kingdom and Europe, Commonwealth, Asia, and North Africa.	Immigrants with US schooling earn higher wages than immigrants with non-US schooling. This wage advantage results from both greater educational attainment and higher returns to education and cannot be attributed to greater English proficiency. Returns to years of non-US schooling are higher for immigrants who complete their schooling in the United States and can be interpreted as US schooling upgrading education received in the source country. For immigrants without US schooling, returns are higher for immigrants from highly developed countries and countries in which English is an official language.	RRI (linear spline function): for non-US schooling: less or equal 11 years of schooling 0.8%, more than 11 years of schooling 8.9%. For US schooling: less or equal 11 years of schooling 4.1%, more than 11 years of schooling 10.2%. RRN: n.a.	RREI: for those with non-US schooling 1.8% (sq0.0003), for those with US schooling 4.2% (sq0.0007). RREN n. a. RRYSM: for those with non-US schooling 2.5% (sq0.0003), for those with US schooling 1.0% (-0.0001).

Authors	Country	Data (DA) and Sample (SA)	Dependent Variable (DV) and Identification Assumption (IA)	Main Immigrant Groups	Main Results	Rates of Return to Schooling: Immigrants (RRI), Natives (RRN)	Rates of Return to Experience: Immigrants (RREI), Natives (RREN); Rate of Return to Years Since Migration (RRYSM)
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Chiswick (1978)	United States	DA: 1970 US Census. SA: White men (natives and immigrants) aged 25–64 years.	DV: Log annual earnings. IA: Equal cohort effects.	Mexican, Cuban, Asian/African Immigrants.	Earnings of foreign-born 14.9% lower after 1 year in the country, 9.5% lower after 5 years, equal after 13 years, and 6.4% higher after 20 years.	RRI: overall 5.7%, to foreign schooling 5.8%, to domestic schooling 5.0%. RRN: 7.2%.	RREI: 2.0% (Experience squared: -0.0003). RREN: 3.2% (sq: -0.0005). RRYSM: 1.5% (sq: -0.0009)
Clark and Lindley (2009)	United Kingdom	DA: UK Labour Force Survey 1993–2004. SA: Men aged 16–64 years.	DV: Employment rate (excl. inactive) and log real gross weekly earnings. IA: Equal time effects.	White immigrants from old Commonwealth (Australia, New Zealand, Canada) and European Union and the rest of the world. Nonwhites from Britain's former colonies in Asia, the Caribbean, and Africa and immigrants from the rest of the world.	Differentiate results by ethnicities (whites/nonwhites) and by whether the immigrant has completed his or her education (labor market entrant) or whether he or she still has to complete his or her education in the United Kingdom system (education entrants). Among whites, education entrants perform better in comparison to white natives than labor market entrants. Among labor market entrants, whites do better than nonwhites, while among education entrants, highly qualified prime-aged nonwhites perform, as well as both white immigrants and natives. Patterns of labor market assimilation are found to be diverse depending on ethnicity and immigrant type. Labor market outcomes for all immigrant groups have a tendency to decline with age relative to white natives.	Labor market entrants: RRI: 6.3% for white immigrants, 5.6% for nonwhite immigrants. RRN: 7.9%. Education entrants: RRI for white immigrants: 67.1% for university degree, 30.7% for A-levels, 19.4% for O-levels. RRI for nonwhite immigrants: 77.6% for university degree, 31.4% for A-levels, 23.6. for O-levels. RRN for white natives: 69.8% for university degree, 27.8% for A-levels, 19.9% for O-levels.	n.a.
Dustmann (1993)	Germany	DA: German Socio-Economic Panel (GSOEP). SA: Men aged 16+ years in 1984, who were full- time employed at time of interview (exclude self- employed, civil servants, individuals in education or apprenticeships).	DV: Log monthly gross earnings. IA: Equal cohort effects.	Turkey and Southern Europe.	The author shows within a human capital framework that in the case of temporary migration, the optimal investment into country-specific human capital is lower than in the case of permanent migration. The empirical results indicate that foreign workers in the German labor market receive lower wages than their native counterparts throughout their working history, and that there is no earnings crossover between these two groups. Using data on expected length of stay in the country, the empirical results support the hypothesis that total length of stay positively influences host country-specific human capital investment and thus earnings of immigrants. Earning profiles are less concave (i.e., the longer the total intended duration of stay in the host country).	RRI: for schooling: 1.2%, for job-specific training: 1.0%. RRN: for schooling: 5.5%, for job-specific training: 3.4%.	RREI: 1.9% (sq: -0.0005). RREN: 3.7% (sq: -0.0007). RRYSM: 1.4% (sq: -0.0002).

Table 4.L1 Summary of the Literature on Immigrants' Earnings Assimilation—continued

Table 4.L1 Summary of the Literature on Immigrants' Earnings Assimilation-continued

Authors	Country	Data (DA) and Sample (SA)	Dependent Variable (DV) and Identification Assumption (IA)	Main Immigrant Groups	Main Results	Rates of Return to Schooling: Immigrants (RRI), Natives (RRN)	Rates of Return to Experience: Immigrants (RREI), Natives (RREN); Rate of Return to Years Since Migration (RRYSM)
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Edin, LaLonde, and Åslund (2000)	Sweden	DA: LINDA database, registry data for immigrants entering between 1970 and 1990. SA: Individuals aged 18–55 years at the time of immigration.	DV: Log annual earnings. IA: Equal time effects.	Nordic countries, other OECD countries, political migrants from Yugoslavia, Iran, Iraq, Poland, and Chile.	Economic migrants are much more likely to emigrate than political ones, with the least economically successful economic migrants most likely to leave. Immigrant earnings grew on average by around 20,000SEK relative to natives during their first 10 years in Sweden. Controlling for emigration, the relative earnings growth during the first ten years reduces to 13,500SEK on average. For OECD migrants, the authors estimate negative assimilation. For non-OECD migrants, there is still substantial evidence for unconditional convergence, even after accounting for emigration. A failure to adjust for emigration leads to an overestimation of the rate of economic assimilation, especially for Nordic and OECD immigrants.		RRYSM from cross-sectional estimates: apart from Nordic women no positive returns to YSM for economic migrants, but positve RRYSM for immigrants from non-OECD countries. From the longitudinal analysis: growth for non- OECD migrants slows down after a few years and cross-sectional results likely to be driven by changes in cohort quality.
Friedberg (2000)	Israel	DA: Israeli Census of Population and Housing 1972 and 1984. SA: Men aged 25–65 years. Arabs excluded and only full- time, salaried, nonagricultural workers retained.	DV: Log monthly earnings. IA: Equal cohort effects, test in Appendix using two cross sections that one cross section is sufficient to identify assimilation rates.	Asia and Africa (largest emigration country Morocco), Eastern Europe, USSR, and Western Hemisphere, and Western Europe.	The gap in the residual earnings of immigrant and native workers is eliminated once the national origin of individual's human capital is accounted for. Human capital acquired abroad receives a lower return in the host labor market than human capital acquired domestically. Return to schooling obtained abroad is highest for immigrants from the West (7.1%) and lowest for immigrants from Asia and Africa (5.7%). The returns to experience acquired abroad are generally insignificant. The portability of home country education varies significantly with its level (elementary school education equally valued in home and host country).	RRI : 8.0% for domestic schooling and 7.1% for foreign schooling. RRN: 10.0% to domestic schooling.	RREI: 1.1% for each year of domestic experience, 0.1% for each year of foreign experience. RREN: 1.7%, from estmation of restricted model RRYSM: 0.8%.
Funkhouser and Trejo (1995)	United States	DA: Special supplements to Current Population Survey (CPS) 1979, 1983, 1986, 1988 and 1989. SA: Men aged between 18 and 61 years.	DV: Log average hourly earnings. IA: Equal time effects.	Mexico and other Latin America, Europe, Canada, and Oceania, Asia (mainly Japan, Korea, China, and the Philippines).	Tracking the immigrants' skill levels through the 1980s, the authors find that male immigrants who entered during the late 1980s are more skilled than those who arrived earlier in the decade. This represents a break from the steady decline in immigrant skill levels that took place between 1940 and 1980, but the average skill level of recent immigrants remains low by historical standards.	RRI: 5.1%. RRN: 8.2%.	RREI: 3.4% (sq: -0.0005). RREN: 5.1% (sq: -0.0008). RRYSM: 2.1% (sq: -0.0002).

Authors	Country	Data (DA) and Sample (SA)	Dependent Variable (DV) and Identification Assumption (IA)	Main Immigrant Groups	Main Results	Rates of Return to Schooling: Immigrants (RRI), Natives (RRN)	Rates of Return to Experience: Immigrants (RREI), Natives (RREN); Rate of Return to Years Since Migration (RRYSM)
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Green and Worswick (2004)	Canada	DA: Immigrant Database (IMDB) for immigrants, Survey of Consumer Finances (SCF) for natives, for the years 1981, 1982, 1984–1997. SA: Men aged 25–64 years.	DV: Log real annual earnings. IA: Equal time effects, with comparison group being native new entrants into the labor market.	English-speaking countries (United States, United Kingdom, Australia, New Zealand), North-Western Europe, Others.	Based on a life cycle human capital framework, the authors argue for defining cohort quality based on the present value of all future earnings rather than the entry earnings of an immigrant cohort. Compare newly arriving immigrants to the subgroup of natives who enter the Canadian labor market at the same time as the immigrants, arguing that these should be more likely to be affected by the same macroshocks and subject to the same human capital investment conditions. Findings show that changes affecting all new entrants play an important role in understanding the large cross-cohort earnings decline in Canada between the 1980s and 1990s. Present value comparisons show that the cohorts of the 1980s. Shifts in the source country composition and the general new entrant effects account for over 90% of the 1980s decline. Foreign experience of immigrants from non-English- speaking, non-European countries yields zero return.	n.a.	Paper reports full set of estimates for immigrants' earnings profiles over YSM relative to matched native entrants by education group (high school and university education), age at entry (distinguishing four groups), and entry cohort (distinguishing five groups).
Hayfron (1998)	Norway	DA: Population Census of Norway Data Bank for 1980 and 1990 (8.3% sample of the central register). SA: Men aged 17–55 years in 1980 and 27–65 years in 1990 who work full time, and earn positive income (self-employed and students excluded). Immigrants defined by citizenship.	DV: Log earnings (taxable income from work, sickness pay, unemployment benefits, and income when in labor market programs). IA: Equal time effects.	n.a.	The results show that the 1970–1979 cohort experiences a relative earnings growth of about 11% between 1980 and 1990, which is substantially lower than the cross-sectional estimate of 19%. There is rapid earnings divergence across immigrant cohorts and between the 1960–1969 immigrant cohort and natives.	Cross-sectional estimates: RRI: 2.4% (1980), 1.9% (1990). RRN: 3.8% (1980), 6.9% (1990).	Cross-sectional estimates: RREI: 8.8% (sq: -0.0010) (1980), 9.0% (sq: -0.0010) (1990). RREN: 10.7% (sq: -0.0012) (1980), 4.1% (sq: -0.0004) (1990). RRYSM: n.a.

Table 4.L1 Summary of the Literature on Immigrants' Earnings Assimilation—continued

Table 4.L1 Summary of the Literature on Immigrants' Earnings Assimilation-continued

Authors	Country	Data (DA) and Sample (SA)	Dependent Variable (DV) and Identification Assumption (IA)	Main Immigrant Groups	Main Results	Rates of Return to Schooling: Immigrants (RRI), Natives (RRN)	Rates of Return to Experience: Immigrants (RREI), Natives (RREN); Rate of Return to Years Since Migration (RRYSM)
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Husted, Nielsen, Rosholm, and Smith (2001)	Denmark	DA: Registry data 1984–1995, SA: Men aged 20–59 years (self- employed not observed and individuals in education excluded).	DV: Log hourly wage (constructed), employment. IA: Equal time effects.	Other Nordic countries, EU, Turkey, other European countries, Sri Lanka, Irak, Iran, Vietnam, and Pakistan.	Initial employment probability for refugees is much lower, but after 5–10 YSM approaches the level of non-refugees and natives. Refugees from Africa and Palestine have lower initial employment chances than refugess from e.g. Europe and Vietnam. Refugees start at a much lower wage rate than Danish workers, and after 4 years their wage growth is just above the wage growth for Danes. The slow assimilation is partly due to the immigrants' weak attachment to the Danish labor market.	For hourly wages the paper reports estimates for education indicators. For Danish-born workers, hourly wage is 10.2% higher if they completed secondary education (relative to primary education), 9.4% higher if they completed vocational training, and 29.9% higher if they completed the highest education level.	Classical model: RREI: refugees 0.5% (sq: -0.0000), non-refugees 1.6% (sq: 0.0001). RREN: 2.7% (sq: -0.0003). RRYSM (linear spline function): first 5 years -5.5%, 5-10 years 0%, 10 years and more, 2%.
Kee (1995)	The Netherlands	DA: Quality of Life Surveys (QLS) 1984–1985 for immigrants, Labour Supply Panel 1985 for natives. SA: Men aged 18–65 years.	DV: Log hourly wage rate. IA: Equal cohort effects.	Turks, Moroccans, Surinamese, Antilleans.	Focus of the paper is more on potential discrimination of immigrants in the Dutch labor market. Taking account of sample selection, the authors estimate that for Antilleans and Turks about 35% and 15% of their wage gaps relative to natives are due to "discrimination." There is no indication of discrimination against Surinamese and Moroccan immigrants.	RRI for schooling abroad: Antilleans 5.1%, Surinamese 3.6%, Turks -0.2%, Moroccans 0.2%. RRI for schooling in the Netherlands: Antilleans 4.4%, Surinamese 3.2%, Turks 1.3%, Moroccans 3.0%, RRN: 4.0%.	Cross-sectional estimates: RREI for experience abroad: Antilleans 1.3% (sq: -0.0003), Surinamese 2.7% (sq: -0.0004), Turks 0.4% (sq: -0.0002), Moroccans 0.4% (sq: -0.0001). RREI for experience in the Netherlands: Antilleans 4.8% (-0.0015), Surinamese 5.5% (sq: -0.0012), Turks 2.5% (sq: -0.0007), Moroccans 2.2% (sq: -0.0006). RREN: 3.3% (sq: -0.0005). RRYSM: n.a.
Kossoudji (1989)	United States	DA: 1976 Survey of Income and Education (SIE). SA: Native and foreign-born men aged 20–64 years who are in full-time employment.	DV: Occupation- specific earnings. IA: Equal cohort effects; separate sample by ethnic group to account for the fact that ethnic groups were differently affected by immigration law.	Hispanics and East Asians.	The author estimates a simultaneous equations mixed model of occupational choice and earnings, distinguishing between immigrants who migrated as adults and those who migrated as children.	For immigrants, except for higher levels of occupations (professionals), the returns to education are not signifi- cantly different from zero. For natives, the returns to education are statistically significant and positive for all occupation groups. No significant difference in the returns to education by the location of where education was obtained. Education always significant in occupa- tional choice equations.	Results typically exhibit quadratic shape of experience profiles for workers in all occupations. RREI: experience in the United States has a significant positive effect on earnings for all groups except Hispanic managers and craft workers and Asian sales/clerical and service workers. No significant gain from experience accumulated in the home country. RREN: positive returns for all occupations, returns higher than for Hispanics, but lower than for Asians. RRYSM: n.a.

Authors	Country	Data (DA) and Sample (SA)	Dependent Variable (DV) and Identification Assumption (IA)	Main Immigrant Groups	Main Results	Rates of Return to Schooling: Immigrants (RRI), Natives (RRN)	Rates of Return to Experience: Immigrants (RREI), Natives (RREN); Rate of Return to Years Since Migration (RRYSM)
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
LaLonde and Topel (1991)	United States	DA: 1970 and 1980 US Census. SA: Men aged 16–64 years who worked 40 or more weeks during the preceding calendar year.	DV: Log weekly wages (constructed). IA: Equal time effects.	Europeans, Asians, Mexicans, Middle Easteners, and Other Latin American and Caribbean.	Initial relative earnings between immigrants and natives declined between 1970 and 1980 with an initial earnings disadvantage of 20% in 1970 and 35% in 1980. Decline can be attributed to changes in the composition of source countries towards Asian and Latin American countries. Assimilation rates after 10 years since arrival in the United States are positive and large. Assimilation is found to be more rapid for groups who start with lower wages, such as Asians and Middle Easterners.	n.a.	Cumulative effect of 10 YSM (relative to old immigrants of same ethnicity): Europeans 8%, Asians 24%, Middle Easterners 42%, Mexicans 21%, Other Latin Americans and Caribbeans 19%.
Longva and Raaum (2003)	Norway	DA: 1980 and 1990 Norwegian Population Census (supplemented by administrative data for 1990). SA: Men aged 17–55 years in 1980 and 27–65 years in 1990 who work full time, and earn positive income (self-employed and students excluded in 1980). Uses universe of immigrant population defined by country of origin conditional on their presence in Norway in 1992.	DV: Log earnings (taxable income from work, sickness pay, unemployment benefits, and income when in labor market programs). IA: Equal time effects.	OECD countries.	The authors find that the earnings assimilation of immigrants in Norway from 1980 to 1990 differs considerably between cohorts and by country of origin. They estimate the relative earnings growth for the 1970–1979 immigrant cohort to be 6% over the decade (lower than the 11% estimated by Hayfron (1998)). They find earnings of OECD immigrants to be comparable to natives as opposed to non–OECD immigrants, who earn considerably less than natives at the time of entry, but improve gradually over time.	All immigrants: RRI: 3.9% (1980), 4.3% (1990). Separate estimates for OECD and non-OECD immigrants: RRI OECD: 4.4% (1980), 4.7% (1990). RRI non-OECD: 2.4% (1980), 3.3% (1990). RRN: 3.7% (1980), 5.2% (1990).	Cross-sectional estimates. All immigrants: RREI: 6.0% (sq: -0.0007) (1980), 5.3% (sq: -0.0006) (1990). Separate estimates for OECD and non-OECD immigrants: RREI OECD: 8.0% (sq: -0.0009) (1980), 6.1% (sq: -0.0007) (1990). RREI non-OECD: 4.6% (sq: -0.0006) (1980), 4.6% (sq: -0.0005) (1990). RREN: 11.9% (sq: -0.0014) (1980), 4.5% (sq: -0.0005) (1990). RRYSM: n.a.

Table 4.L1 Summary of the Literature on Immigrants' Earnings Assimilation—continued

Continued

Table 4.L1 Summary of the Literature on Immigrants' Earnings Assimilation-continued

Authors	Country	Data (DA) and Sample (SA)	Dependent Variable (DV) and Identification Assumption (IA)	Main Immigrant Groups	Main Results	Rates of Return to Schooling: Immigrants (RRI), Natives (RRN)	Rates of Return to Experience: Immigrants (RREI), Natives (RREN); Rate of Return to Years Since Migration (RRYSM)
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Lubotsky (2007)	United States	DA: Longitudinal data: Social Security Earnings records 1951–1997 linked to 1990 and 1991 Survey of Income and Program Participation (SIPP) and 1994 Current Population Survey (CPS). SA: Men born between 1930 and 1969.	DV: Log annual earnings (social security earnings). IA: Equal time and experience effects; but author only interested in differences in wage growth between cross-sectional data set and longitudinal dataset, which are adjusted in the same way.	n.a.	Actual earnings growth among immigrants who remained in the United States until the 1990s was considerably slower than implied by estimates based on repeated cross-sectional data. Over their first 20 years immigrant earnings in the longitudinal data grew by 10–15% relative to natives, while repeated cross sections suggest a growth about twice as fast, of about 26%. Selective emigration by low-wage immigrants leads to overestimation of economic assimilation when using census data. Back-and-forth migration, which leads to misclassification of many low-wage immigrants as more recent arrivals, has caused typical estimates to overstate the measured decline in the entry level of earnings of immigrants between the 1960s and 1980s by one-third.	n.a.	RRYSM: repeated cross-sectional data suggest immigrants relative earnings gap to narrow by 13% in the first ten years and an earnings growth of 10–20 pp in each successive decade; longitudinal data: relative earnings grow by 12–15% in the first 15 years in the United States and relatively little thereafter.
Sanromá, Ramos, and Simón (2009)	Spain	DA: Spanish National Immigrant Survey 2007. SA: Immigrants aged 15–65 years (working at least 10 hours per week and earning net monthly earnings above 200 Euros), immigrants with Spanish nationality excluded.	DV: Log net monthly wages. IA: Equal cohort effects.	Latin America and Eastern Europe.	With the exception of immigrants from developed countries and immigrants who have studied in Spain, the returns to host country human capital are higher than returns to home country human capital. Having legal status is associated with a wage premium.	RRI: for foreign schooling 1.8%, for schooling in Spain 3.3%. RRN: 4 % (from Wage Structure Survey 2006). Immigrants from developed countries have higher return to home country education (6.0%) than Latin Americans (1.8%) and Eastern Europeans (1.1%).	RREI 0.7% (sq: -0.0002). RREN: n.a. RRYSM: 1.4% (sq: -0.0000).
Authors	Country	Data (DA) and Sample (SA)	Dependent Variable (DV) and Identification Assumption (IA)	Main Immigrant Groups	Main Results	Rates of Return to Schooling: Immigrants (RRI), Natives (RRN)	Rates of Return to Experience: Immigrants (RREI), Natives (RREN); Rate of Return to Years Since Migration (RRYSM)
---	---------------	--	--	--	---	---	---
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Schaafsma and Sweetman (2001)	Canada	DA: 1986, 1991, and 1996 Canadian Census. SA: Men aged 16–64 years on the survey date who worked more than 40 weeks in the previous year.	DV: Log earnings (including self- employment). IA: Effect of age on earnings the same for immigrants and natives in the specifications that include age at immigration.	n.a.	The authors observe a correlation between age at immigration and earnings, which is driven by three main effects: absence of a return to source country work experience, the return to education varying with age at immigration, and an acculturation effect for immigrants who are visible minorities or whose mother tongue is not English. Educational attainment and earnings vary systematically across age at immigration. Immigrants who arrive around age 15–18 complete fewer years of schooling than those who arrive either earlier or later.	RRI for schooling abroad: 5.7% (1986), 5.9% (1991), 6.3% (1996). RRI for schooling in Canada: 5.5% (1986), 6.3% (1991), 7.0% (1996). RRN: 7.3% (1986), 7.6% (1991), 7.7% (1996).	Cross-sectional estimates: RREI for experience abroad: 0.9% (sq: -0.0002) (1986), 0.8% (sq: -0.0001) (1991), 0.9% (sq: 0.0001) (1996). RREI for experience in Canada: 5.1% (sq: -0.0011) (1986), 4.6% (sq: -0.0009) (1991), 4.4% (sq: -0.0008) (1996). RREN: 6.2% (sq: -0.0010) (1986), 6.0% (sq: -0.0009) (1991), 6.1% (sq0.0009) (1996). RRYSM: n.a.
Schoeni (1997)	United States	DA: 1970, 1980 and 1990 US Census. SA: Men aged 25–60 years.	DV: Log weekly wages (constructed), including self- employment, and wage and salary income. IA: Equal time effects.	Europeans, Mexico, Japan, Korea and China, United Kingdom and Canada, Central America, Philippines, Caribbean, Africa, other Hispanics and Middle East/other Asia.	Europeans have entered the US labor market with relatively high wages and have earned wages comparable to natives over their life course. Japanese, Koreans and Chinese had a lower initial wage, but have quickly caught up with US-born workers. Mexicans and Central Americans entered with low wages and the wage gap between them and comparable US workers has not shrunk. Wages are closely linked to education and returns to education are higher if some schooling was obtained in the United States.	RRI: for immigrants without US schooling: 4.5% (1970), 5.1% (1980), 5.3% (1990). For immigrants with some US schooling: 6% (1970), 5.6% (1980), and 5.7% (1990). RRN: 7.9% (1970), 7.1% (1980), and 10.3% (1990). Returns vary substantially by country of origin: in 1990 the RRI was 5.3% for Mexicans, around 8% for most other groups, and 13.1% for Japanese, Koreans, and Chinese.	The author accounts for six 5-year categories for YSM. Detailed results for each country of origin group reported with full interactions of all variables with census year dummies.
Shields and Wheatley Price (1998)	England	DA: UK Quarterly Labour Force Survey 1992–1994, pooled cross section. SA: Men aged 16–64 years, resident in England.	DV: Log gross hourly earnings (constructed). IA: Equal cohort effects.	Irish and other Whites, Indian, Pakistani, Bangladeshi, African, Caribbean.	Native-born nonwhites and whites (other) receive higher returns from schooling obtained in the United Kingdom than native-born whites. All other immigrant groups have lower returns to schooling than native-born whites. For nonwhite natives, UK labor market experience is more beneficial and for all immigrant groups less beneficial than for white natives. For Irish and nonwhite immigrants in England there is not statistically significant return to experience abroad.	RRI for foreign education: white British 3.7%, Irish 4.2%, other whites 7.4%, nonwhites 3.3%. RRI for UK education: white British 3.8%, Irish 4.4%, other whites 10%, nonwhites 4.1%. RRN: whites 4.9%, nonwhites 6.6%.	Cross-sectional estimates: RREI UK experience: white British 3.3% (sq: -0.0007), Irish 2.5% (sq: -0.0006), other whites 2.4% (sq: -0.0003), nonwhites 3.0% (sq: -0.0006). RREI foreign experience: white British 2.4% (sq: -0.0008), Irish 2.0% (sq: -0.0007), other whites 4.5% (sq: -0.0008), nonwhites 0.2% (-0.0001). RREN: whites 3.8% (sq: -0.0006), nonwhites 4.5% (sq: -0.0008). RRYSM: n.a.

Table 4.L1 Summary of the Literature on Immigrants' Earnings Assimilation-continued

Continued

Table 4.L1 Summary of the Literature on Immigrants' Earnings Assimilation-continued

Authors	Country	Data (DA) and Sample (SA)	Dependent Variable (DV) and Identification Assumption (IA)	Main Immigrant Groups	Main Results	Rates of Return to Schooling: Immigrants (RRI), Natives (RRN)	Rates of Return to Experience: Immigrants (RREI), Natives (RREN); Rate of Return to Years Since Migration (RRYSM)
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Stewart and Hyclak (1984)	United States	DA: 1970 US Census. SA: Immigrant males aged 14–65 years.	DV: Log annual earnings. IA: Equal cohort effects.	Mexican, Cuban, Asian/African Immigrants, United Kingdom, and European.	Earnings differ greatly by race and country of origin. Immigrants from Scandinavia and Western Europe earned higher incomes than migrants from the United Kingdom, whereas immigrants from the Orient, South America and the West Indies earned significantly less. Black and Hispanic immigrants were found to have depressed earnings profiles with the differential relative to nonblack immigrants increasing over time.	RRI: overall 3.5%, to home country schooling 3.2%, to host country schooling 3.4%. Returns to having attended vocational training: 3.1%. RRN: n.a.	RREI: overall 2.8% (sq: -0.0006), to experience in home country 1.5% (sq: -0.0003), to experience in host country 3.4% (sq: -0.0007). RRYSM: overall 1.9% (sq: -0.0003).
Venturini and Villosio (2008)	Italy	DA: Work Histories Italian Panel (linked employer employee database), 1990–2003. SA: Full-time male workers aged 18–45 years (public employment, agricultural sector, and housekeeping excluded).	DV: Log weekly wage and number of days worked per year (constructed). IA: Panel data estimated separately for immigrants and natives. For selection correction: GNP in sending country being valid instrument for the probability of staying in the host country.	Eastern Europe (Albania, Romania and Ukraine), North Africa, Asia (mainly Philippines), Latin America.	The results with and without the return intention controls that are observed in the data are very similar. The return to experience on the job is almost the same for natives and immigrants. The return to age is higher among natives than among immigrants. Migrants initially start at similar earnings levels, but their wage growth over time is smaller than for natives. Initial differences in days worked per year between native and immigrant workers persist over time. The relative wage differential over time is increasing faster for Africans than for the other immigrant groups, whereas for Asians and Eastern Europeans the widening of the employment differential stops after 5 years.	n.a.	RREI: 3.1% (sq: -0.0001). RREN: 5.5% (sq: -0.0003).
Wilkins (2003)	Australia	DA: Australian Bureau of Statistics Education and Training Survey 1997. SA: Men aged 15–64 years, employed full time at the time of the survey.	DV: Log hourly wages (constructed). IA: Equal cohort effects.	English-speaking immigrants, non- English-speaking immigrants.	The author accounts for age at migration and potentially different effects of years since migration for different arrival ages. For a given stock of human capital, initial earnings are lower for younger arrivals, but their earnings growth is faster with time in the destination country. The return to eduation for immigrants with language difficulties is significantly lower.	Without control for age at migration: RRI for English speakers 4.6%, non-English speakers 5.1%. With control for age at migration: RRI for English speakers 4.5%, non-English speakers 4.4%. RRN: n.a.	Without control for age at migration: RREI for English speakers 1.7% (sq. -0.0003), for non-English speakers 1.9% (sq0.0003). With control for age at migration: RREI for English speakers 0.3% (sq: -0.0001), for non-English speakers 1.0% (sq: -0.0003), both main effects not statistically signifcant. RREN: n.a. RRYSM: greater for child arrivals than later arrivals.

Note: Main immigrant groups in Column (7) refer to main groups focused on in the analysis, not necessarily the main groups present in the country.

Authors	Country	Data	Definition First Generation	Definiton Second Generation	Dependent Variable	Generational Income Elasticity for Immigrants (And Natives If Available)
(1)	(2)	(3)	(4)	(5)	(6)	(7)
Aydemir, Chen, and Corak (2009)	Canada	1981, 2001 Canadian Census.	Foreign-born men with a foreign-born spouse who have children aged 5–17 years in 1981.	Canadian-born, both parents foreign-born, aged 25–37 years in 2001.	Log weekly earnings, schooling.	Sons: 0.27*. Daughters: -0.048 (insignificant). For annual earnings, Sons: 0.18*. Daughters: -0.093. For natives, 0.19 for annual earnings for the overall population (reported from other study in Table 1).
Borjas (1992)	United States	General Social Surveys (GSS) and the National Longitudinal Surveys of Youth (NLSY) 1977–1989.	Foreign-born men.	US-born, at least one foreign-born parent, aged 18–64 years in the GSS and 22–29 years in the NLSY.	Educational attainment and occupation (using Hodge–Siegel– Rossi prestige score) in GSS. Educational attainment and log wage in NLSY.	GSS education: 0.27*; total effect of parental and ethnic capital: 0.48. GSS occupations: 0.20*; total effect of parental and ethnic capital: 0.64. NLSY education: 0.27*; total effect of parental and ethnic capital: 0.37. NLSY wage: 0.35*; total effect of parental and ethnic capital: 0.61.
Borjas (1993)	United States	1940, 1950, 1960, 1970 US Census.	Foreign-born men aged 25– 64 years.	US-born men, at least one parent foreign-born.	Earnings relative to third-generation Americans.	0.27* (relating 1970 second-generation workers to their presumed 1940 immigrant fathers). Reduces to 0.25* with ethnic capital (the group average in 1970) included.
Borjas (2006b)	United States	1940, 1950, 1960, 1970 US Census.	Foreign-born, aged 18–64 years.	US-born, at least one foreign-born parent, aged 18–64 years.	Log weekly earnings relative to third-generation Americans.	Men: 0.511 for 1940–1970 and 0.560 for 1970–2000. Women: 0.242 for 1940–1970 and 0.280 for 1970–2000.
Card, DiNardo, and Estes (2000)	United States	1940 and 1970 US Census, Pooled 1994–1996 Current Population Survey.	Foreign-born men aged 16– 66 years.	US-born men and women, both parents foreign-born, aged 16–66 years.	Mean log weekly wages and mean years of schooling.	Men: 0.44* for 1940–1970 and 0.62* for 1970–1995. Women: 0.21* for 1940–1970 and 0.50* for 1970–1995.
Carliner (1980)	United States	1970 US Census.	Foreign-born ("earlier immigrants"), aged 18–64 years, distinguish those who arrived in the United States between 1965 and 1970 ("recent immigrants").	US-born men, at least one foreign-born parent, aged 18–64 years. Third generation: US-born, both parents also US-born.	Log hourly wages and log annual earnings.	n.a.

Table 4.L2 Summary of the Literature on Immigrants' Intergenerational Mobility

Table 4.L2 Summary of the Literature on Immigrants' Intergenerational Mobility—continued

Authors	Country	Data	Definition First Generation	Definiton Second Generation	Dependent Variable	Generational Income Elasticity for Immigrants (And Natives If Available)
(1)	(2)	(3)	(4)	(5)	(6)	(7)
Deutsch, Epstein, and Lecker (2006)	Israel	1995 Census of Israel.	Male Jews, older than 10 years of age when they immigrated to Israel between 1948 and 1952, from Asian–African countries.	Foreign-born men, aged 10 years or younger who came between 1948 and 1952, and Israelis, aged between 33 and 53 years in 1995, with foreign-born fathers. Third generation: Israelis younger than 33 years of age in 1995 with immigrant fathers whose age at immigration was 10 years or younger.	Log monthly gross wages.	n.a.
Dustmann (2008)	Germany	GSOEP 1984–2002.	Foreign-born men.	German-born men, father foreign- born, aged 20–34 years.	Log hourly wages (constructed).	Baseline estimate: 0.15 [*] . If at least five wage observations used for average wage: 0.37 [*] . If at least eight wage observations used: 0.41 [*] . Last estimate drops to 0.39 [*] if control for father's permanent migration propensity is included. Baseline estimate for natives: 0.18 [*] . If at least five wage observations used: 0.25 [*] . If at least eight wage observations used: 0.29 [*] .
Gang and Zimmermann (2000)	Germany	GSOEP 1984–2002.	Foreign-born men.	German-born to foreign parents or who arrived before the age of 16, aged 17–38 years in 1984.	Total years of education, categorical schooling levels, and receipt or absence of vocational training.	Migrants' education has no effect on the educational attainment of their children. Natives' education has an effect on the educational attainment of the next generation; father's education has a larger impact than mother's education.
Hammarstedt and Palme (2006)	Sweden	1975, 1980, Swedish Census, foreign-born individuals who immigrated to Sweden between 1916 and 1969 and were gainfully employed in 1970. Data on all biological children for the years 1997, 1998, and 1999.	Foreign-born men aged 20– 64 years in 1975 and 1980.	Swedish-born, father foreign- born, aged 20–64 years in 1997, 1998, and 1999.	Annual earnings.	OLS, average of 1975 and 1980 earnings: 0.207*. IV, using parent's educational attainment: 0.39*. OLS natives: 0.14*. IV natives: 0.22*. Regressions include quadratic polynomial in age for first and second generation on RHS.

Authors	Country	Data	Definition First Generation	Definiton Second Generation	Dependent Variable	Generational Income Elasticity for Immigrants (And Natives If Available)
(1)	(2)	(3)	(4)	(5)	(6)	(7)
Nielsen, Rosholm, Smith, and Husted (2003).	Denmark	Entire population of immigrants and 10% of Danish population for 1985–1997.	Foreign-born men and women from less-developed countries.	Danish-born, both parents foreign-born, aged 18–35 years, left education.	Log hourly wages in the first job after leaving education.	Sons: 0.001. Daughters: -0.003. Native sons: -0.009*. Native Daughters: 0.001*.
Osterberg (2000)	Sweden	Swedish Income Panel (SWIP) data from 1978 to 1997. From 1978, a 1% sample of native- born and a 10% sample of foreign-born were taken. Supplementary 10% samples of people immigrating each year from 1979 until 1997.	Foreign-born individuals aged less than 65 years in 1982. Observed between 1978 and 1982.	Swedish-born, at least one foreign-born parent ("second generation"). Foreign-born who immigrated to Sweden when not older than 16 years of age ("young immigrants"). Foreign-born with both parents Swedish-born ("adopted immigrants"). All individuals aged 25 years and more and observed between 1993 and 1997.	Log of average of son's and daughter's reported annual earnings over the period 1993–1997.	Sons (log of father's earnings): second generation: 0.079*, young immigrant: 0.107*, adopted: 0.007*. Sons (log of mother's earnings): second generation: 0.079*, young: 0.076*, adopted: 0.076*. Daughters (log of father's earnings): second generation: 0.037*, young: 0.068*, adopted: -0.004*. Daughters (log of mother's earnings): second generation: 0.041*, young: 0.045*, adopted: -0.025*. Native sons (log of father's earnings): 0.068*. Native sons (log of mother's earnings): 0.022. Native daughters (log of father's earnings): 0.042*, Native daughters (log of mother's earnings): 0.080*.
Riphahn (2003)	Germany	German Microcensuses for 1989, 1991, 1993, 1995, and 1996.	Foreign citizen with a valid year of entry into Germany.	German-born with foreign citizenship, aged 16–19 years.	Currently attending advanced school (Gymnasium), binary variable.	-0.285* (coefficient for father's lowest schooling degree) and 0.267* (coefficient for father's advanced vocational training)0.442* (coefficient for mother's lowest schooling degree) and 0.367* (coefficient for mother's advanced vocational training).

Table 4.L2 Summary of the Literature on Immigrants' Intergenerational Mobility-continued

Continued

Authors	Country	Data	Definition First Generation	Definiton Second Generation	Dependent Variable	Generational Income Elasticity for Immigrants (And Natives If Available)
(1)	(2)	(3)	(4)	(5)	(6)	(7)
Smith (2003)	United States	Census (1940–1970) and Current Population Survey (CPS); special supplements from 1979, 1983, 1986, and 1988 and CPS from 1994–1998.	Foreign-born men.	US-born men, at least one foreign-born parent. Third generation or more: both parents US-born.	Years of schooling and log wages.	Years of schooling: 0.50* (second generation regressed on first) and 0.22* (third generation regressed on second). Log wages: 0.46* (second generation regressed on first) and 0.27* (third generation regressed on second).
Trejo (2003)	United States	Current Population Survey, 1979 and 1989.	Foreign-born, parents also foreign-born, aged over 16 years.	US-born men, at least one foreign-born parent, aged 18–61 years. Third generation: US-born whose parents are also US-born.	Log hourly earnings (constructed).	n.a.

Table 4.L2 Summary of the Literature on Immigrants' Intergenerational Mobility—continued

Note: A $(^{\ast})$ indicates statistical significance at the 5 percent level.

REFERENCES

- Acemoglu, D., Autor, D., 2010. Skills, Tasks and Technologies: Implications for Employment and Earnings. In: Ashenfelter, O., Card, D. (Eds.), Handbook of Labor Economics. Vol. 4. Elsevier-North Holland, Amsterdam forthcoming.
- Acosta, P., 2006. Labor Supply, School Attendance, and Remittances from International Migration: the Case of El Salvador World Bank Policy Research Working Paper No. 3903.
- Algan, Y., Dustmann, C., Glitz, A., Manning, A., 2010. The Economic Situation of First- and Second-Generation Immigrants in France, Germany, and the UK. Econ. J. 120 (542), F4–F30.
- Altonji, J.G., Card, D., 1991. The Effects of Immigration on the Labor Market Outcomes of Less-skilled Natives. In: Abowd, J.M., Freeman, R.B. (Eds.), Immigration, Trade, and the Labor Market. University of Chicago Press, Chicago, pp. 201–234 (Chapter 7).
- Antman, F., 2009. The Intergenerational Effects of Paternal Migration on Schooling: What Can We Learn from Children's Time Allocations? Department of Economics, University of Colorado at Boulder Working Paper No. 08–03.
- Aydemir, A., Borjas, G.J., 2007. Cross-Country Variation in the Impact of International Migration: Canada, Mexico, and the United States. J. Eur. Econ. Assoc. 5 (4), 663–708.
- Aydemir, A., Chen, W.-H., Corak, M., 2009. Intergenerational Earnings Mobility Among the Children of Canadian Immigrants. Rev. Econ. Stat. 91 (2), 377–397.
- Aydemir, A., Robinson, C., 2008. Global Labour Markets, Return and Onward Migration. Can. J. Econ. 41 (4), 1285–1311.
- Autor, D.H., Levy, F., Murnane, R.J., 2003. The Skill Content of Recent Technological Change: An Empirical Exploration. Q. J. Econ. 118 (4), 1279–1334.
- Baker, M., Benjamin, D., 1994. The Performance of Immigrants in the Canadian Labor Market. J. Labor Econ. 12 (3), 369–405.
- Barrett, A., O'Connell, P.J., 2001. Is There a Wage Premium for Returning Irish Migrants? Econ. Soc. Rev. 32 (1), 1–21.
- Bartel, A., 1989. Where Do the New U.S. Immigrants Live? J. Labor Econ. 7 (4), 371-391.
- Barth, E., Bratsberg, B., Raaum, O., 2004. Identifying Earnings Assimilation of Immigrants under Changing Macroeconomic Conditions. Scand. J. Econ. 106 (1), 1–22.
- Barth, E., Bratsberg, B., Raaum, O., 2006. Local Unemployment and the Relative Wages of Immigrants: Evidence from the Current Population Surveys. Rev. Econ. Stat. 88 (2), 243–263.
- Basilio, L., Bauer, T., 2010. Transferability of Human Capital and Immigrant Assimilation: An Analysis for Germany. IZA Discussion Paper No. 4716.
- Becker, G.S., Tomes, N., 1979. An Equilibrium Theory of the Distribution of Income and Intergenerational Mobility. J. Polit. Econ. 87 (6), 1153–1189.
- Beggs, J.J., Chapman, B.J., 1988a. The International Transferability of Human Capital: Immigrant Labour Market Outcomes in Australia. In: Baker, L., Miller, P. (Eds.), The Economics of Immigration. Australian Government Publishing Service, Canberra, pp. 143–147.
- Beggs, J.J., Chapman, B.J., 1988b. Immigrant Wage Adjustment in Australia: Cross Section and Time-Series Estimates. Econ. Rec. 64 (186), 161–167.
- Behrman, J., Taubman, P., 1985. Intergenerational Earnings Mobility in the United States: Some Estimates and a Test of Becker's Intergenerational Endowment's Model. Rev. Econ. Stat. 67 (1), 141–151.
- Beine, M., Defoort, C., Docquier, F., 2010. A Panel Data Analysis of the Brain Gain. World Dev. forthcoming.
- Beine, M., Docquier, F., Rapoport, H., 2001. Brain Drain and Economic Growth: Theory and Evidence. J. Dev. Econ. 64 (1), 275–289.
- Beine, M., Docquier, F., Rapoport, H., 2008. Brain Drain and Human Capital Formation in Developing Countries: Winners and Losers. Econ. J. 118 (528), 631–652.
- Bell, B.D., 1997. The Performance of Immigrants in the United Kingdom: Evidence from the GHS. Econ. J. 107 (441), 333–344.

Belot, M.V.K., Hatton, T.J., 2008. Immigrant Selection in the OECD. CEPR Discussion Paper No. 6675.

Ben-Porath, Y., 1967. The Production of Human Capital and the Life Cycle of Earnings. J. Polit. Econ. 75 (4), 352–365.

- Berman, E., Lang, K., Siniver, E., 2003. Language-skill Complementarity: Returns to Immigrant Language Acquisition. Labour Econ. 10, 265–290.
- Bertoli, Simone, 2010. The Informational Structure of Migration Decision and Migrants Self-Selection. Economics Letters, Vol. 108 (1), pp. 89–92.
- Betts, J., Fairlie, R., 2003. Does Immigration Induce 'Native Flight' from Public Schools into Private Schools? J. Public Econ. 87 (5–6), 987–1012.
- Bijwaard, G.E., 2004. Dynamic Economic Aspects of Migration. Medium Econometrische Toepassingen 12 (3), 26–30.
- Bishop, K.C., 2008. A Dynamic Model of Location Choice and Hedonic Valuation Unpublished Working Paper.
- Björklund, A., Jäntti, M., 1997. Intergenerational Income Mobility in Sweden Compared to the United States. Am. Econ. Rev. 87 (5), 1009–1018.
- Blanden, J., Goodman, A., Gregg, P., Machin, S., 2004. Changes in Intergenerational Mobility in Britain. In: Corak, M. (Ed.), Generational Income Mobility in North America and Europe. Cambridge University Press, pp. 122–146.
- Bleakley, H., Chin, A., 2004. Language Skills and Earnings: Evidence from Childhood Immigrants. Rev. Econ. Stat. 86 (2), 481–496.
- Bleakley, H., Chin, A., 2008. What Holds Back the Second Generation? The Intergenerational Transmission of Language Human Capital among Immigrants. J. Hum. Resour. 43 (2), 267–298.
- Borjas, G.J., 1982. The Earnings of Male Hispanic Immigrants in the United States. Ind. Labor Relat. Rev. 35, 343–353.
- Borjas, G.J., 1985. Assimilation, Changes in Cohort Quality, and the Earnings of Immigrants. J. Labor Econ. 3 (4), 463–489.
- Borjas, G.J., 1987. Self-Selection and the Earnings of Immigrants. Am. Econ. Rev. 77 (4), 531-553.
- Borjas, G.J., 1992. Ethnic Capital and Intergenerational Mobility. Q. J. Econ. 107 (1), 123–150.
- Borjas, G.J., 1993. The Intergenerational Mobility of Immigrants. J. Labor Econ. 11 (1), 113–135.
- Borjas, G.J., 1994. Long-Run Convergence of Ethnic Skill Differentials: The Children and Grandchildren of the Great Migration. Ind. Labor Relat. Rev. 47 (4), 553–573.
- Borjas, G.J., 1995a. Assimilation and Changes in Cohort Quality Revisited: What Happened to Immigrant Earnings in the 1980s? J. Labor Econ. 13 (2), 201–245.
- Borjas, G.J., 1995b. Ethnicity, Neighborhoods, and Human-Capital Externalities. Am. Econ. Rev. 85 (3), 365–390.
- Borjas, G.J., 1999. The Economic Analysis of Immigration. In: Ashenfelter, O., Card, D. (Eds.), Handbook of Labor Economics Vol. 3. (3), 1697–1760 (Chapter 28).
- Borjas, G.J., 2000. The Economic Progress of Immigrants. In: Borjas, G.J. (Ed.), Issues in the Economics of Immigration, NBER Books, pp. 15–50 (Chapter 1).
- Borjas, G.J., 2003. The Labor Demand Curve Is Downward Sloping: Reexamining the Impact of Immigration on the Labor Market. Q. J. Econ. 118 (4), 1335–1374.
- Borjas, G.J., 2006a. Immigration in High-Skill Labor Markets: The Impact of Foreign Students on the Earnings of Doctorates. NBER Working Paper 12085.
- Borjas, G.J., 2006b. Making It in America: Social Mobility in the Immigrant Population. Future Child 16 (2), 55–71.
- Borjas, G.J., 2007. Do Foreign Students Crowd Out Native Students from Graduate Programs? In: Stephan, P.E., Ehrenberg, R.G. (Eds.), Science and the University (Science and Technology in Society). University of Wisconsin Press.
- Borjas, G.J., Bratsberg, B., 1996. Who Leaves? The Outmigration of the Foreign-Born. Rev. Econ. Stat. 78 (1), 165–176.
- Borjas, G.J., Freeman, R.B., Katz, L.F., 1997. How Much Do Immigration and Trade Affect Labor Market Outcomes? Brookings Pap. Econ. Act 1, 1–67.

- Bound, J., Turner, S., Walsh, P., 2009. Internationalization of U.S. Doctorate Education. Population Studies Center Research Report 09-675, University of Michigan.
- Bratsberg, B., Raaum, O., Sorlie, K., 2007. Foreign-Born Migration to and from Norway. In: Özden, Ç., Schiff, M. (Eds.), International Migration, Economic Development and Policy. Palgrave Macmillan, New York, pp. 259–290 (Chapter 9).
- Bratsberg, B., Ragan, Jr., J.F. 2002. The Impact of Host-Country Schooling on Earnings: A Study of Male Immigrants in the United States. J. Hum. Resour. 37 (1), 63–105.
- Butcher, K.F., Card, D., 1991. Immigration and Wages: Evidence from the 1980s. Am. Econ. Rev. 81 (2), 292–296.
- Card, D., 1990. The Impact of the Mariel Boatlift on the Miami Labor Market. Ind. Labor Relat. Rev. 43 (2), 245–257.
- Card, D., 2001. Immigrant Inflows, Native Outflows, and the Local Labor Market Impacts of Higher Immigration. J. Labor Econ. 19 (1), 22–64.
- Card, D., DiNardo, J., Estes, E., 2000. The More Things Change: Immigrants Children of Immigrants in the 1940s, the 1970s, and the 1990s. In: Borjas, G.J. (Ed.), Issues in the Economics of Immigration. NBER, and University of Press, Chicago.
- Carliner, G., 1980. Wages, Earnings, and Hours of First, Second and Third Generation American Males. Econ. Inquiry 18 (1), 87–102.
- Carliner, G., 1981. Wage Differences by Language Group and the Market for Language Skills in Canada. J. Hum. Resour. 16 (3), 384–399.
- Casey, T., Dustmann, C., 2008. Intergenerational Transmission of Language Capital and Economic Outcomes. J. Hum. Resour. 43 (3), 660–687.
- Checchi, D., De Simone, G., Faini, R., 2007. Skilled Migration, FDI, and Human Capital Investment. IZA Discussion Paper No. 2795.
- Chiquiar, D., Hanson, G.H., 2005. International Migration, Self-Selection and the Distribution of Wages: Evidence from Mexico and the United States. J. Polit. Econ. 113 (2), 239–281.
- Chiswick, B.R., 1977. Sons of Immigrants: Are They at an Earnings Disadvantage? Am. Econ. Rev. Pap. Proc. 67, 376–380.
- Chiswick, B.R., 1978. The Effect of Americanization on the Earnings of Foreign-born Men. J. Polit. Econ. 86 (5), 897–921.
- Chiswick, B.R., 1991. Speaking, Reading, and Earnings among Low-Skilled Immigrants. J. Labor Econ. 9 (2), 149–170.
- Chiswick, B.R., 1998. Hebrew Language Usage: Determinants and Effects on Earnings among Immigrants in Israel. J. Popul. Econ. 11 (2), 253–271.
- Chiswick, B.R., DebBurman, N., 2004. Educational Attainment: Analysis by Immigrant Generation. Econ. Educ. Rev. 23 (4), 361–379.
- Chiswick, B.R., Lee, Y.L., Miller, P.W., 2005. Parents and Children Talk: English Language Proficiency within Immigrant Families. Rev. Econ. Househ. 3 (3), 243–268.
- Chiswick, B.R., Miller, P.W., 1992. Language in the Immigrant Labor Market: The Immigrant Experience in Canada and the United States. In: Chiswick, B.R. (Ed.), Immigration, Language and Ethnicity: Canada and the United States. American Enterprise Institute, Washington, DC, pp. 229–296.
- Chiswick, B.R., Miller, P.W., 1994. The Determinants of Post-Immigration Investments in Education. Econ. Educ. Rev. 13 (2), 163–177.
- Chiswick, B.R., Miller, P.W., 1995. The Endogeneity between Language and Earnings: International Analyses. J. Labor Econ. 13 (2), 246–288.
- Chiswick, B.R., Miller, P.W., 2002. Immigrant Earnings: Language Skills, Linguistic Concentration and the Business Cycle. J. Popul. Econ. 15 (1), 31–57.
- Chiswick, B.R., Miller, P.W., 2003. The Complementarity of Language and Other Human Capital: Immigrant Earnings in Canada. Econ. Educ. Rev. 22 (5), 469–480.
- Chiswick, B.R., Miller, P.W., 2005. Do Enclaves Matter in Immigrant Adjustment? City Community 4 (1), 5–35.

- Chiswick, B.R., Miller, P.W., 2007. The International Transferability of Immigrants' Human Capital Skills. IZA Discussion Paper No. 2670.
- Chiswick, B.R., Miller, P.W., 2008. Why is the Payoff to Schooling Smaller for Immigrants? Labour Econ. 15 (6), 1317–1340.
- Chiswick, B.R., Repetto, G., 2001. Immigrant Adjustment in Israel: The Determinants of Literacy and Fluency in Hebrew and Their Effects on Earnings. In: Djajic, S. (Ed.), International Migration: Trends, Policies and Economic Impact. Routledge, London, pp. 204–228.
- Clark, K., Lindley, J.K., 2009. Immigrant Assimilation Pre and Post Labour Market Entry: Evidence from the UK Labour Force Survey. J. Popul. Econ. 22 (1), 175–198.
- Co, C.Y., Gang, I., Yun, M.-S., 2000. Returns to Returning. J. Popul. Econ. 13 (1), 57-79.
- Cobb-Clark, D.A., 1993. Immigrant Selectivity and Wages—The Evidence for Women. Am. Econ. Rev. 83 (4), 986–993.
- Cohen-Goldner, S., Eckstein, Z., 2008. Labor Mobility of Immigrants: Training, Experience, Language, and Opportunities. Int. Econ. Rev. 49 (3), 837–872.
- Commander, S., Kangasniemi, M., Winters, L.A., 2004. The Brain Drain: Curse or Boon? A Survey of the Literature. In: Baldwin, R., Winters, L.A. (Eds.), Challenges to Globalization. University of Chicago Press, Chicago, pp. 235–272.
- Constant, A.F., Zimmermann, K.F., 2008. Measuring Ethnic Identity and Its Impact on Economic Behavior. J. Eur. Econ. Assoc. 6 (2-3), 424–433.
- Corak, M. (Ed.), 2004. Generational Income Mobility in North America and Europe. Cambridge University Press, Cambridge.
- Corak, M., 2006. Do Poor Children Become Poor Adults? Lessons for Public Policy from a Cross Country Comparison of Generational Earnings Mobility. In: Creedy, J., Kalb, G. (Eds.), Research on Economic Inequality, Vol. 13. Dynamics of Inequality and Poverty. Elsevier Press, Amsterdam.
- Corak, M., Heisz, A., 1999. The Intergenerational Earnings and Income Mobility of Canadian Men: Evidence from Longitudinal Income Tax Data. J. Hum. Resour. 34 (3), 504–533.
- Cortes, K.E., 2004. Are Refugees Different from Economic Immigrants? Some Empirical Evidence on the Heterogeneity of Immigrant Groups in the United States. Rev. Econ. Stat. 86 (2), 465–480.
- Cortes, K.E., 2006. The Effects of Age at Arrival and Enclave Schools on the Academic Performance of Immigrant Children. Econ. Educ. Rev. 25 (2), 121–132.
- Cortés, P., 2010. The Feminization of International Migration and its Effects on the Families Left Behind: Evidence from the Philippines. unpublished manuscript.
- Cortés, P., Pan, J., 2009. Outsourcing Household Production: The Demand for Foreign Domestic Helpers and Native Labor Supply in Hong Kong. Working Paper, University of Chicago, Booth School of Business.
- Cortés, P., Tessada, J., 2009. Low-Skilled Immigration and the Labor Supply of Highly Educated Women. unpublished manuscript.
- Cox Edwards, A., Ureta, M., 2003. International Migration, Remittances and Schooling: Evidence from El Salvador. J. Development Economics 72 (2), 429–461.
- Cunha, F., Heckman, J.J., Lochner, L., Masterov, D.V., 2006. Interpreting the Evidence on Life Cycle Skill Formation. In: Hanushek, E.A., Welch, F. (Eds.), Handbook of the Economics of Education. North-Holland, Amsterdam, pp. 697–812 (Chapter 12).
- Cutler, D.M., Glaeser, E.L., Vigdor, J.L., 2005. Ghettos and the Transmission of Ethnic Capital. In: Loury, G., Modood, T., Teles, S.M. (Eds.), Ethnicity, Social Mobility and Public Policy: Comparing the US and UK. Cambridge University Press, UK, pp. 204–222.
- D'Addio, A.C., 2019. Intergenerational Transmission of Disadvantage: Mobility or Immobility across Generations? A Review of the Evidence for OECD Countries. OECD Social, Employment and Migration Working Paper No. 52.
- Dahl, G., 2002. Mobility and the Return to Education: Testing a Roy Model with Multiple Markets. Econometrica 70 (6), 2367–2420.
- De Brauw, A., Giles, J., 2006. Migrant Opportunity and the Educational Attainment of Youth in Rural China. IZA Discussion Paper No. 2326.

- De Freitas, G., 1980. The Earnings of Immigrants in the American Labor Market. PhD dissertation, Columbia University.
- Deutsch, J., Epstein, G.S., Lecker, T., 2006. Multi-Generation Model of Immigrant Earnings: Theory and Application. In: Polachek, S.W., Chiswick, C., Rapoport, H. (Eds.), Research in Labor Economics. Vol. 24. Elsevier, Amsterdam, pp. 217–234.
- Docquier, F., Marfouk, A., 2006. International Migration by Educational Attainment (1990–2000). In: Özden, Ç., Schiff, M. (Eds.), Remittances and Development. Palgrave Macmillan, New York.
- Docquier, F., Rapoport, H., 2009. Skilled Migration: The Perspective of Developing Countries. In: Baghwati, J., Hanson, G. (Eds.), Skilled Immigration Today: Prospects, Problems and Policies. Russell Sage Foundation, New York, pp. 247–284 (Chapter 9).
- Domingues Dos Santos, M., Postel-Vinay, F., 2003. Migration as a Source of Growth: the Perspective of a Developing Country. J. Popul. Econ. 16 (1), 161–175.
- Domingues Dos Santos, M., Postel-Vinay, F., 2004. The Impact of Temporary Migration on Human Capital Accumulation and Economic Development. Brussels Econ. Rev. 47 (1), 77–88.
- Duleep, H.O., Regets, M.C., 1999. Immigrants and Human-Capital Investment. Am. Econ. Rev. Pap. Proc. 89 (2), 186–191.
- Duleep, H.O., Regets, M.C., 2002. The Elusive Concept of Immigrant Quality. IZA Discussion Paper No. 631.
- Dustmann, C., 1993. Earnings Adjustment of Temporary Migrants. J. Popul. Econ. 6 (2), 153-168.
- Dustmann, C., 1994a. Return Intentions of Migrants: Theory and Evidence. CEPR Discussion Paper No. 906.
- Dustmann, C., 1994b. Speaking Fluency, Writing Fluency and Earnings of Migrants. J. Popul. Econ. 7 (2), 133–156.
- Dustmann, C., 1995. Savings Behavior of Migrant Workers—a Life Cycle Analysis. Zeitschrift fuer Wirtschafts-und Sozialwissenschaften 4, 511–533.
- Dustmann, C., 1997. Differences in the Labour Market Behaviour between Temporary and Permanent Migrant Women. Labour Econ. 4 (1), 29–46.
- Dustmann, C., 1999. Temporary Migration, Human Capital, and Language Fluency of Migrants. Scand. J. Econ. 101 (2), 297–314.
- Dustmann, C., 2000. Temporary Migration and Economic Assimilation. Swedish Econ. Policy Rev. 7, 213–244.
- Dustmann, C., 2003. Return Migration, Wage Differentials and the Optimal Migration Duration. Eur. Econ. Rev. 47 (2), 353–369.
- Dustmann, C., 2008. Return Migration, Investment in Children, and Intergenerational Mobility: Comparing Sons of Foreign- and Native-Born Fathers. J. Hum. Resour. 43 (2), 299–324.
- Dustmann, C., Fabbri, F., 2003. Language Proficiency and Labour Market Performance of Immigrants in the UK. Econ. J. 113 (489), 695–717.
- Dustmann, C., Fabbri, F., Preston, I., 2005. The Impact of Immigration on the British Labour Market. Econ. J. 115 (507), F324–F341.
- Dustmann, C., Fadlon, I., Weiss, Y., 2010. Return Migration, Human Capital Accumulation and the Brain Drain. J. Dev. Econ. forthcoming.
- Dustmann, C., Frattini, T., Halls, C., 2010. Assessing the Fiscal Costs and Benefits of A8 Migration to the UK. Fisc. Stud. 31 (1), 1–41.
- Dustmann, C., Frattini, T., Lanzara, G., 2010. Educational Achievements and Ambitions of Second Generation Immigrants—An International Comparison. mimeo, UCL.
- Dustmann, C., Frattini, T., Preston, I., 2008. The Effect of Immigration along the Distribution of Wages. CReAM Discussion Paper No. 03/08.
- Dustmann, C., Glitz, A., Schönberg, U., 2010. Referral-based Job Search Networks. unpublished manuscript.
- Dustmann, C., Glitz, A., Vogel, T., 2010. Employment, Wages, and the Economic Cycle: Differences between Immigrants and Natives. Eur. Econ. Rev. 54 (1), 1–17.
- Dustmann, C., Machin, S., Schönberg, U., 2010. Ethnicity and Educational Achievement in Compulsory Schooling. Econ. J. 120 (546), F272–F297.

- Dustmann, C., Theodoropoulos, N., 2010. Ethnic Minority Immigrants and Their Children in Britain. Oxf. Econ. Pap. 62 (2), 209–233.
- Dustmann, C., van Soest, A., 2001. Language Fluency and Earnings: Estimation with Misclassified Language Indicators. Rev. Econ. Stat. 83 (4), 663–674.
- Dustmann, C., van Soest, A., 2002. Language and the Earnings of Immigrants. Ind. Labor Relat. Rev. 55 (3), 473–492.
- Dustmann, C., van Soest, A., 2004. An Analysis of Speaking Fluency of Immigrants Using Ordered Response Models With Classification Errors. J. Bus. Econ. Stat. 22 (3), 312–321.
- Dustmann, C., Weiss, Y., 2007. Return Migration: Theory and Empirical Evidence from the UK. Br. J. Ind. Relat. 45, 236–256.
- Eckstein, Z., Weiss, Y., 2004. On the Wage Growth of Immigrants: Israel 1990–2000. J. Eur. Econ. Assoc. 2 (4), 665–695.
- Edin, P.-A., Fredriksson, P., Aslund, O., 2003. Ethnic Enclaves and the Economic Success of Immigrants— Evidence from a Natural Experiment. Q. J. Econ. 118 (1), 329–357.
- Edin, P.-A., LaLonde, R.J., Aslund, O., 2000. Emigration of Immigrants and Measures of Immigrant Assimilation: Evidence from Sweden. Swedish Econ. Policy Rev. (7), 163–204.
- Entorf, H., Minoiu, N., 2005. What a Difference Immigration Policy Makes: A Comparison of PISA Scores in Europe and Traditional Countries of Immigration. German Econ. Rev. 6 (2), 355–376.
- Farré, L., González, L., Ortega, F., 2009. Immigration, Family Responsibilities and the Labor Supply of Skilled Native Women. CReAM Discussion Paper No. 16/09.
- Feliciano, C., 2005. Educational Selectivity in U.S. Immigration: How Do Immigrants Compare to Those Left Behind? Demography 42 (1), 131–152.
- Fernández-Huertas Moraga, J., 2010. New Evidence on Emigrant Selection. Rev. Econ. Stat. forthcoming.
- Finn, M.G., 2007. Stay Rates of Foreign Doctorate Recipients from US Universities, 2005. Oak Ridge Institute for Science and Education, Oak Ridge.
- Friedberg, R.M., 1992. The Labor Market Assimilation of Immigrants in the United States: The Role of Age at Arrival. Unpublished Manuscript, Brown University, Providence, RI.
- Friedberg, R.M., 2000. You Can't Take It with You? Immigrant Assimilation and the Portability of Human Capital. J. Labor Econ. 18 (2), 221–251.
- Friedberg, R.M., 2001. The Impact of Mass Migration on the Israeli Labor Market. Q. J. Econ. 116 (4), 1373–1408.
- Funkhouser, E., Trejo, S.J., 1995. The Labor Market Skills of Recent Male Immigrants: Evidence from the Current Population Survey. Ind. Labor Relat. Rev. 48 (4), 792–811.
- Gang, I.N., Zimmermann, K.F., 2000. Is Child Like Parent? Educational Attainment and Ethnic Origin. J. Hum. Resour. 35 (3), 550–569.
- Ganzeboom, H.B.G., De Graaf, P.M., de Leeuw, J., Treiman, D.J., 1992. A Standard Internacional Socio-Economic Index of Occupational Status. Soc. Sci. Res. 21 (1), 1–56.
- Glitz, A., 2011. The Labor Market Impact of Immigration: A Quasi-Experiment Exploiting Immigrant Location Rules in Germany. J. Labor Econ. forthcoming.
- Gonzalez, A., 2003. The Education and Wages of Immigrant Children: The Impact of Age at Arrival. Econ. Educ. Rev. 22 (2), 203–212.
- Goos, M., Manning, A., 2007. Lousy and Lovely Jobs: The Rising Polarization of Work in Britain. Rev. Econ. Stat. 89 (1), 118–133.
- Gould, E.D., Lavy, V., Paserman, M.D., 2009. Does Immigration Affect the Long-term Educational Outcomes of Natives? Quasi-Experimental Evidence. Econ. J. 119 (540), 1243–1269.
- Gould, E.D., Moav, O., 2010. When is 'Too Much' Inequality Not Enough? The Selection of Israeli Emigrants. CReAM Discussion Paper No. 14/10.
- Green, C., Kler, P., Leeves, G., 2007. Immigrant Overeducation: Evidence from Recent Arrivals to Australia. Econ. Educ. Rev. 26 (4), 420–432.
- Green, D.A., Worswick, C., 2004. Immigrant Earnings Profiles in the Presence of Human Capital Investment: Measuring Cohort and Macro Effects. IFS Working Paper 04/13.

- Grenier, G., 1984. The Effect of Language Characteristics on the Wages of Hispanic American Males. J. Hum. Resour. 19 (1), 35–52.
- Grogger, J., Hanson, G., 2008. Income Maximization and the Selection and Sorting of International Migrants. NBER Working Paper No. 13821, Cambridge, MA.
- Hanson, G., 2007. Emigration, Labor Supply, and Earnings in Mexico. In: Borjas, G.J. (Ed.), Mexican Immigration. University of Chicago Press, Chicago, pp. 117–142.
- Hanson, G., 2010. International Migration and the Developing World. In: Rodrik, D., Rosenzweig, M. (Eds.), Handbook of Development Economics. Vol. 5. North Holland, Amsterdam, pp. 4363–4414 (Chapter 66).
- Hanson, G., Woodruff, C., 2003. Emigration and Educational Attainment in Mexico. Mimeo, University of California at San Diego.
- Hanushek, E.A., Woessmann, L., 2008. The Role of Cognitive Skills in Economic Development. J. Econ. Literature 46 (3), 607–668.
- Hanushek, E.A., Woessmann, L., 2010. The Economics of International Differences in Educational Achievement. In: Hanushek, E.A., Machin, S., Woessmann, L. (Eds.), Handbook of the Economics of Education. Vol. 3. North Holland, Amsterdam forthcoming.
- Hanushek, E.A., Zhang, L., 2009. Quality-Consistent Estimates of International Schooling and Skill Gradients. J. Hum. Capital 3 (2), 107–143.
- Hammarstedt, M., Palme, M., 2006. Intergenerational Mobility, Human Capital Transmission and the Earnings of Second-Generation Immigrants in Sweden. IZA Discussion Paper No. 1943.
- Hammarstedt, M., 2009. Intergenerational Mobility and the Earnings Position of First-, Second-, and Third-Generation Immigrants. Kyklos 62 (2), 275–292.
- Hayfron, J.E., 1998. The Performance of Immigrants in the Norwegian Labor Market. J. Popul. Econ. 11 (2), 293–303.
- Heckman, J.J., 1979. Sample Selection Bias as a Specification Error. Econometrica 47 (1), 153-161.
- Heckman, J.J., Honoré, B., 1990. The Empirical Content of the Roy Model. Econometrica 58 (5), 1121–1149.
- Hicks, J., 1932. The Theory of Wages. MacMillan, New York.
- Hirschman, C., 2001. The Educational Enrollment of Immigrant Youth: A Test of the Segmented-Assimilation Hypothesis. Demography 39 (3), 317–336.
- Hoxby, C.M., 1998. Do Immigrants Crowd Disadvantaged American Natives out of Higher Education? In: Hamermesh, D.S., Bean, F.D. (Eds.), Help or Hindrance? The Economic Implications of Immigration for African Americans. Russell Sage Foundation, New York, pp. 282–321.
- Hunt, J., 2010. Which Immigrants Are Most Innovative and Entrepreneurial? Distinctions by Entry Visa. IZA Discussion Paper No. 4745.
- Hunt, J., Gauthier-Loiselle, M., 2010. How Much Does Immigration Boost Innovation? Am. Econ. J.: MacroEcon. 2 (2), 31–56.
- Husted, L., Nielsen, H.S., Rosholm, M., Smith, N., 2001. Employment and Wage Assimilation of Male First Generation Immigrants in Denmark. Int. J. Manpower 22, 39–68.
- Iara, A., 2006. Skill Diffusion by Temporary Migration? Returns to Western European Working Experience in the EU Accession Countries. A Global Development Network Working Paper.
- Ibarrarán, P., Lubotsky, D., 2007. Immigration and Self-Selection: New Evidence from the 2000 Mexican Census. In: Borjas, G.J. (Ed.), Mexican Immigration to the United States. University of Chicago Press, pp. 159–192 (Chapter 5).
- ICMPD, 2006. Comparative Study on Policies towards Foreign Graduates. Study on Retention Policies towards Foreign Students in Industrialised Countries. ICMPD, Vienna.
- Jackson, O., 2010. Does Immigration Crowd Natives Into or Out of Higher Education? University of Michigan Working Paper.
- Jaeger, D.A., 2007. Green Cards and the Location Choices of Immigrants in the United States, 1971–2000. Res. Labor Econ. 27, 131–183.
- Jensen, P., Pedersen, P.J., 2007. To Stay or Not To Stay? Out-Migration of Immigrants from Denmark. Int. Migr. 45 (5), 87–113.

- Johnson, N., Kotz, S., 1972. Distribution in Statistics: Continuous Multivariate Distributions. John Wiley & Sons, New York.
- Kaestner, R., Malamud, O., 2010. Self-selection and International Migration: New Evidence from Mexico. NBER Working Paper No. 15765.
- Kato, T., Sparber, C., 2010. Quotas and Quality: The Effect of H-1B Visa Restrictions on the Pool of Prospective Undergraduate Students from Abroad. CReAM Discussion Paper No. 10/10.
- Kee, P., 1995. Native-Immigrant Wage Differentials in the Netherlands: Discrimination? Oxf. Econ. Pap. 47 (2), 302–317.
- Kennan, J., 2010. Higher Education Subsidies and Human Capital Mobility. Working Paper.
- Kennan, J., Walker, J.R., 2010. The Effect of Expected Income on Individual Migration Decisions. Econometrica forthcoming.
- Kerr, W.R., Lincoln, W., 2010. The Supply Side of Innovation: H-1B Visa Reforms, U.S. Ethnic Invention. NBER Working Paper No. 15768.
- Khan, A., 1997. Post-migration Investment in Education by Immigrants in the United States. Q. Rev. Econ. Finance 37 (Supplement 1), 285–313.
- Kossoudji, S.A., 1988. English Language Ability and the Labor Market Opportunities of Hispanic and East Asian Immigrant Men. J. Labor Econ. 6 (2), 205–228.
- Kossoudji, S.A., 1989. Immigrant Worker Assimilation: Is It a Labor Market Phenomenon? J. Hum. Resour. 24 (3), 494–527.
- LaLonde, R.J., Topel, R.H., 1991. Immigrants in the American Labor Market: Quality, Assimilation, and Distributional Effects. Am. Econ. Rev. Pap. Proc. 81 (2), 297–302.
- LaLonde, R.J., Topel, R.H., 1992. The Assimilation of Immigrants in the U.S. Labor Market. In: Borjas, G.J., Freeman, R.B. (Eds.), Immigration and the Work Force: Economic Consequences for the United States and Source Areas. University of Chicago Press, Chicago, pp. 67–92.
- Lazear, E.P., 1979. Why Is There Mandatory Retirement? J. Polit. Econ. 87 (6), 1261–1284.
- Lindley, J., Lenton, P., 2006. The Over-Education of UK Immigrants: Evidence from the Labour Force Survey. Sheffield Economic Research Paper Series No. 2006001.
- Longva, P., Raaum, O., 2003. Earnings Assimilation of Immigrants in Norway—A Reappraisal. J. Popul. Econ. 16 (1), 177–193.
- López-Córdova, J.E., 2005. Globalization, Migration, and Development: The Role of Mexican Migrant Remittances. Economía 6 (1), 217–256.
- Lubotsky, D., 2007. Chutes or Ladders? A Longitudinal Analysis of Immigrant Earnings. J. Polit. Econ. 115 (5), 820–867.
- Lucas, R.E.B., 2005. International Migration: Lessons from Recent Data. World Bank. Washington, DC. Presentation in the Migration Seminar Series.
- Manacorda, M., Manning, A., Wadsworth, J., 2011. The Impact of Immigration on the Structure of Male Wages: Theory and Evidence from Britain. J. Eur. Econ. Assoc. forthcoming.
- Marchiori, L., Shen, I.-L., Docquier, F., 2009. Brain Drain in Globalization: A General Equilibrium Analysis from the Sending Countries' Perspective. IZA Discussion Paper No. 4207.
- Mattoo, A., Neagu, I.C., Özden, C., 2008. Brain Waste? Educated Immigrants in the US Labor Market. J. Dev. Econ. 87 (2), 255–269.
- McKenzie, D.J., Rapoport, H., 2007. Self-selection Patterns in Mexico-U.S. Migration: The Role of Migration Networks. CReAM Discussion Paper No. 01/07.
- McKenzie, D., Rapoport, H., 2010. Can Migration Reduce Educational Attainment? Evidence from Mexico. J. Popul. Econ. forthcoming.
- McManus, W., Gould, W., Welch, F., 1983. Earnings of Hispanic Men: The Role of English Language Proficiency. J. Labor Econ. 1 (2), 101–130.
- Mishra, P., 2007. Emigration and Wages in Source Countries: Evidence from Mexico. J. Dev. Econ. 82 (1), 180–199.
- Mountford, A., 1997. Can A Brain Drain Be Good for Growth in the Economy? J. Dev. Econ. 53 (2), 287–303.
- Neal, D., 2006. Why Has Black-White Skill Convergence Stopped? In: Hanushek, E., Welch, F. (Eds.), Handbook of the Economics of Education. Vol. 1. Elsevier, Amsterdam, pp. 511–576 (Chapter 9).

- Nekby, L., 2004. The Emigration of Immigrants, Return vs. Onward Migration: Evidence from Sweden. Research Papers in Economics No. 2004:7, Stockholm University, Department of Economics.
- Neymotin, F., 2009. Immigration and its Effect on the College-Going Outcomes of Natives. Econ. Educ. Rev. 28 (5), 538–550.
- Nielsen, C.P., 2007. Immigrant Overeducation: Evidence from Denmark. World Bank Policy Research Working Paper 4234.
- Nielsen, H.S., Rosholm, M., Smith, N., Husted, L., 2003. The School-to-Work Transition of 2nd Generation Immigrants in Denmark. J. Popul. Econ. 16 (4), 755–786.
- OECD, 2006. Where Immigrant Students Succeed—A Comparative Review of Performance and Engagement in PISA 2003. OECD Publishing, Paris.
- OECD, 2007. PISA 2006-Science Competencies for Tomorrow's World. OECD Publishing, Paris.
- OECD, 2008. International Migration Outlook. OECD Publishing, Paris.
- OECD, 2009. PISA 2006 Technical Report. OECD Publishing, Paris.
- OECD, 2010. International Migration Outlook. OECD Publishing, Paris.
- Okkerse, L., 2008. How to Measure Labour Market Effects of Immigration: A Review. J. Econ. Surv. 22 (1), 1–30.
- Oosterbeek, H., Webbink, D., 2006. Assessing the Returns to Studying Abroad. CPB Working Paper No. 64.
- Orrenius, P.M., Zavodny, M., 2005. Self-selection Among Undocumented Immigrants from Mexico. J. Dev. Econ. 78 (1), 215–240.
- Ortega, F., Peri, G., 2009. The Causes and Effects of International Migrations: Evidence from OECD Countries 1980-2005. NBER Working Paper No. 14833.
- Osterberg, T., 2000. Economic Perspectives on Immigrants and Intergenerational Transmissions. Ekonomiska Studier 102, 115–149.
- Ottaviano, G.I.P., Peri, G., 2011. Rethinking the Effects of Immigration on Wages. J. Eur. Econ. Assoc. forthcoming.
- Palifka, B., 2003. Study Abroad and the Transition from School to Work for Mexican College Graduates. Working Paper.
- Parent, D., 2009. Intergenerational Progress in Educational Attainment When Institutional Change Really Matters: A Case Study of Franco-Americans vs. French-Speaking Quebecers. CLSRN Working Paper No. 2009–36.
- Peri, G., Sparber, C., 2008. Highly Educated Immigrants and Occupational Choice. CReAM Discussion Paper No. 13/08.
- Peri, G., Sparber, C., 2009. Task Specialization, Immigration, and Wages. Am. Econ. J.: Appl. Econ. 1 (3), 135–169.
- Perreira, K.M., Harris, K.M., Lee, D., 2006. Making it in America: High School Completion by Immigrant and Native Youth. Demography 43 (3), 511–536.
- Piil Damm, A., 2009. Ethnic Enclaves and Immigrant Labor Market Outcomes: Quasi-Experimental Evidence. J. Labor Econ. 27 (2), 281–314.
- Ramos, F.A., 1992. Out-Migration and Return Migration of Puerto Ricans. In: Borjas, G.J., Freeman, R.B. (Eds.), Immigration and the Work Force: Economic Consequences for the United States and Source Areas. University of Chicago Press, Chicago, pp. 49–66 (Chapter 2).
- Rampney, B.D., Dion, G.S., Donahue, P.L., 2009. NAEP 2008 Trends in Academic Progress (NCES-2009-479), National Center for Education Statistics, Institute of Education Sciences. U.S. Department of Education, Washington, DC.
- Rapoport, H., Docquier, F., 2006. The Economics of Migrants' Remittances. In: Kolm, S.-C., Ythier, J.M. (Eds.), Handbook on the Economics of Giving, Reciprocity and Altruism. Vol. 2. North Holland, Amsterdam, pp. 1135–1198 (Chapter 17).
- Reinhold, S., Thom, K., 2009. Temporary Migration and Skill Upgrading: Evidence from Mexican Migrants. MEA Discussion Paper No. 09182.
- Riphahn, R.T., 2003. Cohort Effects in the Educational Attainment of Second Generation Immigrants in Germany: An Analysis of Census Data. J. Popul. Econ. 16 (4), 711–737.

Rivera-Batiz, F., 1990. English Language Proficiency and the Economic Progress of Immigrants. Econ. Lett. 34 (3), 295–300.

- Rivera-Batiz, F., 1992. English Language Proficiency and the Earnings of Young Immigrants in U.S.. Labor Markets. Rev. Policy Res. 11 (2), 165–175.
- Rooth, D.-O., Ekberg, J., 2003. Unemployment and Earnings for Second Generation Immigrants in Sweden. Ethnic Background and Parent Composition. J. Popul. Econ. 16 (4), 787–814.
- Rosenzweig, M., 2007. Education and Migration: A Global Perspective. Yale University, Mimeo.
- Roy, A.D., 1951. Some Thoughts on the Distribution of Earnings. Oxf. Econ. Pap. 3 (2), 135-146.
- Sanromá, E., Ramos, R., Simón, H., 2008. The Portability of Human Capital Immigrant Assimilation: Evidence for Spain. IZA Discussion Paper No. 3649.
- Sanromá, E., Ramos, R., Simón, H., 2009. Immigrant Wages in the Spanish Labour Market: Does the Origin of Human Capital Matter? IZA Discussion Paper No. 4157.
- Schaafsma, J., Sweetman, A., 2001. Immigrant Earnings: Age at Immigration Matters. Can. J. Econ. 34 (4), 1066–1099.
- Schiff, M., 2005. Brain Gain: Claims about Its Size and Impact on Welfare and Growth Are Greatly Exaggerated. IZA Discussion Paper No. 1599.
- Schnepf, S.V., 2007. Immigrants' Educational Disadvantage: An Examination across Ten Countries and Three Surveys. J. Popul. Econ. 20 (3), 527–545.
- Schoeni, R., 1997. New Evidence on the Economic Progress of Foreign-Born Men in the 1970s and 1980s. J. Hum. Resour. 32 (4), 683–740.
- Shields, M.A., Wheatley Price, S., 1998. The Earnings of Male Immigrants in England: Evidence from the Quarterly LFS. Appl. Econ. 30 (9), 1157–1168.
- Shortland, P., 2006. People on the Move: A Study of Migrant Movement Patterns to and from New Zealand. New Zealand Department of Labour, Wellington.
- Sjaastad, L.A., 1962. The Costs and Returns of Human Migration. J. Polit. Econ. 70 (5), 80-93.
- Smith, J., 2003. Assimilation across the Latino generations. Am. Econ. Rev. Pap. Proc. 93 (2), 315-319.
- Solon, G., 1992. Intergenerational Income Mobility in the United States. Am. Econ. Rev. 82 (3), 393–408.
- Solon, G., 1999. Intergenerational Mobility in the Labor Market. In: Ashenfelter, O.C., Card, D. (Eds.), Handbook of Labor Economics. Vol. 3A. North-Holland, Amsterdam, pp. 1761–1800 (Chapter 29).
- Solon, G., 2002. Cross-Country Differences in Intergenerational Earnings Mobility. J. Econ. Perspect. 16 (3), 59–66.
- Solon, G., 2004. A Model of Intergenerational Mobility Variation over Time and Place. In: Corak, M. (Ed.), Generational Income Mobility in North America and Europe. Cambridge University Press, Cambridge, England, pp. 38–47 (Chapter 2).
- Stark, O., Helmenstein, C., Prskawetz, A., 1997. A Brain Drain with a Brain Gain. Econ. Lett. 55 (2), 227–234.
- Stark, O., Wang, Y., 2002. Inducing Human Capital Formation: Migration as a Substitute for Subsidies. J. Public Econ. 86 (1), 29–46.
- Stewart, J.B., Hyclak, T., 1984. An Analysis of the Earnings Profiles of Immigrants. Rev. Econ. Stat. 66 (2), 292–296.
- Storesletten, K., 2000. Sustaining Fiscal Policy through Immigration. J. Polit. Econ. 108 (2), 300-323.
- Tainer, E., 1988. English Language Proficiency and the Determination of Earnings among Foreign-born Men. J. Hum. Resour. 23 (1), 108–122.
- Todd, P.E., Wolpin, K.I., 2003. On the Specification and Estimation of the Production Function for Cognitive Achievement. Econ. J. 113 (485), F3–F33.
- Trejo, S., 2003. Intergenerational Progress of Mexican-Origin Workers in the US Labor Market. J. Hum. Resour. 38 (3), 467–489.
- Van Ours, J.C., Veenman, J., 2006. Age at Immigration and Educational Attainment of Young Immigrants. Econ. Lett. 90, 310–316.
- Venturini, A., Villosio, C., 2008. Labour-Market Assimilation of Foreign Workers in Italy. Oxf. Rev. Econ. Policy 24 (3), 517–541.

- Vidal, J.-P., 1998. The Effect of Emigration on Human Capital Formation. J. Popul. Econ. 11 (4), 589–600.
- Wiers-Jenssen, J., Try, S., 2005. Labour Market Outcomes of Higher Education Undertaken Abroad. Studies High. Educ. 30 (6), 681–705.
- Wilkins, R., 2003. Immigrant Earnings Adjustment: The Impact of Age at Migration. Australian Econ. Pap. 42 (3), 292–315.
- Willis, R.J., 1987. Wage Determinants: A Survey and Reinterpretation of Human Capital Earnings Functions. In: Ashenfelter, O., Layard, R. (Eds.), first ed. Handbook of Labor Economics. Vol. 1. first ed. Elsevier, Amsterdam, pp. 525–602 (Chapter 10).
- Willis, R.J., Rosen, S., 1979. Education and Self-Selection. J. Polit. Econ. 87 (5), S7-S36.
- Yang, D., 2008. International Migration, Remittances, and Household Investment: Evidence from Philippine Migrants' Exchange Rate Shocks. Econ. J. 118 (528), 591–630.
- Zimmerman, D.J., 1992. Regression Toward Mediocrity in Economic Stature. Am. Econ. Rev. 82 (3), 409–429.



Inequality, Human Capital Formation, and the Process of Development

Oded Galor

Department of Economics, Brown University, Providence, RI 02912

Contents

1.	Introduction	443
	1.1. From the Classical to the Modern Perspective	443
	1.2. The Credit Market Imperfections Channel	445
	1.3. The Political Economy Channel	446
2.	The Benchmark Model	447
	2.1. Production of Final Output	448
	2.2. Factor Prices	449
	2.3. Individuals	449
	2.3.1 Preferences and Budget Constraint	450
	2.3.2 Optimization	451
	2.4. Fundamental Assumptions	451
	2.5. Occupational Choice	452
	2.5.1 Income of an Unskilled Worker	452
	2.5.2 Income of a Skilled Worker	452
	2.5.3 Parental Transfers and Occupational Choices	453
	2.6. Bequest Dynamics	455
	2.7. Distribution, Skill Composition, and Income	457
	2.7.1 Income Distribution and the Composition of Skills	457
	2.7.2 Persistence of Inequality	458
	2.7.3 Skill Composition and Income Per Capita	459
	2.8. Inequality and Economic Development	460
	2.9. Robustness	463
	2.9.1 Labor-Augmenting Technological Progress	463
	2.9.2 Interactions across Dynasties	465
	2.9.3 Random Shocks	465
	2.9.4 Concave Production of Human Capital and Alternative Utility Functions	466
3.	A Unified Theory of Inequality and Growth	466
4.	Nonfinancial Hurdles for Human Capital Accumulation	468
	4.1. Concentration of Landownership	468
	4.2. Social-Political Transitions	470
	4.3. Gender Inequality	471
5.	Evidence	471
	5.1. Inequality, Human Capital Formation, and Economic Growth	471
	5.2. Industrialization and Human Capital Formation	473

	5.2.1 Industrial Demand for Education	475
	5.2.2 Land Concentration: A Hurdle for Human Capital Formation	481
	5.2.3 Land Reforms and Education Reforms	482
	5.2.4 Political and Education Reforms	485
6.	Concluding Remarks	487
Re	eferences	489

Abstract

Conventional wisdom about the relationship between income distribution and economic development has been subjected to dramatic transformations in the past century. While Classical economists advanced the hypothesis that inequality is beneficial for economic development, the Neoclassical paradigm, which had subsequently dominated the field of macroeconomics, dismissed the Classical hypothesis and promoted the viewpoint that the study of income distribution has no importance for the understanding of macroeconomic activity and the growth process. A metamorphosis in these perspectives has taken place in the past two decades. Theory and subsequent empirical evidence have demonstrated that income distribution has a significant impact on the growth process.

The modern approach has demonstrated that in the presence of credit market imperfections, income distribution has a long-lasting effect on investment in human capital, entrepreneurial activity, aggregate income, and economic development. Moreover, in contrast to the Classical viewpoint, which underscored beneficial effects of inequality for the growth process, the modern perspective advanced the hypothesis that inequality may be detrimental for human capital formation and economic development.

The replacement of physical capital accumulation by human capital accumulation as the prime engine of economic growth has changed the qualitative impact of inequality on the process of development. In early stages of industrialization, as physical capital accumulation was a prime source of economic growth, inequality enhanced the process of development by channeling resources toward individuals whose marginal propensity to save is higher. However, in later stages of development, as human capital has become the main engine of economic growth, a more equal distribution of income, in the presence of credit constraints, has stimulated investment in human capital and economic growth.

While the process of industrialization raised the importance of human capital in the production process, reflecting its complementarity with physical capital and technology, human capital accumulation has not benefited all sectors of the economy. Inequality in the ownership of factors of production has generated an incentive for some better-endowed agents to block the implementation of institutional changes and policies that promote human capital formation, resulting in a suboptimal level of investment in human capital from a growth perspective. The transition from an agricultural to an industrial economy changed the nature of the main economic conflict in society. Unlike the agrarian economy, which was characterized by a conflict of interests between the landed aristocracy and the masses, the process of industrialization has brought about an additional conflict between the entrenched landed elite and the emerging capitalist elite. In light of a lower degree of complementarity between human capital and the agricultural sector, education has increased the productivity of labor in industrial production more than in agricultural and primary good production, inducing rural-to-urban migration and a decline in the return to landowners. Thus, while industrialists have had a direct economic incentive to support education policies that would foster human capital formation, landowners, whose interests lay in the reduction of the mobility of their labor force, have favored policies that deprived the masses of education.

The adverse effect of the implementation of public education on landowners' income from agricultural production has been magnified by the concentration of land ownership. As long as landowners affected the political process and thereby the implementation of growth-enhancing education policies, inequality in the distribution of land ownership has been a hurdle for human capital accumulation, slowing the process of industrialization, and the transition to modern growth. Variation in the distribution of ownership over land and other natural resources across countries has contributed to disparity in human capital formation and the industrial composition of the economy, and thus to divergent development patterns across the globe. Moreover, in some societies, geographical conditions that led to income inequality brought about oppressive institutions designed to maintain the political power of the elite and to preserve the existing inequality.

Keywords

Education Gender Gap Human Capital Income Distribution Inequality Development Unified Growth Theory

1. INTRODUCTION

Conventional wisdom about the relationship between income distribution and macroeconomic activity has been subjected to dramatic transformations in the past century. While Classical economists advanced the hypothesis that inequality is beneficial for economic development, the Neoclassical paradigm, which had subsequently dominated the field of macroeconomics, dismissed the Classical thesis and promoted the viewpoint that the study of income distribution has no significance for the understanding of macroeconomic activity and the growth process.

A metamorphosis in these perspectives has taken place in the past two decades. Theory and subsequent empirical evidence have demonstrated that income distribution does, in fact, have a significant impact on the growth process. Moreover, unlike the Classical viewpoint, which underlined the beneficial effects of inequality for the growth process, the modern theory has highlighted the potential adverse effects of inequality on the process of development.

1.1. From the Classical to the Modern Perspective

The Classical approach advanced the hypothesis that inequality is beneficial for economic development in the postindustrialization period (Kaldor (1955)).¹ It suggests that since

¹ Development economists advanced an additional hypothesis about the relationship between inequality and economic development that is largely tangential to the understating of this association in the modern, post-industrialization era. As argued by Rosenstein-Rodan (1943), Lewis (1954), Baldwin (1956), and North (1959), and formulated by Murphy, Shleifer, and Vishny (1989), in the absence of international demand for domestic industrial goods, a wide distribution of the income generated from the leading agricultural sector may be critical for industrialization.

the marginal propensity to save increases with wealth, inequality channels resources toward individuals whose marginal propensity to save is higher, increasing aggregate savings, capital accumulation, and economic growth.² However, the Classical hypothesis was implicitly dismissed by the representative agent paradigm that had dominated the field of macroeconomics. The influential Neoclassical approach rejected the relevance of heterogeneity, and thus the distribution of income, for macroeconomic analysis. It implicitly interpreted the observed relationship between inequality and economic growth as capturing the effect of the growth process on the distribution of income.³

The Neoclassical viewpoint has been challenged in the past two decades, as both theories and subsequent empirical evidence have demonstrated that income distribution has a significant impact on the growth process. In contrast to the representative agent approach which dominated the field of macroeconomics for several decades, the modern perspective, originated by Galor and Zeira (1988, 1993), has underlined the role of heterogeneity in the determination of macroeconomic activity. It has advanced a novel viewpoint that heterogeneity, and thus income distribution, plays an important role in the determination of aggregate economic activity and economic growth in the short run as well as in the long run.

Galor and Zeira have demonstrated that in the presence of credit market imperfections income distribution has a long-lasting effect on investment in human capital, aggregate income, and economic development. Moreover, in contrast to the Classical hypothesis, which underscored the virtues of inequality for economic growth, their research advanced the hypothesis that inequality may be detrimental for human capital formation and economic development.

The modern perspective about the relationship between inequality and economic development has subsequently emerged, resulting in a voluminous body of research that have highlighted the adverse effect of inequality on the process of development.⁴ The initial research has been widely classified into two broad approaches for the examination of the relationship between inequality and growth: the credit market imperfection approach and the political economy approach.⁵

² Echoing the insight of the Classical economists, it was established that within a Neoclassical growth model with a convex saving function, the distribution of income might lead to either an equalitarian or unequal distribution of income in the long run (Stiglitz (1969)), where the less-egalitarian equilibrium is superior (Bourguignon (1981)).

³ This viewpoint can be traced to the hypothesis advanced by Kuznets (1955), according to which, the inverted U relationship between inequality and economic development that he found reflects a causation from the process of development to the distribution of income.

⁴ This chapter, written from a macro-growth perspective, focuses on the literature that explores the effect of inequality on the development process, rather than on the forces that prevent (Loury (1981)) or generate persistent inequality within an economy (Benabou (1996); Durlauf (1996a); Fernández and Rogerson (1996), Mookherjee and Ray (2003)) or across economies (Galor and Mountford (2008); Galor (2010)).

⁵ An additional line of research that has generated less attention examined the effect of inequality on aggregate demand, innovations, and growth, in the presence of non-homothetic preferences (Chou and Talmain (1996); Matsuyama (2000); Foellmi and Zweimuller (2006)).

1.2. The Credit Market Imperfections Channel

The credit market imperfection approach for the study of income distribution and economic growth has explored the implications and the robustness of the effect of inequality on the process of development in the presence of credit market imperfections.

Galor and Zeira have demonstrated that in the presence of credit market imperfections and fixed costs associated with investment in education, occupational choices (and thus the efficient segmentation of the labor force between skilled and unskilled workers) are affected by the distribution of income. In particular, if the interest rate for borrowers is higher than that for lenders, inequality may result in an underinvestment in human capital.⁶ Inequality may therefore adversely affect macroeconomic activity and economic development in the short-run, and due to intergenerational transfers and their effect on the persistence of inequality, it may generate a detrimental effect on economic development in the long run as well.⁷

The credit market imperfection approach for the study of the effects of income distribution on economic growth, which has subsequently emerged, maintained the two fundamental assumptions of the Galor–Zeira model (i.e., credit market imperfections and fixed costs associated with individual-specific investment projects), establishing the robustness of the main hypothesis.⁸ Notably, Banerjee and Newman (1993) examine the effect of inequality on a different type of occupational choices—the choice between becoming an entrepreneur or a worker. They demonstrate that if credit markets are imperfect and fixed costs are associated with entrepreneurial activities, inequality may result in an underinvestment in entrepreneurial activity and may therefore be harmful for economic development.⁹ Furthermore, they establish that as long as wages are endogenous, the main hypothesis of the credit market imperfection approach is robust to

⁶ Although the provision of public education mitigates the effect of inequality on human capital formation, the adverse effect is still maintained due to the differential effect of inequality on: (i) the importance of forgone earnings in education decisions, (ii) the allocation of parental inputs in the production of the children's human capital (Galor and Tsiddon (1997b)), (iii) the ability of parents to optimally select the schooling environment for their children (Benabou (1996); Durlauf (1996a); Fernández and Rogerson (1996)), and (iv) assortative mating (Fernández and Rogerson (2001); Fernandez, Guner, and Knowles (2005)). Moreover, it should be noted that while the positive effect of wage inequality on the incentive to invest in human capital may counterbalance the adverse effect of limited parental resources on investment in human capital, other forms of inequality (e.g., wealth inequality and inequality in the distribution of income between capital, land, and labor) do not affect the incentive to invest in human capital.

⁷ In contrast, in sufficiently poor economies, where the fixed cost of education is high in comparison to the level of income per capita, inequality may permit at least members of the upper tail of the income distribution to undertake investment in human capital. Hence, higher inequality would be expected to be associated with higher investment in education.

⁸ The Galor–Zeira setup was further exploited by Quah (1996) to shed light on the emergence of convergence clubs (and thus persistent inequality) in the world economy, in the presence of imperfect capital mobility across economies.

⁹ Aghion and Bolton (1997) further demonstrate that redistribution improves the efficiency of the economy, because it enhances equality of opportunity and the trickle-down process from the rich to the poor.

the introduction of random shocks to the outcome of investment (in human capital or entrepreneurial activities).¹⁰

The interplay between income inequality and equality of opportunities that has been underlined by Galor and Zeira led to an additional strand of research within the credit market imperfection approach. This research examines the effect of inequality on the degree of intergenerational mobility and thus the efficiency in the allocation of talents across occupations (Fershtman, Murphy, and Weiss (1996); Owen and Weil (1998); Maoz and Moav (1999); Checchi, Ichino, and Rustichini (1999); Hassler, Rodriguez Mora, and Zeira (2007)).¹¹

Furthermore, the interaction between income inequality and credit market imperfections was placed at the center of an important literature that examines the relationship between segregation and persistent inequality.¹² These studies have demonstrated that in the presence of credit market imperfections, inequality enhances segregation across communities and thus, in the light of local externalities in the production of human capital, it may generate persistent education and income gaps (Benabou (1996); Durlauf (1996a, 1996b); Fernández and Rogerson (1996)).¹³

1.3. The Political Economy Channel

The political economy approach for the study of the relationship between inequality and economic growth further advanced the viewpoint that inequality is harmful for economic development. Earlier studies have argued that inequality generates a pressure to adopt redistributive policies, and the distortions associated with these policies adversely affect investment in physical and human capital and thus the growth process.

These studies have suggested that in societies that are characterized by inequality, distributional conflicts may bias political decisions in favor of appropriation. Hence, since the incentives for productive accumulation of physical capital, human capital, and know-ledge hinge on the ability of individuals to privately appropriate the return on their investment, inequality may diminish investment and economic growth. In particular, using the median voter paradigm, it was hypothesized that in a more equitable society, taxation on physical capital (Alesina and Rodrik (1994)) and human capital (Persson

¹⁰ See also Piketty (1997).

¹¹ The adverse effect of inequality on occupational choices and intergenerational mobility is robust to the removal of credit market imperfections as long as parental human capital and social background are introduced into the analysis (Galor and Tsiddon (1997a, 1997b); Hassler and Mora (2000); Zilcha (2003); Mejía and St-Pierre (2008); Brezis and Temin (2008)).

¹² In an influential study, Loury (1981) has underlined the lack of intergenerational persistence of inequality, despite credit market imperfections.

¹³ Eicher, García-Peñalosa, and van Ypersele (2009) examine the interaction between inequality, corruption, and education.

and Tabellini (1994)) is lower, limiting the degree of distortions in investment decisions, and promoting economic growth.

In light of the inconsistency of this mechanism with empirical evidence (Perotti (1996)), subsequent theories advanced the thesis that inequality may in fact generate an incentive for better-endowed agents to lobby against redistribution, preventing efficient redistribution policies from being implemented (Saint-Paul and Verdier (1996); Benabou (2000, 2002)). Moreover, others have examined the long-run effects of inequality in the ownership of factors of production on the incentive for better-endowed agents to block the implementation of institutional changes and policies that promote human capital formation and thus economic growth (Engerman and Sokoloff (2000); Galor, Moav, and Vollrath (2009)).

2. THE BENCHMARK MODEL

The basic framework of analysis for the effect on inequality on income per capita is the Galor–Zeira model. The model demonstrates that under plausible conditions (i.e., credit market imperfections and fixed costs in the acquisition of human capital), income distribution has a long-lasting effect on investment in human capital, aggregate income, and the development process.¹⁴ In particular, if the interest rate for borrowers is higher than that for lenders, as is universally the case, the distribution of income affects occupational choices (and thus the efficient segmentation of the labor force between skilled and unskilled workers), and it may result in an underinvestment in human capital. Inequality may therefore adversely affect macroeconomic activity and economic development in the short-run, and due to intergenerational transfers and their effect on the persistence of inequality, it may adversely affect economic development in the long run as well.

Consider a small open overlapping-generations economy in which economic activity extends over infinite discrete time. In every period, the economy produces a single homogeneous good that can be used for consumption and investment. The good is produced in two sectors using capital, skilled labor, and unskilled labor in the production process. The stock of physical capital in every period is formed by aggregate domestic saving in the preceding period and net of international borrowing, whereas the segmentation of the labor force between skilled and unskilled labor in every period is the outcome of individuals' education decisions in an environment characterized by credit market imperfections.

¹⁴ The main hypothesis of the credit market imperfection approach is robust, however, to the removal of the assumption of a fixed cost of education or investment projects, as long as savings are an increasing function of wealth (Moav (2002); Galor and Moav (2004)).

2.1. Production of Final Output

Production occurs within a period. The output produced in the domestic economy at time t, Y_t , is the sum of the output produced in the skilled labor-intensive sector, Y_t^s , and the unskilled labor-intensive sector, Y_t^u . Namely,

$$Y_t = Y_t^s + Y_t^u. ag{5.1}$$

The output produced in the skilled labor-intensive sector is governed by a Neoclassical constant-returns-to-scale production technology,

$$Y_t^s = F(K_t, L_t^s) \equiv L_t^s f(k_t); \quad k_t \equiv K_t / L_t^s,$$
 (5.2)

where K_t and L_t^s are the quantities of physical capital and skilled labor employed in production at time *t*. Capital depreciates fully within a period.¹⁵ The intensive production function, $f(k_t)$, is monotonically increasing, strictly concave in k_t , and satisfies the Neoclassical boundary conditions that assure the existence of an interior level of k_t that maximizes profit.

The output produced in the unskilled-intensive sector is governed by a linear production technology that converts the input of unskilled labor into final output. In particular,

$$Y_t^u = aL_t^u, (5.3)$$

where a > 0 is the marginal productivity of unskilled labor in the production of final output.

Producers operate in a perfectly competitive environment. Given the interest rate, r_t , and the wage rate of skilled labor, w_t^s , producers in the skilled labor-intensive sector in period t choose the level of employment of capital, K_t , and the skilled labor, L_t^s , so as to maximize profits. That is,

$$\{K_t, L_t^s\} = \arg\max[L_t^s f(k_t) - w_t^s L_t^s - r_t K_t].$$
(5.4)

The producers' inverse demand for factors of production is therefore

$$r_t = f'(k_t) \equiv r(k_t);$$

$$w_t^s = f(k_t) - f'(k_t)k_t \equiv w^s(k_t).$$
(5.5)

Similarly, producers in the unskilled labor-intensive sector demand labor as long as the wage of an unskilled laborer does not exceed its productivity, *a*. The demand for unskilled labor in this sector in period *t* is therefore perfectly elastic at the wage level $w_t^s = a$.

¹⁵ Imperfect capital depreciation has no effect on the qualitative results.

2.2. Factor Prices

Suppose that capital is perfectly mobile internationally and the world interest rate is constant over time at level r > 0. Producers can borrow and individuals can lend unlimited funds at this rate at the world market.

The interest rate in the domestic economy in period t, r_t , is therefore equal to the constant world interest rate r, that is,

$$r_t = r. (5.6)$$

In particular, if the entire aggregate saving in the domestic economy would have been channeled toward the domestic production and would have generated a marginal productivity of capital that exceeds the world interest rate, international capital would flow into the domestic economy till the marginal productivity of capital in the domestic economy would be equal to the world interest rate. However, if the entire aggregate saving in the domestic economy would have been channeled toward the domestic production and would have generated a marginal productivity of capital that would be lower than the world interest rate, domestic savings would flow into the world economy until the marginal productivity of capital in the domestic economy would be equal to the world interest rate.

International capital mobility implies therefore that the ratio of capital to skilled labor employed in production, k_t , is constant over time. In particular, as follows from (5.5) and (5.6),

$$k_t = f'^{-1}(r) \equiv k,$$
 (5.7)

and thus the wage of an skilled worker,

$$w_t^s = w^s(k) \equiv w^s, \tag{5.8}$$

is constant over time at a level w^s .

Furthermore, the perfectly elastic demand for unskilled labor in the unskilledintensive sector implies that as long as unskilled labor is present in the economy,

$$w_t^u = a \equiv w^u. \tag{5.9}$$

2.3. Individuals

In every period, a generation that consists of a continuum of individuals of measure one is born. Each individual has a single parent and a single child. Individuals, within as well as across generations, are identical in their preferences and innate abilities. However, they may differ in their family wealth and thus, due to imperfect capital markets, in their investment in human capital.¹⁶

¹⁶ The introduction of population growth does not affect the qualitative results.

Individuals live for two periods. In the first period of life (childhood), an individual can either join the labor force as an unskilled worker or devote time to a costly acquisition of human capital. Their consumption in this period is an integral part of parental consumption.

Individuals receive a parental transfer (bequest) toward the end of the period. Those who choose to become skilled workers, channel the parental transfer toward the cost of education. If parental transfer is insufficient to cover the entire cost of education, they can access an imperfect capital market and borrow the remaining part at the borrowers' interest rate. If parental transfer exceeds the cost of education, the excess is saved for the second period of life at the lender interest rate. In contrast, individuals who choose to join the labor market directly as unskilled workers save their parental transfer and their wage income for adulthood.

In the second period of their lives (adulthood), individuals who did not acquire education in the first period of life continue to work as unskilled workers, whereas those who acquired education join the labor force as skilled workers. All individuals allocate their wage income and the returns on their savings between family consumption and capital transfers to their children.

2.3.1 Preferences and Budget Constraint

Preferences of an individual who is born in period t (a member of generation t) are defined over household consumption in adulthood, c_{t+1} , and over the intergenerational transfer (bequest) to the offspring, b_{t+1} . The preferences are represented by a log-linear utility function,

$$u_t = \alpha \log c_{t+1} + (1 - \alpha) \log b_{t+1}, \tag{5.10}$$

where $\alpha \in (0, 1)$.¹⁷

The budget constraint of a member of generation t during adulthood is therefore

$$c_{t+1} + b_{t+1} \le \omega_{t+1}, \tag{5.11}$$

where the level of wealth of individual *t* in the second period of life, ω_{t+1} , reflects the parental transfer and occupational decisions made in the first period of life. It consists of the individual's wage income in the second period of life net of loan repayments, capital income on savings, and wealth carried from the first period of life.

¹⁷ This utility function reflects the joy of giving. As discussed in the robustness section, the qualitative analysis will not be affected if preferences are defined over the utility of the offspring.

2.3.2 Optimization

In the second period of life (adulthood), individuals allocate their second-period wealth, ω_{t+1} , between consumption, c_{t+1} , and bequest, b_{t+1} , so as to maximize their utility function subject to their second-period budget constraint.

$$\{c_{t+1}, b_{t+1}\} = \operatorname{argmax}[\alpha \ln c_{t+1} + (1-\alpha) \ln b_{t+1}]$$

subject to: $c_{t+1} + b_{t+1} \le \omega_{t+1}.$ (5.12)

Hence, a fixed fraction, α , of their second-period wealth, ω_{t+1} , is devoted to consumption, and the remaining fraction, $(1 - \alpha)$, is devoted to bequest;

$$c_{t+1} = \alpha \omega_{t+1};$$

$$b_{t+1} = (1 - \alpha) \omega_{t+1}.$$
(5.13)

Moreover, the indirect utility function of members of generation t, v_t (i.e., the level of utility generated by the optimal choices of c_{t+1} and b_{t+1}) is monotonically increasing in their second-period wealth, ω_{t+1} ;

$$v_t = [\alpha \ln \alpha + (1 - \alpha) \ln (1 - \alpha)] + \ln \omega_{t+1}.$$
(5.14)

Thus, an occupational choice in the first period of life that maximizes the individual's second-period wealth, ω_{t+1} , maximizes the individual's utility.

2.4. Fundamental Assumptions

The effect of the distribution of income on occupational choices and macroeconomic activity in the short run is generated by the presence of credit market imperfections, whereas the long-run effects of income distribution on macroeconomic activity is triggered by a fixed cost associated with the acquisition of human capital.

Suppose that credit markets are imperfect. While individuals can lend unlimited funds at the world interest rate, r, the interest rate for individuals who wish to borrow in order to invest in human capital is higher than r, reflecting monitoring cost designed to avoid default and the inability of human capital to serve as a tangible collateral for the loan.¹⁸ Hence,

$$r < i, \tag{A1}$$

where r is the interest rate paid to lenders, and i is the interest rate on loans that are designed to finance investment in human capital.

¹⁸ Due to reputation and the cost of mobility, firms are assumed to be unable to evade debt payment and thus they can borrow at the world interest rate *r*. This simplifying assumption has no bearing on the qualitative results.

Suppose further that the acquisition of education is associated with a fixed cost, $c^{h} = h$.¹⁹ This fixed cost may reflect the indivisibility of human capital formation in general and of academic degrees in particular.²⁰ The fixed cost of education can be viewed as a weighted average of the payments to teachers, administrators, and maintenance workers in the school system (i.e., a weighted average of the wages of skilled and unskilled workers). In particular,

$$c^{h} = \theta w^{s} + (1 - \theta) w^{u} \equiv h > 0, \tag{A2}$$

for some $\theta \in [0, 1]$.

2.5. Occupational Choice

In the first period of life, individuals make an occupational choice. They can either acquire education and work in adulthood as skilled workers or join the labor force directly as unskilled worker and remain unskilled in adulthood.

2.5.1 Income of an Unskilled Worker

An individual *t* who decides to join the labor force directly as an unskilled worker earns in the first period of life the wage of an unskilled worker, w^{u} . In addition, in the end of the first period, the individual receives a bequest of b_t . Since consumption in childhood is an intrinsic part of the household consumption, these resources are saved for adulthood. In the second period of life (adulthood), the individual's wealth consists of first-period saving, $w^{u} + b_t$, capital income on the saving, $(w^{u} + b_t)r$, in addition to their second-period wage income, w^{u} . Hence, the second-period wealth, ω_{t+1}^{u} , of an unskilled member of generation *t* who receive an inheritance, b_t , is

$$\omega_{t+1}^{u} = w^{u}(2+r) + (1+r)b_{t} \equiv \omega^{u}(b_{t}).$$
(5.15)

2.5.2 Income of a Skilled Worker

An individual *t* who decides to acquire education and to join the labor force in the second period of life as a skilled worker earns in the second period of life the wage of a skilled worker, w^s . The wealth of the individual in period t + 1, ω_{t+1}^s , depends on whether parental transfers in the first period of life, b_t , are sufficient to cover the cost of education, *h*. If $b_t < h$, the individual borrows in the first period of life the additional required funds, $(h - b_t)$, and repays the loan along with the interest rate for borrowers, *i*, from the wage

¹⁹ As underlined in section 2.9, the main hypothesis is robust, however, to the removal of the assumption of a fixed cost of education or investment projects, as long as savings are an increasing function of wealth.

²⁰ This indivisibility is reflected in a discrete jump in the return to high school graduates versus high school dropouts, or in the return to college graduates versus college dropouts.

income in the second period of life. However, if $b_t \ge h$, the individual finances the entire cost of education using the parental transfer, saving the excess funds, $(b_t - h)$. In the second period of life, the individual wealth consists of wage income, w_s , saving, $b_t - h$, and capital income, $(b_t - h)r$.

Hence, the second-period wealth, ω_{t+1}^{s} , of a skilled member of generation t who receives an inheritance, b_t , is

$$\omega_{t+1}^{s} = \omega^{s}(b_{t}) \equiv \begin{cases} w^{s} - (h - b_{t})(1 + i) & \text{if } b_{t} \le h \\ w^{s} + (b_{t} - h)(1 + r) & \text{if } b_{t} \ge h, \end{cases}$$
(5.16)

or equivalently,

$$\omega_{t+1}^{s} = \omega^{s}(b_{t}) \equiv \begin{cases} w^{s} - (1+i)h + (1+i)b_{t} & \text{if} \quad b_{t} \le h \\ w^{s} - (1+r)h + (1+r)b_{t} & \text{if} \quad b_{t} \ge h. \end{cases}$$
(5.17)

2.5.3 Parental Transfers and Occupational Choices

A member of generation t who receives an inheritance, b_t , acquires education if

$$\omega_{t+1}^{s} = \omega^{s}(b_{t}) > \omega_{t+1}^{u} = \omega^{u}(b_{t}).$$
(5.18)

Hence, the desirability of investment in human capital for a member of generation t depends on the individual's level of inheritance, b_t .

Since individuals are identical in their abilities and the only ex-ante source of heterogeneity among individuals is parental income, the presence of skilled and unskilled workers in society in every time period would require additional assumptions. In particular, if investment in human capital is profitable even for individuals who ought to finance the entire cost of education via borrowing, then counterfactually, all individuals would invest in human capital. Furthermore, if investment in human capital is not profitable even for those who can finance the entire cost of education from parental transfer, then counterfactually, no individual will invest in human capital.

Thus, suppose that investment in human capital is beneficial for individuals who can finance the entire cost of education without borrowing, that is,

$$w^{s} - (1+r)h > w^{u}(2+r), \tag{A3}$$

and suppose it is detrimental for individuals who must finance the entire cost of education via borrowing, that is,²¹

$$w^{s} - (1+i)h < 0.$$
 (A4)

²¹ At this stage of the analysis, it is sufficient to assume that $u^{\delta} - (1+i)h < u^{\mu}$ (2+r). However as will become apparent, the existence of multiple steady-state equilibria in the dynamics of bequests necessitates a stronger assumption, that is, $w^{\delta} - (1+i)h < 0$.



Figure 5.1 The Threshold Level of Bequest, f, above Which Investment in Human Capital Is Profitable.

As follows from (5.15) and (5.17) and as depicted in Fig. 5.1, $\omega^{s}(0) < 0 < \omega^{u}(0)$ (Assumption A4), $\omega^{s}(b_{t}) > \omega^{u}(b_{t})$ for all $b_{t} \ge h$ (Assumption A3), and there exists a level of bequest, f, such that

$$\omega_{t+1}^{s} = \omega^{s}(f) = \omega_{t+1}^{u} = \omega^{u}(f), \qquad (5.19)$$

where

$$f \equiv \frac{w^{u}(2+r) - [w^{s} - (1+i)h]}{i-r} > 0.$$
(5.20)

Hence, members of generation t choose to acquire education if they receive parental transfer, b_t , that exceeds that threshold level f. Namely,

$$\omega_{t+1}^s = \omega^s(b_t) > \omega_{t+1}^u = \omega^u(b_t) \quad \text{if and only if} \quad b_t > f.$$
(5.21)

Thus, income distribution affects occupational choices in the short run. Let the distribution of inheritance at time t be $D_t(b_t)$, that is,

$$\int_0^\infty D_t(b_t) \mathrm{d}b_t = L_t \equiv 1, \qquad (5.22)$$

where $L_t \equiv 1$ is the size of the adult generation in period *t*. It follows that the fractions of the adult generation in period t + 1 that choose to become unskilled workers, l_{t+1}^u , and the fraction that choose to become skilled workers, l_{t+1}^s , are

$$l_{t+1}^{u} = \int_{0}^{f} D_{t}(b_{t}) db_{t},$$

$$l_{t+1}^{s} = \int_{f}^{\infty} D_{t}(b_{t}) db_{t}.$$
(5.23)

Hence, the distribution of income in period *t* directly determines the segmentation of the adult generation in period t + 1 between skilled and unskilled workers, affecting the level of income per capita in this period.²²

2.6. Bequest Dynamics

The long-run effects of the distribution of income on the process of development and macroeconomic activities are determined by the interaction between occupational choices and the evolution of bequest.

As follows from the solution to the individual's optimization (5.13), members of generation t transfer a fraction $(1 - \alpha)$ of their wealth, ω_{t+1} , to their offspring. That is,

$$b_{t+1} = (1 - \alpha)\omega_{t+1}, \tag{5.24}$$

where

$$\omega_{t+1} = \begin{cases} \omega_{t+1}^{u} = \omega^{u}(b_t) & \text{if} \quad b_t \le f \\ \omega_{t+1}^{s} = \omega^{s}(b_t) & \text{if} \quad b_t > f. \end{cases}$$
(5.25)

Hence, the inheritance received by members of generation t determines their occupational choices, wealth, and their level of bequest to their offspring.

The evolution of bequest is determined by the sequence $\{b_t\}_{t=0}^{\infty}$ such that, as follows from the definition of $\omega^{"}(b_t)$ and $\omega^{s}(b_t)$ given by (5.15) and (5.17),

$$b_{t+1} = \phi(b_t) \equiv \begin{cases} (1-\alpha)[w^u(2+r) + (1+r)b_t] & \text{if } 0 \le b_t \le f \\ (1-\alpha)[w^s - (1+i)h + (1+i)b_t] & \text{if } f \le b_t \le h \\ (1-\alpha)[w^s - (1+r)h + (1+r)b_t] & \text{if } h \le b_t. \end{cases}$$
(5.26)

Hence, the dynamical system is piecewise linear. In particular, if

$$(1-\alpha)(1+r) < 1$$

 $(1-\alpha)(1+r) > 1,$ (2.A5)

then

$$\phi'(b_t) \equiv \begin{cases} (1-\alpha)(1+r) < 1 & \text{if } 0 \le b_t < f\\ (1-\alpha)(1+i) > 1 & \text{if } f < b_t < h\\ (1-\alpha)(1+r) < 1 & \text{if } h < b_t. \end{cases}$$
(5.27)

 $^{^{22}}$ In addition, it affects that fraction of the younger generation that joins the labor force as unskilled workers at time t.



Figure 5.2 Bequest Dynamics: Multiple Steady-State Equilibria in Intergenerational Transfers.

The dynamical system is characterized by multiple locally stable steady-state equilibria, as depicted in Fig. 5.2, under additional restrictions on the parameters of the model.²³

Dynasties whose initial levels of intergenerational transfers are below g converge in the long run to the lower steady-state equilibrium level, \overline{b}^{μ} , where the level of parental transfers is insufficient to permit investment in human capital by offspring. In contrast, dynasties whose initial level of intergenerational transfers is above g permit investment in human capital by offspring and the levels of intergenerational transfer among members of those dynasties converge in the long run to the higher steady-state equilibrium level, \overline{b}^{i} . In particular,

$$\lim_{t \to \infty} b_t \begin{cases} = \overline{b}^u \equiv \frac{(1-\alpha)w^u(2+r)}{1-(1-\alpha)(1+r)} & \text{if } b_t < g; \\ = \overline{b}^s \equiv \frac{(1-\alpha)[w^s - (1+r)h]}{1-(1-\alpha)(1+r)} & \text{if } b_t > g. \end{cases}$$
(5.28)

The level of bequest that determines the segmentation of society between educated and uneducated individuals, as well as the segmentation of the labor force in the long run, is

$$g \equiv \frac{(1-\alpha)[(1+i)h - w^s]}{(1-\alpha)(1+i) - 1},$$
(5.29)

as can be derived from (5.26), where g > 0 as follows from Assumptions A4 and A5.

²³ Since $\phi(0) = (1 - \alpha)w^{\mu}(2 + r) > 0$ and since $\phi'(b_t) < 1$ for $b_t \ge h$; multiplicity of locally stable steady-state equilibria is guaranteed if $\phi(f) < f$ and $\phi(h) > h$. Namely, if $w^{\mu}(2 + r)[(1 - \alpha)(1 + i) - 1] < [(1 + i)h - w^{\epsilon}][1 - (1 - \alpha)(1 + r)]$ and $(1 - \alpha)w^{\epsilon} > h$.

The threshold level of bequest, g, above which investment in human capital is beneficial for members of the dynasty in the long run is lower if (1) the cost of education, h, is lower, (2) the wage of a skilled worker, w^s , and thus the incentive to become a skilled worker is higher, (3) the interest rate for borrowers, i, is lower, or (4) the propensity of individuals to bequeath, $(1 - \alpha)$, is higher.²⁴

2.7. Distribution, Skill Composition, and Income

2.7.1 Income Distribution and the Composition of Skills

Given the distribution of inheritance at time t, $D_t(b_t)$, the critical level of bequest, g, determines the long-run composition of the labor force. As depicted in Fig. 5.3, the fraction of each generation that in the long run becomes unskilled workers, \overline{l}^{μ} , and skilled workers, \overline{l}^{i} , is

$$\lim_{t \to \infty} l_{t+1}^{u} = \int_{0}^{g} D_{t}(b_{t}) db_{t} \equiv \overline{l}^{u};$$

$$\lim_{t \to \infty} l_{t+1}^{s} = \int_{g}^{\infty} D_{t}(b_{t}) db_{t} \equiv \overline{l}^{s},$$
(5.30)



Figure 5.3 Income Distribution and Skill Composition.

²⁴ Note that $\partial g/\partial (1 - \alpha) < 0$ if Assumption A4 is satisfied.

where

$$\partial \bar{l}^s / \partial g < 0. \tag{5.31}$$

Thus, the distribution of income determines not only the composition of skills and macroeconomic activities in the short run, but via its effect on future intergenerational transfers, it also affects the composition of skills and macroeconomic activity in the long run. Moreover, inequality persists over time, and the distribution of income tends toward bimodality.

2.7.2 Persistence of Inequality

The initial distribution of income affects occupational choices and the distribution of income in the short run. However, as demonstrated in Fig. 5.4, the interaction between occupational choices and intergenerational transfers affects occupation choices and the distribution of income in the long run as well.

While heterogeneity of ability may permit upward mobility of high-ability, lowincome individuals and downward mobility of low-ability, high-income inequality,



Figure 5.4 Persistence of Inequality.

income inequality nevertheless operates toward the segmentation of society into two clubs: a club of poor, uneducated individuals and a club of rich, educated individuals. This outcome, which may raise important social and economic concerns, has significant direct effects on aggregate economic activity and economic growth.

2.7.3 Skill Composition and Income Per Capita

Income distribution affects the composition of skills in the labor force and thus has an impact on the level of income per capita in the short run as well as the long run.

In the long run, the labor force consists of $\overline{l}^{"}$ young unskilled workers, $\overline{l}^{"}$ adult unskilled workers, and \overline{l}° skilled workers. The steady-state level of income in the economy is given therefore by (1) the wage income of unskilled individuals in the first period of life, (2) wage and capital income of unskilled individuals in the second period of life, and (3) wage and capital income of skilled individuals in the second period of life.

The steady-state level of income of a skilled individual in the second period of life, I_2^s , consists of wage income, w^s , and capital income, $(\overline{b}^s - h)r$, reflecting the return on saving in the first period of life. Hence,²⁵

$$I_{2}^{s} = w^{s} + (\overline{b}^{s} - h)r.$$
(5.32)

The steady-state level of income of an unskilled individual in the second period of life, I_2^u , consists of a wage income, w^u , and capital income, $(\overline{b}^u + w^u)r$, reflecting the return on the saving, that is,

$$I_2^u = w^u + (\overline{b}^u + w^u)r.$$
(5.33)

Finally, the income of an unskilled individual in the first period of life, I_1^{μ} , consists only of the wage income, w^{μ} :²⁶

$$I_1^u = w^u. (5.34)$$

The aggregate level of income in the domestic economy in the steady state, \overline{Y} , is therefore

$$\overline{Y} = I_2^s \overline{l}^s + I_2^u \overline{l}^u + I_1^u \overline{l}^u.$$
(5.35)

²⁵ In the steady state the level of intergenerational transfer among skilled dynasties exceeds the cost of education, *h*. Hence skilled individuals have a positive level of saving on which they receive the interest rate for lenders, *r*.

²⁶ Note that individuals in the first period of their lives do not have capital income. They receive an inheritance that constitutes their wealth in the first period, but this is not earned income and is thus not relevant for the calculation of an economy's aggregate income.
Using the fact that $\overline{l}' + \overline{l}'' = 1$, it follows that

$$\overline{Y} = [w^{u}(2+r) + r\overline{b}^{u}](1-\overline{l}^{s}) + [w^{s} + r(\overline{b}^{s} - h)]\overline{l}^{s}$$

= $w^{u}(2+r) + r\overline{b}^{u} + [(w^{s} - rh) - w^{u}(2+r) + r(\overline{b}^{s} - \overline{b}^{u})]\overline{l}^{s}.$ (5.36)

The steady-state level of income per capita is $\overline{\gamma} = \overline{Y}/2$, noting that the size of the population is equal to 2 in every time period.

An increase in the fraction of skilled workers increases therefore income per capita. Namely,

$$\frac{\partial \overline{\gamma}}{\partial \overline{l}^s} = \left[(w^s - rh) - w^u (2 + r) + r(\overline{b}^s - \overline{b}^u) \right] / 2 > 0, \tag{5.37}$$

noting that $\overline{b}^s > \overline{b}^u$ (as follows from (5.28)) and $(w^s - rh) - w^u(2+r) > 0$ (as a result of Assumption A3).

Moreover, a society characterized by a lower threshold of bequest above which individuals invest in human capital (i.e., a lower level of g) has a higher level of income per capita in the steady state. Namely, as follows from (5.31) and (5.37),

$$\frac{\partial \overline{\gamma}}{\partial g} = \frac{\partial \overline{\gamma}}{\partial \overline{l}^s} \frac{\partial \overline{l}^s}{\partial g} < 0.$$
(5.38)

Thus, the initial distribution of income affects income per capita in the short run as well as in the long run. If the distribution of income across members of society is characterized by a lower fraction of individuals that cannot invest in human capital, then income per capita increases in the long run.

For a given distribution of income, income per capita in the long run is higher the lower is the threshold level of bequest, g, above which investment in human capital is beneficial for members of the dynasty in the long run. Namely, income per capita in the long run is higher if (1) the cost of education, h, is lower, (2) the wage of a skilled worker, w^s , and thus the incentive to become a skilled worker is higher, (3) the interest rate for borrowers, i, is lower, or (4) the propensity of individuals to bequeath, $(1 - \alpha)$, is higher.

2.8. Inequality and Economic Development

Income distribution affects the growth process and the level of income per capita in the long run. Inequality in the distribution of income may have an adverse effect on the growth process in a nonpoor economy, whereas inequality may have a beneficial effect on the growth process in poor economies.

Consider an economy in period t where income per capita is sufficiently large relative to the cost of education. In particular, suppose that the average level of bequest in period t, \hat{b}_t , exceeds the critical level, g, above which investment in human capital is beneficial for members of the dynasty in the long run, that is,

$$\hat{b}_t \equiv b_t^s l_t^s + b_t^u l_t^u > g.$$
 (5.39)

As is illustrated in Fig. 5.5, for a given average level of bequest in period t, \hat{b}_t , an increase in inequality (for a wide class of measures of inequality) will be associated with an increase in the number of individuals below the critical level g. Thus, inequality in nonpoor economies is likely to reduce investment in human capital and may thus decrease the long-run level of income per capita.

Consider an economy in period t where income per capita is sufficiently low relative to the cost of education. In particular, suppose that the average level of bequest in period t, \hat{b}_t , is lower than the critical level, g, above which investment in human capital is beneficial for members of the dynasty in the long run, that is,

$$\hat{b}_t \equiv b_t^s l_t^s + b_t^u l_t^u < g. \tag{5.40}$$



Figure 5.5 The Adverse Effect of Inequality on the Process of Development: A Nonpoor Economy.

As is illustrated in Fig. 5.6, for a given average level of bequest in period t, \hat{b}_t , an increase in inequality (for a wide class of measures of inequality) will be associated with an increase in the number of individuals above the critical level g. Hence, inequality in poor economies may induce investment in human capital and may thus increase the long-run level of income per capita.

Thus the model generates the following testable predictions. Among economies that are identical in their structural characteristics (i.e., production technologies, preferences, the cost of education, and the degree of credit market imperfections) and therefore in the threshold level of bequest above which investment in human capital is beneficial,

- 1. Higher inequality in the distribution of income will be associated with higher income per capita across poor economies.
- **2.** Higher inequality in the distribution of income will be associated with lower income per capita across nonpoor economies.



Figure 5.6 The Positive Effect of Inequality on the Process of Development: A Poor Economy.

2.9. Robustness

2.9.1 Labor-Augmenting Technological Progress

Suppose that the economy experiences an exogenous *labor-augmenting* technological progress that transforms the labor force and increases the productivity of workers in both the skilled labor-intensive and the unskilled labor-intensive sector.

The output produced in the skilled labor-intensive sector in period t is

$$Y_{t}^{s} = F(K_{t}, A_{t}L_{t}^{s}) \equiv A_{t}L_{t}^{s}f(k_{t}); \quad k_{t} \equiv K_{t}/A_{t}L_{t}^{s},$$
(5.41)

where A_t is the level of technology in period t, and $A_t L_t^s$ is the number of efficiency units of skilled labor employed in production at time t. Similarly, the output produced at the unskilled labor-intensive sector in period t is

$$Y_t^u = A_t a L_t^u. ag{5.42}$$

Technology evolves over time at a constant exogenous rate λ .

$$A_{t+1} = (1+\lambda)A_t, (5.43)$$

where $\lambda > 0$ is the rate of labor-augmenting technological progress, and the level of technology at time 0, A_0 , is exogenously given.

As follows from the producer's profit maximization and the presence of perfect international capital mobility,

$$w_t^s = A_t[f(k) - f'(k)k] \equiv A_t w^s$$

$$w_t^u = A_t a \equiv A_t w^u$$

$$r_t = r.$$
(5.44)

Suppose further that the acquisition of education is associated with a fixed cost, c_t^h , that reflects the indivisibility of human capital formation in general and of academic degrees in particular. The fixed cost of education can be viewed as a weighted average of the payments to teachers, administrators, and maintenance workers in the school system (i.e., a weighted average of the wages skilled and unskilled workers):

$$c_t^h = \theta A_t w^s + (1 - \theta) A_t w^u \equiv A_t h \quad \text{for some } \theta \in [0, 1].$$
(5.45)

Hence, the second period wealth, ω_{t+1}^{u} , of an unskilled member of generation *t* who receive an inheritance, b_t , is

$$\omega_{t+1}^{u} = A_{t} \omega^{u} (2 + r + \lambda) + (1 + r) b_{t} \equiv \omega^{u} (b_{t}, A_{t}), \qquad (5.46)$$

whereas the second-period wealth, ω_{t+1}^{s} , of a skilled member of generation t who received an inheritance, b_t , is

$$\omega_{t+1}^{s} = \omega^{s}(b_{t}, A_{t}) \equiv \begin{cases} A_{t+1}w^{s} - (A_{t}h - b_{t})(1+i) & \text{if} \quad b_{t} \le A_{t}h \\ A_{t+1}w^{s} + (b_{t} - A_{t}h)(1+r) & \text{if} \quad b_{t} \ge A_{t}h, \end{cases}$$
(5.47)

or equivalently,

$$\omega_{t+1}^{s} = \omega^{s}(b_{t}, A_{t}) \equiv \begin{cases} A_{t}[w^{s}(1+\lambda) - (1+i)h] + (1+i)b_{t} & \text{if} \quad b_{t} \le A_{t}h \\ A_{t}[w^{s}(1+\lambda) - (1+r)h] + (1+r)b_{t} & \text{if} \quad b_{t} \ge A_{t}h. \end{cases}$$
(5.48)

Modifying Assumptions A3 and A4 and assuming that

$$w^{s}(1+\lambda) - (1+i)h < 0$$

$$w^{u}(2+r) > [w^{s} - (1+i)h] + \lambda(w^{s} - w^{u}),$$
(5.49)

it follows from (5.46) and (5.48) that there exists a level of bequest, f_t , such that

$$\omega_{t+1}^{s} = \omega^{s}(f_{t}) = \omega_{t+1}^{u} = \omega^{u}(f_{t}), \qquad (5.50)$$

where

$$f_t = \frac{A_t \{ w^u (2+r) - [w^s - (1+i)h] - \lambda (w^s - w^u) \}}{(i-r)} = f(A_t).$$
(5.51)

Moreover,

$$\frac{f_t}{A_t} = \frac{w^u (2+r) - [w^s - (1+i)h] - \lambda(w^s - w^u)}{(i-r)} \equiv \hat{f} > 0.$$
(5.52)

The evolution of bequest is given therefore by

$$b_{t+1} = \begin{cases} (1-\alpha)\{A_t w^u (2+r+\lambda) + (1+r)b_t\} & b_t \in [0, f_t] \\ (1-\alpha)\{A_t [w^s (1+\lambda) - (1+i)h] + (1+i)b_t\} & b_t \in [f_t, A_t h] \\ (1-\alpha)\{A_t [w^s (1+\lambda) - (1+r)h] + (1+r)b_t\} & b_t \in [A_t h, \infty). \end{cases}$$
(5.53)

Let $\hat{b}_{t+1} \equiv b_{t+1}/A_{t+1}$, then

$$\hat{b}_{t+1} \equiv \psi(b_t) = \begin{cases} \left[\frac{1-\alpha}{1+\lambda}\right] \{w^u(2+r+\lambda) + (1+r)\hat{b}_t\} & \hat{b}_t \in [0,\hat{f}] \\ \left[\frac{1-\alpha}{1+\lambda}\right] \{[w^s(1+\lambda) - (1+i)h] + (1+i)\hat{b}_t\} & \hat{b}_t \in [\hat{f},h] \\ \left[\frac{1-\alpha}{1+\lambda}\right] \{[w^s(1+\lambda) - (1+r)h] + (1+r)\hat{b}_t\} & \hat{b}_t \in [h,\infty]. \end{cases}$$
(5.54)

Hence as long as

$$(1 - \alpha)(1 + r) < (1 + \lambda); (1 - \alpha)(1 + i) > (1 + \lambda),$$
 (5.55)

 $\psi(\hat{f}) < \hat{f}$ and $\psi(h) > h$, the dynamical system is characterized by multiple steady states, where the unstable equilibrium is

$$\hat{g} = \frac{(1-\alpha)[(1+i)h - w^{s}(1+\lambda)]}{[(1-\alpha)(1+i) - (1+\lambda)]} > 0.$$
(5.56)

Thus, the qualitative analysis is unaffected by labor-augmenting technological progress. Moreover, if technological progress is a function of the skilled composition of the labor force, inequality would have an effect on the growth rate of the economy in the steady state.

2.9.2 Interactions across Dynasties

The basic Galor–Zeira model establishes the potential adverse effect of inequality on economic growth in an economy in which wages, for simplicity, are unaffected by the composition of the labor force. The structure of the basic model is designed to assure that factor prices are constant over time, permitting a simple characterization of the dynamics of income distribution and its implication for aggregate economic activities and economic development. However, as established in the second part of Galor and Zeira (1993), the main hypothesis is robust to the endogenization of wages and thus to the incorporation of interdependence in investment decisions across dynasties.

The robustness of the main insights of the theory in an environment where factor prices are endogenously determined and investment decisions across dynasties are interdependent is established in a large number of studies (e.g., Banerjee and Newman (1993); Galor and Moav (2004)).

2.9.3 Random Shocks

The persistent effect of inequality is immune to shocks to the outcome of investment in human capital, as long as wages are endogenous. Notably, Banerjee and Newman (1993) examine the effect of inequality on a different type of occupational choice (i.e., the choice between becoming an entrepreneur or a worker, rather than the choice between becoming either a skilled or an unskilled worker). They demonstrate that if credit markets are imperfect and fixed costs are associated with entrepreneurial activities, inequality may result in an underinvestment in entrepreneurial activity and may therefore be harmful for economic development. Their study establishes that as long as wages are endogenous, the main hypothesis of the credit market imperfection approach as a whole is robust to the introduction of random shocks to the outcome of investment (in human capital or entrepreneurial activities).

2.9.4 Concave Production of Human Capital and Alternative Utility Functions

The qualitative impact of income distribution is unaffected by the incorporation of a divisible, concave production function of human capital, as long as the saving rate is an increasing function of income (Moav (2002); Galor and Moav (2004)).

Moreover, as shown in earlier versions of the Galor–Zeira model, the results are robust to alternative forms of intergenerational altruism in which the utility function is defined over the utility of the offspring rather than the level of intergenerational transfer to the offspring.

3. A UNIFIED THEORY OF INEQUALITY AND GROWTH

The modern perspective on the relationship between inequality and economic development has been initially segmented. It lacked a unified hypothesis regarding the relationship between inequality and the growth process, particularly in light of the (seemingly) contrasting predictions generated by the classical approach and the modern approach. The development of a unified theory of inequality and growth that captures that changing role of inequality in the process of development has provided a needed intertemporal reconciliation between the Classical viewpoint and the modern perspective, while permitting the dominating theories within the modern perspective to be placed within a broader framework.

The theory advanced by Galor and Moav (2004) suggests that the replacement of physical capital accumulation by human capital accumulation as the prime engine of economic growth has changed the qualitative impact of inequality on the process of development. In early stages of industrialization, as physical capital accumulation was a prime source of economic growth, inequality enhanced the process of development by channeling resources toward individuals whose marginal propensity to save is higher. However, in later stages of development, as physical capital accumulated and the demand for human capital increased, human capital has become the main engine of economic growth.²⁷ A more equal distribution of income, in the presence of credit constraints, has stimulated investment in human capital and promoted economic growth.

The central hypothesis of this unified approach stems from the recognition that human capital accumulation and physical capital accumulation are fundamentally asymmetric. In contrast to physical capital, human capital is inherently embodied in humans, and the existence of physiological constraints subjects its accumulation at the individual level to diminishing returns. The aggregate stock of human capital would therefore be larger if its accumulation would be widely spread among individuals in society, whereas

²⁷ The rise in the demand for skilled labor may be viewed as an outcome of: (i) capital–skill complementarity, (ii) a skillbiased technological change, or (iii) an unbiased technological acceleration, reflecting the comparative advantage of educated individuals in coping with a changing technological environment (Nelson and Phelps (1966); Schultz (1975); Foster and Rosenzweig (1996)).

the aggregate productivity of the stock of physical capital is largely independent of the distribution of its ownership in society. This asymmetry between the accumulation of human and physical capital suggests therefore that as long as credit constraints are largely binding, a more equal distribution of income is conducive for human capital accumulation, whereas, provided that the marginal propensity to save increases with income, inequality is conducive for physical capital accumulation.

Therefore, the theory provides a reconciliation between conflicting viewpoints about the effect of inequality on economic growth. It suggests that the Classical viewpoint, regarding the positive effect of inequality on the process of development, reflects the state of the world in early stages of industrialization, when physical capital accumulation is the prime engine of economic growth. In contrast, the central hypothesis of the credit market imperfection approach, regarding the negative effect of inequality on economic growth, reflects the later stages of development when human capital accumulation is the prime engine of economic growth and credit constraints are largely binding.

In early stages of industrialization, physical capital is scarce, the rate of return to human capital is lower than the rate of return to physical capital, and the process of development is fueled by capital accumulation. The positive effect of inequality on aggregate saving dominates therefore the negative effect on investment in human capital and inequality raises aggregate savings and capital accumulation and enhances the process of development. In later stages of development, as physical capital accumulates, the complementarity between capital and skills increases the rate of return to human capital. Investment in human capital accumulation increases and the accumulation of human capital as well as physical capital fuels the process of development. Since human capital is embodied in individuals and individuals' investment in human capital is subjected to diminishing marginal returns, the aggregate return to investment in human capital is maximized if investment in human capital is widely spread among individuals in society. Equality alleviates the adverse effect of credit constraints and has therefore a positive effect on the aggregate level of human capital and economic growth. Moreover, the differences in the marginal propensities to save across individuals narrow as wages increase, and the negative effect of equality on aggregate saving subsides. Therefore, in later stages of development, as long as credit constraints are sufficiently binding, the positive effect of inequality on aggregate saving is dominated by the negative effect on investment in human capital, and equality stimulates economic growth. However, as wages further increase, credit constraints become less binding, differences in the marginal propensity to save further decline, and the aggregate effect of income distribution on the growth process becomes less significant.²⁸

²⁸ Inequality may widen once again due to skilled- or ability-biased technological change induced by human capital accumulation. This line of research was explored theoretically by Galor and Tsiddon (1997b), Acemoglu (1998), Caselli (1999), and Galor and Moav (2000), among others. It is consistent with recent evidence provided by Berman, Bound, and Machin (1998), Goldin and Katz (1998), and Autor, Katz, and Kearney (2008), among others.

Although the replacement of physical capital accumulation by human capital accumulation as a prime engine of economic growth in currently developed economies is instrumental for the understanding of the role of inequality in the process of development of these advanced economies, this unified theory generates an insight about the role of inequality in the growth process of less developed economies as well. The presence of international capital inflows has diminished the role of inequality in stimulating physical capital accumulation in less-developed economies. Moreover, the adoption of skill-biased technologies by some of these economies has increased the return to human capital and has strengthened the positive effect of a more equal distribution of income on human capital formation and economic growth.

The unified theory of inequality and growth may provide greatly needed theoretical guidance for empirical research in this field. In contrast to the credit market imperfection approach, which suggests that the effect of inequality depends on the country's level of income (i.e., inequality is beneficial for poor economies and harmful for others), the unified theory of inequality and growth suggests that the effect of inequality on growth depends on the relative return to physical and human capital. In economies in which the return to human capital is relatively lower, inequality is beneficial for economic growth, whereas in economies in which the return to human capital is relatively higher and credit constraints are largely binding, equality is beneficial for the development process.

4. NONFINANCIAL HURDLES FOR HUMAN CAPITAL ACCUMULATION

While the process of industrialization raised the importance of human capital in the production process, reflecting its complementarity with physical capital and technology, human capital accumulation has not benefited all sectors of the economy. Inequality in the ownership of factors of production has generated an incentive for some betterendowed agents to block the implementation of institutional changes and policies that promote human capital formation, resulting in a suboptimal level of investment in human capital from a growth perspective. In particular, variation in the distribution of ownership over land and other natural resources across countries has contributed to the observed disparity in human capital formation and to the divergent development patterns across the globe.

4.1. Concentration of Landownership

The transition from an agricultural to an industrial economy changed the nature of the main economic conflict in society. Unlike the agrarian economy, which was characterized by a conflict of interests between the landed aristocracy and the masses, the process of industrialization has brought about an additional conflict between the entrenched landed elite and the emerging capitalist elite. In light of a lower degree of complementarity between human capital and the agricultural sector, education has increased the productivity

of labor in industrial production more than in agricultural and primary good production, inducing rural-to-urban migration and thus a decline in the return to landowners. Thus, while industrialists have had a direct economic incentive to support education policies that would foster human capital formation (Galor and Moav (2006)), landowners, whose interests lay in the reduction of the mobility of their labor force, have favored policies that deprived the masses of education (Galor, Moav, and Vollrath (2009)).²⁹

The adverse effect of the implementation of public education on landowners' income from agricultural production has been magnified by the concentration of land ownership. Thus, as long as landowners affected the political process and thereby the implementation of growth-enhancing education policies, inequality in the distribution of land ownership has been a hurdle for human capital accumulation, slowing the process of industrialization and the transition to modern growth.³⁰

Economies in which land and other natural resources have been more equally distributed have implemented earlier public education campaigns and have benefited from the emergence of a skill-intensive industrial sector and a rapid process of development. In contrast, among economies marked by a more unequal distribution of ownership over land and other natural resources, resource abundance that was a source of richness in the early stages of development has led in later stages to underinvestment in human capital, an unskilled labor-intensive industrial sector, and a slower growth process. Thus, variation in the distribution of ownership over land and other natural resources across countries has contributed to disparity in human capital formation and the industrial composition of the economy, and thus to divergent development patterns across the globe.³¹

An alternative mechanism that underlines the adverse effect of inequality on human capital formation and economic development has been advanced by Engerman and

²⁹ In accordance with the unified approach for the study of inequality and economic development, this line of research suggests that capital accumulation in the process of industrialization gradually intensified the relative scarcity of skilled labor and generated an incentive for human capital accumulation. Investment in human capital, however, has been suboptimal due to credit market imperfections, and public investment in education has been therefore growth enhancing. Due to the complementarity between physical and human capital in production, the capitalists were among the prime beneficiaries of the accumulation of human capital by the masses. They therefore had the incentive to support the provision of public education that improved their economic well-being and contributed significantly to the demise of the capitalists-workers class structure and to changes in the nature of inequality in society that were conducive to economic development. Mutually beneficial reforms are also considered by Lizzeri and Persico (2004) and Doepke and Zilibotti (2005).

³⁰ Interestingly, during the 19th century, the emergence of a *broad-based* demand for human capital-intensive services by the landowners in land-rich economies in Latin America (e.g., Argentina) triggered the establishment of an extensive public education system prior to the onset of significant manufacturing activities (Galiani, Heymann, Dabús, and Tohmé (2008)). Thus, lack of concentration of land ownership (that was conducive for a *broad-based* demand for human capital-intensive services by the landowners) had a positive effect on human capital formation even prior to industrialization.

³¹ Rajan (2009) reinforces this thesis, suggesting that rent preservation and its interaction with inequality in ownership over factor endowment is a recipe for paralysis and poverty.

Sokoloff (2000) and Acemoglu, Johnson, and Robinson (2005). They argued and provided evidence that geographical conditions that led to income inequality brought about oppressive institutions (e.g., restricted access to the democratic process and to education) designed to maintain the political power of the elite and to preserve the existing inequality between the elite and the masses.³² Thus, Engerman and Sokoloff (2000) underlined the role of the sustained conflict between the elite and the masses in the delay in the implementation of growth-enhancing educational policies and thus in the adverse effect of inequality on the process of development, suggesting the perpetual desirability of extractive institutions for the ruling elite in the absence of changes in the political structure.³³

4.2. Social-Political Transitions

Inequality and its association with sociopolitical instability have been identified as an additional adverse force in the process of development. In particular, the effect of inequality on social conflict and on political and educational reforms was examined by Alesina and Perotti (1996), Acemoglu and Robinson (2000), Bourguignon and Verdier (2000), and Gradstein (2007).³⁴ These studies suggest that reforms and redistribution from the elite to the masses diminish the tendency for sociopolitical instability and may therefore stimulate investment and economic growth. In particular, Acemoglu and Robinson (2000) argue that the extension of the franchise during the nineteenth century can be viewed as a commitment device to ensure future income redistribution from the elite to the masses.³⁵

In contrast, Galor and Moav (2006) have argued that the transformation in class structure and inequality can be viewed as a byproduct of a productive cooperation between capitalists and workers, rather than an outcome of a divisive class struggle. In accordance with the unified approach for the study of inequality and economic development, this line of research suggests that capital accumulation in the process of

³² Acemoglu, Johnson, and Robinson (2005) maintain that economic performance across countries have a colonial origin, reflecting the institutional quality that were introduced by European colonialism across the globe. They have argued that historical reversals in the economic performance of societies have a colonial legacy that reflects the imposition of extractive institutions by European colonizers in affluent regions that benefited from favorable geographical conditions in the pre-colonial era and the implementation of growth-enhancing institutions in poorer regions.

³³ In contrast, Galor, Moav, and Vollrath (2009) demonstrate that even if the political structure in the economy remains unchanged, economic development and a gradual diversification of the assets held by the landed aristocracy may ultimately trigger the implementation of growth-promoting institutions once the stake of the landed aristocracy in the efficient operation of the industrial sector dominates their overall economic interest.

³⁴ See also Bowles and Gintis (1975).

³⁵ Mejía and Posada (2007) identify conditions under which a social conflict lead to the transition to democracy and those under which purely economic forces lead to the transition, underlying the relative role of inequality, the importance of a human capital externalities in production, and the feasibility of redistribution by the masses.

industrialization gradually intensified the relative scarcity of skilled labor and generated an incentive for human capital accumulation. However, investment in human capital has been suboptimal due to credit market imperfections, and public investment in education has been therefore growth enhancing, as per Galor and Zeira (1993). Due to the complementarity between physical and human capital in production, the capitalists were among the prime beneficiaries of the accumulation of human capital by the masses. They therefore had the incentive to support the provision of public education that improved their economic well-being and contributed significantly to the demise of the capitalists–workers class structure and to changes in the nature of inequality in society that were conducive to economic development.

4.3. Gender Inequality

The decline in gender inequality, which was brought about by the rise in the demand for human capital in the process of development, reinforced the positive association between a more-egalitarian distribution of income and economic growth. The decline in gender inequality contributed to the onset of the demographic transition, as well as to the rise in female labor force participation, fostering the growth process as a whole. The decline in the gender wage gap has affected household fertility decisions, female labor force participation, and thus the growth process.

As suggested by Galor and Weil (1996, 1999), technological progress and capital accumulation complemented mentally intensive tasks and substituted for physically intensive tasks in industrial production. In light of the comparative physiological advantage of men in physically intensive tasks and of women in mentally intensive tasks, the demand for women's labor input gradually increased, inducing a decline in fertility rates, a significant increase in labor force participation, and a transition from stagnation to growth.³⁶

5. EVIDENCE

5.1. Inequality, Human Capital Formation, and Economic Growth

Several attempts have been made to examine the theoretical predictions of the credit market imperfections approach and the political economy approach about the effect of inequality and heterogeneity on economic growth. Consistent with the hypothesis advanced by the theories, early cross-country analyses by Alesina and Rodrik (1994),

³⁶ The decline in the overall level of inequality that was associated with the emergence of human capital has been linked theoretically, empirically and quantitatively to the reduction in fertility and therefore in light of the quantity-quality trade-off (e.g., Rosenzweig and Wolpin (1980); Hanushek (1992)) to lower levels of investment in human capital and income. See, Galor and Zang (1997), Dahan and Tsiddon (1998), Kremer and Chen (2002), de la Croix and Doepke (2003), and Hassler, Jose V. Rodriguez Mora, and Zeira (2007).

Persson and Tabellini (1994), and Perotti (1996) have established a negative association between the level of inequality and economic growth.

Importantly, Perotti (1996) conducted an examination of the various channels through which inequality may affect economic growth, as proposed by the modern theoretical perspective. His study provides support for the validity of the human capital formation channel, showing that inequality is indeed associated with lower level of human capital formation, and lower human capital formation is associated with lower levels of economic growth.³⁷ Further support for the main predictions of the education channel, advanced in the context of the credit market imperfection approach, has been generated by Deninger and Squire (1998). Utilizing the distribution of land as a proxy for the distribution of assets, they find that initial inequality has a significant adverse effect on education and economic growth. Moreover, consistent with the predictions of the credit market imperfections approach that credit constraints ought to have a larger effect on the investment decisions of individuals with lower income, they find that initial inequality primarily hurts the poor.³⁸

In contrast to the human capital channel, Perotti's examination of the political economy channel was not favorable to the theories advanced by Alesina and Rodrik (1994) and Persson and Tabellini (1994). His findings refute this early hypothesis of the political economy approach, demonstrating that in contrast to their proposed channel, inequality is in fact associated with lower levels of taxation, while lower levels of taxation, contrary to the theories, are associated with lower levels of economic growth.

Later studies have deviated from the desirable examination of the channels through which inequality may affect growth and restricted their attention to the reduced form relationship between inequality and growth. Notably, Forbes (2000) and Barro (2000) examined the effect of inequality on economic growth in a panel of countries. They find a positive and zero effect, respectively, of an increase in inequality on economic growth.

However, these findings ought to be interpreted very cautiously. They appear to have no bearing on the validity of the theories and are not very informative about the overall effect of inequality. First, these studies examine the effect of inequality beyond its effects through education, fertility, and investment. For instance, Barro (2000) has found that, once controls for education, fertility, and investment are introduced, there is no relationship between inequality and economic growth in the entire sample. Therefore, his findings suggest that inequality does not have a direct effect on growth beyond its effects through education, fertility, and investment (i.e., the dominating channels through which

³⁸ The adverse effects of financial constraints on economic development are well established (Levine (2005)).

³⁷ In line with related theoretical arguments that human capital formation and fertility are negatively related and thus inequality would be expected to have contrasting effects on these two variables (e.g., Galor and Zhang (1997)), Perotti (1996) suggests that the human capital channel is reinforced by the introduction of fertility. Inequality is associated with higher fertility rates and a lower level of investment in human capital, which are in turn associated with lower economic growth.

inequality operates), implying perhaps that the dominating channels through which inequality operates are those proposed in the literature. In particular, if the control for fertility is dropped in Barro (2000), the effect of inequality on growth is significantly negative, as predicted by the theory. Moreover, these studies examine the effect of inequality in the short run (i.e., the effect of inequality on the average growth rate in the subsequent 5–10 years), while as suggested by the theories, inequality is likely to have mostly longer-run effects (e.g., via the formation of human capital).

Moreover, even within the context of the limited scope of the studies of Forbes (2000) and Barro (2000), their econometric methodology and their findings have been challenged. Banerjee and Duflo (2003) argued that the linear regression structure imposed in these and in earlier empirical studies is inconsistent with the predictions of the theories, and the qualitative findings may be an artifact of the imposed linearity. They find that changes in inequality (in any direction) are associated with lower growth rates. Moreover, in line with the adverse long-run impact of inequality proposed by the theories, they find a negative relationship between growth rates and lagged inequality.

Recently, Easterly (2007) has reaffirmed the hypothesis advanced by the modern theories that inequality has an adverse effect on human capital formation and economic development. Using agricultural endowments as an instrument for inequality in order to overcome concerns about measurement errors and the endogeneity of inequality, his cross-country analysis suggests that inequality has been a barrier to schooling and economic prosperity.

5.2. Industrialization and Human Capital Formation

The process of industrialization was characterized by a gradual increase in the relative importance of human capital in the production process. As underlined by Unified Growth Theory (Galor (2011)), this important development was triggered by acceleration in the rate of technological progress and the role of human capital in adapting to a rapidly changing technological environment.

In the first phase of the Industrial Revolution, human capital played a limited role in the production process. Education was motivated by a variety of factors, including religion, enlightenment, social control, moral conformity, sociopolitical stability (i.e., the shadow of rebellion of the masses), social and national cohesion, and military efficiency. The extent to which public education was provided was not correlated with industrial development, and it differed across countries due to political, cultural, social, historical, and institutional factors. Human capital had a limited role in the production process; education instead served religious, social, and national goals. As argued by Landes (1969), although certain workers—supervisory and office personnel in particular—were required to be able to read and do the elementary arithmetical operations in order to perform their duties, a large fraction of the work of industry was performed by illiterates, especially in the early days of the Industrial Revolution.

In contrast, during the second phase of the Industrial Revolution, the demand for skilled labor in the growing industrial sector markedly increased. Human capital formation was designed primarily to satisfy the increasing skill requirements in the process of industrialization, and industrialists became involved in shaping the educational system. Moreover, the reversal of the Malthusian relationship between income and population growth during the demographic transition corresponded to a further increase in the level of resources invested in each child.

Evidence relating to the evolution of the return on human capital during this period is scarce and controversial.³⁹ One can mistakenly argue that the lack of clear evidence about the increase in the return on human capital during this period indicates the absence of a significant increase in the demand for human capital. However, this partial equilibrium argument is flawed. The return on human capital is affected by the demand and supply of human capital. Technological progress in the second phase of the Industrial Revolution brought about an increase in demand for human capital, and indeed, in the absence of a supply response, one would have expected an increase in the return on human capital. However, the significant increase in schooling that took place during the nineteenth century (in particular, the introduction of public education), which lowered the cost of education, generated a significant increase in the supply of educated workers. Some of this supply response was a direct reaction to the increase in demand for human capital and thus may have only operated to partially offset the increase in the return on human capital. However, the removal of the adverse effect of credit constraints on the acquisition of human capital (as reflected by the introduction of public education) generated an additional force that increased the supply of educated labor and operated to reduce the return on human capital.

Reassuringly, Becker, Hornung, and Woessmannand (2011), in the first rigorous attempt to examine empirically the role of education in the process of industrialization, provide evidence about the important role that education played in the process of industrialization. Using variation in pre-industrial education across counties in Prussia in 1816 as instrument for later education levels in these counties, they find that education contributed significantly to industrialization in two industrial phases, in 1849 and in 1882. Moreover, as implied by Unified Growth Theory, they find that the role of education has been intensified in the second phase of Prussia's industrialization.⁴⁰

³⁹ Not surprisingly, existing evidence focusing on the return on *old* skills (e.g., construction) does not find that the return on such skills increased in England over the course of the nineteenth century (Clark (2005)).

⁴⁰ The rise in the demand for education in the process of industrialization (prior to the demographic transition) is underlined in the theories of Galor and Weil (2000) and Galor and Moav (2002). Moreover, the contribution of education to industrialization is consistent with both theories, although Galor and Moav (2002) attribute a more significant role for education in the first phase of industrialization.

5.2.1 Industrial Demand for Education

Education reforms in developed countries in the eighteenth and nineteenth centuries are indicative of the significance of industrial development in the formation of human capital during the second half of the nineteenth century. In particular, differences in the timing of the establishment of a national system of public education between England and continental Europe are instrumental in isolating the role that industrial forces played in human capital formation.

England: During the first phase of the Industrial Revolution (1760–1830), capital accumulation increased significantly without a corresponding increase in the supply of skilled labor. The investment–output ratio increased from 6% in 1760 to 12% in 1831 (Crafts (1985), p. 73), whereas literacy rates remained largely unchanged, and the state devoted virtually no resources to raising the level of literacy among the masses (Mokyr (2001)). Literacy was largely a cultural skill or a hierarchical symbol and was of limited use in the production process. For instance, in 1841, only 5% of male workers and only 2% of female workers were employed in occupations in which literacy was strictly required (Mitch (1992)). Furthermore, an illiterate labor force could operate the existing technology, and economic growth was not impeded by educational retardation.⁴¹ Workers developed skills primarily through on-the-job training, and child labor was highly valuable.

The development of a national public system of education in England lagged behind other Western European countries by nearly half a century (Sanderson (1995)).⁴² England's early industrialization occurred without direct state intervention in the development of the minimal skills required for industrial production (Green (1990)). England initiated a sequence of reforms in its educational system after the 1830s, and literacy rates gradually increased. The process was initially motivated by nonindustrial reasons, such as religion, social control, moral conformity, enlightenment, and military efficiency, as was the case in other European countries (e.g., Germany, France, Holland, and Switzerland) that had supported public education much earlier. However, in light of the modest demand for skills and literacy by the capitalists, the level of governmental support was rather small.⁴³

As the Industrial Revolution progressed to its second phase, the demand for skilled labor in the growing industrial sector markedly increased, and the proportion of

⁴¹ Some have argued that the low skill requirements even declined over this period. For instance, Sanderson (1995) suggests that the emerging economy created a whole range of new occupations that required even less literacy and education than the old ones.

⁴² For instance, in his parliamentary speech in defense of his 1837 education bill, Whig politician Henry Brougham reflected on this gap: "It cannot be doubted that some legislative effort must at length be made to remove from this country the opprobrium of having done less for education of the people than any of the more civilized nations on earth" (Green (1990), pp. 10–11).

⁴³ Even in 1869, the English government funded only one-third of school expenditure (Green (1990)).

children aged 5–14 in primary schools rose from 11% in 1855 to 25% in 1870 (Flora, Kraus, and Pfenning (1983)). Literacy became an increasingly desirable characteristic for employment, as indicated by job advertisements of the period (Mitch (1993)). In light of industrial competition from other countries, capitalists started to recognize the importance of technical education for the provision of skilled workers. As noted by Sanderson (1995, pp. 10–13), "Reading... enabled the efficient functioning of an urban industrial society laced with letter writing, drawing up wills, apprenticeship indentures, passing bills of exchange, and notice and advertisement reading." Moreover, manufacturers argued that "universal education is required in order to select, from the mass of the workers, those who respond well to schooling and would make a good foreman on the shop floor." (Simon (1987), p. 104).

As it became apparent that skills were necessary for the creation of an industrial society, replacing previous concerns that the acquisition of literacy would make the working classes receptive to radical and subversive ideas, capitalists lobbied for the provision of public education.⁴⁴ The pure laissez-faire policy failed to develop a proper educational system, and capitalists demanded government intervention in the provision of education. As Leeds iron-master and advocate of technical education, James Kitson explained to the Select Committee on Scientific Instruction (1867–1868): "[T]he question is so extensive that individual manufacturers are not able to grapple with it, and if they went to immense trouble to establish schools they would be doing it in order that others may reap the benefit" (Green (1990), p. 295).⁴⁵ An additional turning point in the attitude of English capitalists toward public education was the Paris Exhibition of 1867, where the limitations of English scientific and technical education became evident. Unlike the 1851 exhibition in which England won most of the prizes, the English performance in Paris was rather poor, and of the 90 classes of manufacturers, Britain dominated only in 10.⁴⁶

In 1868, the government established the parliamentary Select Committee on Scientific Education. This was the origin of nearly 20 years of various parliamentary investigations into the relationship between the sciences, industry, and education designed to address the capitalists' outcry over the necessity of universal public education. A sequence of reports by the Committee in 1868, by the Royal Commission on Scientific Instruction

⁴⁴ There was a growing consensus among workers and capitalists about the virtues of reform. The labor union movement was increasingly calling for a national system of nonsectarian education. The National Education League (founded in 1869 by radical Liberals and Dissenters) demanded a free, compulsory, nonsectarian national system of education (Green (1990)).

⁴⁵ Indeed, the Factory Act of 1802 required owners of textile mills to provide elementary instruction for their apprentices, but the law was poorly enforced (Cameron (1993)).

⁴⁶ Lyon Playfair, who was one of the jurors, reported that "a singular accordance of opinion prevailed that our country has shown little inventiveness and made little progress in the peaceful arts of industry since 1862." The cause of this lack of progress "upon which there was most unanimity conviction is that France, Prussia, Austria, Belgium and Switzerland possess good systems of industrial education and that England possesses none" (Green (1990), p. 296).

and the Advancement of Science during 1872–1875, and by the Royal Commission on Technical Education in 1882 underlined the inadequate training for supervisors, managers, proprietors, and workers. They argued that most managers and proprietors did not understand the manufacturing process and thus failed to promote efficiency, investigate innovative techniques or value the skills of their workers (Green (1990)). In particular, W. E. Forster, the vice president of the committee of the Council of Education, told the House of Commons: "Upon the speedy provision of elementary education depends our industrial prosperity ... if we leave our work-folk any longer unskilled ... they will become overmatched in the competition of the world" (Hurt (1971), pp. 223–224). The reports made various recommendations that highlighted the need to redefine elementary schools, to revise the curriculum throughout the entire school system (particularly with respect to industry and manufacturing), and to improve teacher training.

In addition, in 1868, the Schools Inquiry Commission investigated the secondary schools. It found that the level of instruction in the vast majority of schools was very unsatisfactory, reflecting the employment of untrained teachers and the use of antiquated teaching methods. Its main proposal was to organize a state inspection of secondary schools and provide efficient education geared to the specific needs of its consumers. In particular, the Royal Commission on Technical Education of 1882 confirmed that England was being overtaken by the industrial superiority of Prussia, France, and the United States and recommended the introduction of technical and scientific education to secondary schools.

It appears that the English government gradually yielded to the capitalists and increased contributions to elementary as well as higher education. In the 1870 Education Act, the government assumed responsibility for ensuring universal elementary education. In 1880, prior to the significant extension of the franchise in 1884—which made the working class the majority in most industrial counties—education was made compulsory throughout England. The 1889 Technical Instruction Act allowed the new local councils to set up technical instruction committees, and the 1890 Local Taxation Act provided public funds that could be spent on technical education (Green (1990)). Finally, the 1902 Balfour Education Act marked the establishment of a national education system that provided free compulsory elementary education (Ringer (1979); Green (1990)).

School enrollment of 10-year-olds increased from 40% in 1870 to 100% in 1900. The literacy rate among men, which was stable at around 65% during the first phase of the Industrial Revolution, increased significantly during the second phase reaching nearly 100% at the end of the nineteenth century (Cipolla (1969)). Also, the proportion of children aged 5–14 in primary schools increased significantly in the second half of the nineteenth century, from 11% in 1855 to 74% in 1900 (Flora, Kraus, and Pfenning (1983)).

Continental Europe: The early development of public education occurred in the western countries of continental Europe (e.g., Prussia, France, Sweden, and the Netherlands) well before the Industrial Revolution and was motivated by social, religious, political, and national factors. However, as was the case in England, massive educational reforms occurred in the second half of the nineteenth century due to the rising demand for skills in the process of industrialization. As noted by Green (1990, pp. 293–294), "In continental Europe industrialization occurred under the tutelage of the state and began its accelerated development later when techniques were already becoming more scientific; technical and scientific education had been vigorously promoted from the center as an essential adjunct of economic growth and one that was recognized to be indispensable for countries which wished to close Britain's industrial lead."

In France, the initial development of the education system occurred well before the Industrial Revolution, but the process was intensified and transformed to satisfy industrial needs during the second phase of industrialization. The early development of elementary and secondary education in the seventeenth and eighteenth centuries was dominated by the church and religious orders. Some state interventions in technical and vocational training were designed to reinforce development in commerce, manufacturing, and military efficiency. After the French Revolution, the state established universal primary schools, but enrollment rates remained rather low. The state concentrated on the development of secondary and higher education, with the objective of producing an effective elite to operate the military and governmental apparatus. Secondary education remained highly selective, offering general and technical instruction largely to the middle class (Green (1990)). Legislative proposals during the National Convention quoted by Cubberley (1991) are revealing about the underlying motives behind education in this period: "[C]hildren of all classes were to receive education, physical, moral and intellectual, best adapted to develop in them republican manners, patriotism, and the love of labor ... They are to be taken into the fields and workshops where they may see agricultural and mechanical operations going on."

The process of industrialization in France, the associated increase in the demand for skilled labor, and the breakdown of the traditional apprenticeship system significantly affected the state's attitude toward education. State grants for primary schools gradually increased in the 1830s, and some legislation was introduced to provide primary education in all regions, extend higher education, and provide teacher training and school inspections. The number of communities without schools fell by 50% from 1837 to 1850 and, as the influence of industrialists on the structure of education intensified, education became more stratified according to occupational patterns (Anderson (1975)). This legislation reflected the increasing need for skilled labor in the economic environment of the period (Green (1990)). The eagerness of capitalists for rapid education reforms was reflected by the organization of industrial societies that financed

schools specializing in chemistry, design, mechanical weaving, spinning, and commerce (Anderson (1975)).

As was the case in England, competition led industrialists to lobby for the provision of public education. The Great Exhibition of 1851 and the London Exhibition of 1862 created the impression that the technological gap between France and other European nations was narrowing and that French manufacturers should invest in the education of their labor force to maintain their technological superiority. Subsequently, reports on the state of industrial education by commissions established in 1862-1865 reflected the pleas of industrialists for the provision of industrial education on a large scale and for the implementation of scientific knowledge in the industry. "The goal of modern education ... can no longer be to form men of letters, idle admirers of the past, but men of science, builders of the present, initiators of the future."47 Education reforms in France were extensive in the second phase of the Industrial Revolution, and by 1881, a universal, free, compulsory, and secular primary school system had been established, and technical and scientific education was further emphasized. Illiteracy rates among conscripts tested at the age of 20 declined gradually from 38% in 1851-1855 to 17% in 1876–1880 (Anderson (1975)), and the proportion of children aged 5–14 in primary schools increased from 52% in 1850 to 86% in 1901 (Flora, Kraus, and Pfenning (1983)).

In Prussia, as in France, where the initial steps toward compulsory education took place at the beginning of the eighteenth century, well before the Industrial Revolution, education was viewed primarily as a way to unify the state. In the second part of the eighteenth century, education was made compulsory for all children aged 5–13. Never-theless, these regulations were not strictly enforced partly due to the lack of funding (reflecting the difficulty of taxing landlords for this purpose) and partly due to their adverse effect on child labor income. At the beginning of the nineteenth century, motivated by the need for national cohesion, military efficiency, and trained bureaucrats, the education system was further reformed. Provincial and district school boards were established, education became compulsory (and secular) for a 3-year period, and the gymnasium was reconstituted as a state institution that provided 9 years of education for the elite (Cubberley (1991); Green (1990)).

Similarly to England and France, industrialization in Prussia triggered the implementation of universal elementary schooling. Taxes were imposed to finance the school system, and teachers' training and certifications were established. Secondary schools started to serve industrial needs as well; the *Realschulen*—which emphasized the teaching of mathematics and science—were gradually adopted, and vocational and trade schools were founded. Total enrollment in secondary school increased

⁴⁷ L'enseignement Professionnel, ii (1864, p. 332), quoted in Anderson (1975, p. 194).

sixfold from 1870 to 1911 (Flora, Kraus, and Pfenning (1983)). Furthermore, the Industrial Revolution significantly affected the nature of education in German universities. German industrialists, who perceived advanced technology as a competitive advantage, lobbied for reforms in the operation of universities and offered to pay to reshape their activities toward technological training and industrial applications of basic research (McClelland (1980)).

The evolution of education in the Netherlands also reflected the interest of capitalists in the skill formation of the masses. In particular, as early as the 1830s, industrial schools were established and funded by private organizations, representing industrialists and entrepreneurs. Ultimately, in the latter part of the nineteenth century, the state urged by industrialists and entrepreneurs—started to support these schools (Wolthuis (1999)).

United States: The process of industrialization in the United States also increased the importance of human capital in the production process. Evidence provided by Abramovitz and David (2000) and Goldin and Katz (2001) suggests that during 1890–1999 the contribution of human capital accumulation to the growth process of the United States nearly doubled.⁴⁸ As argued by Goldin (1998), the rise of the industrial, business, and commerce sectors in the late nineteenth and early twentieth centuries increased the demand for managers, clerical workers, and educated sales personnel who were trained in accounting, typing, shorthand, algebra, and commerce. Furthermore, in the late 1910s, technologically advanced industries demanded blue-collar craft workers who were trained in geometry, algebra, chemistry, mechanical drawing, and related skills. The structure of education was transformed in response to industrial development and the increasing importance of human capital in the production process, and American high schools adapted to the needs of the modern workplace of the early twentieth century. Total enrollment in public secondary schools increased seventyfold from 1870 to 1950 (Kurian (1994)).⁴⁹

⁴⁹ As noted by Galor and Moav (2006), due to differences in the structure of education finance in the United States in comparison to European countries, capitalists in the United States had only limited incentives to lobby for the provision of education and to support it financially. Unlike the central role that government funding played in the provision of public education in European countries, the evolution of the educational system in the United States was based on local initiatives and funding. The local nature of education initiatives induced community members, in urban as well as rural areas, to play a significant role in advancing their schooling systems. American capitalists, however, faced limited incentives to support the provision of education within a county in an environment where labor was mobile across counties and the benefits from educational expenditure in one county could be reaped by employers in other counties.

⁴⁸ Literacy rates in the United States were rather high prior to this increase in the demand for skilled labor. Literacy rates among the white population were already 89% in 1870, 92% in 1890, and 95% in 1910 (Engerman and Sokoloff (2000)). Education in earlier periods was motivated by social control, moral conformity, and social and national cohesion, as well as by skills required for trade and commerce. In particular, Bowles and Gintis (1975) and Field (1976) argue that educational reforms are designed to *sustain* the existing social order by displacing social problems into the school system.

5.2.2 Land Concentration: A Hurdle for Human Capital Formation

The transition from an agricultural to an industrial economy altered the nature of the conflict among interest groups in society. The conflict of interest between the elite and the masses which had characterized the agricultural stage of development was transformed into a conflict between the entrenched landed elites and emerging capitalist elites. As documented in Section 5.2.1, the capitalists who were striving for an educated labor force supported policies that promoted public education. In contrast, as documented in this section, landowners, whose interest lay in the reduction of the mobility of the rural labor force, favored policies that deprived the masses of education (Galor, Moav, and Vollrath (2009)).⁵⁰

Anecdotal evidence suggests that the degree of concentration of land ownership across countries and regions is inversely related to education expenditure and attainment. North and South America provide the most distinctive set of suggestive evidence regarding the relationship among the distribution of land ownership, education reforms, and the process of development. The original colonies in North and South America had vast amounts of land per person and levels of income per capita that were comparable to those of Western Europe. However, North and Latin America differed in the distribution of land and resources. While the United States and Canada have been characterized by a relatively egalitarian distribution of land ownership, in the rest of the New World, land and resources have been persistently concentrated in the hands of the elite (Deninger and Squire (1998)).

Persistent differences in the distribution of land ownership between North and Latin America were associated with a significant divergence in education and income levels across these regions (Maddison (2001)). Although all economies in the Western hemisphere were developed enough in the early nineteenth century to justify investment in primary schools, only the United States and Canada were engaged in the education of the general population (Coatsworth (1993); Engerman and Sokoloff (2000)).⁵¹

Variations in the degree of inequality in the distribution of land ownership among Latin American countries were reflected in differences in investment in human capital as well. In particular, Argentina, Chile, and Uruguay, in which inequality in the distribution of land ownership was less pronounced, invested significantly more in education (Engerman and Sokoloff (2000)). Similarly, Nugent and Robinson (2002) show that in Costa Rica and Colombia, where coffee is typically grown on small farms (reflecting

⁵⁰ Interestingly, during the 19th century, the emergence of a *broad-based* demand for human capital-intensive services by the landowners in land-rich economies in Latin America (e.g., Argentina) triggered the establishment of an extensive public education system prior to the onset of significant manufacturing activities (Galiani, Heymann, Dabús, and Tohmé (2008)).

⁵¹ One may view the conflict that led to the Civil War in the United States as a struggle between the industrialists in the North, who were striving for a large supply of (educated) workers, and the landowners in the South, who wanted to sustain the existing system and to ensure a large supply of cheap (uneducated) labor.

lower inequality in the distribution of land), income and human capital are significantly higher than in Guatemala and El Salvador, where coffee plantations are rather large.⁵² Moreover, one of the principles championed by the progressives during the Mexican Revolution of 1910 was compulsory free public education. However, the achievement of this goal varied greatly by state. In the north, where land distribution was more equitable, enrollment in public schools increased rapidly as industrialization advanced following the revolution. In contrast, the south, which was dominated by *haciendas*, employing essentially slave labor, there was virtually no increase in school enrollment following the revolution (Vaughan (1982)). Similarly, rural education in Brazil lagged behind some other Latin American countries due to the immense political power of the local landlords. Hence, in 1950, 30 years after the Brazilian government had instituted an educational reform, nearly 75% of the nation was still illiterate (Bonilla (1965)).

Moreover, the adverse effect of the concentration of land ownership on education expenditure has been established empirically. Galor, Moav, and Vollrath (2009) exploit exogenous source of variation in the concentration of land ownership across states in the United States using data from the beginning of the twentieth century, and they find that inequality in the distribution of land ownership indeed had an adverse effect on public investment in education. Becker, Cinnirella, and Woessmann (2010, 2012) exploit variation in the concentration of land ownership across counties in Prussia and show that landownership was negatively related to educational enrolment in Prussia in 1816 and in 1849.

5.2.3 Land Reforms and Education Reforms

Evidence from Japan, Korea, Russia, and Taiwan indicates that land reforms were followed by, or occurred simultaneously with, significant education reforms. There are two interpretations for these historical episodes. First, land reforms could have diminished the economic incentives of landowners to block education reforms. Second, an unfavorable shift in the balance of power from the viewpoint of the landed aristocracy could have brought about the implementation of both land and education reforms, consistent with the basic premise that landowners opposed spending on education, whereas others (e.g., the industrial elite) favored it.

Japan and the Meiji Restoration: Toward the end of the Tokugawa regime (1600– 1867), although the level of education in Japan was impressive for its time, the provision of education was sporadic and had no central control or funding, reflecting partly the resistance of the landholding military class to education reforms (Gubbins (1973)).

⁵² In contrast to the proposed theory, Nugent and Robinson (2002) suggest that a holdup problem generated by the monopsonistic power in large plantations prevents commitment to reward investment in human capital, whereas smallholders can capture the reward to human capital and therefore have the incentive to invest. This mechanism does not generate the economic forces that permit the economy to escape this institutional trap.

The opportunity to modernize the educational system arrived following the overthrow of the traditional feudal structure shortly after the Meiji Restoration of 1868. In 1871, an imperial decree initiated the abolishment of the feudal system. In a sequence of legislation during 1871–1883, decisions regarding land utilization and the choice of crops were transferred from landlords to farmers, prohibitions on the sale and mortgage of farmland were removed, a title of ownership was granted to the legal owners of the land, and communal pasture and forest land were transferred from the ownership of wealthy landlords to the ownership of the central government. This legislation resulted in the distribution of land among small family farms, a structure that persisted until the rise of a new landlord system during the 1930s (Hayami (1975), chapter 3).

Education reform and land reform evolved simultaneously. In 1872, the Educational Code established compulsory and locally funded education for all children between the ages 6 and 14 (Gubbins (1973)). In addition, the central government funded a secondary school and university system. The Education Code of 1872 was refined in 1879 and 1886, setting the foundations for the structure of Japanese education until World War II. The progress in education attainment following the land reforms of the Meiji government was substantial: while in 1873 only 28% of school-age children attended schools, this ratio increased to 51% by 1883 and to 94% by 1903 (Passin (1965)).

Russia before the Revolution: Education in tsarist Russia lagged well behind comparable European countries at the close of the nineteenth century. Provincial councils dominated by wealthier landowners were responsible for their local school systems and were reluctant to favor the education of the peasants (Johnson (1969)). Literacy rates in rural areas were 21% in 1896, and the urban literacy rate was 56%. As the tsar's grip on power weakened during the early 1900s, the political power of the wealthy landowners gradually declined, leading to a sequence of agrarian reforms that were initiated by the Premier Pyotr Stolypin in 1906. Restrictions on the mobility of peasants were abolished, fragmented landholdings were consolidated, and the formation of individually owned farms was encouraged and supported through the provision of government credit. Stolypin's reforms accelerated the redistribution of land to individual farmers, and landholdings of the landed aristocracy declined from about 35–45% in 1860 to 17% in 1917 (Johnson (1969)).

Following the agrarian reforms and the declining influence of the landed aristocracy, the provision of compulsory elementary education was proposed. The initial effort of 1906 languished, but the newly created representative body, the Duma, continued to pressure the government to provide free compulsory education. During the period 1908–1912, the Duma approved a sequence of significant increases in expenditures for education (Johnson (1969)). The share of the provincial council's budget that was allocated to education increased from 20% in 1905 to 31% in 1914 (Johnson (1969)), the share of the central government's budget devoted to the Ministry of Public Education increased threefold from 1.4% in 1906 to 4.9% in 1915, and the share of the entire

population that was actively attending schools increased threefold from 1.7% in 1897 to 5.7% in 1915 (Dennis (1961)).

South Korea and Taiwan: The process of development in Korea was marked by major land reform followed by a massive increase in governmental expenditure on education. During the Japanese occupation in 1905–1945, land distribution in Korea became increasingly skewed, and by 1945, nearly 70% of Korean farming households were simply tenants (Eckert (1990)). During the period 1948–1950, the Republic of Korea instituted the Agricultural Land Reform Amendment Act that drastically affected landholdings.⁵³ The principle of land reform was enshrined in the constitution of 1948 and the actual implementation of the Agricultural Land Reform Amendment Act began on March 1950.⁵⁴ This act prohibited tenancy and land renting, put a maximum on the amount of land any individual could own, and dictated that individuals could only own land if they actually cultivated it. Owner cultivated farm households increased sixfold from 349,000 in 1949 to 1,812,000 in 1950, and tenant farm households declined from 1,133,000 in 1949 to essentially 0 in 1950 (Yoong-Deok and Kim (2000)).

Land reforms were accompanied by soaring expenditures on education. In 1949, a new Education Law was passed in South Korea that focused specifically on transforming the population into a technically competent workforce capable of industrial work. This legislation led to dramatic increases in the number of schools and students at all levels of education. Between 1945 and 1960, the number of elementary schools increased by 60%, and the number of elementary students went up by a staggering 165%. In secondary education, the growth was even more dramatic, with both the number of schools and the number of students growing by a factor of ten. The number of higher-education institutions quadrupled, and that of higher education students increased from only 7,000 in 1945 to more than 100,000 in 1960. In 1948, Korea allocated 8% of government expenditures to education. Following a slight decline due to the Korean War, educational expenditure increased to 9% in 1957 and to 15% in 1960, remaining at that level thereafter (Sah-Myung (1983)).

Taiwan experienced similar reforms during the same period, once Japanese colonization ended. The government of Taiwan implemented reforms during 1949–1953, enforcing rent reductions, selling public land to individual farmers who had previously been tenants, and permitting the purchase of rented land. In 1948, prior to these reforms, 57% of farm families were full or part owners, and 43% were tenants or hired

⁵³ A major force behind this land reform was the aim of the post–World War II U.S. provisional government to remove the influence of the large landowners (who were either Japanese or collaborators with the Japanese).

⁵⁴ Formally, education reform took place prior to the land reforms, but the provision for land reform was enshrined in the constitution prior to educational reform. The imminent land reform could have reduced the incentives for the landed aristocracy to oppose this education reform.

hands; by 1959, the share of full or part owners had increased to 81%, and the share of tenants had dropped to 19% (Chen (1961)).

A massive education reform accompanied these land reforms. The number of schools in Taiwan grew by 5% per year between 1950 and 1970, while the number of students grew by 6% per year. The pattern of growth mirrors that of South Korea, with an especially impressive growth of 11% per year in the number of secondary students and 16% per year in higher education students. Funding for education grew from 1.8% of GNP in 1951 to 4.1% in 1970 (Lin (1983)).

In 1950, South Korea and Taiwan were primarily agricultural economies with a GDP per capita (measured in 1990 international dollars) of about \$770 and \$940, respectively. South Korea and Taiwan's GDP per capita lagged well behind many countries in Latin America, such as Colombia (\$2150) and Mexico (\$2360), despite sharing with these countries a legacy of vast inequality in the distribution of agricultural land. In contrast to the Latin American countries, the implementation of land reforms in South Korea and Taiwan and its association with education reforms contributed to their tremendous growth performance in the postwar period. With a level of income per capita in 1950 that placed them not only far behind the nations of Latin America but also behind Congo, Liberia, and Mozambique, these two countries have each grown at an average rate of nearly 6% per year between 1950 and 1998, leaving behind the countries of sub-Saharan Africa and overtaking the Latin American countries. In 1998, South Korea and Taiwan had GDP per capita levels that were 150% higher than that of Colombia and 100% higher than in Mexico (Maddison (2001)).

5.2.4 Political and Education Reforms

The nineteenth century was marked by significant political reforms along with the previously described education reforms and impressive human capital formation. One could therefore challenge the significance of the industrial motive for educational reform, suggesting that political reforms during the nineteenth century shifted the balance of power toward the working class and enabled workers to implement education reforms independently of the interests of the industrial elite. Have political institutions, rather than changes in economic incentives in the process of development, been the prime force behind the formation of human capital during this period?

Political reforms that took place in the nineteenth century had no apparent effect on education reforms during this period, strengthening the hypothesis that industrial development and the increasing demand for human capital were indeed the trigger for human capital formation and the subsequent onset of the demographic transition. Education reforms took place in autocratic states that did not relinquish political power throughout the nineteenth century, and major reforms occurred in societies in the midst of the process of democratization well before the stage at which the working class constituted the majority of voters. In particular, as depicted in Fig. 5.7, the most significant education reforms in England were completed *before* the voting majority shifted to the working class. The Reform Act of 1832 nearly doubled the total electorate, but only 13% of the voting-aged population was enfranchised. Artisans, the working classes, and some sections of the lower middle classes remained outside the political system. The franchise was extended further by the Reform Acts of 1867 and 1884, and the total electorate nearly doubled in each of these episodes. However, working-class voters did not become the majority in all urban counties until 1884 (Craig (1989)).

Figure 5.7 shows that a trend of significant increases in primary education was established well before the extension of the voting rights in the context of the 1867 and 1884 Reform Acts. In particular, the proportion of children aged 5–14 in primary schools increased fivefold (and surpassed 50%) over the three decades prior to the extension of the franchise in 1884 in which the working class was granted a majority in all urban counties. Furthermore, the political reforms do not appear to have had an effect on the pattern of educational reform. In fact, the average growth rate of school attendance from decade to decade during 1855–1920 reached a peak around the Reform Act of 1884 and started declining thereafter. However, it is interesting to note that the abolition of education fees in nearly all elementary schools occurred only in 1891, after the Reform Act of 1884, suggesting that the political power of the working class may have affected the distribution of education cost across the population, but the decision to educate the masses appears to be taken independently of the political power of the working class.





Figure 5.7 Evolution of Voting Rights and School Enrollment.

Source: Flora, Kraus, and Pfenning (1983).

reforms of 1884 that made the working class the majority in most counties. Moreover, while the shadow of the rebellion of the masses that may have prompted political reforms may have also contributed to education reforms, industrial demand for human capital in the second phase of industrialization dominated sociopolitical concerns in human capital formation.

In France, the trend of expanding education also preceded the major political reforms that gave the voting majority to the working class (Fig. 5.7). Prior to 1848, restrictions limited the electorate to less than 2.5% of the voting-aged population. The 1848 revolution led to the introduction of universal voting rights for nearly all adult males and resulted in a majority for working class voters. Nevertheless, the proportion of children aged 5–14 in primary schools doubled (and exceeded 50%) during the two decades prior to the extension of the franchise in 1848. Furthermore, the political reforms of 1848 did not appear to have an effect on the pattern of education expansion.

A similar pattern occurred in other European countries. Political reforms in the Netherlands did not affect the trend in education expansion, and the proportion of children aged 5–14 in primary schools exceeded 60% well before the major political reforms of 1887 and 1897. Similarly, the trends of political and education reforms in Sweden, Italy, Norway, Prussia, and Russia do not lend credence to the alternative hypothesis.⁵⁵

6. CONCLUDING REMARKS

Conventional wisdom about the relationship between income distribution and economic development has been subjected to dramatic transformations in the past century. While Classical economists advanced the hypothesis that inequality is beneficial for economic development, the Neoclassical paradigm, which had subsequently dominated the field of macroeconomics, dismissed the Classical hypothesis and promoted the viewpoint that the study of income distribution has no significance for the understanding of macroeconomic activity and the growth process.

A metamorphosis in these perspectives has taken place in the past two decades. Theory and subsequent empirical evidence have demonstrated that income distribution does, in fact, have a significant impact on the growth process. Moreover, unlike the Classical viewpoint, which underlined the beneficial effects of inequality for the growth process, the modern perspective has highlighted the potential adverse effects of inequality on the process of development.

The replacement of physical capital accumulation by human capital accumulation as the prime engine of economic growth has changed the qualitative impact of inequality

⁵⁵ Relatedly, Galor, Moav, and Vollrath (2009) find a positive effect of education on political reforms across countries in the period 1960–2000.

on the process of development. In early stages of industrialization, as physical capital accumulation was a prime source of economic growth, inequality enhanced the process of development by channeling resources toward individuals whose marginal propensity to save is higher. However, in later stages of development, as human capital has become the prime engine of economic growth, a more equal distribution of income, in the presence of credit constraints, has stimulated investment in human capital and promoted economic growth.

While the process of industrialization raised the importance of human capital in the production process, reflecting its complementarity with physical capital and technology, human capital accumulation has not benefited all sectors of the economy. Inequality in the ownership of factors of production has generated an incentive for some betterendowed agents to block the implementation of institutional changes and policies that promote human capital formation, resulting in a suboptimal level of investment in human capital from a growth perspective.

The transition from an agricultural to an industrial economy changed the nature of the main economic conflict in society. Unlike the agrarian economy, which was characterized by a conflict of interests between the landed aristocracy and the masses, the process of industrialization has brought about an additional conflict between the entrenched landed elite and the emerging capitalist elite. In light of a lower degree of complementarity between human capital and the agricultural sector, education has increased the productivity of labor in industrial production more than in agricultural and primary good production, inducing rural-to-urban migration and thus a decline in the rental rate. Thus, while industrialists have had a direct economic incentive to support education policies that would foster human capital formation, landowners, whose interests lay in the reduction of the mobility of their labor force, have favored policies that deprived the masses of education, as long as their stake in the productivity of the industrial sector was insufficient. The adverse effect of the implementation of public education on landowners' income from agricultural production has been magnified by the concentration of land ownership. Thus, as long as landowners affected the political process and thereby the implementation of growth-enhancing education policies, inequality in the distribution of land ownership has been a hurdle for human capital accumulation, slowing the process of industrialization and the transition to modern growth.

Economies in which land and other natural resources have been more equally distributed have implemented earlier public education campaigns and have benefited from the emergence of a skill-intensive industrial sector and a rapid process of development. In contrast, among economies marked by a more unequal distribution of ownership over land and other natural resources, resource abundance that was a source of richness in the early stages of development has led in later stages to underinvestment in human capital, an unskilled labor-intensive industrial sector, and a slower growth process. Thus, variation in the distribution of ownership over land and other natural resources across countries has contributed to disparity in human capital formation and the industrial composition of the economy, and thus to divergent development patterns across the globe. Moreover, geographical conditions that led to income inequality brought about oppressive institutions designed to maintain the political power of the elite and to preserve the existing inequality.

REFERENCES

- Abramovitz, M., David, P.A., 2000. American Macroeconomic Growth in the Era of Knowledge-Based Progress: The Long-Run Perspective. In: Engerman, S.L., Gallman, R.E. (Eds.) The Cambridge Economic History of the United States. vol. 2. Cambridge University Press, New York, 1–92.
- Acemoglu, D., 1998. Why Do New Technologies Complement Skills? Directed Technical Change and Wage Inequality. Q. J. Econ. 113 (4), 1055–1089.
- Acemoglu, D., Johnson, S., Robinson, J.A., 2005. Institutions as the Fundamental Cause of Long-Run Growth. In: Aghion, P., Durlauf, S.N. (Eds.) Handbook of Economic Growth. vol. IA. Elsevier North-Holland, Amsterdam, The Netherlands.
- Acemoglu, D., Robinson, J.A., 2000. Why Did the West Extend the Franchise? Democracy, Inequality, and Growth in Historical Perspective. Q. J. Econ. 115 (4), 1167–1199.
- Aghion, P., Bolton, P., 1997. A Theory of Trickle-Down Growth and Development. Rev. Econ. Stud. 64 (2), 151–172.
- Alesina, A., Perotti, R., 1996. Income Distribution, Political Instability, and Investment. Eur. Econ. Rev. 40 (6), 1203–1228.

Alesina, A., Rodrik, D., 1994. Distributive Politics and Economic Growth. Q. J. Econ. 109 (2), 465–490. Anderson, R.D., 1975. Education in France 1848–1870. Clarendon Press, Oxford.

- Autor, D.H., Katz, L.F., Kearney, M.S., 2008. Trends in US Wage Inequality: Revising the Revisionists. Rev. Econ. Stat. 90 (2), 300–323.
- Baldwin, R.E., 1956. Patterns of Development in Newly Settled Regions. Manchester Sch. Econ. Soc. Stud. 24 (2), 161–179.
- Banerjee, A.A., Duflo, E., 2003. Inequality and Growth: What Can the Data Say? J. Econ. Growth 8 (3), 267–299.
- Banerjee, A., Newman, A., 1993. Occupational Choice and the Process of development. J. Polit. Economy 101 (2), 274–298.
- Barro, R.J., 2000. Inequality and growth in a panel of countries. J. Econ. Growth 5 (1), 5-32.
- Becker, S.O., Cinnirella, F., Woessmann, L., 2010. The Trade-Off between Fertility and Education: Evidence from before the Demographic Transition. J. Econ. Growth 15 (3), 177Ű–204.
- Becker, S.O., Cinnirella, F., Woessmann, L., 2012. The Effect of Investment in Children's Education on Fertility in 1816 Prussia. Cliometrica 6 (1).
- Becker, S.O., Hornung, E., Woessmannand, L., 2011. Education and Catch up in the Industrial Revolution. Am. Econ. J. Macroecon 3.
- Benabou, R., 1996. Equity and Efficiency in Human Capital Investment: the Local Connection. Rev. Econ. Stud. 63, 237–264.
- Benabou, R., 2000. Unequal Societies: Income Distribution and the Social Contract. Am. Econ. Rev. 90 (1), 96–129.
- Benabou, R., 2002. Tax and Education Policy in a Heterogeneous-Agent Economy: What Levels of Redistribution Maximize Growth and Efficiency? Econometrica 70 (2), 481–517.
- Berman, E., Bound, J., Machin, S., 1998. Implications of Skill-Biased Technological Change: International Evidence. Q. J. Econ. 113 (4), 1245–1279.
- Bonilla, F., 1965. Brazil. In: Coleman, J.S. (Ed.) Education and Political Development. Princeton University Press, Princeton.

- Bourguignon, F., 1981. Pareto Superiority of Unegalitarian Equilibria in Stiglitz's Model of Wealth Distribution with Convex Saving Function. Econometrica 49 (6), 1469–1475.
- Bourguignon, F., Verdier, T., 2000. Oligarchy, Democracy, Inequality and Growth. J. Dev. Econ. 62 (2), 285–314.
- Bowles, S., Gintis, H., 1975. Capitalism and Education in the United States. Social. Revolut. 5 (25), 101–138.
- Brezis, E.S., Temin, P., 2008. Elites and Economic Outcomes. In: Durlauf, S.N., Blume, L.E. (Eds.) New Palgrave Dictionary of Economics. Macmillan, Basingstoke.
- Cameron, R., 1993. A Concise Economic History of the World: From Paleolithic Times to the Present. Oxford University Press, USA.
- Caselli, F., 1999. Technological Revolutions. Am. Econ. Rev. 89 (1), 78-102.
- Checchi, D., Ichino, A., Rustichini, A., 1999. More Equal but Less Mobile? Education Financing and Intergenerational Mobility in Italy and in the US. J. Public. Econ. 74 (3), 351–393.
- Chen, C., 1961. Land Reform in Taiwan. China Publishing Company, Taiwan.
- Chou, C., Talmain, G., 1996. Redistribution and Growth: Pareto Improvements. J. Econ. Growth 1 (4), 505–523.
- Cipolla, C.M., 1969. Literacy and Development in the West. Penguin Books, Harmondsworth, Middlesex.
- Clark, G., 2005. The Condition of the Working Class in England 1209–2004. J. Polit. Economy 113 (6), 1307–1340.
- Coatsworth, J.H., 1993. Notes on the Comparative Economic History of Latin America and the United States. In: Bernecker, W., Tobler, H.W. (Eds.) Development and Underdevelopment in America. Walter de Gruyter, New York.
- Crafts, N., 1985. British Economic Growth during the Industrial Revolution. Oxford University Press, Oxford.
- Craig, F.W.S., 1989. British Electoral Facts, 1832–1987. Gower Press, Brookfield.
- Cubberley, E.P., 1991. The History of Education. Reverside Press, Cambridge, MA.
- Dahan, M., Tsiddon, D., 1998. Demographic Transition, Income Distribution, and Economic Growth. J. Econ. Growth 3 (1), 29–52.
- de la Croix, D., Doepke, M., 2003. Inequality and Growth: Why Differential Fertility Matters. Am. Econ. Rev. 93 (4), 1091–1113.
- Deninger, K., Squire, L., 1998. New Ways of Looking at Old Issues: Inequality and Growth. J. Dev. Econ. 57 (2), 259–288.
- Dennis, G., 1961. Education. In: Florinsky, M.T. (Ed.) Encyclopedia of Russia and the Soviet Union. McGraw-Hill, New York.
- Doepke, M., Zilibotti, F., 2005. The Macroeconomics of Child Labor Regulation. Am. Econ. Rev. 95 (5), 1492–1524.
- Durlauf, S.N., 1996a. A Theory of Persistent Income Inequality. J. Econ. Growth 1 (1), 75–93.
- Durlauf, S.N., 1996b. Neighborhood Feedbacks, Endogenous Stratification, and Income Inequality. In: Barnett, W., Gandolfo, G., Hillinger, C. (Eds.) Dynamic Disequilibrium Modeling. Cambridge University Press, Amsterdam.
- Easterly, W., 2007. Inequality Does Cause Underdevelopment: Insights from a New Instrument. J. Dev. Econ. 84 (2), 755–776.
- Eckert, C., 1990. Korea Old and New: A History. Ilchokak, Seoul.
- Eicher, T.S., García-Peñalosa, C., van Ypersele, T., 2009. Education, Corruption, and the Distribution of Income. J. Econ. Growth 14 (3), 205–231.
- Engerman, S.L., Sokoloff, K.L., 2000. History Lessons: Institutions, Factor Endowments, and Paths of Development in the New World. J. Econ. Perspect. 14 (3), 217–232.
- Fernandez, R., Guner, N., Knowles, J., 2005. Love and Money: A Theoretical and Empirical Analysis of Household Sorting and Inequality*. Q. J. Econ. 120 (1), 273–344.
- Fernández, R., Rogerson, R., 1996. Income Distribution, Communities, and the Quality of Public Education. Q. J. Econ. 111 (1), 135–164.
- Fernández, R., Rogerson, R., 2001. Sorting and Long-Run Inequality*. Q. J. Econ. 116 (4), 1305–1341.

- Fershtman, C., Murphy, K.M., Weiss, Y., 1996. Social Status, Education, and Growth. J. Polit. Economy 104 (1), 108–132.
- Field, A., 1976. Educational Reform and Manufacturing Development in Mid-Nineteenth Century Massachusetts. J. Econ. Hist. 36 (1), 263–266.
- Flora, P., Kraus, F., Pfenning, W., 1983. State Economy and Society in Western Europe 1815–1975, Vol. 1., Palgrave Macmillan, London.
- Foellmi, R., Zweimuller, J., 2006. Income Distribution and Demand-Induced Innovations. Rev. Econ. Stud. 73 (4), 941–960.
- Forbes, K.J., 2000. A Reassessment of the Relationship between Inequality and Growth. Am. Econ. Rev. 90 (4), 869–887.
- Foster, A.D., Rosenzweig, M.R., 1996. Technical Change and Human-Capital Returns and Investments: Evidence from the Green Revolution. Am. Econ. Rev. 86 (4), 931–953.
- Galiani, S., Heymann, D., Dabús, C., Tohmé, F., 2008. On the Emergence of Public Education in Land-Rich Economies. J. Dev. Econ. 86 (2), 434-446.
- Galor, O., 2010. The 2008 Lawrence, R. Klein Lecture—Comparative Economic Development: Insights from Unified Growth Theory. Int. Econ. Rev. 51 (1), 1–44.
- Galor, O., 2011. Unified Growth Theory. Princeton University Press, Princeton, NJ.
- Galor, O., Moav, O., 2000. Ability-Biased Technological Transition, Wage Inequality, and Economic Growth. Q. J. Econ. 115 (2), 469–497.
- Galor, O., Moav, O., 2002. Natural Selection and the Origin of Economic Growth. Q. J. Econ. 117 (4), 1133–1191.
- Galor, O., Moav, O., 2004. From Physical to Human Capital Accumulation: Inequality and the Process of Development. Rev. Econ. Stud. 71 (4), 1001–1026.
- Galor, O., Moav, O., 2006. Das Human-Kapital: A Theory of the Demise of the Class Structure. Rev. Econ. Stud. 73 (1), 85–117.
- Galor, O., Moav, O., Vollrath, D., 2009. Inequality in Landownership, the Emergence of Human-Capital Promoting Institutions, and the Great Divergence. Rev. Econ. Stud. 76 (1), 143–179.
- Galor, O., Mountford, A., 2008. Trading Population for Productivity: Theory and Evidence. Rev. Econ. Stud. 75 (4), 1143–1179.
- Galor, O., Tsiddon, D., 1997a. Technological Progress, Mobility, and Economic Growth. Am. Econ. Rev. 87 (3), 363–382.
- Galor, O., Tsiddon, D., 1997b. The Distribution of Human Capital and Economic Growth. J. Econ. Growth 2 (1), 93–124.
- Galor, O., Weil, D.N., 1996. The Gender Gap, Fertility, and Growth. Am. Econ. Rev. 86 (3), 374-387.
- Galor, O., Weil, D.N., 1999. From Malthusian Stagnation to Modern Growth. Am. Econ. Rev. 89 (2), 150–154.
- Galor, O., Weil, D.N., 2000. Population, Technology, and Growth: From Malthusian Stagnation to the Demographic Transition and beyond. Am. Econ. Rev. 90 (4), 806–828.
- Galor, O., Zang, H., 1997. Fertility, Income Distribution, and Economic Growth: Theory and Cross-Country Evidence. Japan. World. Econ. 9 (2), 197–229.
- Galor, O., Zeira, J., 1988. Income Distribution and Investment in Human Capital: Macroeconomics Implication. Department of Economics, Hebrew University, Working Paper No. 197.
- Galor, O., Zeira, J., 1993. Income Distribution and Macroeconomics. Rev. Econ. Stud. 60 (1), 35-52.
- Goldin, C., 1998. America's Graduation from High School: The Evolution and Spread of Secondary Schooling in the Twentieth Century. J. Econ. Hist. 58 (2), 345–374.
- Goldin, C., Katz, L.F., 1998. The Origins of Technology-Skill Complementarity. Q. J. Econ. 113 (3), 693-732.
- Goldin, C., Katz, L.F., 2001. The Legacy of US Educational Leadership: Notes on Distribution and Economic Growth in the 20th Century. Am. Econ. Rev. 91 (2), 18–23.
- Gradstein, M., 2007. Inequality, Democracy and the Protection of Property Rights. Econ. J. 117 (516), 252–269.
- Green, A., 1990. Education and State Formation. Macmillan, Palgrave, Hampshire, UK.
- Gubbins, J.H., 1973. The Making of Modern Japan, Scholarly Resources. Wilmington, DE.

- Hanushek, E.A., 1992. The Trade-Off between Child Quantity and Quality. J. Polit. Economy 100 (1), 84–117.
- Hassler, J., Mora, J.V.R., 2000. Intelligence, Social Mobility, and Growth. Am. Econ. Rev. 90 (4), 888–908.
- Hassler, J., Jose V. Rodriguez Mora, Zeira, J., 2007. Inequality and Mobility. J. Econ. Growth 12 (3), 235–259.
- Hayami, Y., 1975. A Century of Agricultural Growth in Japan. University of Minnesota Press, Minneapolis. Hurt, J., 1971. Education in Evolution. Paladin, London.
- Johnson, W., 1969. Russia's Educational Heritage. Octagon Books, New York.
- Kaldor, N., 1955. Alternative Theories of Distribution. Rev. Econ. Stud. 23 (2), 83-100.
- Kremer, M., Chen, D., 2002. Income Distribution Dynamics with Endogenous Fertility. J. Econ. Growth 7 (3), 227–258.
- Kurian, G., 1994. Datapedia of the United States, 1790–2000: America Year by Year. Bernan Press, Lanham, MD.
- Kuznets, S., 1955. Economic Growth and Income Inequality. Am. Econ. Rev. 45 (1), 1–28.
- Landes, D.S., 1969. The Unbound Prometheus: Technical Change and Industrial Development in Western Europe from 1750 to the Present. Cambridge University Press, Cambridge, UK.
- Levine, R., 2005. Finance and Growth: Theory and Evidence. Handb. Econ. Growth 1, 865–934.
- Lewis, A.W., 1954. Economic Development with Unlimited Supplies of Labour. Manchester Sch. 22 (2), 139–191.
- Lin, C.J., 1983. The Republic of China (Taiwan). In: Thomas, R.M., Postlethwaite, T.N. (Eds.) Schooling in East Asia. Pergamon Press, Oxford.
- Lizzeri, A., Persico, N., 2004. Why Did the Elites Extend the Suffrage? Democracy and the Scope of Government, with an Application to Britain's Hage of Reformt'*. Q. J. Econ. 119 (2), 707–765.
- Loury, G.C., 1981. Intergenerational Transfers and the Distribution of Earnings. Econometrica 49 (4), 843–867.
- Maddison, A., 2001. The World Economy: A Millennial Perspective. OECD, Paris.
- Maoz, Y.D., Moav, O., 1999. Intergenerational Mobility and the Process of Development. Econ. J. 109 (458), 677–697.
- Matsuyama, K., 2000. A Ricardian Model with a Continuum of Goods under Nonhomothetic Preferences: Demand Complementarities, Income Distribution, and North-South Trade. J. Polit. Economy 108 (6), 1093–1120.
- McClelland, C., 1980. State, Society, and University in Germany. Cambridge University Press, Cambridge, UK, pp. 1700–1914.
- Mejía, D., Posada, C.-E., 2007. Populist Policies in the Transition to Democracy. Eur. J. Polit. Economy 23 (4), 932–953.
- Mejía, D., St-Pierre, M., 2008. Unequal Opportunities and Human Capital Formation. J. Dev. Econ. 86 (2), 395–413.
- Mitch, D.F., 1992. The Rise of Popular Literacy in Victorian England: The Influence of Private Choice and Public Policy. University of Pennsylvania Press.
- Mitch, D.F., 1993. The Role of Human Capital in the First Industrial Revolution. In: Mokyr, J. (Ed.) The British Industrial Revolution: An Economic Perspective. Westview Press, Boulder.
- Moav, O., 2002. Income Distribution and Macroeconomics: The Persistence of Inequality in a Convex Technology Framework. Econ. Lett. 75 (2), 187–192.
- Mokyr, J., 2001. The Rise and Fall of the Factory System: Technology, Firms, and Households Since the Industrial Revolution. Carnegie-Rochester Conf. Ser. Public Policy 55 (1), 1–45.
- Mookherjee, D., Ray, D., 2003. Persistent Inequality. Rev. Econ. Stud. 70 (2), 369-393.
- Murphy, K.M., Shleifer, A., Vishny, R., 1989. Income Distribution, Market Size, and Industrialization. Q. J. Econ. 104 (3), 537–564.
- Nelson, R.R., Phelps, E.S., 1966. Investment in Humans, Technological Diffusion, and Economic Growth. Am. Econ. Rev. 51 (2), 69–75.
- North, D.C., 1959. Agriculture in Regional Economic Growth. J. Farm Econ. 41 (5), 943–951.
- Nugent, J., Robinson, J., 2002. Are Endowments Fate? CEPR Discussion Paper No. 3206.

- Owen, A.L., Weil, D.N., 1998. Intergenerational Earnings Mobility, Inequality and Growth. J. Monet. Econ. 41 (1), 71-104.
- Passin, H., 1965. Japan. In: Coleman, J.S. (Ed.) Education and Political Development. Princeton University Press, Princeton.
- Perotti, R., 1996. Growth, Income Distribution, and Democracy: What the Data Say. J. Econ. Growth 1 (2), 149–187.
- Persson, T., Tabellini, G., 1994. Is Inequality Harmful for Growth? Am. Econ. Rev. 84 (3), 600-621.
- Piketty, T., 1997. The Dynamics of the Wealth Distribution and the Interest Rate with Credit Rationing. Rev. Econ. Stud. 64 (2), 173-189.
- Quah, D.T., 1996. Convergence Empirics Across Economies with (some) Capital Mobility. J. Econ. Growth 1 (1), 95–124.
- Rajan, R.G., 2009. Rent Preservation and the Persistence of Underdevelopment. Am. Econ. J. Macroeconomics 1 (1), 178-218.
- Ringer, F., 1979. Education and Society in Modern Europe. Indiana University Press, Bloomington.
- Rosenstein-Rodan, P.N., 1943. Problems of Industrialisation of Eastern and South-Eastern Europe. Econ. J. 53 (210), 202–211.
- Rosenzweig, M.R., Wolpin, K.I., 1980. Testing the Quantity-Quality Fertility Model: The Use of Twins as a Natural Experiment. Econometrica 48 (1), 227-240.
- Sah-Myung, H., 1983. The Republic of Korea (South Korea). In: Thomas, R.M., Postlethwaite, T.N. (Eds.) Schooling in East Asia. Pergamon Press, Oxford.
- Saint-Paul, G., Verdier, T., 1996. Inequality, Redistribution and Growth: A Challenge to the Conventional Political Economy Approach. Eur. Econ. Rev. 40 (3-5), 719-728.
- Sanderson, M., 1995. Education, Economic Change, and Society in England. Cambridge University Press, Cambridge, UK, pp. 1780-1870.
- Schultz, T.W., 1975. The Value of the Ability to Deal with Disequilibria. J. Econ. Lit. 13 (3), 827-846.
- Simon, B., 1987. Systematization and Segmentation in Education: the Case of England. In: Muller, D., Ringer, F., Simon, B. (Eds.) The Rise of the Modern Educational System. Cambridge University Press, Cambridge.
- Stiglitz, J., 1969. Distribution of Income and Wealth among Individuals. Econometrica 37 (3), 382-397.
- Vaughan, M., 1982. The State, Education, and Social Class in Mexico, 1880-1928. Northern Illinois University Press, DeKalb, IL.
- Wolthuis, J., 1999. Lower Technical Education in the Netherlands, 1798-1993: The Rise and Fall of a Subsystem. Garant, Apeldoorn, The Netherlands.
- Yoong-Deok, J., Kim, Y., 2000. Land Reform, Income Redistribution, and Agricultural Production in Korea. Econ. Dev. Cult. Change 48 (2), 253-268.
- Zilcha, I., 2003. Intergenerational Transfers, Production and Income Distribution. J. Public Econ. 87 (3-4), 489-513.

493



The Design of Performance Pay in Education

Derek Neal

Department of Economics and the Committee on Education, University of Chicago and NBER

Contents

1.	Screening Teachers	499
2.	Moral Hazard	504
	2.1. A Simple Treatment of the Multi-Tasking Problem	506
	2.2. Is Incentive Pay Efficient in Education?	508
	2.3. Generalizability	515
	2.4. Other Hidden Actions and the Contamination of Information	521
	2.5. Choosing Targets and Prizes	523
	2.6. Common Design Flaws	527
	2.6.1 Setting Standards Too High	527
	2.6.2 Incorrect Handicapping	529
	2.6.3 Using the Wrong Sample	530
	2.6.4 Holding the Line	531
	2.6.5 Subjective Targets	532
	2.6.6 The Value of Relative Performance Systems	534
	2.6.7 Aggregation	536
	2.7. Steps Forward	538
	2.7.1 Team Competition	539
	2.7.2 Limitations of Assessment-Based Incentives	541
	2.7.3 Heterogeneity	541
	2.7.4 Back to Screening	543
3.	Markets	543
4.	Conclusion	546
Acknowledgments		548
References		548

Abstract

This chapter analyzes the design of incentive schemes in education while reviewing empirical studies that evaluate performance pay programs for educators. Several themes emerge. First, education officials should not use one assessment system to create both educator performance metrics and measures of student achievement. To mitigate incentives for coaching, incentive systems should employ assessments that vary in both format and item content. Separate no-stakes assessments provide more reliable information about student achievement because they create no incentives for educators to take hidden actions that contaminate student test scores. Second, relative performance

schemes are too rare in education. These schemes are more difficult to manipulate than systems built around psychometric or subjective performance standards. Third, assessment-based incentive schemes are mechanisms that complement rather than substitute for systems that promote parental choice, e.g. vouchers and charter schools.

Keywords

Alignment Standards Relative Performance Hidden Action

In most countries, the vast majority of elementary and secondary schools are publicly funded. Therefore, the study of the design of incentive systems employed in education is primarily the study of how governments structure the institutions and mechanisms used to procure a specific type of service. In democratic countries, the study of incentive design in education may also include explorations of agency problems that exist between voters and the elected officials that design these mechanisms.¹ However, in this chapter, I ignore these latter issues and focus on the design challenges that face a benevolent public education authority that seeks to maximize the welfare generated from public expenditures on education.

In broad terms, the personnel policies in most firms are designed, at least in part, to solve two problems. Firms must select and assign workers appropriately and then elicit efficient effort from them. If one treats schools as firms, the first problem involves the screening and assignment of teachers. Systems that create links between measures of teacher performance and the retention of teachers or the granting of tenure influence the distribution of talent among persons who occupy teaching jobs in equilibrium. Because the personnel policies employed in most public school systems create only weak links between teacher performance and retention or tenure decisions, scholars and reform advocates often cite existing tenure and retention policies as a potential source of inefficiency in personnel policies.² Nonetheless, systems that directly link retention and tenure decisions to measures of teacher performance have not yet been implemented on a large scale. Thus, my discussion below is restricted to simulation results that may help shed light on the potential effects of various reforms to tenure and retention policies. To date, there are no empirical studies that evaluate the effects of performance based promotion and retention systems.

A large section of this chapter examines the effects of performance pay systems that are designed to induce more effort from teachers. Although few incentive systems in education are currently employed as screening devices, many existing systems represent

¹ See Dixit (2002).

² See Ballou and Podgursky (1997) for an extensive treatment of the features of hiring and tenure processes in public schools that appear inefficient.
attempts to solve moral hazard problems between teachers and education authorities. Teachers typically work in a setting where the majority of their actions are hidden from their supervisors and where the contextual information that determines the efficient choice of actions at any point in time is also often hidden from their supervisors. In this setting, it is prohibitively expensive to write forcing contracts that specify desired actions for each potential classroom setting and then provide the monitoring required to make sure that these desired actions are taken.

Faced with these monitoring problems, education authorities may pursue one of two strategies. They can pay educators flat salaries and seek to shape the effort choices of educators through professional development and the processes used to screen and recruit teachers, or they can link incentive payments to some measure of educator performance. In recent decades, education authorities throughout the world have begun to experiment with the latter approach, and the existing literature contains many papers that evaluate the impacts of various teacher incentive programs. However, at this point, the literature contains few papers that formally explore the design of incentive systems for educators. I argue here that many performance pay schemes in education are poorly designed, and a careful review of the empirical literature on the results of various performance pay schemes reveals that poor design yields poor results in predictable ways.

Most performance pay schemes in education are constructed as contests against predetermined performance standards, in which teachers receive rewards if their measured performance exceeds a specified target. A large literature notes that such schemes are problematic when teachers can take actions that inflate the measured achievement of their students relative to their students' true skill levels, and I devote considerable attention to this issue. However, I also note that the tasks of choosing the psychometric performance standards used in such contests and maintaining the integrity of these standards over time are difficult ones. Variation in student assessment results reflects not only variation in educator performance but also variation in the backgrounds and aptitudes of students. Systems that do not correctly control for student characteristics in the creation of performance targets for educators create incentives for educators to avoid certain types of students or schools. In addition, no existing contest schemes contain procedures that adjust performance standards over time to reflect secular progress in available teaching methods. Finally, there is considerable suggestive evidence that performance standards can be compromised by testing agencies that make changes to assessment content or the scaling of assessments over time that compromise the meaning of psychometric scales.

Performance pay schemes built around subjective performance evaluation avoid the technical problems involved in setting statistical performance standards, but these systems have not worked well in practice. Two recent studies suggest that when one group of government educators evaluates the performance of another group of government educators subjectively, performance pay schemes may well morph into increases in base pay for teachers that are not accompanied by improvements in teacher performance.

Even if we assume that an education authority has access to a set of performance metrics that do isolate variation in educator performance relative to the current pedagogical frontier, a simple model of pay for performance contests shows that education authorities can waste resources by setting performance standards either too low or too high. Systems that set standards too low either pay more in prize money than is required given the effort they elicit or elicit less than efficient effort levels or both. Further, systems that set standards too high can be especially wasteful because some educators respond by keeping effort levels low and treating the incentive system as a lottery. Circumstantial evidence suggests that designers of recent systems have, in some cases, set performance standards well above efficient levels, and in other cases, set them far below efficient levels.

In contrast to performance pay schemes built around fixed performance standards, relative performance schemes often elicit more useful information for two reasons. First, the evolution of the distribution of measured performance among educators over time can provide information about how the education production frontier is evolving over time. Further, systems that involve competition among educators for a fixed pool of reward money cannot easily be manipulated into a means of raising base pay for an entire population of teachers who make no changes in their effort levels. Nonetheless, relative performance incentive schemes are rare in education and thus far have been mostly confined to the realm of short-term experiments. Although these experiments produced some encouraging results, there are no large scale relative pay for performance schemes in operation at this time.

Many accountability and performance pay systems employ test scores from assessment systems that produce information used not only to determine rewards and punishments for educators but also to inform the public about secular progress in student learning. As long as education authorities keep trying to accomplish both of these tasks with one set of assessments, they will continue to fail at both tasks. If the goal of assessing students is to measure trends in secular achievement, separate no-stakes assessments provide information that is not likely to be contaminated by hidden actions. However, when authorities use one set of assessment results for both incentive pay and student assessment, educators face incentives to take numerous hidden actions that simultaneously inflate their own measured performance and contaminate information about levels of student achievement.

If education authorities implement separate assessment systems for performance incentives and student assessment, they still face the possibility that educators will engage in wasteful hidden actions that manipulate the results of tests used to determine performance pay, but authorities can mitigate some of these concerns by linking performance pay to the results of assessments that contain no year to year overlap in item content or format. This design eliminates the incentive for teachers to engage in coaching behaviors that do not build lasting skills but simply prepare students for a particular set of questions or test formats. Although assessments without repeated items and common formats cannot be readily placed on a common psychometric scale, the ordinal content of these assessment results can be used to implement performance contest schemes that elicit efficient effort from teachers.

Throughout much of the chapter, I consider an education authority that employs many teachers in many different schools and seeks to design personnel policies that screen and motivate teachers. However, the final sections of the chapter consider the design of incentive systems that operate at the school level, and I discuss how governments can design systems that require schools to compete for public support. Seen through this lens, voucher systems, charter schools, and other systems that expand school choice are complements to and not substitutes for incentive systems built around assessment results.

1. SCREENING TEACHERS

A large empirical literature documents the fact that measured teacher productivity varies greatly among teachers, holding constant observed characteristics of their students and their school environments. However, it is difficult to use observed characteristics of candidate teachers to predict who will actually perform well in the classroom.³ This later finding is consistent with two different views of the information structure in teacher labor markets. The first view contends that candidate teachers know whether or not they will be effective teachers, but they cannot directly reveal this information to prospective employers in a credible way. In this asymmetric information scenario, personnel policies must be designed in ways that induce teachers to reveal their ability type indirectly. A second view is that neither a new teacher or her principal knows how effective she will be and that both parties learn about her effectiveness as she gains experience. In this symmetric learning scenario, personnel policies dictate how teacher compensation evolves as new information about her productivity is revealed and also whether or not she will be allowed to continue teaching given her performance record.

For the purpose of this chapter, I adopt the second view and consider the design of policies that maximize the output of teachers employed at a point in time as well as the sequence of teachers who will occupy a given position over time. I return below to the question of how pay should vary with measured performance. For now, I focus on the issue of whether or not teachers should be allowed to continue teaching based on their past record.

Rockoff and Staiger (2010) make the first formal attempt to derive firing rules that maximize the steady-state output of teachers in a school district. They note that the

³ See Aaronson, Barrow, and Sander (2003); Rockoff (2004); Rivkin, Hanushek, and Kain (2005); and Hanushek and Rivkin (2006).

measured productivity of teachers varies greatly among teachers and that existing research suggests that current hiring and screening procedures in public schools may do little to narrow the dispersion of productivity among new teachers.⁴ They also note that the most important cost of replacing teachers after one or two years on the job is that new teachers typically perform significantly worse than teachers with one or two years of experience.

Using a set of assumptions about the reliability of measured teacher performance, the dispersion in teacher performance, the returns to early experience, and the rate of exogenous exits from teaching, Rockoff and Staiger derive optimal firing rules under various assumptions about how many years teachers teach before school districts make an up or down decision on their retention. They choose rules that maximize the steady-state average productivity per teacher, which is equivalent to maximizing the steady-state total output of the school system since they are holding constant the number of teaching positions in their hypothetical school system.

The policy recommendations that Rockoff and Staiger produce are quite different from current practice in modern school systems. They consider a number of different scenarios that involve different tenure clocks, variances of measurement error in teacher productivity, and hiring costs. However, they always conclude that school systems should dismiss at least two thirds of each new cohort of teachers during their first few years of experience.

The Rockoff and Staiger approach is based on a steady-state analysis that involves the following thought experiment: for any retention policy that describes when teachers will be retained or dismissed based on their history of measured performance, derive the steady-state distribution of teacher quality in the system. Then, choose the policy that maximizes average steady-state teacher quality.

It is not clear that this exercise is the most relevant for policy analysis. If a given school system adopted a Rockoff and Staiger retention policy today that applied to new hires but existing teachers continued to enjoy the same employment protection they enjoy now, it could easily take 20 years for the system to approach the steady-state that Rockoff and Staiger describe. A different and possibly more relevant approach is to consider the policy that maximizes the expected discounted value of teacher quality generated by the sequence of teachers who will occupy a position that is open today. Further, because it is standard in the literature to assume that individual teacher productivity is not influenced by the quality of her co-workers, the optimal rule for one position is the optimal rule for all positions.

Rockoff and Staiger note that most of the returns to experience among teachers come, on average, quite early in their careers. They conclude that the existing literature

⁴ See Ballou and Podgursky (1997).

implies that the performance of first year teachers is on average roughly .47 standard deviations below the average quality of experienced teachers,⁵ but teachers with one year of experience perform almost as well as more experienced teachers, and the returns to experience appear to be roughly zero after two years of experience. Here, I ignore the small returns to experience in year two and focus on the larger returns to experience in the first year of teaching. Under the assumption that the education authority is risk neutral and that the authority is maximizing the present discounted value of teacher quality measured in standard deviation units, the assumption that new teachers are, on average .47 standard deviations less productive than experienced teachers is equivalent to the assumption that the authority must pay a search cost of .47 to fire an experience the first year, the retention policy problem facing an education authority can be described using a well-known model of job matching.

Let θ denote the true productivity of a given teacher. θ is not observed directly, but each period t that a teacher works, the education authority observes a productivity signal, x_t . In year one, $x_1 = -.47 + \theta + \varepsilon_1$. For years t > 1, $x_t = \theta + \varepsilon_t$. Here, ε_t represents measurement error or some other transitory component of measured productivity. For all $t = 1, 2, ..., \infty$, ε_t is drawn identically and independently over time and teachers. The model is denominated in standard deviation units of teacher quality. Assume that $\theta \sim N(0, 1)$ and that $\varepsilon_t \sim N(0, \sigma_{\varepsilon}^2) \forall t$. Let m_t be the posterior belief about expected productivity of a given teacher based on the history of her measured performance, $(x_{t-1}, x_{t-2}, ..., x_1)$, and let ρ_t equal the precision of the authority's beliefs about teacher quality at the beginning of year t of her career. Teachers never die in this model, but there is an exogenous probability, δ , that a teacher leaves teaching in a given period for reasons unrelated to her productivity. Finally, let β be the authority's discount rate.

The timing of events is as follows: the authority hires a new teacher. At the end of the teacher's first period of work, the authority observes x_1 and forms (m_2, ρ_2) . The authority then either allows the teacher to work another period or fires the teacher and hires a new teacher. If the authority retains the teacher, the authority repeats the same review and retention decision process at the end of the teacher's second period of work using both signals, (x_1, x_2) , and the same process repeats in future periods. At the beginning of each period, the education authority is trying to maximize the expected present value of teacher productivity generated by the teachers who fill a particular slot. The Bellman equation that describes the problem facing the education authority is:

$$V(m_t, \rho_t) = \max[V_0, m_t + \beta(1 - \delta)E[V(m_{t+1}, \rho_{t+1}) | m_t] + \beta\delta V_0]$$

⁵ According to Rockoff and Staiger (2010), a one standard deviation improvement in teacher quality is associated with roughly a .15 standard deviation increase in expected student achievement, and on average, the students of rookie teachers perform about .07 standard deviations below students of experienced teachers.

Here, $V(m_t, \rho_t)$ is the value of having a current teacher with t - 1 periods of experience and a history of productivity such that $m_t = E(\theta | x_{t-1}, x_{t-2}, ..., x_1)$, and $V_0 = V(0, 1)$ is the expected value of hiring a new teacher.⁶

Many readers may recognize that I have characterized the education authority's problem using Jovanovic's (1979) model of job matching. Jovanovic describes how a worker optimally searches for a job when he believes that his potential match with each new job comes from the same distribution. I use the model to describe how an education authority optimally fills a vacancy when the authority believes that each new teacher is drawn from the same productivity distribution.⁷ The Jovanovic model is well-known in labor economics, and it is well established that the optimal policy for the authority is to choose a set of cutoff values, $(r_1, r_2, r_3, ...)$, such that teachers are dismissed at the beginning of period *t* if $m_t < r_t$. As long as one assumes that a teacher's actions only affect output in her own classroom, the authority can maximize the expected present value of total productivity in the school system by using this same policy to fill all teaching positions.

I have solved this model using the parameters for δ and σ_{ε}^2 that Rockoff and Staiger employ, and to simplify the numerical analysis, I assume no teacher works more than thirty years.⁸ Given the exogenous quit rate of $\delta = .05$, this assumption has virtually no affect on the optimal cutoffs early in a teacher's career.

The Jovanovic approach differs conceptually from Rockoff and Staiger's steady-state analysis because it explicitly incorporates discounting and because it imposes no tenure clock. As Rockoff and Staiger acknowledge, policies that maximize steady-state payoffs do not properly discount the returns that occur in steady-state. The main cost of firing a teacher is the poor expected performance of the new replacement teacher. This cost is paid today. However, if we assume that existing teachers would continue to enjoy their current employment protections following any changes to the tenure system for new teachers, the benefits of a higher steady-state average teacher quality would come decades from now. Further, rules that force up or out tenure decisions early in a teacher's career raise optimal promotion standards because the education authority cannot correct the mistake of giving tenure to a candidate who is later revealed to be less than deserving.

Thus, it is not surprising that the Jovanovic simulations yield much more conservative dismissal policies than those produced by the Rockoff and Staiger simulations. Exact dismissal rates vary with parameter choices, but the typical set of rules implies that roughly fifty percent of new teachers should be dismissed after one year and small

⁶ ρ_t is only a function of t because this is a normal learning problem.

⁷ Jovanovic (1979) assumed that workers receive all the surplus for employer-employee matches. I am assuming that there is a fixed wage for teachers that the authority must pay to any teacher that fills a slot. Thus, the authority simply wants to maximize the expected present value of productivity generated by each teaching slot.

⁸ $\delta = .05$ and $\sigma_e^2 = 1.5$, which implies a reliability ratio of .4.

fractions of new teachers should be dismissed in years two through six of their tenure with roughly forty percent of new teachers never facing dismissal. Nonetheless, both simulations suggest more stringent firing rules than we currently observe in most public school systems.⁹ Thus, it is important to consider whether or not these simulations form a solid basis for considering drastic changes in personnel practices.

Although both exercises provide interesting starting points for broader research on retention policy, both also share important shortcomings. First, if public schools adopt aggressive firing policies, schools may have to raise salaries to maintain the current quality of their applicant pool. It is not clear how elastic the quality constant supply of potential teachers is, but it is certainly a key consideration for any policy makers who contemplate following Rockoff and Staiger's advice. Second, the more important assumption built into both sets of simulations is that teacher productivity is a fixed trait that does not vary with teacher effort other than through mechanical learning by doing. The simulation exercises described here help us think about some of the costs of the current hiring and firing procedures in public schools, but those who take the resulting dismissal rules seriously as viable policy prescriptions are implicitly or explicitly embracing the view that differences in measured teacher productivity are entirely due to differences in teacher talent and not differences in teacher effort.

Given this starting point, the only way to deal with low performing teachers is to terminate them. Better incentive provision has no value. However, this view of personnel policy is rather extreme given the existing literature on incentives in professional labor markets, and it also reflects a false interpretation of some well known results from the empirical literature on teacher productivity.

The fact that teachers vary in terms of their measured productivity does not imply anything about whether or not most teachers provide socially efficient levels of effort given their talent or whether or not it is possible to improve the entire distribution of teacher productivity through the use of incentives. Further, while the evidence on heterogeneous teacher productivity surely reflects a degree of true talent heterogeneity among teachers, it may also reflect differences among teachers in their own personal effort norms. Given the absence of incentive pay and the level of job security protections in many public school settings, these differences in personal norms could be an important source of ex post differences in teacher performance.

The distinction between talent heterogeneity and norm heterogeneity is important when one is trying to forecast the expected benefits from better incentive provision

⁹ Work by Adams (1997) implies that total separations among young teachers are likely around half the levels of dismissals implied by the rules generated by the Jovanovic simulations. Thus, even if one assumes that all current teachers who quit are being forced out, the implied dismissal rates in the data are quite different than those implied by either set of simulations.

for teachers. If bad teachers are simply teachers who are not able to learn how to teach well, then better performance pay schemes should yield negligible improvements in the distribution of teacher performance. On the other hand, if bad teachers are teachers who are not motivated to take the steps required to teach well, then improvements in the design of incentives may generate significant improvements in the distribution of teacher performance without significant increases in total teacher compensation.

Finally, the types of firing rules discussed here can never operate only as screening mechanisms. Policies that link retention decisions to measures of teacher performance should induce more effort from teachers,¹⁰ and if differences in effort norms are important ex ante, the introduction of these policies should alter the ex post distribution of teacher productivity. In fact, it seems reasonable to conjecture that, if a school system announced even the Jovanovic style dismissal rules that I describe above, the administrators of this system would observe that the threat of dismissal alters the distribution of teacher productivity by compressing differences in teacher effort levels among teachers who share the same talent level. In this scenario, the dismissal rules announced ex ante would no longer be optimal ex post because key parameters in the simulation would be influenced by the change in policy.¹¹

There is little evidence that existing hiring procedures in public schools work well as mechanisms for identifying candidates who will perform well in the classroom. Further, many public school teachers receive tenure in almost a perfunctory manner quite early in their careers.¹² These observations give credence to the notion that better screening and retention polices could yield large gains in teacher productivity. However, the combination of perfunctory tenure, civil service employment protections, and civil service salary schedules also suggest that the dead weight loss associated with inefficient effort allocation among existing teachers is a first order concern as well, regardless of whether or not one contends that many existing teachers should not be allowed to continue teaching.

2. MORAL HAZARD

The literature on the use of assessment based incentive schemes in education often draws a distinction between accountability systems and performance pay systems. Assessment-based accountability systems are promoted as vehicles for holding public schools accountable

¹⁰ The analyses presented here rest on the assumption that teachers earn more than they could in other jobs requiring the same effort levels. If no teachers are earning rents, then it is hard to imagine how any changes in personnel policies could improve teacher performance without spending more money on teacher salaries.

¹¹ Technically, the simulations that I conducted and that Rockoff and Staiger conducted suffer from the same problem because estimates of the variance of teacher value-added are taken from the existing stock of current teachers. However, the Rockoff and Staiger agenda is motivated by the view that there is now a weak correlation at best between being a poor performing teacher and a teacher that leaves teaching.

¹² See Ballou and Podgursky (1997).

for their use of public funds. These systems define achievement standards for students and then measure the performance of schools using metrics that describe the degree of discrepancy between the standards set by the accountability systems and the achievement of the student populations in various schools. Further, these systems often also include a set of sanctions that school administrators and teachers face if their students fail to meet the performance targets set by the accountability system.

In sum, accountability systems typically seek to accomplish two tasks. They attempt to measure the performance of schools relative to a set of public standards in a manner that is consistent over schools and over time. Further, they create incentives for educators to provide effective instruction for their students. Thus, the paradigm that dominates accountability system design involves a two-step procedure. First, measure performance relative to public standards. Then, reward or punish schools based on success or failure to meet these standards.

Because accountability systems typically contain rewards and sanctions that are either not spelled out in detail or less than credible because they cannot be enforced ex post,¹³ the primary function of most accountability systems is performance measurement. In contrast, performance pay systems are more explicitly focused on incentive provision and often contain precise mappings between the performance of students and the compensation and employment status of educators.

In this chapter, I focus most of my attention on performance pay systems for several reasons. To begin, the purpose of this chapter is to explore theory and evidence concerning the design of incentive systems for educators, and performance pay systems are explicit incentive schemes. Further, one of my main conclusions will be that accountability systems should not be used as incentive systems. Systems that serve as mechanisms for providing public information about the achievement of students and the performance of schools relative to public education standards should not contain rewards or sanctions that provide incentives for educators.

Donald Campbell (1976) offered the following summary of common patterns he observed in case studies of the use of performance metrics for incentive provision in government agencies,

"I come to the following pessimistic laws (at least for the U.S. scene): The more any quantitative social indicator is used for social decision-making, the more subject it will be to corruption pressures and the more apt it will be to distort and corrupt the social processes it is intended to monitor."

Campbell, (1976)

¹³ One of the most infamous examples of a system that contains incredible threats of sanctions is the No Child Left Behind Act of 2001 (NCLB). Neal (2010) provided a discussion of how the school re-organization threats attached to the 100% proficiency requirement create confusion concerning how the law will be enforced in future years when this requirement becomes binding and tens of thousands of schools have failed to meet it.

A key component of social decision making in Campbell's analyses is resource allocation among government workers or their units. Thus, one way to understand Campbell's Law is that, when government tries to pursue two missions, e.g. incentive provision and performance measurement, with one system, it fails at both missions. Campbell offered this observation as an empirical law. I will use a simple model of worker responses to incentive schemes to explain why Campbell observed what he observed. In section 2.7, I will discuss how education officials can improve both performance measurement and incentive provision by developing separate systems that address these goals independently.

2.1. A Simple Treatment of the Multi-Tasking Problem

The multi-tasking model of Holmstrom and Milgrom (1991) is the tool that economists most often employ to organize their thoughts about the responses of teachers to various merit pay schemes. Here, I use a special case of their model to build intuition concerning the forces that shape the optimal design of incentives in education.¹⁴

Consider an education authority that hires one teacher to teach one student. The teacher allocates her effort among two different tasks. Let t_1 be the time that the teacher devotes to task one, and let t_2 denote the time she devotes to task two. The human capital production technology is such that

$$h = f_1 t_1 + f_2 t_2 + e$$

where (h - e) is the human capital acquired by the student as a result of the teacher's efforts. Here, *h* is an addition to the value of a skill set, and it is measured in dollars. f_1 and f_2 are constants, and *e* is a random shock to the learning process that captures factors beyond the teacher's control that affect the student's rate of learning. The authority cannot observe *h*, t_1 , or t_2 . However, the authority can observe a statistical measure of teacher performance, *p*, where

$$p = g_1 t_1 + g_2 t_2 + v$$

 g_1 and g_2 are constants, and v is a random shock that influences measured performance. Here, we assume that both e and v are shocks drawn independently from distributions with mean zero that do not depend on the actions of the teacher, (t_1, t_2) . We also assume that the teacher's utility function can be described by

$$U = X - C(t_1, t_2)$$

¹⁴ Here, I follow the Gibbons (2010) exposition of the model. See Baker (2002) for a related treatment.

where X is the teacher's expected income and $C(t_1, t_2)$ describes the teacher's cost of effort for any pair (t_1, t_2) . Now, suppose the education authority seeks to design an optimal compensation contract of the form

$$w = s + bp$$

where s is a base salary, and b is a bonus rate associated with the performance measure p. The base salary s is not interesting for our purposes because it is only a mechanism for dividing surplus between the teacher and the authority. Given any choice of b, one can choose a base salary large enough to elicit the teacher's participation given some outside utility option U_0 .

The optimal bonus rate b is the solution to the following problem:

$$\max_{b} f_1 t_1(b) + f_2 t_2(b) - C(t_1(b), t_2(b)) \quad s.t.$$
$$[t_1(b), t_2(b)] = \arg\max_{t_1, t_2} s + b(g_1 t_1 + g_2 t_2) - C(t_1, t_2)$$

In words, the optimal bonus rate maximizes the difference between the expected value of the human capital created by the teacher's action and the cost of the teacher's actions taking into account that the teacher's response to any bonus rate, b, will be to chose actions that maximize her utility given b. Assume the following cost function for teacher effort,

$$C(t_1, t_2) = .5(t_1 - \overline{t_1})^2 + .5(t_2)^2$$

where \bar{t}_1 is a norm for time devoted to effective instruction. The education authority may have established this norm through previous expenditures devoted to screening potential teachers or professional development activities. The key is that, for the purposes of this analysis, the norm is fixed and not affected by the incentive scheme. The education authority chooses the optimal incentive structure taking \bar{t}_1 as given.

Given this simple setup, a few calculations reveal that the optimal bonus rate is

$$b^* = \frac{f_1 g_1 + f_2 g_2}{g_1^2 + g_2^2} = \frac{\|\mathbf{f}\|}{\|\mathbf{g}\|} \cos \theta$$

where θ is the angle between the vectors (f_1, f_2) and (g_1, g_2) . See Figure 6.1.¹⁵

By assuming that workers and firms are risk neutral and that costs are quadratic, I have made sure that the formula for *b* is simple. Nonetheless, this formula highlights two factors that shape the optimal strength of incentives in more general settings. To begin, $\cos \theta$ is an alignment factor. If the vectors are orthogonal, e.g. $(f_1 = 0, f_2 > 0)$ and

¹⁵ The points [(0, 0), (f_1, f_2) , (g_1, g_2)] form a triangle that can be split into two right triangles. Based on the right triangle that includes the origin, it is easy to show that $\cos \theta = \frac{f_1g_1 + f_2g_2}{\|\mathbf{f}\| \|\mathbf{g}\|}$.



Figure 6.1 Alignment ard Incentives.

 $(g_1 > 0, g_2 = 0)$, then $\cos \theta = \cos 90 = 0$, and $b^* = 0$. If the performance measure is aligned perfectly with true output, the two vectors fall on the same ray from the origin, $\cos \theta = \cos 0 = 1$. The ratio preceding $\cos \theta$ is a scale factor. Note that when the performance metric is perfectly aligned, the ratio of the two vector lengths simply transforms the units of the performance metric into dollars, i.e. the monetary value of the human capital created by the teacher's actions.

2.2. Is Incentive Pay Efficient in Education?

I return specifically to the topics of alignment and scale below, but I first want to discuss what this model teaches concerning whether or not the presence of at least some form of incentive pay is optimal, i.e. $b^* > 0$. This version of the multi-tasking model implies that positive incentives are optimal, $b^* > 0$, as long as $f_1g_1 + f_2g_2 \neq 0$. To see this, note that in cases where $b^* < 0$, the authority can always implement $b^* > 0$ by defining a new performance metric p'=-p.

The condition, $f_1g_1 + f_2g_2 > 0$, always holds if at least three of the four constants, (f_1, f_2, g_1, g_2) , are strictly positive and none are negative, i.e. if one task contributes to both output and the performance measure, the other task contributes to one or both, and neither task is detrimental to either real output or the performance measure, then some positive level of incentive pay is optimal. Define t_1 as time spent teaching the curriculum using teaching techniques deemed to be best practice, and then note that t_2 may be given many different labels, e.g. coaching students regarding test-taking strategies, changing student answers before assessments are graded, etc. Most discussions of educators' responses to high-stakes testing regimes implicitly assume that $f_1 > 0$, $g_1 > 0$, $g_2 > 0$, and

thus, if the gaming activities, t_2 , do not harm students, $f_2 \ge 0$, then optimal policy in this framework must involve some positive level of performance pay, $b^* > 0$.

On its surface, the condition $f_2 \ge 0$ seems like a fairly weak requirement, and thus the formula for b^* above seems to indicate that some positive level of incentive pay is always efficient. However, this feature of b^* is not a robust feature of the multi-tasking model because it hinges on separability in the cost function.

Since education requires both teacher and student time and since students have finite energy and attention, cost functions of the following form may be just as interesting to consider:

$$C(t_1, t_2) = .5(t_1 + t_2 - \overline{t})^2$$

Here, t_1 and t_2 are perfect substitutes, and \overline{t} is a norm for total effort that influences teacher costs. I assume that in the absence of any incentive pay, teachers choose $t_1 = \overline{t}$ and $t_2 = 0$.¹⁶ Given this setting, if the education authority chooses b > 0, then teachers choose $t_1 = 0$ as long as $g_2 > g_1$ and there are many combinations of f_1 , f_2 , g_2 , and \overline{t} such that the authority should optimally choose to have no incentive pay and accept $f_1\overline{t}$ as classroom output. When $f_1 > f_2$ and $f_1\overline{t}$ represents baseline output, any incentive scheme that causes teachers to substitute small amounts of t_2 for t_1 lowers output while holding costs constant.

Nonetheless, if \overline{t} is low enough, incentive pay may still increase total surplus. Since t_1 cannot be negative, the global benefits of increasing t_2 well beyond \overline{t} may compensate for the loss of $t_1 = \overline{t}$. Thus, whether or not $b^* = 0$ hinges on the technology of classroom instruction and the norm, \overline{t} , that exists in a school system.

A key consideration in the literature on responses to incentive schemes in education has been the precise nature of the activities, t_2 , induced by the incentive scheme and the relative values of f_2 and g_2 given the maintained assumption that $g_2 > g_1$. However, it is equally important to consider whether or not t_2 actions represent an increase in total teacher effort or simply substitution away from t_1 . In the latter scenario, effort norms and the nature of teacher cost functions are key. Schemes that induce teachers to devote more effort to coaching and test preparation may improve overall performance among teachers who were devoting little time to effective teaching ex ante. However, these same schemes are harmful when they induce effective teachers to replace effective instruction with significant amounts of coaching.

Now that I have presented a basic model that can serve as a guide to interpreting empirical results on incentive schemes in education, I draw your attention to Table 6.1. This table contains a summary of pay for performance schemes that are either ongoing

¹⁶ Teachers are indifferent among any combinations (t_1, t_2) such that $t_1 + t_2 = \overline{t}$, but here I assume that they use the students' best interests as a tie breaker.

Program	Place/Time	Description	Study	Results
Career Ladder Evaluation System	Tennessee, 1985–1989	5-stage career ladder for teachers that awarded bonuses after reaching the third stage. Bonuses ranged from \$1000 for the third certification level to \$7000 for the fifth certification level.	Dee and Keys (2004)	Math scores increased by 3%, reading scores by 2%, but only increases in math were statistically significant. Teachers on the lower 3 rungs were more effective at promoting math achievement, and teachers at the higher rungs were more effective at promoting reading achievement.
CIS	Kenya, 1997–1998	School based program that awarded bonuses to schools for either being the top-scoring school or for showing the most improvement. Bonuses were divided equally among all teachers in a school, who were working with grades 4–8.	Glewwe, Ilias, and Kremer (2010)	The program increased government exam participation. It did not increase scores in the first year, but treatment scores rose by .14 SDs relative to controls in the second year. However, this improvement did not persist after the completion of the program, and there were no improvements on parallel low stakes NGO exams.
ABC	North Carolina 1996–Present	School based program that awards bonuses to all teachers if school-wide scores meet statistical target. \$1500 maximum bonus. Part of the state accountability system.	Vigdor (2009)	Large Gains in Math and Reading Proficiency on the State Test. NAEP trends suggest that reading gains are suspect, but math gains may reflect real improvement.
DSAIP	Dallas, 1991–1995	Schools were ranked based on gains in student learning. Approximately the top 20% of schools received awards for	Ladd (1999)	Pass rates on standardized tests of Reading and Math increased significantly, but only for white and Hispanic students.

Continued

Program	Place/Time	Description	Study	Results
		each member of their staff. Principals and teachers received \$1000 bonuses, and other staff received \$500.		Black students did not exhibit significant gains relative to other cities. The dropout rate decreased more in Dallas relative to other cities from 1991 to 1994.
KIRIS	Kentucky, 1992–1996	Schools could earn bonus money if they achieved growth targets for school-wide performance on assessments as well as other objectives.	Koretz and Barron (1998), Koretz (2002)	Scores on KIRIS assessments rose dramatically in all subjects, but Kentucky students showed modest gains or no improvement on many comparable NAEP or ACT tests.
Teachers' Incentive Intervention	Israel, 1995–1997	Schools were ranked based on their relative performance adjusted for student background characteristics. Credits hours, matriculation exam pass rates, and dropout rates served as performance criteria. The top 1/3 of schools received awards. 75% of the award went to bonuses for teachers, 25% of the award went to facilities improvements.	Lavy (2002)	Clear evidence of improved outcomes on most dimensions with larger impacts observed in religious schools. Matriculation certificates did not increase in secular schools, but average test scores increased in both secular and religious schools.
PRP	England, 1999–Present	Teachers submit applications for bonus pay and provide documentation of better than average performance in promoting student achievement. Teachers who are promoted become eligible for future raises if they meet documented criteria.	Atkinson et al. (2009)	No clear evidence of improvement. Given one strategy that sought to adjust for experience differences between treatment and controls, English and Science teachers showed modest improvement. Math teachers did not show improvement.

Program	Place/Time	Description	Study	Results
TAP 1999-	17 states, 227 schools	Statistical VAM method produces teacher performance indices of 1 to 5. Teachers with scores of 3 or greater earn a bonus that increases with their score.	Hudson (2010)	Introduction of TAP raises math achievement relative to samples in a synthetic control group by .15 SDs. Reading impacts positive but smaller and imprecisely estimated.
Israel (experiment)	Israel, 2000–2001	A rank-order tournament among teachers of each subject, with fixed rewards of several levels. Teachers were ranked based on how many students passed the matriculation exam, as well as the average scores of their students.	Lavy (2009)	There were overall improvements in pass rates in Math and English due to an overall change in teaching methods, increased after school teaching, and increased responsiveness to student needs among teachers. Increased exam participation rates also played a role in test score gains.
Andhra Pradesh (Randomized Evaluation Study)	India, 2005–2007	100 schools got group bonuses based on school performance, and 100 got individual bonuses based on teacher performance. Bonuses were awarded based on how much the percentage gain in average test scores exceeded 5%.	Muralidharan and Sundararaman (2010)	After 2 years, students in incentive schools scored better than the control group by .28 SDs in math, and .17 SDs in language. These students also tended to do better on questions of all difficulty. Students at incentive schools also did better in non- incentive subjects.
Achievement Challenge Pilot Project (ACPP)	Little Rock, Arkansas, 2004–2007	Individual teachers were awarded bonuses based on their students' improvement on the Iowa Test of Basic Skills. Awards were determined by the level of growth and number of students a teacher had.	Winters (2008)	There was statistically significant improvement in all three subjects (math, reading, language) tested. Students increased 3.5 Normal Curve Equivalent (NCE) points in math (.16 SDs), 3.3 NCE points in reading (.15 SDs), and 4.6 NCE points in language (.22 SDs).

Continued

Program	Place/Time	Description	Study	Results
POINT	Nashville, TN, 2006–2009	Teachers volunteered to participate in a performance pay experiment. Bonuses of 5 K, 10 K, and 15 K were awarded for surpassing the 80%, 90%, and 95% threshold in the historic distribution of value-added.	Springer et al. (2010)	Program involved 5th- through 8th-grade math teachers. Some evidence of achievement gains in 5th- grade math in years two and three, but these gains did not persist over the next school year. No evidence of positive program impacts in other grades. Attrition rates from the study were high years two and three. Attrition is concentrated among inexperienced teachers.
NYC School- Wide Bonus Program	New York City, 2007–2011	Random sample of "high-need" schools participated in a bonus pay scheme. The scheme involved team incentive pay at the school level linked to growth targets, but school compensation committees distributed the bonus money among teachers. The two bonus levels were \$3000 per teacher and \$1500 per teacher. The program was added on top of an accountability program that already put performance pressure on schools.	Goodman and Turner (2010)	Performance scores were weighted averages of improvements in test score performance and inspections of school environment. Target scores required lower-performing schools to make greater improvements. 2008–2009 was the only full year of implementation. 89% of eligible schools won the maximum bonus. There is no clear evidence that the program improved student achievement.

Continued

Program	Place/Time	Description	Study	Results
Portugal's Scale Reform	Portugal, 2007–Present	Abandoned single pay scale in favor of two scale system. Promotion to higher pay scale involved a level jump of about 25% of monthly salary. Teachers in the same school who already worked on the higher pay scale performed the performance assessments for junior teachers.	Martins (2009)	Using schools in the Azores and Madeira as well as private schools as controls, there is no evidence of achievement gains induced by the program and consistent evidence that the program harmed achievement on national exams.
MAP	Florida	Districts choose their own method for measuring teacher contribution to achievement.	No Independent Study	
Procomp	Denver, 2006–Present	Teachers and principals negotiate achievement targets for individual students. Teachers can also earn bonuses for meeting state growth expectations for their students based on statistical targets. Finally, teachers may earn bonuses for serving in a "distinguished" school. School ratings are determined by test scores, parent surveys, and attendance. Maximum performance bonus is 5%.	No Independent Study	
Qcomp	Minnesota, 2007–Present	Much of performance pay linked to evaluations of lesson plans and their implementation. Schools or districts develop their own plans for measuring teacher contributions to measured students achievement.	No Independent Study	

or have been implemented in the recent past. The table devotes particular attention to schemes that have been evaluated by teams of independent scholars, and it covers performance pay schemes from several different countries. Most of these studies address schemes implemented in the developed world, but a few address performance pay in developing countries. As I work through various implications of agency theory for the design of incentive schemes, I will refer back to related empirical results in Table 6.1.

2.3. Generalizability

Table 6.1 shows that many assessment-based performance pay schemes do generate noteworthy increases in student performance on the particular assessment used for incentive provision. Thus, Table 6.1 provides much evidence against the notion that educators simply do not respond to incentives. The exceptions are the PRP system in England, the recent pay scale reform in Portugal, and two recent experiments in New York City and Tennessee. I will comment below concerning the unique features of these schemes that may have muted incentives.

I begin my review of the empirical studies in Table 6.1 by asking how many studies provide evidence that a particular incentive scheme induced changes in teacher effort allocations that improved results on a particular assessment but did not improve students' actual skill levels. In the framework set out above, it seems natural to assume that, given most interpretations of t_2 , $f_1 > 0$, $f_2 \ge 0$, $g_1 \ge 0$, $g_2 > 0$, and this implies that any incentive scheme with b > 0 will induce teachers to supply more total effort $t_1 + t_2 > \overline{t}$, since the marginal costs of both efforts are zero given $t_1 + t_2 = \overline{t}$. However, the choice of b > 0 is clearly not welfare improving if the increased total effort by teachers improves measured performance, p, without generating improvements in actual student human capital, h. This combination of outcomes implies that teachers are expending more effort in response to the incentive scheme without improving student learning, which suggests that improvements in measured performance are coming through increases in t_2 that crowd out time devoted to t_1 and result in lower student human capital.

If the ex post evaluation of a given incentive scheme reveals that student learning did improve, this is not clear evidence that the introduction of the scheme improved welfare. Such a finding constitutes evidence that the scheme created real benefits for students, but these benefits may or may not be greater than the costs of the program. These costs include not only the resources required to implement the program but also any losses of student skill in areas that are not assessed and therefore given less attention after such schemes are implemented.

On the other hand, if studies that evaluate the effects of a given incentive plan reveal no real improvements in student skill, then there is good reason to suspect that the plan is not efficient. Implementing incentive schemes usually requires new resources, and schemes that do not generate real improvements in student skills that are the targets of assessments are not likely sources of improvements in skills that are not directly assessed.

Empirical research on these questions is fundamentally difficult because neither policy makers or researchers observe true skill, h. Nonetheless, a significant body of research on high-stakes testing systems attempts to make inferences about changes in h by exploring whether or not assessment gains induced by particular incentive systems generalize to other assessments. For example, assume that a school district decides to implement a performance pay system for fifth-grade math teachers, and the district links teacher pay to student results on assessment A. Further, assume that following the introduction of this program, student results on assessment A improve. The generalizability issues that interests many researchers in educational statistics are variations on the following counterfactual question;

"Suppose that in every period, the fifth-grade students in this district had also taken a second math assessment, B, and teachers were not rewarded or punished as a result of student outcomes on assessment B. Would one have observed gains on assessment B following the introduction of incentive pay that were comparable to the gains observed on assessment A?"

In sum, do gains measured on the assessments used to determine incentive payments reflect increases in skill that create general improvements in math assessment results or only improvements specific to one assessment format or a particular set of questions?

If gains on a particular high-stakes assessment do not generalize, this is not clear evidence that the incentive system induced no changes in teacher behavior that created real increases in skill. Assessments differ in terms of their relative focus on various topics and assessment B may simply not cover the skills assessed on A that improved. Nonetheless, if one finds that gains on a particular high-stakes assessment do not generalize at all to other assessments that are designed to cover the same curriculum, it is possible that the gains on the high-stakes assessment represent no lasting contributions to student skill. In this case, the district likely induced socially wasteful allocations of teacher effort that improved high-stakes assessment results without improving student skills.

Koretz (2002) summarizes results from several studies of generalizability, and he discusses three different types of teacher behavior that could generate gains on high-stakes tests that do not generalize to other assessments of the same subject matter. To begin, teachers may narrow their instructional focus. If teachers respond to incentives by devoting more class time to topics listed in the curriculum and stressed on a related high-stakes assessments, then scores on these assessments may rise substantially while scores on broader assessments of the same subject may show only modest improvements. This scenario is a plausible explanation for the results found in some generalizability studies, but it seems far fetched as an explanation for why some studies document large improvements on high-stakes assessments while scores on contemporaneous low-stakes assessments of the same domain remain flat or even fall.

Here, it seems more likely that teachers are engaging in what Koretz calls coaching. Coaching involves activities that improve scores on a given assessment without improving student mastery of a subject. Stecher (2002) reviews observational studies of coaching behaviors and cites a striking example of such behavior in two Arizona school districts that introduced high-stakes assessment systems. Shephard and Dougherty (1991) report that teachers in these districts reduced the number of writing assignments they gave to students and increased the number of assignments that involved having students find mistakes in prepared passages. This change in teaching practice likely harmed the development of writing skill among students, but it makes sense as a strategy for raising test scores on standardized tests.¹⁷

Koretz also noted that some educator responses to high-stakes assessment systems go beyond coaching and constitute cheating. I discuss specific examples of cheating in the next section, but for now, I note that both coaching and cheating should generate measured achievement gains that do not generalize to other assessments.

Clean evidence on the generalizability of assessment gains is rare, and the existing literature does not speak with one voice. Some studies provide fairly persuasive evidence that the measured gains induced by a particular performance pay program represented little or no improvement in actual subject mastery. Others provide suggestive evidence that at least a portion of the measured gains induced by particular programs reflects real skill gains.

I begin by considering two programs in Table 6.1 that both involve performance pay that is determined by assessments results collected within state accountability programs. The ABC program in North Carolina allows all teachers in a given school to earn a bonus of up to \$1,500 per teacher based on the test score performance of all the students in the school relative to targets determined by a statistical model that conditions on historical performance in the particular school in question and in the state as a whole. The KIRIS system in Kentucky began in 1992. This system also provided bonus pay for teachers based on team performance. All teachers in a school could earn bonuses if the overall performance of students in their school surpassed targets determined by KIRIS formulas.

Koretz and Barron (1998) examine the effects of KIRIS on achievement during the period 1992–1996. Vigdor (2009) examines the effects on ABC of student achievement in North Carolina. Both studies compare trends in NAEP scores with trends in scores on the state specific assessments used to create school accountability measures and

¹⁷ Stecher (2002) reviewed several related practices that have been documented in other states. In math, a related practice involves working only on math problems that follow a format or rubric know to be present on a particular high-stakes assessment.

determine bonus payments. Koretz and Barron (1998) report results in standard deviation units. Vigdor (2009) reports trends in proficiency rates. These studies provide evidence that KIRIS and ABC produced noteworthy gains in reading and math scores on state assessments. Further, in some subjects and grades, the improvements on the KIRIS exams were extremely large.

Nonetheless, NAEP scores in Kentucky improved by only modest amounts and at rates no greater than one would have expected based on national trends, and reading proficiency rates in North Carolina follow a similar pattern. In fact, eighth grade reading proficiency levels on NAEP in North Carolina have been lower than for most of the past decade than they were in the late 1990s when the state introduced the ABC system. Still, since the introduction of ABC, proficiency rates in math on both the state assessment and the NAEP have risen steadily, and although Vigdor does not compare North Carolina NAEP trends in math with trends in other states, the math results from ABC are at least consistent with the hypothesis that ABC generated gains in math achievement that are not entirely specific to the ABC assessment system.

The ABC and KIRIS programs are of particular interest here because they involved cash payments to educators and independent researchers have explored the generalizability of the gains induced by these systems. However, there is a larger literature on the generalizability of gains induced by high-stakes accountability systems generally. Jacob (2005) concludes that an accountability system introduced in the Chicago Public Schools in 1996 generated noteworthy gains in scores on high-stakes assessments, but he reports that scores on low stakes assessments did not improve among third and sixth grade students relative to what one would have expected based on pre-existing trends in Chicago test scores. Jacob finds that both high and low stakes scores rose sharply among eighth graders, and he concludes that the Chicago accountability program generated increases in general skills among older students but not among younger students.

Klein et al (2000) examine data from Texas during the 1990s. The Texas Assessment of Academic Skills (TASS) program began in the 1990s, and this accountability program received considerable attention because scores on TASS exams rose dramatically following the introduction of state wide accountability measures. However, Klein et al demonstrate that, between 1992 and 1996, changes in NAEP reading and math tests did not always square with corresponding changes in TASS scores. Fourth grade math scores on the NAEP rose sharply in Texas relative to scores in other states, but changes in NAEP fourth-grade reading scores and changes in NAEP eighth grade math scores in Texas followed the same pattern in Texas that one observes nationwide.

Hanushek and Raymond (2005) analyze differences among states in NAEP scores trends during the 1990s and conclude that accountability systems improve student learning if they contain real threats of sanctions for educators when students perform poorly. They reach this conclusion by comparing the time pattern of state-level changes in NAEP scores with the timing of the introduction of state level accountability systems

of different types. They conclude that accountability systems that only create public report cards for schools generate at most small gains in achievement but systems that contain real sanctions for poor educator performance generate noteworthy gains in NAEP scores.

In 2001, the No Child Left Behind Act (NCLB) forced all states to adopt accountability systems that contained, at the least, threats of serious sanctions for poorly performing schools. Because NCLB is a nationwide program, it is nearly impossible to precisely assess its impact on general skill development, but it is clear that measured achievement gains on most state assessments have greatly exceeded gains on the NAEP since 2001.

In sum, the United States literature suggests that assessment based incentives schemes typically generate measured improvements on the assessments used to determine rewards and sanctions for educators, and in some cases but not all, these gains appear to reflect improvements in general student skills. Readers may be less than surprised to learn that results from generalizability studies outside the United States also provide mixed results.

Glewwe (2009) argues that agency problems between public school teachers and education authorities are often much more severe in developing countries than in the developed world. In many developing countries, teachers earn wages that are many times greater than per capita GDP, yet teachers are often absent from school and often absent from their classrooms even when they attend school. He summarizes evidence from a number of developing countries and makes a compelling case that public school teachers in many developing countries perform poorly while earning large wage rents.

Given these stylized facts, policy makers and researchers are interested in learning more about the potential benefits of performance pay schemes for educators in lessdeveloped countries. Two recent studies employ data from field experiments in Kenya and India. These settings are interesting because, in both countries, teachers earn much more than the typical worker and also work within civil service systems that provide extraordinary job security and few performance pressures. The high wages offered to teachers in these countries permit both governments to fill teaching positions with well educated people, but the civil service systems in both countries create widespread concern about teacher effort. As in many other developing countries, absenteeism is a significant problem, and policy makers have concerns about the effort level of teachers who do show up for work. Given the status quo in both countries, some may conjecture that the introduction of incentive pay should create real benefits in both countries. However, the results from these experiments are quite mixed.

Glewwe, Ilias, and Kremer (2010) evaluate an incentive pay program run as an experiment in Kenya by International Child Support (ICS).¹⁸ The program began by

¹⁸ ICS is a Dutch organization that funds education and health interventions that seek to help children in developing countries.

selecting 100 schools that appeared to be performing less than optimally. From these schools, the program administrators chose 50 schools to participate in a program that awarded prizes to teachers based on student test score performance in their schools. The plan involved team incentives since all teachers who worked with students in grades 4 through 8 received common prizes based on an aggregate measure of the performance of all of their students.

The prizes ranged in value from 21 to 43 percent of the monthly earnings of a typical teacher. Students took two types of exams. The program linked teacher prizes to scores on government exams, but ICS also created another set of exams that involved no stakes for teachers. The program generated little evidence of test score improvements during the first year or the program. In the second year, the program created large score gains on government tests but no improvements in scores on the low stakes exams.

Glewwe et al conclude that teachers responded to the program by increasing the number of test preparation sessions held for students. They find no evidence of improvements in teacher attendance or classroom practice. Further, they report that even the improvements on government exams did not persist in year three after the incentive program ended.

The fact that the ICS experiment generated measured improvements in student achievement that did not generalize to a parallel low stakes assessment is not shocking given the results reviewed above. However, it is noteworthy that relative student performance on high-stakes exams returned to pre-program levels when the incentive experiment ended. Thus, the test preparation sessions and other activities that generated the measured improvements in high-stakes test performance during the program did not even generate lasting improvements in test taking skills or knowledge specific to the government exams.

While the Glewwe et al results provide suggestive evidence that the Kenyan program was socially wasteful, a recent incentive pay program in India appears to have generated some real gains for students. The Andhra Pradesh Randomized Evaluation Study (APRES) is a large study of experimental interventions in government primary schools located in the Indian state of Andhra Pradesh. School years begin in the middle of June in Andhra Pradesh, and the program began in late June of 2005 with baseline testing of students in treatment and control schools. Two of the treatments specified in this project involved bonus pay schemes based on student test score outcomes in future years.

Let $\Delta \overline{s}$ equal the percentage point increase in the average score of the students in a given classroom or school. Teachers received bonus pay equal to $500 * \max[0, \Delta \overline{s} - 5]$. Teachers who participated in the group incentive plan received bonuses based on school-level average improvements, so, if the average score in a team-incentive school increased by .07, all teachers received a bonus of 1,000 rupees. Teachers in the

individual incentive program received bonuses according to the same formula based on the performance of their own students.

Muralidharan and Sundararaman (2010) estimate the impacts of these two incentive programs by comparing test score outcomes in treatment schools over the subsequent two years to outcomes in a group of control schools. The APRES design randomly assigned 100 schools to each of the two treatments and the control sample. Both incentive programs generated significant improvements in student tests scores. Taken as a whole, the incentive programs raised scores over a two-year period relative to the control group by .27 standard deviations in math and .17 standard deviations in language. The measured impacts in year two are somewhat larger in schools treated with the individual incentive program.

The APRES experiment did not collect test score data from any parallel set of low-stakes math and reading exams. Thus, it is not possible to perform a direct analysis of the generalizability of these gains. However, Muralidharan and Sundararaman provide much suggestive evidence that these gains do reflect at least some real contributions to students' subject mastery. Scores on social studies and science tests also rose significantly in incentive schools relative to control schools even though teachers received bonus pay based only on the math and language results. Further, there is evidence that teachers in incentive schools assigned more work and conducted classes beyond regular school hours. On the other hand, there is evidence that part of the extra class time was devoted to taking practice tests, which may have involved some coaching behaviors. Further, there is no evidence that teachers in incentive schools improved their attendance rates, which remained far below levels found in developed countries.

The contrast between the results from Kenya and India points to the need for more research on what features of the design and implementation of incentive programs improve outcomes. One obvious difference between the two programs is that the Kenyan program tied reward pay to results on national examinations that had been in place for a long time while the APRES experiment developed their own exams for the program. The greater apparent prevalence of coaching as opposed to improved teaching in Kenya may signal that familiarity with the national exam system greatly raised the relative returns to coaching. This conjecture is quite speculative at this point, but I argue below that coaching is less of a concern if education authorities implement assessment systems such that the specific item content and format of each assessment is not predictable.

2.4. Other Hidden Actions and the Contamination of Information

In the previous section, I discussed implicit and explicit evidence that teachers coach students for specific exam questions and question formats in response to high-stakes assessments. Although coaching is typically not an optimal allocation of teacher effort, some forms of coaching may generate some lasting human capital gains for students, and if coaching activities reflect reduced leisure on the part of teachers rather than reductions in effective teaching time, it is possible that these incentive schemes are improving educator performance relative to the performance one expects given public sector monitoring alone. Nonetheless, a different literature documents other ways that some teachers respond to assessment based incentive schemes that are almost certainly wasteful from a social perspective. In Koretz's (2002) taxonomy, these activities constitute cheating.

Jacob and Levitt (2003) provide clear and compelling evidence that some teachers or principals in Chicago responded to the introduction of high-stakes accountability in 1996 by simply changing their students' answer sheets before returning them. It is worth noting that these cheaters were not terribly sophisticated. Jacob and Levitt found that some classes got entire blocks of questions correct even though their performance on the remaining questions implies that it should have been almost impossible for the whole class to get any one set of even two or three questions correct. The scores for students linked with cheating often reflect large increases from the previous year, and these same students experience small improvements or declines in the following year. Jacob and Levitt conclude that cheating took place in between three and five percent of Chicago classrooms following the introduction of high-stakes testing.

Figlio and Winicki (2005) present evidence that schools in Virginia that faced the most serious threats of sanctions under a state accountability system responded by increasing the sugar content of student meals on the day the state administered high-stakes tests. They also cite several media reports of similar behavior in response to high-stakes assessment systems in other areas of the country. School officials appear to be responding to a literature that links test score performance to glucose level in test takers, and these actions represent a textbook example of how agents may respond to the presence of an incentive system by taking hidden actions that inflate their measured performance but contribute nothing to their actual performance.

Jacob and Levitt (2003) and Figlio and Winicki (2005) show that high-stakes assessment systems induce some educators to engage in behaviors that are socially wasteful. These socially wasteful behaviors as well as the coaching activities described above contaminate public information about school performance in two ways. First, since these types of manipulations inflate assessment results, these behaviors contaminate measures of how student achievement is evolving over time on average in a state, district, or school. Second, because some educators are likely more prone to engage in these manipulations than others, these manipulations also distort our understanding of the relative performance of different districts, schools, and teachers. This second point is often missed in current policy debates. The case studies that Campbell (1976) reviewed involve scenarios in which gaming behaviors contaminate information about the performance of some unit or agency over time. However, if the teachers and principals in various schools differ in their personal norms concerning their distaste for coaching or cheating behaviors, then heterogeneity in coaching or cheating contaminates the information that assessments provide concerning relative performance levels in a cross-section of teachers or schools at a point in time.

Suppose school A has higher measured performance than school B under a low stakes assessment regime, but the measured performance of school B exceeds that of school A after the introduction of a high-stakes assessment program. There are two possibilities. School B may have instituted real improvements and now is more effective than school A, or the staff of school B may simply be more willing than the staff of school A to engage in coaching or cheating behaviors that inflate measured performance. This last possibility may be thought of as Campbell's Law turned on its side because it points to the possibility that hidden responses to incentive schemes may contaminate not only time-series information concerning the true evolution of average performance but also cross-sectional information about the true relative performance of various units at a point in time.

Some policy makers may have their own preferred strategies for minimizing cheating or coaching through the use of independent testing agencies or other devices. However, if the assessment system used to measure student performance or educator performance relative to public standards is a no stakes system that is completely separate from any system of incentives or sanctions for educators, there is no reason for educators to engage in coaching or cheating in the first place. Any assessment-based performance pay system must contain strategies for minimizing gaming behaviors, but the best strategy for making sure that public measurement systems actually produce reliable measurements is to make these systems separate from any systems that reward or punish educators.¹⁹

2.5. Choosing Targets and Prizes

In section 2.1 above, I presented a model where the education authority takes the performance metric as given and must choose an optimal linear piece-rate given this performance metric. Most assessment based incentive programs in education do not resemble piece rate schemes where educators earn bonus pay as a linear function of some scaled performance metric. Instead, most incentive programs for educators are contest schemes, and more often than not, these contests do not involve competition among educators but rather competition to surpass a performance target.

In section 2.1, I described a human capital production function in which teacher actions are the only source of growth in human capital or increase in measured performance, and I used this model to discuss the alignment of incentive schemes when teachers can take multiple hidden actions. Now, I want to set aside the issue of alignment and focus on the choice of performance targets and prizes given a well aligned

¹⁹ Cullen and Reback (2006) showed that schools may also alter the results of assessment-based accountability systems by manipulating the distribution of students who are tested.

performance metric. The existing literature contains little discussion of these issues. The papers summarized in Table 6.1 contain no formal analyses of how to set performance targets for contests or prizes given certain performance targets in order to maximize some clearly defined social objective function.

I assume that teachers can engage in only one action, *t*, which one can think of as time spent employing optimal teaching practices. Further, because I am concerned with how the authority sets performance targets, I also model changes in student human capital and measured academic performance that do not result from teacher effort but rather from baseline learning that reflects activities directed by a student's parents or the student himself.

For now, I continue to assume that each teacher teaches one student and specify the model as follows:

$$h = \gamma(h_0) + ft + e$$

$$p = \varphi(h_0) + gt + v$$

$$U = X - \frac{c}{2} (t - \overline{t})^2$$

Here, as in section 2.1, h is the human capital the student possesses at the end of the period, and p is the measured performance of the student at the end of the period. But now, the educational production function includes $\gamma(h_0)$, which captures baseline learning that is not attributed to teacher effort, and $\varphi(h_0)$, which captures the effect of baseline learning on measured achievement. Both of these baseline learning factors are functions of the student's human capital stock at the beginning of the period, h_0 . The parameters f and g capture the effects of t on human capital growth and changes in measured performance respectively. The terms e and v are mean zero error terms that reflect shocks to the creation and measurement of human capital. Both are drawn identically and independently over all student-teacher pairs, and both distributions are unimodal and symmetric around zero. Let $\Phi(.)$ denote the cumulative distribution function of v. Realizations of v determine the outcomes of contests in equilibrium.

As before, U is teacher utility and X denotes expected teacher income. The cost of effort function is quadratic around the effort norm \overline{t} . Given this setup, the condition $(t^* - \overline{t}) = \frac{f}{c}$ defines socially optimal teacher effort. To keep things simple, I have chosen a setting such that optimal teacher effort is the same for all teachers regardless of the levels of h_0 their students possess. However, performance standards in this setting will vary with h_0 .²⁰

²⁰ In versions where students learn at different rates given the same instruction, the efficient level of instruction will vary among students even when all teachers are homogeneous. If both teachers and students are heterogeneous, the social optimum also involves not only a specification of instruction levels for each student but also the assignment of students to teachers.

I begin by discussing the design of optimal contests against performance standards in a setting where the authority understands teacher preferences and the technology of instruction, knows the quantity $\varphi(h_0)$ for each student and observes p but not h at the end of the period. Let the authority define $\hat{p}(h_0, t^*)$ as the expected measured performance for a student who begins the period with human capital h_0 and receives efficient instruction from his teacher. Assume the authority knows $\hat{p}(h_0, t^*)$ for each student and let the authority announce the following contest scheme. Teachers receive their base salary independent of their effort choice. They also receive bonuses, π , if the measured performance of their students is greater than or equal to the relevant values of $\hat{p}(h_0, t^*)$. The problem facing each teacher is

$$\max_{t} \pi [1 - \Phi(\hat{p}(h_0, t^*) - \varphi(h_0) - gt)] - \frac{c}{2}(t - \overline{t})^2$$

and the teacher's first order condition is

$$\pi g \phi(\hat{p}(h_0, t^*) - \varphi(h_0) - gt) = c(t - \overline{t})$$

Now, suppose that the authority chooses $\pi = \frac{f}{g\phi(0)}$, then the solution to this first-order condition becomes $(t^* - \overline{t}) = \frac{f}{c}$, and it is straightforward to show that the second order condition for a local maximum is also satisfied at t^* . However, more work is required to demonstrate that t^* is a global solution to the teacher's problem. If the density $\phi(v)$ falls too quickly as |v| increases, the teacher may find that the total cost of choosing t^* is greater than the expected return.²¹

Nonetheless, for reasonable parameterizations of this model, when the authority sets the performance standard for a given teacher at $\hat{p}(h_0, t^*)$, there is a prize associated with this standard that elicits efficient effort from the teacher, and the teacher will win the prize with probability one half. This contest scheme is a rather obvious place to begin, but there may be many other combinations of prizes and targets that also elicit efficient effort from teachers. Consider changing the performance standard by an amount Δ while choosing a new prize level $\frac{f}{g\phi(\Delta)}$, the first order condition for a teacher choosing optimal effort when facing a contest of the form $\left(\frac{f}{g\phi(\Delta)}, \hat{p}(h_0, t^*) + \Delta\right)$ is satisfied at t^* , and for many values of Δ , t^* may remain the teacher's optimal effort choice.

Let Ω denote the set of all values of Δ such that teachers choose t^* when facing the contest $\left(\frac{f}{g\phi(\Delta)}, \hat{p}(h_0, t^*) + \Delta\right)$. Given $\Delta \in \Omega$, each contest is associated with an expected

²¹ It is well established that one can design two-person contests such that both workers chose efficient effort as part of a pure strategy Nash equilibrium. See Lazear and Rosen (1981). However, these equilibria require that chance play a sufficient role in the outcome of these contests. A contest against the standard $\hat{p}(h_0, t^*)$ is analogous to a game against a machine that always chooses efficient effort, and by making sure that chance plays a sufficient role in determining the outcome of such contests, the authority can ensure that the teacher's best response is to also choose efficient effort.

payoff $\frac{f}{g\phi(\Delta)}[1 - \Phi(\Delta)]$. Based on the history of performance pay experiments in public education, I assume that the base salary of teachers is fixed and that any prize money teachers earn from the introduction of performance pay systems is additional income. Further, I assume that the education authority's goal is to minimize the additional payroll cost of introducing a contest scheme that induces efficient effort. Thus, the optimal Δ minimizes the expected prize payoff $\frac{f}{g\phi(\Delta)}[1 - \Phi(\Delta)]$, subject to the constraint that $\Delta \in \Omega$. A complete characterization of the solution to this problem is rather tedious, but several features of the optimal solution are worth noting.

To begin, the optimal prize involves a scaling factor, $\frac{J}{g}$, that parallels the scale factor in our optimal piece rate formula in section 2.1. The issue of scaling is front and center in any piece rate scheme, but the issue must also be confronted in contest schemes. The authority needs to understand how to translate the scale of the performance metric into values of student skill stocks in order to choose prizes correctly.²²

Turning to the choice of performance standard, the optimal Δ cannot be negative. Since the authority is considering only contests, $\Delta \in \Omega$, that elicit efficient effort, $\Delta < 0$ implies that teachers win more often than in the $\Delta = 0$ contest. Further, the prize $\frac{f}{g\phi(\Delta)}$ is larger than the prize in the $\Delta = 0$ contest because $\phi(\cdot)$ is maximal at zero. These results imply that the expected cost of the $\Delta = 0$ contest is less than the expected cost for any $\Delta < 0$ contest that ellicits efficient effort.

Although the optimal contest involves $\Delta \ge 0$, the authority must be careful not to choose a Δ that is too large. If Δ is too demanding, teachers may find it optimal to choose some $t < t^*$ because the total cost of choosing t^* exceeds the expected increase in prize winnings from choosing t^* . For example, let $\Phi(v)$ represent a normal distribution with variance σ_v^2 . Then, it is straightforward to show that t^* is not an optimal response to any contest $\left(\frac{f}{g\phi(\Delta)}, \hat{p}(h_0, t^*) + \Delta\right)$ if $\Delta > \frac{c\sigma_v^2}{fg} = \frac{q_v^2}{g(t^* - \bar{t})}$.²³

To provide some insight into this condition, note that if f = 1 and g = 1, then the unit of time used to measure *t* is the unit such that teachers raise the expected value of a student's human capital by one dollar when they allocate one more unit of effective instruction to the student. Further, the units of *v* are such that one can think of these shocks to measured

²² Cunha and Heckman (2008) discuss methods that allow researchers to map test scores for youth into expected values of future adult outcomes like earnings. These methods cannot provide direct evidence on the meaning of scales associated with new assessments unless there are ways to equate the new assessment scales to the scales of tests taken when the current generation of adults was in school. More work is needed in this area to provide better guidance concerning the correct pricing of the psychometric performance measures used in performance pay schemes.

²³ To see that the second order condition is violated for these values of Δ when $t = t^*$ use the fact that $\phi'(\Delta) = -\frac{\Delta}{\sigma_v^2}\phi(\Delta)$. The optimal choice for *t* may be greater or less than t^* , but cases involving inefficiently low levels of effort are the main concern here.

human capital as the equivalent of deletions or additions to the total amount of instruction they receive. Here, in contests where $\Delta > \frac{\sigma_v}{(t^*-t)}\sigma_v$, the teacher's second order condition is violated at t^* . Thus, if one is willing to assume that the effort innovation $(t^* - \overline{t})$ offsets a one standard deviation shock of bad luck, then the optimal $\Delta \in [0, \sigma_v]$.

Even though I have modeled a rather simple contest, a full characterization of the optimal Δ is beyond the scope of this chapter. The point of deriving bounds on the optimal Δ for this case is to demonstrate that it takes little effort to construct environments in which education authorities can easily make one of two mistakes by choosing performance standards in an ad hoc way. First, since $\pi(\Delta)$ is the only prize such that the teacher's first order condition is satisfied when she chooses effort t^* in response to a contest against the performance standard $\hat{p}(h_0, t^*) + \Delta$, an authority that began by choosing $\Delta > \frac{c\sigma_{\nu}^2}{\ell_{\theta}}$ would find that no prize exists such that both the first and second order conditions of the teacher's problem are satisfied at t^* . It is possible to set standards that are too high in the sense that, given such standards, there are no prizes that elicit efficient effort. The typical outcome in these cases is that the authority chooses a prize level that elicits less than efficient effort, but it is possible that the authority could set a prize so large that teachers supplied more than efficient effort.²⁴ Second, authorities can set standards that are clearly too low and waste resources relative to the $\Delta = 0$ benchmark. Any contest that results in significantly more than half of the contestants winning a prize is either not eliciting efficient effort or is wasting prize money.

2.6. Common Design Flaws

Most performance pay programs adopt prizes and performance standards without conducting any formal analyses of expected responses by teachers, and the prevailing view seems to be that simply providing incentives through standards and prizes should improve effort allocation among teachers. However, the model outlined here raises concerns about ad hoc approaches to the design of performance contests. Contests that may seem reasonable to many can actually be wasteful.

2.6.1 Setting Standards Too High

Political forces often create pressure for "high standards" in education, but these pressures can be counterproductive. Although it is clearly wasteful to set standards too low, standards well beyond $\hat{p}(h_0, t^*)$ may induce no additional effort from teachers.

²⁴ Further, although teachers in the model above never respond to incentives by choosing $t=\bar{t}$ because the quadratic cost function assumed here imposes a marginal effort cost of zero at \bar{t} , any fixed cost associated with adjusting effort away from the norm, \bar{t} , introduces the possibility that teachers would respond to excessively demanding standards by staying at \bar{t} .

The POINT program (2006–2009) allowed math teachers in grades 5 through 8 in Nashville, TN to volunteer for a performance pay program. The volunteers were randomly assigned to treatment and control groups. Those in the treatment group were eligible for three levels of bonus pay: \$5,000, \$10,000, and \$15,000, dollars. The reward levels were linked to value-added performance targets associated with the 80th, 90th, and 95th percentiles in the historical distribution of student gains on the Tennessee Comprehensive Assessment Program (TCAP). Although these prizes are significant relative to base levels of teacher pay, Springer et al (2010) report that the students in treatment classrooms typically performed no better than students in control classrooms.²⁵ Further, Springer et al (2010) are not reporting that the gains induced by the program did not generalize. Rather, they report no clear pattern of gains on the high-stakes assessment, TCAP.

It is tempting to say that the POINT results are quite puzzling, given the size of the prizes involved in POINT and the balance of the existing literature on educator responses to incentive schemes. However, it is just as important to note that POINT may have set targets so high that teachers responded optimally by doing roughly what they had done before, $t \approx \overline{t}$. Springer et al (2010) provide an appendix which claims that the expected marginal gains from more teacher effort were likely significant for many teachers. However, their figures suggest that roughly one half of the teachers in the experiment faced less than a twenty percent chance of winning a bonus based on their past record of performance. Although it is quite difficult to determine what the marginal gains to effort were for any of the POINT teachers, it takes little creativity to choose cost and density functions such that the one half of teachers who faced less than a twenty percent chance of winning based on their past performance would have found it optimal to remain at or near \overline{t} .

In the model above, all teachers are equally talented and thus share the same cost of effort, and this is likely not true of teachers in the POINT project. However, the presence of teacher heterogeneity only increases the likelihood that *at least some* teachers responded to the system with no change in effort. The estimated treatment effects in the POINT project almost certainly reflect a weighted average of many different changes in teacher effort, but researchers who work with POINT data in the future should carefully investigate the possibility that a significant portion of POINT teachers *optimally* chose not to change their effort levels.

This observation is closely related to the literature on educational triage in accountability systems. Many systems, including the implementation of NCLB in many states, hold all students to a single proficiency standard. However, this "high standards" for all approach often induces teachers to divert resources away from some students who are currently in great need of special attention and who also have no realistic chance of

²⁵ There was some indication of improvement among fifth graders after year one. However, these impacts did not persist into sixth grade.

reaching proficiency in the near term. Gillborn and Youdell (2000) began this literature with work on the responses of English schools to the structure of national exam systems. Neal and Schanzenbach (2010) document this behavior among Chicago teachers following the introduction of NCLB.

2.6.2 Incorrect Handicapping

The POINT system used a simple value-added approach to transform test scores into performance metrics for teachers. In terms of the model above, this procedure is an attempt to condition on $\varphi(h_0)$ when setting standards for teacher performance. It is obvious that all performance pay schemes based on targets must solve this measurement problem. However, it is surprising how often policy makers have adopted rather heuristic approaches that produced less than desirable results.

The MAP system now in place in Florida replaced an earlier system called STAR. The STAR system attempted to assign points to teachers for different possible innovations in reading levels that their students might experience during a given year. These point allocations formed a Value Table with rows for each initial reading level, columns for each terminal reading level, and entries that specified performance points for each possible outcome. The Value Table methodology represented an attempt to make sure that all teachers competed "on a level playing field" as the law required.²⁶ However, Neal (2009a) reports that the initial results from Hillsborough County provided strong suggestive evidence that the point allocations overstated the relative performance of teachers who worked in affluent schools, and the STAR system was altered and then replaced shortly after its introduction.

The ABC system in North Carolina sets performance targets at the school level and also uses rather ad hoc statistical procedures to attempt to control for baseline differences in school characteristics. Vigdor (2009) reports that this system may also be biased against schools that serve economically and academically disadvantaged students, and Clotfelter et al (2004) report that the introduction of ABC created a dramatic relative decline in the retention rates of faculty in schools serving disadvantaged student populations. These changes in retention rates are quite large, and there is no evidence that these departures were concentrated among weak teachers.

Systems that employ statistical procedures to set performance targets must be implemented with care. Any performance pay scheme that employs a statistical procedure to set performance targets will create incentives for even good teachers to leave their current students if the procedures set performance standards for these students that are too demanding relative to the standards set for others.

²⁶ The method assigned positive points to student improvements and assigned more points to improvements that are less common. Teachers received point deductions for students who regressed.

2.6.3 Using the Wrong Sample

The response of many in the education research community to these observations is that the STAR scheme, the ABC formula, and other ad hoc adjustment schemes are transparently flawed. Advocates of value-added models (VAM) contend that performance metrics without these flaws are available. Consider the following regression model:

$$\gamma_{ijt} = x_{ijt}\beta + D_{ijt}\theta + \varepsilon_{ijt}$$

Here, γ_{ijt} is the test score of student *i* in classroom *j* at time *t*, and x_{ijt} is a set of student, peer, and resources variables that serve as controls for the baseline growth expected from student *i*. The matrix D_{ijt} includes a set of dummy variables that indicate the assignment of students to classrooms at time *t*, and ε_{ijt} is an error term that captures shocks to measured performance. Regression models of this form produce vectors $\hat{\theta} = (\hat{\theta}_1, \hat{\theta}_2 \dots \hat{\theta}_j)$ that contain metrics of classroom performance for all classrooms, $j = 1, 2, \dots J$. Although I use *j* to index classrooms, *j* can also index schools in systems where teachers receive bonus pay for team performance. In either case, performance pay systems built around the VAM approach award prizes to the teachers who work in schools or classrooms associated with values of $\hat{\theta}_i$ that exceed some target level.

VAM advocates contend that this approach is the best way to produce performance metrics for educators that correctly control for differences among students in the expected growth in measured achievement attributable to differences in baseline growth among students, $\varphi(h_0)$. However, in order to set appropriate performance targets, policy makers also need to control for the expected measured gains from efficient instruction, gt^* , and many implementations of VAM fail to address this second issue.

The most widespread and statistically sophisticated assessment based incentive program in the United States is TAP. TAP involves several components, but the assessment based component involves running a regression like the one above and giving teachers a score of one through five based on their rank in the vector $\hat{\theta} = (\hat{\theta}_1, \hat{\theta}_2 \dots \hat{\theta}_j)$. Teachers are then rewarded if they earn a score of three or more. If one ignores the rounding procedure, TAP is paying a bonus to teachers with measured performance above the median, and thus some may see this system as analogous to the $\Delta = 0$ contest above, i.e. the scheme that employs $\hat{p}(h_0, t^*)$ as the performance standard. Note that, in the $\Delta = 0$ contest, all teachers choose efficient effort and win the bonus with probability .5.

However, TAP is not analogous to the $\Delta = 0$ contest. The VAM models TAP uses to produce performance metrics for TAP teachers employ data from both TAP and non-TAP schools. Because non-TAP teachers typically work in schools without performance pay systems and because TAP addresses the widely held belief that teacher effort is not efficient in many traditional public schools, it makes sense to assume that many of the teachers in the VAM samples are not supplying efficient effort. Thus, TAP is using VAM in a manner that sets performance standards below the expected value of measured performance given efficient effort, and the analysis above shows that contest schemes built around standards set below $\hat{p}(h_0, t^*)$, i.e. $\Delta < 0$, waste resources.

Nonetheless, the simple model above does predict that teachers will respond to systems with "low" standards if the prizes are high enough, and in a recent study, Hudson (2010) reports that TAP does improve measured student performance. Hudson compares school-wide improvements in test score performance following the introduction of TAP to test score changes in a composite set of control schools that did not introduce TAP, and she finds that TAP does raise math scores and may improve reading scores as well.

2.6.4 Holding the Line

So far, I have focused on how difficult it may be for education authorities to specify an efficient system of performance standards and prize payments using standard psychometric performance metrics. However, even if an education authority were endowed with an efficient system at a point in time, the authority would find it difficult to maintain the integrity of its performance standards over time. We have already discussed how coaching on the part of teachers can inflate assessment results, but even in a world with no coaching or gaming, placing the results of different assessment forms on a common scale over time is technically quite difficult, which implies that it is difficult to verify the integrity of psychometric scales as well as any performance metrics derived from them.

Suppose that a political organization representing teachers put hidden pressure on testing agencies to make assessments less challenging over time while scoring them in the same manner. If this organization were successful, the scores associated with various performance targets would correspond to lower levels of teacher effort and actual student skill, and the fact that the performance standards had been compromised would be hard for the public or the education officials that represent them to detect. Those who think this concern is far-fetched should consult the literature on the integrity of proficiency standards under NCLB. A detailed review of this literature is beyond the scope of this survey, but there is considerable evidence that political pressures have compromised the meaning of proficiency scores over time under NCLB.²⁷

In addition, the School-Wide Bonus Program in New York City may be an example of problems that arise when changes in exam difficulty compromise performance standards over time. Goodman and Turner (2010) describe an experiment in New York City that began during the 2007–2008 school year and continued through the 2008–09, school year. Schools could earn bonuses of either \$1,500 per teacher or

²⁷ See Cronin et al (2007) and Neal (2010).

\$3,000 per teacher if they met targets for school improvement scores. The improvement scores were weighted averages of measures of progress in student achievement and measures of school environment factors such as attendance and safety. The program required schools that began at lower performance levels to meet higher improvement targets in order to win a bonus. The program involved competition at the school level, but compensation committees at each winning school divided the bonus monies among teachers, and in many instances, the committees exercised their discretion and deviated from an equal sharing rule.

Because the program was announced in the middle of the 2007–2008 school year, 2008–2009 was the first year that the program was in place before the school year began. In 2008–2009, 135 of 152 treatment schools (89%) won the maximum prize of \$3,000 per teacher, and there is scant evidence that the program had any positive impacts on student achievement. The program is quite complicated, and it was layered on top of the New York City accountability system. The failure of the program to impact student achievement may reflect confusion about exactly how the program worked as well as the fact that many schools in the treatment and control samples already faced significant performance pressures from the accountability system. However, there is another possibility. It is widely believed that the state assessments used to generate student achievement and school performance measures became easier over time starting in 2007 and that the scoring and scaling of these assessments did not reflect these changes in exam difficulty. It is almost impossible to know the extent to which teachers were aware of this trend, but ex post, the program operated almost like a change in base pay. Although treatment schools did not perform better than control schools, more than 91% of treatment schools won a bonus, and 89% of treatment schools won the maximum prize.

2.6.5 Subjective Targets

Private-sector firms also face difficult performance measurement issues, and these firms rarely rely solely on statistical methods to solve these problems. Instead, schemes that link rewards to subjective performance evaluations are common in the private sector. Several of the entries in Table 6.1 describe systems that link performance pay for educators to subjective evaluation schemes, and despite their prominence in private industry, the results of these schemes in public education are not impressive.

The Performance Related Pay (PRP) system in England involves two forms of bonus pay for teachers who have already reached the maximum pay level in the standard pay scale. The first is a permanent increase in base salary. The second involves future opportunities to move up to even higher levels of base pay dictated by an extended salary schedule. Atkinson et al. (2009) examine the performance of eligible versus ineligible teachers following the introduction of this system in 1999, and they find no significant effects of eligibility on student performance. When they attempt
to correct these estimates for experience differences between eligible teachers and non-eligible teachers, their results imply that the program increased teacher performance in science but may have harmed teacher performance in math and English.

In sum, there is little clear evidence that the PRP system improved instruction in English schools, and given the process that determined awards, some may not be surprised by this result. The initial cohort of eligible teachers were those who were already at the top level of the standard pay scale. These teachers applied for a permanent increase in base pay and movement to a new promotion and salary schedule by submitting cases which contained evidence that their "students are making progress as good or better than similar students nationally." Wragg, et al. (2001) report that 88% of eligible teachers applied and 97% of those who applied received the award. Unless the returns to teacher experience in England are quite exceptional, the officials who reviewed these cases adopted a lenient interpretation of "as good or better."

Martins (2009) describes a similar performance pay program implemented in Portugal in 2006–2007. This program linked promotion to a higher salary schedule and one-time cash prizes to individual teacher performance evaluations. These evaluations were supposed to consider the performance of students on internal and external exams, feedback from parents on teacher performance, attendance records, and participation in research and professional development activities. However, these evaluations were not conducted by independent third party inspectors. Martins writes that "criteria for progression (promotion and prizes) were to be assessed at each school, by those teachers (already) in the higher pay scale."

Using private school students and students on Portuguese Islands,²⁸ Martins finds that student exam scores on internal tests remained flat or fell slightly following the reform, and scores on national exams fell substantially. Martins does not have an experimental control sample, but the results he reports are so negative that it is difficult to believe that the Portuguese system produced any real achievement gains for students, and students may have been harmed.

In private firms, the person who evaluates a worker's performance is either an owner of the firm or an agent of the owner. In public education, subjective performance evaluation is more problematic because many principals and administrators work under employment and salary rules that create only weak links between the quality of their personnel decisions and their own compensation. Thus, some may not be surprised that performance pay systems that involve one group of public employees making subjective determinations about the bonus payments given to another group of public employees did not generate noteworthy gains in student achievement.²⁹

²⁸ Azores and Madeira implemented weaker versions of the performance pay reforms.

²⁹ See Prendergast (1999) for a discussion of problems that may arise in subjective evaluation systems within large private organizations if agency problems exist between managers and owners.

Still, the English and Portuguese systems are not unique. Many of the entries in Table 6.1 involve systems in which educators are involved in creating the performance standards that they or their coworkers are required to meet in order to earn a bonus. The ProComp system in Denver, the Qcomp system in Minnesota, and the MAP system in Florida all involve district or school level discretion in defining the performance standards that determine performance pay. These programs have not been formally evaluated, but one must worry that these systems may morph into vehicles for raising the base pay of most or all teachers whether or not these teachers improve their performance.

2.6.6 The Value of Relative Performance Systems

Education officials can avoid some of the problems highlighted thus far in section 2.6 if they commit to performance pay schemes that are true relative performance systems. In relative performance schemes, there is a fixed amount of prize money set aside, and all of the prize money is distributed to some worker or workers ex post based on relative performance comparisons among the workers. The reliance on relative performance measures means that some teachers will win and others will lose by construction. Thus, there is no way to manipulate these systems so that every worker receives a bonus even if no worker improved their performance. It is quite difficult to convert relative performance schemes into changes in base pay through corruption activities, whether the activities involve corruption of psychometric standards or manipulation of subjective performance evaluations.³⁰

Further, relative performance schemes can provide information that the education authority needs to maintain the value of incentive schemes over time. Even if an authority knew the level of measured performance associated with efficient effort at a point in time, developments in pedagogy, changes in assessments, or contamination of performance metrics may cause this level to rise or fall over time. In some environments, the authority can use movements in average measured performance to infer how levels of measured performance associated with efficient effort are moving over time. Competition among teachers in relative performance that are associated with efficient classroom effort. Thus, VAM methods on samples of teachers who all face the same incentive system may create adequate control for both student differences in expected baseline achievement growth and the efficient levels of instruction that the system induces teachers to allocate to students.³¹

³⁰ If all teachers could collude on low effort, then the prizes would be handed out each period based on measurement error and each teacher would enjoy an increase in expected base pay without changing their effort. However, it seems unlikely that teachers in an entire school district or state could maintain such collusion.

³¹ See Holmstrom (1982). Barlevy and Neal (2011) describe specifically how this insight applies to the design of incentive systems for educators.

Table 6.1 describes three systems that involve both competition among educators for a fixed set of prizes and the use of VAM methods to rank schools or teachers. Ladd (1999), Lavy (2002), and Lavy (2009) all contain evaluations of experimental performance pay schemes. Ladd (1999) describes a system implemented in Dallas in 1991. The Dallas system was a tournament among schools. Schools received performance scores based on estimates of average value added in the school as well as measures of attendance and dropout rates. The VAM estimates of school performance employed scores from several different assessment systems, and the procedure produced measures of relative school improvement in performance. Each year, about 20 percent of the schools won performance bonuses. All staff in winning schools received a bonus. Principals and teachers received one thousand dollars.

Lavy (2002) describes a tournament among secondary schools in Israel that took place in 1995–1997. Here, secondary schools received performance scores based on estimates of their contributions to improvement in three areas: credit units per student, the fraction of students receiving a matriculation certificate, and the school dropout rate. The top one third of schools received awards that varied with the overall performance ranking of the schools. The largest prize resulted in bonuses for teachers that equaled roughly five percent of the starting salary for a new teacher. The smallest prize generated bonuses that were one fourth as large.

Lavy (2009) describes a tournament among individual Israeli secondary school teachers in 2000–2001. Individual teachers received performance scores based on the average score of their students on the matriculation exam and their students' pass rate. Teachers who taught the same subject competed against each other. Further, because the regression models used to produce relative performance measures included school fixed effects, teachers were competing against other teachers in their school and were rewarded for being exceptional relative to their peers. The program ranked teachers according to pass rate performance and average score performance and used a point system to form an aggregate ranking. The pass rate score contributed more to the overall teacher ranking. Winners received performance pay bonuses based on their total performance index, and the top performers received large bonuses.

None of these programs involved random assignment of schools or teachers to treatment. Thus, the authors employ several empirical strategies that attempt to pin down the causal impacts of these programs. Although some may quibble with the details of any one of these three papers, the results taken as a whole paint a fairly consistent picture. All three papers find that these programs generated significant increases in measured achievement among students, but all three also report significant heterogeneity in estimated treatment effects for different sub-populations. Ladd (1999) reports that the Dallas program generated large gains for white and Hispanic students but not for Black students. Lavy (2002) and Lavy (2009) find that both Israeli programs generated larger improvements among students with lower baseline performance as well as students from less educationally advantaged families, but Lavy also notes that both programs included design features that generated stronger incentives for teachers to direct relatively more attention to weak students.

None of these three papers have access to the type of low-stakes assessment data required to make definitive statements about the generalizability of the measured gains induced by these programs. However, the Dallas system may have been more difficult than many to game because it involved test data from several different assessment systems as well as measures of attendance and dropout rates. Further, Lavy (2009) presents evidence that the Israeli program induced substantial changes in teacher effort and pedagogy.

All of these systems represent components of experimental programs. I know of no ongoing large scale performance pay systems in education that are true relative performance pay schemes. This outcome may reflect the fact that teachers and their unions recognize that relative performance schemes cannot be manipulated into systems that simply change base pay for all teachers.

2.6.7 Aggregation

Although the programs described in Ladd (1999), Lavy (2002), and Lavy (2009) appear to have worked fairly well, the tournament structure of these programs raises important implementation questions. In a world where each teacher has only one student, tournaments would be relatively easy to implement. One could define leagues based on baseline student characteristics, and the within-league rank of each student would determine whether or not his teacher won a prize.

However, because teachers and schools work with many students at one time, the construction of performance rankings based on assessment data is not so straightforward. Imagine a setting with assessments that produced perfectly reliable measures of student skill. Further, suppose one teacher had two students who both began the year with a math score of 150 and then ended the year with scores of 155 and 160. Finally, suppose another teacher had two students who began the year with scores of 100 and 200 respectively and ended the year with scores of 110 and 205 respectively. Based on such data, how could one rank the performance of the two teachers without understanding the values to society of bringing students from 100 to 110, 150 to 155, 150 to 160, or 200 to 205?

The VAM methods used in all three experiments assert that our two hypothetical teachers performed equally well simply because the average score improvement in both hypothetical classrooms was 7.5. The experiments in Dallas and Israel took the average of VAM residuals to create performance ranks for classrooms and sometimes schools, and one must ask when averages that are expressed in units of a particular psychometric scale provide valid rankings of total performance for schools or teachers. These averages provide valid rankings if the VAM model is correctly specified and if scores on a given psychometric scale are a fixed affine transformation of the social value of the underlying

skill levels associated with various scores. Put differently if $p_{ijt} = \gamma_{ijt} = ah_{ijt} + c$, where γ_{ijt} is the test score for student *i* in class *j* in period *t*, h_{ijt} is the social value of this student's skills at the end of period *t*, and a > 0 and *c* are constants, then VAM rankings of class-room of school performance will be accurate.

Yet, if an education authority could create a psychometric scale with these magical properties, then pay for performance schemes based on piece rates must be considered as serious policy options.³² The absence of piece rate schemes in practice may reflect many factors, but I conjecture that a key factor is that the use of piece rates would focus attention on the fact that education authorities do not know whether or not a teacher who moves a child from 150 to 155 on a given developmental scale is creating greater, lesser, or equal social value than a teacher who moves a child from 200 to 205. But, if this is one reason that we do not observe piece-rates schemes based on VAM estimates of teacher performance metrics, there is no reason to accept VAM rankings as ex post performance rankings that determine the allocations of prizes in a tournament. Many VAM estimators are quite complex, and the literature contains lengthy debates about the relative value of different VAM approaches, but the results from all VAM models are sensitive to the psychometric scaling of assessment results, and this fact should give advocates of these models pause.³³

Further, in some contexts, the literal interpretation of VAM performance rankings indicts the whole enterprise. Imagine two fifth grade math teachers in a large district. Both are supposed to take their students as far as they can through a common curriculum, but one teacher works with children in a disadvantaged school who began elementary school not knowing how to count and the other teaches in a selective magnet school designed for gifted children. Now, assume that the test score results from both teachers' classes are part of a state or district wide sample used as inputs into a VAM model that produces a vector $\hat{\theta}$ which contains a performance measure for all fifth grade math teachers in the district. The elements of θ associated with our two hypothetical teachers are supposed to tell us which teacher performed better during the year or at least which teacher one should expect to have performed better. However, these two teachers did not do the same job because they worked with students who were at completely different places in their academic development, and thus it seems almost nonsensical to ask which teacher did better. Functional form assumptions and the assumption that the units of a given psychometric scale serve as a welfare index allow VAM to rank the performances of these two teachers, but the fact that some

³² Many tournament schemes, like those in the Israeli and Dallas experiments, cannot elicit first best effort from all participants unless all teachers are equally talented, but piece-rate systems are efficient even in the presence of worker heterogeneity.

³³ See Briggs and Betebenner (2009), Briggs and Weeks (2009), Reardon and Raudenbush (2009) for more on this issue.

applications of VAM provide clear answers to nonsensical questions should be a source of concern for VAM advocates and not a selling point for VAM methods.

2.7. Steps Forward

In the previous sections, I described how hidden actions like coaching contaminate the information in high-stakes assessments, and I also discussed how hidden manipulations or subjective determinations of performance targets may transform performance pay schemes into increases in expected base pay for teachers without commensurate changes in teacher effort. Finally, I discussed the benefits of performance pay schemes based on measures of relative performance but noted the problems that may arise when policy makers create performance metrics that depend on the implicit assumption that particular psychometric scales serve as proxies for welfare indices.

In recent work with Gadi Barlevy, Barlevy and Neal (2011), we describe a performance pay scheme for educators with the following properties: (i) educators compete against each other for a fixed set of prize money (ii) reward pay is based on rankings of individual student outcomes. No measure of classroom or school output is involved, and no composite ranking of educator performance is created (iii) the mapping between student assessment results and the performance pay given to specific teachers is invariant to the scale used to report assessment results, and (iv) because the system is scale invariant, it can be implemented using a series of assessments that contain no repeated items and no common format, which removes opportunities for teachers to coach students concerning particular formats or items used in previous assessments.

The system we propose is called "pay for percentile" and it works as follows. Consider the population of students taking fifth grade math in a state or a large school district. At the beginning of the year, place each of these students in a comparison set that contains other students with similar records of academic achievement, common family backgrounds, and similar peers. Then, at the end of the school year, give each student a percentile score that describes the fraction of students in his comparison set that performed less well than he did. Average these percentile scores over all the fifth grade math students in a given classroom or school and call this average a percentile performance index. This index is a winning percentage. It tells us how often students in a given unit perform better than students in other units who began the year at the same achievement levels. Finally, pay educators bonuses that are proportional to their percentile performance indices.³⁴

³⁴ Classroom size and the efficient prize in a standard two-person contest determine this constant of proportionality. The Barlevy and Neal (2011) framework extends the two-contestant, single-output tournament model of Lazear and Rosen (1981) to a setting with many contestants and many distinct but jointly produced outputs. In the context of education, the human capital acquired by each student is a distinct output, but the set of outputs produced in the classroom are produced jointly by choosing a vector of time allocations to different tasks, e.g., lesson planning, lecturing, small group instruction, and individual tutoring.

Note that this system relies only on the ordinal information in assessment results, and because only ranks within comparison sets matter, this system does not require and never produces a measure or ranking of overall educator performance. All students compete in seeded contests against students in other schools, and performance pay for educators is determined by the overall winning percentage of their students in these contests. Even though some teacher actions, e.g. lesson planning, group tutoring, classroom lectures, simultaneously affect the expected contest outcomes for many of their students, we show that such a scheme can elicit efficient effort from all teachers on all tasks that create human capital in their students.

Because pay for percentile employs only information concerning relative ranks, it provides no information that allows education authorities to understand how student performance is evolving over time or how the performance of a school is evolving over time. However, as I argue above, separating incentive provision and performance measurement eliminates incentives for educators to take actions that contaminate performance measurements. Education authorities can always measure progress in student achievement using parallel assessment systems that involve no stakes for educators and also contain the overlap in item content and format that make proper equating possible. By making this system a no-stakes system, education authorities remove incentives for educators to engage in the coaching and manipulation activities that currently contaminate the information produced by many accountability systems.

2.7.1 Team Competition

Lavy (2009) reports some positive effects of an incentive scheme that forces teachers to compete against other teachers in the same school, and Muralidharan and Sundararaman (2010) report that the incentive scheme that linked piece rates bonuses to individual teacher performance in India generated larger measured achievement gains than the scheme that paid team piece-rates. While some may be tempted to conclude that individual incentives are important as a means for overcoming free rider problems, there are benefits from implementing pay for percentile as a team competition rather than competition among individual teachers. Although the experimental results appear positive, systems like the one Lavy (2009) describes could create serious problems if implemented as permanent policies.

The presence of school fixed effects in the Israeli VAM models used to create teacher performance measures implies that the performance of each teacher is being measured relative to the average performance of teachers in her school. This convention creates a clear incentive for teachers to sabotage the work of their peers. Sabotage may not have been a problem in a short-lived experiment where teachers may or may not have fully understood the construction of performance metrics. However, the Jacob and Levitt (2003) results suggest that one should not assume that teachers are unwilling to engage in such behaviors when permanent incentive schemes create clear incentives for such malfeasance.

Systems that involve individual incentive pay but no direct competition among teachers working in the same school are less problematic, but education authorities may still prefer to have teachers compete in teams. The persons who may possess the best information about how a particular fifth grade math teacher in a given school can improve are the other fifth grade math teachers in the same school. Incentive systems should encourage these teachers to share this information rather than withhold it. Thus, it makes sense to allow all the teachers who teach a given subject in a particular grade to compete as a team against teachers in other schools that serve similar communities and students. These teams are often so small that free riding should not be a huge concern and peer monitoring should be quite effective. The majority of incentive schemes described in Table 6.1 are team-incentive schemes, and all of the team incentive plans did generate improvements in measured achievement.

There are also statistical reasons to prefer inter-school rather than intra-school competition. Barlevy and Neal (2011) discuss how existing methods in educational statistics can be adapted to estimate percentile performance indices, and a key assumption in these methods, and other methods used to create educational performance metrics, is that the conditioning sets that define league competition are so rich that one can treat the assignment of teachers to students as random given these conditioning variables. It may be easier to satisfy this requirement when performance pay contests involve only inter-school competition. Rothstein (2010) presents evidence from North Carolina data that, within schools, unobserved dimensions of student aptitude affect the allocation of student among classrooms, and it makes sense that this would be the case. In order to maximize the human capital created in their schools, principals must use all the information at their disposal to make optimal matches between students and teachers. Furthermore, any system that asks teachers within the same school to compete against each other may create resistance from some teachers to accept the students who should optimally be assigned to them. However, at the school level or grade level within a school, every student must be assigned to some teacher, and inter-school competition for team bonuses creates incentives for teachers and principals to make sure that students are assigned optimally among teachers.

While it is true that there may still be concerns about selection among schools by parents, it may be possible when implementing performance pay schemes at the level of a state or country to form leagues for schools to compete in such that schools are well matched on the measured characteristics of students, communities, and parents, and no two schools in the same league serve geographic areas that intersect. Given this arrangement, no parents would have chosen their child's particular school over any of the other schools in their school's league, and concerns about selection into schools on unobserved family traits may be less severe.

2.7.2 Limitations of Assessment-Based Incentives

The design of pay for percentile removes opportunities for teachers to coach students for upcoming assessments based on the specific items and format found in previous assessment. Further, this scheme avoids many thorny issues that arise when education authorities attempt to build performance pay systems that are dependent on the scaling of psychometric performance measures. However, any assessment-based performance pay scheme for educators will create alignment problems, and pay for percentile is no exception. Educators still benefit from cheating, e.g. giving students answers during the exam. Further, assessment-based schemes do not reward teachers for building non-cognitive skills that are not assessed.

Concerns about cheating can potentially be addressed by mandating that all assessments be monitored by third party testing agencies, but concerns about teachers diverting effort away from the development of important social and emotional skills must be addressed by building systems that reward teachers for contributing to their students' non-cognitive development. Many of the systems described in Table 6.1 are systems involving multiple components, and while I have focused on the assessment-based components of each program, the presence of other components is an important design issue. Many reasonable social welfare functions imply that the optimal set of personnel policies for educators should create incentives for teachers to foster both the cognitive and non-cognitive development of their students. In the next section, I will discuss a strategy for eliciting information from parents concerning the performance of educators with regard to the social and emotional development of children.

2.7.3 Heterogeneity

All incentive pay schemes in education that are built around statistical performance metrics appear to be designed as mechanisms for eliciting effort from a homogeneous group of teachers. The objective incentive schemes described in Table 6.1 involve statistical targets that are the same for all teachers holding constant the characteristics of their students. Further, the tournament schemes employed in Israel and Dallas involve no handicapping. Given student characteristics, all teachers compete on equal footing. Pay for percentile is similar.

However, if teachers differ in the talent levels, one common set of performance standards cannot elicit efficient effort from all teachers. Further, simple tournament schemes typically do not elicit efficient effort from heterogeneous contests without some handicapping system. Thus, if the education authority can observe teacher characteristics that serve as exogenous proxies for effective talent, then the authority can improve efficiency by seeding contests based not only on student characteristics but also on these teacher characteristics as well as measures of resources within the classroom that may affect teacher effective-ness.³⁵ If this seeding process creates competition among teams of teachers such that teams who compete against each other have symmetric beliefs about their true talent levels, then there will exist prizes such that these seeded contests elicit efficient effort from all teachers. However, if some teams of teachers know that they are either better or worse than the typical team of teachers that shares their characteristics, then more elaborate mechanisms are required.³⁶

Some may advocate piece-rate schemes as a strategy for inducing efficient effort from heterogeneous teachers. While I have already noted that this approach requires that education authorities translate an entire psychometric scale into monetary units, another implementation concern may be even more important. Tournament schemes can be implemented using a fixed amount of money that the authority introduces as an addition to total teacher compensation. Thus, tournaments allow existing teachers to know that they will not receive wage cuts following the introduction of incentive pay, and they allow education authorities to know ex ante exactly how much the incentive scheme will cost.

These features are attractive politically, but no piece-rate scheme can provide both of these features at once. In piece-rate schemes that involve relative pay for performance, teachers who perform well below average must receive salary reductions, and it is possible that those who perform at the lowest levels would owe performance fines in excess of their base salaries. This observation may offer insight into the fact that none of the systems described in Table 6.1 involve piece rates linked to relative performance measures.

The two piece rate schemes in India and Arkansas link performance pay to absolute measures of teacher output. These schemes guarantee non-negative bonuses for all teachers. However, these programs create the possibility that total prize winnings will exceed the budget an authority has set aside ex ante. Further, although both programs were experiments that lasted only a few years, any absolute piece rate scheme implemented as a permanent policy would invite the corruption and cheating activities expected in all scale dependent incentive systems, and these activities could generate significant growth in total bonus pay over time even if the distribution of teacher performance remained fixed over time.

³⁵ Examples include class size, the presence of a teacher's aide, teacher experience, computer resources, etc.

³⁶ Barlevy and Neal (2011) discuss how heterogeneity in teacher talent affects the properties of pay for percentile and other tournament schemes. Several authors have proposed more complex tournament schemes that address heterogeneity directly but are also more difficult to implement. O'Keeffe et al (1984) and Bhattacharya and Guasch (1988) present contest schemes that involve heterogenous contestants selecting the measurement rules and payoff rules that they will compete under.

2.7.4 Back to Screening

I began section 2 by looking at models of screening in which teachers supplied effort inelastically but enjoyed different levels of talent, but most of section 2 implicitly addresses settings where teachers are homogeneous with respect to their talent levels, or at least homogeneous given a set of observed characteristics, and the goal is to design performance schemes that elicit efficient effort. The agenda for future research in this area should be the design of systems that dictate seeded relative performance contests at each stage of a teacher's career while permitting the entire history of winning percentages in these contests to affect not only performance bonuses but also base pay, pension benefits, retention decisions, and the seeding of future contests among remaining teachers. It is not clear how well education authorities can do if they seek to design systems that both screen and provide incentives. The dynamic aspects of such systems create new complications because teachers know that performance today may not only affect compensation today but also whom they compete against in the future. Further, team incentive schemes are useful for encouraging effective co-operation within schools, but measures of individual teacher performance may be most useful for retention policies. In sum, the existing economics of education literature contains considerable research on the construction of methods for evaluating the impacts of performance pay systems or other incentive systems in education, but the literature on the design of these systems remains quite small and limited in scope, and there is much work to be done.

3. MARKETS

I note above that, even if pay for percentile or some other assessment based incentive scheme can be used to induce all teachers in publicly funded schools to teach their students in ways that promote mastery of the topics specified in a common curriculum, most parents and public officials want teachers to be more than conduits of academic information. Parents want their children to feel safe at school, and they want their children to develop emotionally and socially as well as cognitively. Thus, even if education officials develop an assessment based incentive scheme that induces teachers to teach well, they must also address the concern that schools will spend too much time on academics at the expense of the social and emotional development of children.

This observation implies that assessment based incentive schemes can never be more than one component of the incentive systems that publicly funded schools face. However, it is not obvious how education officials should develop incentive schemes that direct the efforts of educators regarding the non-cognitive development of children. It is not at all clear that education officials will ever be able to design assessments of non-cognitive skills that are both extensive enough and reliable enough to use as a basis for incentive pay. In the absence of systems that directly assess non-cognitive skills, education authorities need to consider indirect mechanisms. Although many education policy debates frame assessment based accountability and expansions of parental choice as opposing alternative mechanisms for eliciting better performance from publicly funded schools, I have written in Neal (2009a) that these policies are best seen as complements. Once policy makers recognize that assessment-based accountability proposals, almost by definition, ignore non-cognitive skill development, it is natural to consider these questions: Who possesses good information about the non-cognitive development of children, and who faces strong incentives to truthfully report information they possess about the non-cognitive development of children? "Parents" is a good answer to both questions, and the value of voucher systems, charter school expansions, and other policies that expand school choice is that they provide a means of enlisting millions of parents as performance monitors. Further, education officials can induce these performance monitors to reveal what they are observing using relative simple market mechanisms.

Three recent papers, Barrow and Rouse (2009), Figlio (2009), and Neal (2009b), review the literature on the effects of private schooling and the effects of access to private schools through voucher programs in particular. Three important conclusions stand out as themes concerning the impacts of vouchers in developed countries. First, the measured cognitive benefits of access to private schools through voucher programs are often modest. Second, the effects of voucher access on parental and student satisfaction are often large. Third, access to private schools often creates substantial gains in total education attainment.

Given the existence of at least three recent survey papers on this topic, I will not provide another literature review here. However, I do note that the literature as a whole implies that vouchers often allow parents to find schools for their children that are better matches on dimensions other than academic quality, and better matches apparently lead to more attainment. If parents do possess the ability to evaluate important non-academic aspects of school performance, then it makes sense to consider mechanisms that provide incentives and opportunities for parents to use their evaluations in ways that shape the behavior of educators who receive public funds.³⁷

³⁷ Further, there is evidence that private schools offer an even broader set of benefits for students in developing countries. Andrabi et al (2010) examine outcomes for private school children in Pakistan. They do not have a voucher experiment that generates random variation in private school access, but they do build an instrumental variables strategy by exploiting interactions between the location of families, the location of public schools, and the historic pattern of settlement in rural villages. They find enormous positive effects of private schooling on achievement even though public schools are funded at much higher levels. Angrist et al (2006) report results from a voucher experiment in Colombia. The vouchers covered roughly half of the cost of private schooling and were assigned by lottery. The study used comparisons between lottery winners and losers to estimate the impacts of being offered access to private schooling. The implied achievement gains associated with private school access were large, and the authors conclude that the implied increase in expected adult earnings among recipients likely exceed the cost of the program.

Neal (2009a) outlines a framework for designing systems that distribute public funds among schools that combines features of assessment-based accountability systems and voucher systems. In this framework, all schools, both private and public, compete on multiple dimensions for public funding. Student assessment results, the results of school inspections, feedback from parents, and parental choices affect whether or not a given school is eligible to receive funding and the level of funding it receives in a given year. Much more work is required before researchers can offer specific guidance concerning the optimal mapping from these varied signals of school performance into the funding levels enjoyed by schools, but assessment based performance pay and vouchers may work well together in systems that require schools to compete for public resources on all relevant dimensions of school performance.

By creating competition among schools for students and public resources, such a system also creates competition among schools for teachers. I noted above that subjective performance pay schemes have produced questionable results in public education, and this presumably reflects the fact that educational administrators are not always penalized when they give raises or promotion to undeserving teachers. However, in a managed competition framework, all the teachers in a school as well as the administrators in the school know that the future capacity of the school to provide higher salaries for its employees is directly influenced by the quality of its personnel policies. The best solutions to the screening and incentive provision problems described above may arise as byproducts of a system that forces schools to compete for the public support they receive. A competitive market for teachers allows schools to build reputations as employers that reward teachers for excellent performance on all dimensions and also allows teachers to benefit from building their own personal reputations.³⁸

Nonetheless, Neal (2009b) points out that, while many countries now have systems that operate like voucher systems and force schools to compete for students, no developed country with a large voucher system allows schools to compete for teachers by following different personnel policies. Systems that force schools to compete for public funding but also force all schools that receive public funding to hire, train, reward, and fire teachers according to a fixed set of personnel policies are incoherent from a design perspective. In any industry, increased competition among firms offers the possibility that the firms which remain in the market going forward will be those who have successfully adopted new and more efficient means of production. Teachers are the key input in educational production. Thus, policies that govern the hiring, training, retention, and motivating of teachers should have large impacts on the efficiency of schools. It makes no sense to promote competition among schools for students while restricting how schools may compete for teachers.

³⁸ See Hoxby (2002) for more on how competition for teachers could affect who teaches and how.

4. CONCLUSION

Current research in the economics of education devotes considerable attention to the methods that researchers use to evaluate the impacts of various innovations in public education policy. It is appropriate that researchers devote great energy to the tasks of discovering what works best and developing methods that actually help us discern what works best. However, economists should begin contributing more to debates among scholars and policy makers concerning how performance pay programs are designed before they are ever implemented and evaluated.

Most of the programs reviewed here provide some evidence that teachers responded to performance pay schemes by changing their effort allocations in some way, and in many cases, there is at least strong suggestive evidence that total teacher effort rose following the introduction of performance pay. Two of the exceptions to this rule are the bonus schemes in England and Portugal that relied on subjective assessments made by either education officials or peer teachers. Ex post, these programs appear to have been vehicles for increasing the baseline pay scale of experienced teachers without requiring improved teacher performance. Whether or not this outcome was anticipated by the political champions of these programs, the lesson taught by these programs, and a larger literature on performance pay in other organizations, is that subjective bonus schemes should not be expected to work well unless they are part of a larger incentive system that provides incentives for those who make subjective performance evaluations to make these evaluations accurately.

The POINT program also stands out as a program that generated few measurable impacts, but the lesson that POINT teaches is different. The performance standards in POINT are completely objective. However, it is not clear that these standards were set at levels that make efficient incentive provision possible. While there may be other plausible explanations for the POINT results, the simple model developed in section 2.5 highlights the possibility that POINT simply set the performance standards too high. Further, whether or not this is true in the case of POINT, the theoretical results from section 2.5 provide an important warning for those who design incentive schemes around psychometric performance targets. It is simply not true that education authorities can choose performance standards in an ad hoc manner and then experiment with different prize levels until they discover a prize level that will elicit efficient effort given their initial choice of standard. Given some performance targets, there is no prize level that would induce efficient effort.

Concerns about the choice of performance standards as well as the manipulation of performance standards can be mitigated to some extent if education authorities require that all incentive schemes involve pay for relative performance. When authorities force educators to compete for a fixed amount of reward money, well designed contests can reveal the expected level of measured performance that is associated with efficient effort levels among teachers. When officials allow competition to determine standards endogenously, they make it difficult for educators or their representatives to compromise performance standards or prevent standards from rising over time as new technologies and teaching methods make better performance possible.

Nonetheless, even in relative performance schemes, manipulation of performance metrics remains a concern. Although relative performance schemes weaken incentives for collusion among teachers, they may induce wasteful forms of competition. Educators in relative performance contests may take actions that are privately beneficial because they raise measured relative performance but socially wasteful because they crowd out teaching activities that create more lasting skills among students.³⁹ The literature suggests that educators often respond to assessment-based incentives by coaching students for specific assessment items or item formats. In fact, studies that examine scores on both high and low stakes assessments for the same population of students offer no evidence that any incentive scheme induced changes in measured performance on high-stakes assessments that even come close to fully generalizing to low stakes assessments of the same material.

Thus, it seems obvious that a key task for those who design future performance pay schemes for teachers is the creation of a series of assessments that consistently cover a well specified curriculum but vary substantially in terms of specific item content and format. Put more pointedly, the designers of assessment-based incentive schemes must take seriously the challenge of designing a series of assessments such that the best response of educators is not to coach but to teach in ways that build true mastery of the intended domain.

Many existing performance pay schemes cannot employ results from such a series of assessments because these systems are built around a particular psychometric scale, and it would typically not be possible to place results from assessments of varying formats on a common psychometric scale. However, ordinal contests like the pay for percentile scheme described in section 2.7 can employ the results from such assessments, and a commitment to ordinal contests and tests without repeated items and formats could go a long way toward eliminating the coaching and test preparation responses that appear to plague many current and previous systems.

This observation is related to the most obvious lesson generated by the material presented in section 2. Education authorities cannot reasonably expect to obtain reliable information about secular trends in performance from assessment series that are part of

³⁹ Further, Barlevy and Neal (2011) point out that although general score inflation does not benefit teachers who compete in a relative performance scheme, teachers as a group can still benefit from manipulating the dispersion of scales. If teachers can collectively pressure testing agencies to compress the distribution of performance metrics, after piece-rates have been set, the contaminated system will provide weaker incentives but pay out the same total prize money to teachers.

incentive systems. Systems that provide reliable information about secular trends in performance must involve assessments that can be properly equated over time, but the overlap in content and format that makes proper equating possible creates opportunities for the coaching behaviors that inflate scores and compromise the meaning of assessment scales. If education officials desire credible measures of secular progress, they must obtain these measures from a series of assessments that contain no stakes for educators.

Finally, because taxpayers and their representatives want schools to build noncognitive as well as cognitive skills, assessment based incentive schemes can never be more than one component of a broad system of incentives for educators. From this starting point, it is clear that assessment based incentive schemes and voucher systems should not be seen as policy substitutes but rather policies that may work well together as part of a broader system that requires schools to compete on several dimensions for access to government funds. By fostering competition among schools that rewards schools for fostering both the cognitive and non-cognitive skills of children, education authorities may create competition among schools for effective teachers that spurs innovation in the creation of new methods for screening, developing, and rewarding teachers.

ACKNOWLEDGMENTS

I thank participants at the 2009 Handbook of Economics of Education Conference in Munich for helpful feedback. I thank Gadi Barlevy, Robert Gibbons, Canice Prendergast, and Phil Reny for useful conversations and guidance. I thank Luca Maini, Eric Nielsen, Armin Rick, and Richard Olson for excellent research assistance. I thank the Searle Freedom Trust for research support. I also thank Lindy and Michael Keiser for research support through a gift to the University of Chicago's Committee on Education.

REFERENCES

- Adams, G.J., 1996. Using a Cox Regression Model to Examine Voluntary Teacher Turnover. J. Exp. Educ. 64 (3), 267–285.
- Andrabi, T., Bau, N., Das, J., Khwaja, A.I., 2010. Test Scores and Civic Values in Public and Private Schools. mimeo.
- Angrist, J., Bettinger, E., Kremer, M., 2006. Long-Term Consequences of Secondary School Vouchers: Evidence from Administrative Records in Colombia. Am. Econ. Rev. 96 (3), 847–862.
- Atkinson, A., Burgess, S., Croxson, B., Gregg, P., Propper, C., Slater, H., Wilson, D., 2009. Evaluating the Impact of Performance-related Pay for Teachers in England. Labour Econ. 16 (3), 251–261.
- Baker, G., 2002. Distortion and Risk in Optimal Incentive Contracts. J. Human Resour. 37 (4), 728-751.
- Ballou, D., Podgursky, M., 1997. Teacher Pay and Teacher Quality. W.E. Upjohn Institute for Employment Research, Kalamazoo, Michigan.
- Barlevy, G., Neal, D., 2011. Pay for Percentile. University of Chicago, forthcoming American Economic Review.
- Barrow, L., Rouse, C.E., 2009. School Vouchers and Student Achievement: Recent Evidence, Remaining Questions. Ann. Rev. Econ. 1, 17–42.

- Bhattacharya, S., Luis Guasch, J., 1988. Heterogeneity, Tournaments, and Hierarchies. J. Polit. Econ. 96 (4), 867–881.
- Briggs, D., Betebenner, D., 2009. Is Growth in Student Achievement Scale Dependent? presented at the annual meetings of the National Council on Measurement in Education, April, 2009.
- Briggs, D.C., Weeks, J.P., 2009. The Sensitivity of Value-Added Modeling to the Creation of a Vertical Score Scale. Educ. Finance Policy 4 (4), 384–414.
- Campbell, D.T., 1976. Assessing the Impact of Planned Social Change. Occasional Working Paper 8. Dartmouth College, Public Affairs Center, Hanover, N.H.
- Clotfelter, C.C., Ladd, H.F., Vigdor, J.L., Allan Diaz, R., 2004. Do School Accountability Systems Make it More Difficult for Low-Performing Schools to Attract and Retain High-Quality Teachers? J. Policy Anal. Manage 23 (2), 251–271.
- Cronin, J., Dahlin, M., Adkins, D., Gage Kingsbury, G., 2007. The Proficiency Illusion. Thomas, B. Fordham Institute, Washington, DC.
- Cunha, F., Heckman, J., 2008. Formulating, Identifying and Estimating the Technology of Cognitive and Noncognitive Skill Formation. J. Hum. Resour. 43, 739–782.
- Cullen, J., Reback, R., 2006. Tinkering Towards Accolades: School Gaming under a Performance Accountability System. In: Advances in Applied Microeconomics. Vol. 14. Elsiever, pp. 1–34.
- Dee, S.T., Keys, B.J., 2004. Does Merit Pay Reward Good Teachers? Evidence from a Randomized Experiment. J. Policy Anal. Manage. 23 (3), 471–488.
- Dixit, A., 2002. Incentives and Organizations in the Public Sector: An Interpretative Review. J. Hum. Resour. 37 (4), 696–727.
- Figlio, D.N., 2009. Voucher Outcomes. In: Mark Berends, Matthew G. Springer, Dale Ballou, Herbert J. Walberg, (Eds.), Handbook of Research on School Choice. Lawrence Erlbaum Associates/Taylor & Francis Group.
- Figlio, D.N., Winicki, J., 2005. Food for Thought: The Effect of School Accountability Plans on School Nutrition. J. Public Econ. 89, 381–394.
- Gibbons, R., 2010. Inside Organizations: Pricing, Policies, and Path Dependence. Annu. Rev. Econom. 2, 337–365.
- Gillborn, D., Youdell, D., 2000. Rationing Education: Policy, Practice, Reform, and Equity. Open University Press, Buckingham.
- Glewwe, P., 2009. Teacher Incentives in the Developing World. In: Springer, M.G. (Ed.), Performance Incentives: Their Growing Impact on American K-12 Education. Brookings.

Glewwe, P., Ilias, N., Kremer, M., 2010. Teacher Incentives. Am. Econ. J. Appl. Econ. 2 (3), 205-227.

- Goodman, S., Lesley. T., Teacher Incentive Pay and Educational Outcomes: Evidence from the New York City Bonus Program, mimeo, Columbia University.
- Hanushek, E.A., Raymond, M.E., 2005. Does School Accountability Lead to Improved Student Performance? J. Policy Anal. Manage. 24 (2), 297–327.
- Hanushek, E.A., Rivkin, S.G., 2006. Teacher Quality. In: Hanushek, E.A., Welch, F. (Eds.), Handbook of the Economics of Education. Vol. 2. Elsevier, B. V.
- Holmstrom, B., 1982. Moral Hazard in Teams. Bell J. Econ. 13 (2), 324-340.
- Holmstrom, B., Milgrom, P., 1991. Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership and Job Design. J. Law Econ. Organiz. 7, 24–52.
- Hoxby, C.M., 2002. Would School Choice Change the Teaching Profession. J. Hum. Resour. 37 (4), 846-891.
- Hudson, S. 2010. The Effects of Performance-Based Teacher Pay on Student Achievement. SIEPR Discussion Paper No. 09-023.
- Jacob, B., 2005. Accountability Incentives and Behavior: The Impact of High Stakes Testing in the Chicago Public Schools. J. Public Econ. 89 (5), 761–796.
- Jacob, B., Levitt, S., 2003. Rotten Apples: An Investigation of the Prevalence and Predictors of Teacher Cheating. Q. J. Econ. 118 (3), 843–877.
- Jovanovic, B., 1979. Job Matching and the Theory of Turnover. J. Polit. Econ. 87 (5), 972–990.
- Klein, S., Hamilton, L., McCaffrey, D., Stecher, B., 2000. What Do Test Scores in Texas Tell Us? Rand Issue Paper 202.

- Koretz, D.M., Barron, S.I., 1998. The Validity of Gains on the Kentucky Instructional Results Information System (KIRIS). RAND, Santa Monica.
- Koretz, D.M., 2002. Limitations in the Use of Achievement Tests as Measures of Educators' Productivity. J. Hum. Resour. 37 (4), 752–777.
- Ladd, H.F., 1999. The Dallas School Accountability and Incentive Program: an Evaluation of its Impact on Student Outcomes. Econ. Educ. Rev. 18 (1), 1–16.
- Lavy, V., 2002. Evaluating the Effect of Teacher's Group Performance Incentives on Pupil Achievement. J. Polit. Econ. 110 (6), 1286–1317.
- Lavy, V., 2009. Performance Pay and Teacher's Effort, Productivity, and Grading Ethics. Am Econ. Rev. 99 (5), 1979–2021.
- Lazear, E., Rosen, S., 1981. Rank Order Tournaments as Optimum Labor Contracts. J. Polit. Econ. 89 (5), 841–864.
- Martins, S.P., 2009. Individual Teacher Incentives, Student Achievement and Grade Inflation. IZA Discussion Paper No. 4051.
- Muralidharan, K., Sundararaman, V., 2010. Teacher Performance Pay: Experimental Evidence from India. UCSD, mimeo.
- Neal, D., 2009a. Designing Incentive Systems for Educators. In: Springer, M. (Ed.), Performance Incentives: Their Growing Impact on American K-12 Education. Brookings.
- Neal, D., 2009b. The Role of Private Schools in Education Markets. In: Berends, M., Springer, M.G., Ballou, D., Walberg, H.J. (Eds.), Handbook of Research on School Choice. Lawrence Erlbaum Associates/Taylor & Francis Group.
- Neal, D., 2010. Aiming for Efficiency Rather than Proficiency. J. Econ. Perspect. 24 (3), 119-131.
- Neal, D., Schanzenbach, D.W., 2010. Left Behind by Design: Proficiency Counts and Test-Based Accountability. Rev. Econ. Stat. 92 (2), 263–283.
- O'Keeffe, M., Viscusi, K.W., Zeckhauser, R.J., 2019. Economic Contest: Comparative Reward Schemes. J. Labor Econ. 2 (1), 27–56.
- Prendergast, C., 1999. The Provision of Incentives in Firms. J. Econ. Lit. 37 (1), 7-63.
- Reardon, S.F., Raudenbush, S.W., 2009. Assumptions of Value-Added Models for Estimating School Effects. Educ. Finance Policy 4 (4), 492–519.
- Rivkin, S.G., Hanushek, E.A., Kain, J.F., 2005. Teachers, Schools, and Academic Achievement. Econometrica 73 (2), 417–458.
- Rockoff, J.E., 2004. The Impact of Individual Teachers on Student Achievement: Evidence from Panel Data. Am. Econ. Rev. 94 (2), 247–252.
- Rockoff, J.E., Staiger, D.O., 2010. Searching for Effective Teachers with Imperfect Information. J. Econ. Perspect. 24 (3), 97–118.
- Rothstein, J., 2010. Teacher Quality in Educational Production: Tracking, Decay, and Student Achievement. Q. J. Econ. 125 (1), 175–214.
- Shepard, L.A., Dougherty, K.C., 1991. Paper presented at the annual meeting of the American Educational Research Association, Chicago, IL. Effects of High Stakes Testing on Instruction.
- Springer, M.G., Ballou, D., Hamilton, L., Le, V., Lockwood, J.R., McCafrey, D., Pepper, M., Stecher, B., 2010. Teacher Pay for Performance: Experimental Evidence from the Project on Incentives in Teaching. National Center on Performance Incentives at Vanderbilt University, Nashville, TN.
- Stecher, B.M., 2002. Consequences of Large-Scale, High-Stakes Testing on School and Classroom Practice. In: Hamilton, L.S., Stecher, M.B., Klein, S.P. (Eds.), Making Sense of Test-Based Accountability in Education. National Science Foundation.
- Vigdor, J.L., 2009. Teacher Salary Bonuses in North Carolina. In: Springer, M. (Ed.), Performance Incentives: Their Growing Impact on American K-12 Education. Brookings.
- Winters, M., Greene, J.P., Ritter, G., Marsh, R., 2008. The Effect of Performance-Pay in Little Rock, Arkansas on Student Achievement. National Center on Performance Incentives, Peabody College of Vanderbilt University, Working Paper 2008-02.
- Wragg, E., Haynes, G., Wragg, C., Chamberlin, R., 2001. Performance Related Pay: The Views and Experiences of 1000 Primary and Secondary Headteachers. Teachers' Incentives Pay Project Occasional Paper 1. School of Education, University of Exeter.



Educational Vouchers in International Contexts

Eric Bettinger

Stanford University

Contents

1.	The Context for Vouchers Abroad	555
	1.1. Link between Public and Private Sectors	556
	1.2. Utilization of Private Sector	556
	1.3. Geographic Links between Public School Attendance	556
	1.4. Unionization and Exit	556
	1.5. Finance	557
2.	Voucher Programs in Colombia, Chile, and Sweden	557
	2.1. Colombia's Voucher Program	557
	2.2. Chilean Voucher Program	558
	2.3. Sweden's Voucher Program	559
3.	Effects of Vouchers on Voucher Users	560
	3.1. Evidence from Colombia	560
	3.2. Evidence from Other Countries	564
4.	Effects of Vouchers on Overall School System	565
	4.1. Evidence from Chile	566
	4.2. Evidence from Sweden	567
	4.3. Evidence from Other Countries	568
5.	Conclusion and Suggested Directions for Future Study	568
References		570

Abstract

International evidence on school choice largely focuses on educational voucher or voucher-like systems. The research to date primarily focuses on two complementary questions: what are the effects of school choice on students who exercise school choice? and what are the effects of school choice on the overall system that allows choice? In this chapter, we review the educational voucher focusing on these two research questions. We primarily focus on educational voucher programs in Chile, Colombia, and Sweden. We discuss each of these programs and the accompanying literature in depth. We briefly discuss research from other countries, especially ongoing research in India, which may provide key insights into voucher and school-choice debates. Although there are a number of similarities between research on school choice in the United States and abroad, research on school choice abroad presents an entirely different set of political circumstances, institutions, and funding schemes. We discuss these issues and their impact on generalizeability of international research. We also recommend future directions for voucher research, particularly, in identifying key components of voucher systems that have led to the observed effects to date.

Keywords

Voucher School Choice International Education

International evidence on school choice largely focuses on educational voucher or voucher-like systems. Worldwide, voucher or voucher-like programs exist in a wide array of developed and developing countries, including programs in Sweden, India, the Netherlands, Chile, Belize, Japan, Canada, Colombia, and Poland. Chile's educational voucher program is the largest voucher program in the world, and Colombia's PACES voucher program serviced almost 144,000 students in its peak. Voucher-like programs in Sweden and the Netherlands provide a pseudo market where public and private schools vie for the top students.

Research on educational choice and vouchers outside the United States typically focuses on two complimentary research questions. The first question focuses on the effects of school choice on students who exercise school choice. For example, in a voucher program, this type of research focuses on how educational vouchers affect the outcomes of students who use the vouchers. Are the users better off than they would have been in the absence of the voucher?

The main focus of this first research question is on the productivity of different schooling options. Since vouchers enable students to attend private schools, scholars have often linked the voucher literature to the literature on the effects of private schooling. For example, Neal (1998) focuses extensively on the relationship between private schooling and educational vouchers in his review of the literature on the effects of private schooling. Similarly, Rouse (1998) measured the effect of vouchers on Milwaukee students who use the voucher. She used voucher assignment as an instrument for the likelihood that students attend private school and reported estimates of the effects of private schooling.

However, the effects of school choice on those who exercise choice need not be solely about the productivity of different schooling options. Bettinger, Kremer, and Saavedra (2010) identified several plausible channels by which the Colombia voucher program has affected students' educational outcomes. These channels include such channels as income effects, changes in peers, changes in incentives, and changes in school resources.

The key empirical challenge in measuring the effects of school choice on those who exercise school choice is that students who exercise school choice often differ systematically from those students who do not. For example, literature on the effects of charter schools in the United States (or similarly contract schools in Colombia) often compares students who attend charter schools with students who do not (e.g., Bifulco and Ladd (2006); Bonilla (2011)). Hoxby et al. (2009) has criticized this research since students who attend charter schools may systematically differ from other students.¹ To overcome these biases, researchers must devise empirical strategies that exploit exogenous variation in the degree to which students exercise choice in order to estimate the effects.

The second key research question in the research on school choice aims at examining the effects of school choice on the overall system. For example, in a voucher program, this research focuses on the improvement of outcomes for all students in the system as a result of the increased opportunities for attendance. Does the emergence of new options improve the overall menu of schools and the students' outcomes within those schools?

The idea that the available set of schools may influence the overall productivity of the system and variations of the theme have long been discussed in economics (e.g., Friedman (1962)). If education was the only public good in the economy, Tiebout's model of residential choice (Tiebout (1956)) would be a model of school choice, and the efficient allocation in the economy would be the one in which students choose an educational provider among a large number of different and competing educational providers.

The link between residential choice and school choice is the key empirical problem for researchers who want to isolate the impact of school choice. Residential choice and school options may be endogenously determined making it difficult to estimate causal relationships. For example, do families with high-ability children choose to live in areas with many school options and hence more school choice? Or does the wider menu of school choice lead to better educational systems and hence high-achieving students? Hence, to identify the impacts of choice, researchers must identify variation in the menu of school choice, which is uncorrelated with students' underlying abilities.

Often, it is difficult to find variation given that many communities see little change in the schooling options available over time. We often lack data on how these communities formed and how educational institutions have evolved. Hence, both in the United States and abroad, researchers largely focus on cases in which the menu of choices available to a set of students has abruptly changed. There have been many "abrupt" changes in schooling options. The emergence of charter and contract schools in the United States, the United Kingdom, and elsewhere has modified the educational markets and introduced changes to the set of schools in local markets. However, generally speaking, the educational policy most responsible for changing the primary and secondary markets outside the United States is educational vouchers. In terms of evidence from international

¹ Hoxby et al. (2009) uses random assignment of charter school openings to create treatment and control groups of students. This random assignment occurs in cases in which schools are oversubscribed. Comparing treatment and control groups that are randomly assigned can create internally valid estimates of the effects of charter schools on students who apply to oversubscribed charter schools. This may not have external validity across all charter schools (Raymond 2009, Reardon 2009).

school choice programs, the most important educational research to date focuses on educational vouchers.

The three educational voucher programs outside the United States that have received the most attention from academics are the voucher programs in Chile, Colombia, and Sweden. There are a number of other voucher programs located throughout the world (e.g., targeted voucher programs in Puerto Rico and Bangladesh); however, there is little evidence on these programs to date. We briefly discuss research from other countries, especially ongoing research in India, which may provide key insights into voucher and school choice debates. Given that most of the evidence on vouchers focuses on Colombia, Chile, and Sweden, much of our discussion focuses on these three programs.

There are a number of similarities and differences between research on school choice in the United States and abroad. In terms of similarities, the economic principles affecting schooling markets abroad are similar to those in the United States. Core economic principles such as competition, information accessibility, marginal costs, marginal benefits, public goods, externalities, and production functions dictate equilibrium in the primary and secondary school markets. Debates on school choice abroad shed light on all of these economic principles; hence, research on school choice abroad builds our knowledge more generally about the school choice and the economics of education.

On the other hand, research on school choice abroad presents an entirely different set of political circumstances, institutions, and funding schemes. The fundamental differences in the schooling markets abroad and United States may limit the generalizeability of this research to situations in the United States, and often, the educational markets are defined such that school choice plays a much more important role than in the United States. In addition, the long-run potential of school choice to influence educational quality may be greater outside the United States (and in particular, in developing countries) making research on school choice important not only for academic reasons but crucial for policy making.

For the interested researcher, there are a number of other reviews of educational vouchers each of which sheds light on features of voucher and voucher-like systems. Woessman (2007) largely focused on the role of autonomy and accountability in the development of schooling markets. West (1996) and Zimmer and Bettinger (2007, 2008) focused more extensively on inventorying existing evidence in the United States and abroad.

This review is organized as follows. First, we discuss some of the key differences between the setting in developing countries and the United States. We do this so that the reader can keep in mind the generalizeability of research abroad on school choice issues inside the United States. Second, we discuss the origin and structure of the Colombian, Chilean, and Swedish voucher systems. Third, we present evidence focusing on the effects of vouchers on voucher recipients. Fourth, we present evidence on the effects of vouchers on the overall educational system. Finally, we suggest some additional lines of research that might shed additional light on school choice abroad.

1. THE CONTEXT FOR VOUCHERS ABROAD

In the United States, the debates around school choice center on both equity and efficiency. From an equity perspective, voucher advocates claim that vouchers provide more and perhaps better options to students stuck in failing schools. From an efficiency perspective, voucher advocates claim that vouchers may improve the overall effectiveness of the educational system through competitive effects.

Outside the United States, there may be several motivations for voucher policies. For example, in Colombia's PACES voucher program, the primary motivation for adopting the policy was rapid overcrowding in Colombia's secondary schools. In 1991, many public high schools hosted multiple sessions of high school per day. One high school would use the facility in the early morning, another would use the facility in the afternoon, and yet another in the evening. Colombia lawmakers saw school vouchers as a means to increase capacity without significant capital expense. In Sweden, by contrast, decentralization of educational decision making allowed more freedom for municipalities to allocate funds between public and private schools according to local needs (Klitgaard (2007)).

Competitive educational markets and equality of opportunity are additional motivations frequently cited in Chile. The Chilean program was greatly influenced by Friedman (1962), and the primary motivations were to improve the overall quality of the system through competition and to provide greater access to schooling among the poor.

The structure of voucher programs and schooling systems also varies significantly from program to program. For example, the value of the voucher also varies dramatically across voucher programs. In Sweden, the subsidy started at 85% and was eventually dropped to 75% of the average public school funding level. In addition, Swedish private schools are heavily regulated by municipalities, and local governments have banned private schools from charging additional fees on top of the voucher. In Colombia, the voucher program originally provided subsidies that were equivalent to private school tuition at "voucher schools," yet the Colombian government allowed inflation to erode the value of the voucher. In the last year of the voucher, it covered about 50% of overall tuition, yet private schools experience much greater autonomy than schools in the public sector.

The variance in motivations and structures of voucher programs makes it so that research across programs and systems must be carefully compared and assimilated. In other words, there is no "pure" voucher system, and research has yet to distinguish what design features of educational voucher systems has led to the observed effects to date. In addition, there are also a few key differences in educational markets abroad that greatly shape the types of research and the potential findings of such research. Although this list is not exhaustive, it provides some basic reasons why voucher research in one country may not be generalizeable in other settings, particularly in the United States:

1.1. Link between Public and Private Sectors

In the United States, the definition of a public school is clearly delineated.² Public schools must be managed and run by government entities. Outside the United States, this division is less clear. Government funds often directly subsidize the private sector. Except in a few voucher locations (e.g., Milwaukee, Cleveland), it is rare in the United States to see public dollars used for religious private schools. In other countries, the link between religious private schools and the public sector is much more prevalent.

1.2. Utilization of Private Sector

In the United States, about 11% of students attend private schools. Outside the United States, this number is much higher (Psacharapoulos (1987)). Private school plays a much more significant role outside the United States than it does inside the United States.

1.3. Geographic Links between Public School Attendance

Worldwide, families tend to show a preference for the school located near them. In the United States, the public school nearest to a family is also the school where parents are obligated to send their child unless they choose not to use a traditional public school. Outside the United States, this geographic link is less powerful. Parents can often choose a public school among a number of options that need not be close.

1.4. Unionization and Exit

Teachers' unions play a strong role in the adoption or failure to adopt voucher programs. They also can play an important part in negotiating the terms by which schools can compete. For example, in Chile, the powerful teachers' union was able to negotiate job security for teachers at schools that have experienced substantial attrition due to voucher competition. Even if all of the students exit the school, the teachers still have a guaranteed paycheck from the government, and the school will remain open. Such an arrangement, negotiated by a powerful union, changes the nature and consequences of competition between the public and private sectors.

² Charter schools may be the only "gray" area between private and public schools. Often, these public schools resemble private schools in their management.

1.5. Finance

The United States has a largely decentralized educational system. Tax collection and expenditure for schools are largely based on the school district level. Outside the United States, especially in school systems where vouchers are present, there is a mix between the degree of local autonomy and centralization. Typically, taxes and revenue for schools are collected at a centralized level, and then local units receive block grants that are used to finance schools. The amount of these block grants is often tied directly to enrollment (e.g., Sweden, Chile).

2. VOUCHER PROGRAMS IN COLOMBIA, CHILE, AND SWEDEN

By 2010, there were at least 15 countries that had adopted educational voucher programs. These voucher programs ranged from small programs in Bangladesh, Belize, Canada, Guatemala, Japan, Lesotho, Poland, and Puerto Rico to large programs in Colombia, Chile, the Netherlands, and Sweden. Significant research has also been conducted in the voucher-like programs in New Zealand, the United Kingdom, and the United States.

We focus our attention on the Colombian, Chilean, and Swedish voucher systems. These three educational voucher programs account for most of the evidence on international educational voucher systems. To understand the evidence and the types of research questions asked in these sites, we should understand the details of the programs.

2.1. Colombia's Voucher Program

Colombia's Programa de Ampliación de Cobertura de la Educación Secundaria (PACES) provided vouchers to students in Colombia. From 1992 to 1997, PACES provided educational vouchers to over 144,000 students. Colombia's voucher program began in 1991 as part of a larger effort to decentralize public services and to expand private sector provision (King, Rawlings, Gutierrez, Pardo, and Torres (1997)). The Colombian government advertised the program in local newspapers and through radio ads, and the program immediately became popular. In order to improve enrollment rates among the poorest families in Colombia, PACES targeted low-income families (King, Orazem, and Wohlgemuth (1998)). To qualify for the voucher, parents had to present a utility bill proving that they lived in one of the two lowest socioeconomic strata (out of six possible strata). Research by Morales-Cobo (1993) suggests that this targeting was effective in Bogotá.

To be eligible for the voucher, children had to be entering the sixth grade, the start of Colombian secondary schools, and be under the age of 16. Children were also only eligible if they had been attending public school in the previous year and had already arranged admission at a participating private secondary school. Colombia has a threetier system of schools with public schools, normal private schools, and elite private schools. Low-income students typically do not attend elite private schools, and the elite private schools did not participate in the voucher program. King, Rawlings, Gutierrez, Pardo, and Torres (1997) investigated differences between public secondary schools and participating private schools. They found that pupil-teacher ratios, test scores, and access to technology were similar across schools.

The voucher allowed students to attend private schools. As we mentioned above, the voucher paid all of students' private school tuitions in the initial years; however, over time, the value of the voucher did not keep pace with inflation. The voucher program started in students' sixth grade years and continued through the end of high school (eleventh grade).

Colombia's voucher program was oversubscribed from the beginning. Local educational leaders used lotteries to assign vouchers. These lotteries also facilitate research on the effects of the vouchers. The lotteries can aid researchers in creating a comparable control group to students who participated in the voucher program. We discuss the importance of the lottery in the next section.

As long as students were promoted at the end of a grade, they could automatically renew their voucher through eleventh grade, the end of Colombian high school. Students failing a grade were supposed to be dropped from the PACES program. Calderon (1996) shows that about 77% of recipients renewed their vouchers. In addition, the rules of the voucher allowed students to transfer to other schools with the voucher; however, few students who transferred schools kept their vouchers.

2.2. Chilean Voucher Program

Chile's voucher program began in 1980 when the Chilean government, in consultation with Milton Friedman, initiated a series of educational reforms designed to decentralize and privatize education (Rounds (1996)). Chile's program provides tuition subsidies to private schools. As in the case of Colombia, Chile's schooling system has three tiers, and the highest, most elite tier of private schools did not accept the voucher. However, with the allowance of new entrants into the marketplace, Chile has had a significant increase in the number of new private schools that accept the voucher. Private school attendance increased from 15 to 42% from 1981 to 2005 (Bravo, Mukhopadhyay, and Todd (2009)).

The voucher program was designed to reward schools for enrollment. Prior to the reform, the federal government provided direct funding to both private and public schools; however, the subsidies were insensitive to enrollments. Private schools generally received significantly less than public schools. Once the reform was established, both public and private schools received equal subsidies that were directly tied to student enrollment (Hsieh and Urquiola (2003)). Besides introducing enrollment-based funding, the federal government also decentralized schooling allowing individuals schools to set their budgets and to manage their curricula.

The Colombian voucher program targeted poor students. The Chilean system, by contrast, allowed for universal take-up. In addition, voucher schools could exercise selective admission. Rounds (1996) finds that the poorest families were less likely to attend voucher schools as a result of selective admissions. The biggest change in the program in recent years was the increase in the size of the subsidies to schools that admit low-income families. Beginning in 2008, schools received additional monies for admitting low-income students.

The fact that Chile's program has run for almost 30 years might provide the most concrete evidence on the long-run effects of a large-scale educational voucher program.

2.3. Sweden's Voucher Program

Sweden's voucher program began in 1992. Up until that point, less than 1% of all students were attending private schools (Klitgaard (2008)). The program was a seeming departure from the ruling Social Democratic Party's ideological core issue that there should be one school system for all students; however, increasing public pressure for school choice and decentralization led to the adoption of Sweden's voucher program (Bunar (2010)). Klitgaard (2008) explains that the public sector reforms in Sweden aimed to grant more freedom of choice between alternative service providers, influence the quality of the public services, and increase the economic efficiency within the public sector.

Sweden's voucher reform really consisted of three complementary reforms. First, financial and managerial reforms were decentralized. Local districts and boards began to govern schools with significant input from parents. Local municipalities also had some responsibility for revenue. Second, educational vouchers allowed students to transfer both between public and private schools and between public schools. The value of the voucher for a private school was equivalent to the average per-pupil expenditure in the public sector, and private schools were unable to increase tuition above the level of the voucher. Finally, within a couple of years of the 1992 reform, the national government centralized guidelines for new schools with greater emphasis on evaluation and inspection (Bunar (2010)). New private schools had to adopt the national curriculum and be subject to the oversight of the local municipalities.

Prior to the reform, students attended neighborhood schools, but as a result of the reform, students could attend any public or private school in the municipality. Schools had a legislated priority list for determining the students they could admit if oversub-scribed. Proximity to the school is the primary criterion. Because neighborhood schools remain the most popular choice and that choice is limited by the availability of new slots, choice has not been very common in some municipalities; however, this reform led to a significant increase in the quantity of private schools and in the share of pupils attending them (Bohlmark and Lindahl (2007)). Although most of the literature in Sweden focuses on the emerging private sector, little attention has been paid to the

degree of school choice between public schools. Soderstrom and Uusitalo (2010), for example, argues that for most students, the choice between different public schools is far more important than the choice between public and private schools.

3. EFFECTS OF VOUCHERS ON VOUCHER USERS

As we mentioned above, there are two complementary questions on which the research on school choice focuses. The first question focuses on the effects of the voucher on voucher users. The central question asked in this research is whether students are better off using the voucher than they would have been in the absence of the voucher.

Research on the effects of vouchers can potentially shed light on a number of persistent questions in the economics of education. For example, families offered a voucher can elect to stay in the current schools in which they are enrolled or transfer to new schools. The decision to transfer to new schools tells us more about parents' overall objective functions. It helps us to understand to what extent academic achievement may be valued above other schools' characteristics including specific services and resources, peers, safety, and so on. Other economists have argued that measuring the effects of vouchers on users can shed light on the nature of peer effects (e.g., Bettinger, Kremer, and Saavedra (2010)) and school production (e.g., King, Orazem, and Wolgemuth (1998)).

To measure the effects of vouchers on students using the voucher, researchers exploit the randomization present in many voucher assignment mechanisms. Randomization is compelling in that it allows for an unbiased estimate of the counterfactual outcome for students chosen to receive a voucher. As we mentioned in the previous section, this is a key feature of the Colombian voucher program, and hence, the most influential studies in the international school choice literature to date focus on this voucher program.³

3.1. Evidence from Colombia

There have been three influential papers focusing on the effects of the Colombian voucher program on student achievement. The first, Angrist, Bettinger, Bloom, Kremer, and King (2002), exploits the local voucher lotteries as a means of identifying the effects of the Colombian vouchers on students who use them.

Angrist, Bettinger, Bloom, Kremer, and King (2002) relies on survey data from the cohort of voucher applicants from the 1995 Bogotá lottery. These students applied for the voucher before the start of secondary school, and Angrist et al. interviewed these

³ Several prominent studies from the United States suggest modest if any effects of educational vouchers on voucher users. Krueger and Zhu (2004), for example, suggests that vouchers increase test scores by about 2.08 percentile points (t = 1.24).

students 3 years after their initial voucher application. Angrist et al. then compared the lottery winners and lottery losers to identify the effects of the voucher.

They find that voucher students had completed about 0.1 years of schools more than their peers and had test scores about 0.2 standard deviations higher. This increased attainment was also manifest in significant reductions in the likelihood that students repeated grades within 3 years of the voucher lottery. Students having five percentage points were less likely to have repeated a grade within the 3 years of the voucher lottery. Angrist, Bettinger, Bloom, Kremer, and King (2002) also find that the incidence of child labor and teen age marriage was lower as a result of the educational voucher.

Although the paper by Angrist et al. paints a very optimistic case for the effects of the voucher, there are a few weaknesses of the paper to be mentioned. For example, their survey response rate was very low relative to other studies. Angrist et al. interviewed 2985 students who had originally applied to the voucher program. They were only able to reach 54% of these individuals. Their argument that the response rates were balanced across lottery winners and losers suggests that there should be no bias to estimated results. In addition, one of the key results might not have represented a "good" outcome. Students who won the voucher lottery had completed more schooling than their counterparts who lost the lottery, but they were not more likely to be attending school 3 years after the voucher lottery. The decreased grade completion could explain these differences; however, one cynical interpretation of the results was that schools might have promoted students simply to keep the voucher monies coming to the school. Students lost the vouchers if they did not successfully complete a school year. Administrators may have had incentives to promote more voucher winners from year to year than they did voucher lottery losers. Administrators may have promoted students in order to keep the voucher monies flowing to the school. Finally, another concern in the Colombian voucher lottery was whether any of the observed effects after 3 years would translate into long-term differences in outcomes. It was not clear whether the observed differences would lead to long-run differences in educational outcomes. Given these concerns, Angrist, Bettinger, and Kremer (2006) attempted to measure more long-run effects of the educational vouchers.

Angrist, Bettinger, and Kremer (2006) attempted to follow up with students to investigate outcomes in students' high-school careers. They investigated whether voucher students were more likely to take the ICFES (Instituto Colombiano Para El Fomento De La Educación Superior), Colombia's primary college entrance exam. Since 90% of students who graduate high school take this exam (although only 75% of these attend college), Angrist, Bettinger, and Kremer (2006) interprets students' taking the ICFES exam as an indicator of high school attendance.

This second voucher study improved on the first in three notable ways. First, because the authors used administrative data representing all students who entered the

voucher lottery, they did not have the attrition in the sample that the Angrist, Bettinger, Bloom, Kremer, and King (2002) paper had. The only students who attrited from the sample were students whose identification numbers were invalid so that they could not be matched to administrative records. Only 12% of the sample had invalid identification numbers, and this was balanced across voucher lottery winners and losers, and once the analysis was restricted to individuals with valid ages, only 3% of students had invalid ages.

The second improvement was the ability to focus on long-run learning outcomes. The prior study was only able to present test score outcomes for the sample of students who attended a special testing session. The second study, by contrast, was able to examine the college entrance exams of the voucher lottery participants. Finally, Angrist, Bettinger, and Kremer (2006) were also able to examine long-run outcomes of the students.

In comparing voucher lottery winners and losers using these administrative data, Angrist, Bettinger, and Kremer (2006) finds that voucher lottery winners were 20% more likely to graduate from high school. This is robust over a number of different specifications. Given that the voucher winners were more likely to have taken the college entrance exam, comparisons of students' entrance exam scores may be difficult to interpret. For example, if the marginal student affected by the voucher was from the bottom of the test score distribution, then their exam score would lower the average of voucher winners, and given that their counterpart voucher lottery losers did not take the exam, the average score of voucher lottery losers would not be affected.

To overcome these issues, Angrist, Bettinger, and Kremer (2006) used both parametric and nonparametric bounding strategies. In the parametric strategies, they used censored models to compare voucher lottery winners and losers. Their nonparametric strategies examined different scenarios that could lead to upward or downward biases in the voucher comparisons. They find that voucher lottery winners achieved higher test scores on the ICFES exam. Their estimated effects on reading and math scores range between a lower bound of 0.09–0.13 standard deviations and an upper bound of 0.50 standard deviations.

Although the papers by Angrist, Bettinger, Bloom, Kremer, and King (2002) and Angrist, Bettinger, and Kremer (2006) suggested strong evidence of the overall effect of educational vouchers, they did not shed light on the underlying mechanisms through which the voucher effects emerged. As we mentioned above, no voucher programs are perfectly identical, and for formulating policy recommendations, it would be important to know how these voucher effects arise.

For example, suppose that private schools are better than public schools. If this is the case, then one might expect voucher winners to have better outcomes since the voucher winners were more likely to attend private schools. However, research by King, Rawlings, Gutierrez, Pardo, and Torres (1997) shows the private schools that voucher

winners attended were very similar to the public schools they would have attended in the absence of the voucher. In addition, recent research by Bettinger, Kremer, and Saavedra (2010) shows that a subset of the voucher lottery winners, who applied to vocational schools, actually attended schools that had worse academic outcomes, more behavioral problems, and fewer resources than the academic schools that the voucher losers attended. Nonetheless, the voucher lottery winners in this group had better academic outcomes.

The improved outcomes among these vocational lottery applicants could have resulted because of changes in student incentives rather than changes in schooling. Students kept their voucher only if they successfully promoted to the next grade (Angrist, Bettinger, Bloom, Kremer, and King (2002)). Repeating a grade is fairly common in Colombia, and considering the behavior of voucher lottery losers, voucher winners who repeated a grade would likely have had to enter the workforce prematurely. The prospect of keeping the voucher and avoiding the labor market may have been enough of an incentive to encourage voucher winners to work harder than other students.

There are other plausible mechanisms. For example, given that 94% of voucher lottery losers sent their children to voucher schools in the initial year of the voucher, the voucher may have only represented an income shock.

Without knowing the mechanism, it is difficult to know whether the results will generalize to other voucher programs. If the mechanism is private schooling and the quality of private schooling is better than public schools in other locations, then vouchers may have positive effects in other settings. However, a peculiar feature of Colombia's program is the fact that retention of students' scholarships depends on students' promotion from grade to grade. If this changes students' incentives and if this is the channel by which vouchers affect students' outcomes, then the voucher effect seen in the Colombian program may not extend to other voucher programs that lack a similar feature. Given our lack of knowledge as to the correct mechanism, the identification of voucher mechanisms is a place where future voucher research could greatly improve on existing research.

Finally, another theme in the voucher literature is whether the voucher leads to stratification. In Colombia, the evidence on the impacts of stratification is less developed and inconclusive. Colombia's educational vouchers targeted low-income families living in the poorest neighborhoods, and Ribero and Tenjo (1997) report that the targeting of vouchers was largely effective in reaching this population. However, applicants to Colombia's voucher program were not a random sample of families from these neighborhoods. Voucher applicants came from families with higher educational levels than other families in the same neighborhoods (Angrist, Bettinger, Bloom, Kremer, and King (2002)).

The Colombian voucher program was instituted to increase access to students. Information and access to the vouchers led to some stratification by income and education; however, no systematic research has evaluated whether Colombia's voucher program has actually improved access to education, and as a result, it is unclear whether the voucher program increased overall enrollments among the most disadvantaged families.

3.2. Evidence from Other Countries

There are three reasons why research on the Chilean voucher program has been less conclusive than the research in the Colombian voucher program. First, Chilean private schools can use selective admissions in accepting voucher students. As a result, Chilean voucher schools could conceivably admit students with better academic qualifications. Second, given that the program has existed for over 30 years and that there has been significant entry of new private schools, it is quite difficult to identify counterfactual outcomes of students who are attending schools that did not exist prior to the reform. Third, given that the voucher program could have impacted the overall quality of the system, the "control" group in any regression specification may have also been influenced by the vouchers.

Some of the early evidence suggested that voucher schools modestly outperformed public schools. This finding was common among many papers (e.g., Sappelli and Vial (2002); Bravo, Contreras and Sanhueza (1999); McEwan and Carnoy (2000)) but was sensitive to the types of controls included in the model, the specific municipalities included in the sample, and the statistical methods used. McEwan (2001), for example, found that Catholic voucher schools tended to be more effective and productive than other schools. McEwan uses the 1997 *Sistema de Medición de Calidad de la Educación* (SIMCE) to estimate the differences between similar students attending public and various types of voucher schools. The effects are somewhat sensitive to specification; for example, the effects of Catholic schools on math scores are 0.257 standard deviations (with a standard error of 0.041) when McEwan controls for individual and peer socioeconomic characteristics. When McEwan includes controls for selectivity as well, the effect drops to 0.103 standard deviations (with a standard error of 0.185).

McEwan and Carnoy (2000) also use the SIMCE to assess the impact of voucher schools. They show that academic achievement is slightly lower (about 0.07 standard deviations) in nonreligious voucher schools, but considering that these voucher schools have about 13% less funding than public schools, McEwan and Carnoy suggest that they may be more cost effective than public schools.

In Sweden, students in independent schools appear to have slightly higher average grades than their counterparts in public schools although the gap has diminished lately. However, it is unclear if this reflects a causal effect of attending a voucher school, as students in independent schools in general come more from families with academic backgrounds (Bunar (2010)). Moreover, the sorting of students across voucher schools and public schools has created additional separation between native Swedes and second-generation immigrants (Bohlmark and Lindahl (2007)).

Bohlmark and Lindahl (2007) use national registry data in Sweden to track the outcomes of all students. They decompose changes arising from school choice into those effects caused by exercising choice and those effects on the system caused by choice. To estimate the effects of choice on families who exercise choice, they look at differences in educational outcomes within siblings, which are correlated with changes in schooling options. However, this effect is small, and Bohlmark and Lindahl find that the individual gain from attending a private school is roughly one-tenth of the size of their estimated impacts on the overall system. They estimate that a 1% increase in private schooling leads to roughly a 0.1 percentile rank increase in academic outcomes.

4. EFFECTS OF VOUCHERS ON OVERALL SCHOOL SYSTEM

The second question that has been extensively evaluated in the educational choice literature is whether school choice has improved the outcomes of the entire educational system. This research question is arguably more important than the former in systems in which significant fractions of students attend both public and voucher schools.

Research on the effects of vouchers on entire systems can potentially shed light on a number of significant questions in the economics of education. For example, economists have long posited that educational vouchers can increase competition between all schools in an educational system. The increased competition may reveal itself as increased cost-effectiveness in terms of either or both reduced costs and improved student achievement. In addition, many economists have questioned the extent to which student resources and achievement are linked. Since the marginal costs of attending students are generally smaller than the average costs, schools that "lose" enrollments through increased competition experience dramatic differences in their financial resources. Changes in their expenditure and outcomes may shed light on the relationship of these variables.

Because we generally do not randomly assign vouchers to different educational systems, it is difficult to identify the effects of vouchers on educational systems. Today, most research attempts to identify variation across geographic areas in the amount of competition at a point either in time or over time. In the TIMSS test data, for example, correlation between the degree of public funding going to private schooling and student test scores is positive (Woessman (2007)). Test scores increase by 9–10 points in math and science as the share of private schooling increases by 14 percentage points (Woessman (2003)). In the international literature linking school choice and its impact

on the overall educational outcomes of students in a country, the most influential studies to date focus on the Chilean and Swedish voucher programs.⁴

4.1. Evidence from Chile

There are a number of papers that analyse the effects of Chile's voucher program. One of the primary difficulties of research on the Chilean voucher system is its impact on the composition of both public and private schools. For example, Hsieh and Urquiola (2003) show that after the reform began in 1981, public school enrollment dropped, yet as enrollments dropped, the average income of families in public school and the average achievement of students in public schools seemed to fall. Hsieh and Urquiola (2003) argue that public school families who were relatively wealthy transferred to the private sector following the reform. They suggest that the finding that voucher schools were more effective than public schools arises from the shift in the types of families at both private and public schools.

McEwan and Carnoy (2000) and Carnoy (1998) showed similar evidence that parents in voucher schools have higher incomes and greater levels of education than parents in public schools. In addition, in the early 1990s, many voucher schools began charging tuition in addition to the voucher, and the difference in parents' incomes and education levels between these tuition-charging voucher schools and the other voucher schools was significant (Anand, Mizala, and Repetto (2006)).

Even if vouchers increase sorting across voucher and nonvoucher schools, this sorting could still lead to improved academic achievement for voucher users and nonusers depending on the nature of peer effects. If improvements in peer quality lead to better educational outcomes for voucher users, then the voucher could improve their outcomes through increased sorting. On the other hand, if the exit of high-quality students reduces peer quality in public schools, then the students left in the public system may have systematically worse outcomes. The aggregate effect of the voucher depends on the strength of these two effects.

Hsieh and Urquiola (2003) explored the aggregate effects of Chile's voucher program. They looked at changes in aggregate test scores throughout Chile and failed to find that test scores had increased as a result of vouchers. They argue that the only way to identify the overall effects of the program is to focus on aggregate outcomes. This is because it is difficult if not impossible to remove the selection bias inherent in comparisons of different schools. These overall effects arise from both the direct effects on voucher recipients and the indirect effects of the voucher program on other students who did not have vouchers. In terms of identifying the effects of the voucher on users,

⁴ Evidence from the United States is positive. For example, Chakratabi (2007) shows that the increase in competition in Milwaukee following the supreme court decision allowing sectarian schools led to improved test scores. Hoxby et al. (2009) finds positive effects of charter school competition in Michigan.

Hsieh and Urquiola (2003) accentuated the problems of selection bias and the difficulties in identifying the effects of the Chilean voucher program on users.

Research by Gallego (2005) provided new evidence on the Chilean voucher program. Gallego noted that the number of Catholic priests existing in a community in the early 1950s can predict the number of and overall enrollment in voucher schools. This arises since most voucher schools were affiliated with Catholicism both before and after the reform. Gallego (2005) uses this fact to identify the effects of the voucher program. He uses the number of priests as an instrument for the penetration of the voucher program in a specific market. His findings suggest positive effects of the voucher program on the academic outcomes of students throughout municipalities where the voucher program had more penetration. Although the result may be indicative of competitive effects, it is driven in part by the effects of the program on voucher recipients. It echoes previous research (McEwan (2001); McEwan and Carnoy (2000)), which suggested that voucher schools affiliated with Catholicism had better outcomes than other voucher schools, public or private. As shown by McEwan and Carnoy (2000), Catholic schools produced better students at a lower cost than other public or private, voucher schools.

One interesting feature of the Chilean system is the inability of public schools to "exit" the system. In a purely competitive voucher program, both public and private schools should be able to exit. However, in Chile, public schools cannot close, and if all of the students from a public school were to exit that school, the municipality would still be forced to pay teachers' salaries and keep the school open.

A recent paper by Bravo, Mukhopadyay, and Todd (2009) provides a different way of estimating the long-run effects of Chile's educational voucher program. They estimate a life-cycle model of earnings and schooling decisions. Using labor force data for individuals educated before and after the voucher reform, they estimate impact of choosing private school and then simulate what educational attainment would have been in absence of the program. They find that educational vouchers increased educational attainment, high school graduation, college attendance and graduation, and wages.

4.2. Evidence from Sweden

The evidence from Sweden focuses on whether the program has improved outcomes throughout the entire educational system or it has exacerbated inequalities by increasing the amount of stratification between schools. As Bunar (2010) summarizes, even the positive estimates of the voucher program suggest that the voucher program may have had only a small impact on the overall educational quality and equity in Sweden.

Three influential studies suggest that the voucher program has improved overall performance throughout Sweden's system. Bohlmark and Lindahl (2007) find that an increase in the private school share by 10 percentage points increases average pupil achievement by almost 1 percentile rank point. Sandstrom and Bergstrom (2002) and

Björklund, Edin, Fredriksson, and Krueger (2004) both conclude that public schools improved as a result of competition from privately operated schools in Sweden.

Although the effects of the voucher program might be positive in the aggregate, one study suggests that this comes at a cost. According to Bohlmark and Lindahl (2007), an increase in the private school share by 10% points leads to a 2% increase in the average school costs throughout a municipality. Since Sweden's reform made school finance the municipality's responsibility, the increased cost strained some municipalities' budgets. Bunar (2010) explains that the higher costs arise for a number of reasons: first, public schools are expected to provide a constant level of educational quality; second, they must accept every child living in a certain attendance zone; third, public schools have long-term financial commitments to specific buildings; and finally, the contracts between teacher unions and municipalities do not make it so easy to lay off the public staff.

4.3. Evidence from Other Countries

There are other studies that focus on the effects of school choice on educational systems (e.g., Lavy (2006) in Israel, Gibbons, Machin and Silva (2008) in the United Kingdom). Ongoing work by Kremer and Muralidharan (2010) might provide the most definitive evidence on the effects of educational vouchers on overall systems.

Kremer and Muralidharan approached a sample of communities in India. In each community, they informed families about a potential voucher program and invited families to apply for this voucher program. They matched the participating communities to other communities and then randomized which of these communities received voucher programs. If the voucher program was oversubscribed within a community selected for vouchers, they conducted a voucher lottery.

Within the communities that received educational vouchers, they could exploit the voucher lotteries to estimate the effects of the voucher on individuals who use the voucher. In addition, within each community, they identified individuals who did not apply for the voucher. These individuals would only be impacted by the voucher if the voucher affected the overall quality of all of the schools in that community. They could compare these students across both voucher and nonvoucher communities and use the randomization to create an unbiased comparison. Their preliminary evidence suggests positive effects for both users and for the overall system.

5. CONCLUSION AND SUGGESTED DIRECTIONS FOR FUTURE STUDY

The evidence to date on the effects of educational vouchers on both voucher users and the overall system is far from conclusive. In the case of the effects of the voucher on voucher users, the most conclusive evidence comes from Colombia; however, it is unclear whether that evidence will generalize to other settings and populations.
The program itself was discontinued in 1998 due to perceived ineffectiveness. In other parts of the world, the impact of receiving a voucher is much less conclusive even within the respective countries. It is difficult to form an accurate counterfactual, and the results are often sensitive to the econometric specifications. For example, the perceived impact of the voucher changes dramatically as additional variables are included in the econometric specification (McEwan (2001)).

In the case of the effects of vouchers on the overall system, the results are also far from conclusive. Although many studies have suggested positive effects in both Sweden and Chile, these studies employ identification schemes that are not perfect. Some of the identification schemes (e.g., Gallego (2005)) are quite convincing, but even the most tightly identified studies could be criticized. For example, Gallego (2005) uses the strength of the Catholic church as an instrument for Catholic private school strength; however, if the strength of the Catholic church is correlated with other characteristics of towns, then the results may be biased. In addition, some careful studies (e.g., Hsieh and Urquiola (2003)) have shown that voucher programs may had no impact on the overall quality of the educational system.

Studies that propose new identification strategies may move the voucher literature forward in significant ways. For example, the ability of Swedish researchers to examine differences in outcomes within families over time represents a substantial improvement over studies that compare different individuals at different points in time. New ideas on identification and carefully designed studies such as the work by Kremer and Muralidharan in India can greatly improve international evidence on the efficacy of educational vouchers and their impact on the overall educational system.

Even if the research questions discussed in this paper were definitively resolved, there would still be significant research to understand the lessons from the voucher literature. As mentioned above, there are significant puzzles to the specific features of voucher programs, which lead to voucher effects. Most studies take place in a specific setting, and identifying voucher effects that are generalizeable to other settings is important for understanding the strengths and weaknesses of voucher systems. Given the heterogeneity in voucher programs across countries, it is important to know which specific features of vouchers might be driving the specific effects observed in different programs. This is an important direction for future research. In addition, given the lack of data on many voucher programs, there are ample opportunities to investigate and contrast existing voucher programs.

More generally, puzzles remain as to the mechanisms by which vouchers affect students have not received significant attention in the international voucher literature. On the question of the effects of the voucher on voucher users, several mechanisms have been explored including peer effects, the quality of private schooling, the salience of vocational education, and incentives. None of these mechanisms have been concretely identified as the cause of the observed voucher effects. On the question of the effects of the voucher on educational systems, competition, sorting, and resource drain are the most frequently mentioned mechanisms; however, there is little known about how schools compete. In Sweden, for example, the "competition" also involved an increase in overall expenditure by the municipality. We might have expected competition to lead to reductions in costs rather than increases. Work by Fiske and Ladd (2000) explores potential ways in which schools might compete (or fail to compete). Understanding the mechanisms and the underlying costs of voucher systems will yield more insights into cost-effective policies and the underlying economics of education in voucher settings.

Finally, additional evidence on other outcomes would expand our knowledge of educational vouchers' impacts. Parents choose schools for both academic and nonacademic reasons, yet most research focuses only on the academic outcomes. Understanding which outcomes matter to parents might yield additional lessons about the productivity of vouchers. Moreover extending voucher research into long-run outcomes (e.g., wages in Bravo, Mukhopadhyay, and Todd (2009)) might further help us understand the scope of vouchers. Then, the importance of the teaching profession and the effects of vouchers on the overall quality of the teaching force may provide important lessons on both mechanisms and implications of voucher programs.

REFERENCES

- Anand, P., Mizala, A., Repetto, A., 2006. Using School Scholarships to Estimate the Effect of Government Subsidized Private Education on Academic Achievement in Chile. University of Chile, Center of Applied Economics. Working Paper Number 220.
- Angrist, J., Bettinger, E., Bloom, E., Kremer, M., King, E., 2002. The Effects of School Vouchers on Students: Evidence from Colombia. Am. Econ. Rev. 92 (5), 1535–1558.
- Angrist, J., Bettinger, E., Kremer, M., 2006. Long-Term Educational Consequences of Secondary School Vouchers: Evidence from Administrative Records in Colombia. Am. Econ. Rev. 96 (3), 847–862.
- Bettinger, E., Kremer, M., Saavedra, J.E., 2010. Are Educational Vouchers Only Redistributive? Econ. J. R. Econ. Soc. 120 (546), F204–F228, 08.
- Bifulco, R., Ladd, H.F., 2006. The Impacts of Charter Schools on Student Achievement: Evidence from North Carolina. Educ. Finance Policy 1 (1), 5090.
- Björklund, A., Edin, P., Fredriksson, P., Krueger, A., 2004. Education, Equality and Efficiency: An Analysis of Swedish School Reforms during the 1990s (IFAU Report 2004:1). Institute for Labour Market Policy Evaluation, Uppsala, Sweden.
- Bohlmark, A., Lindahl, M., 2007. The Impact of School Choice on Student Achievement, Segregation and Costs: Swedish Evidence. Discussion Paper Series. IZA DP No. 2786.
- Bonilla, J., 2011. Contracting Out Public Schools on Academic Achievement: Evidence from Colombia. University of Maryland mimeo, College Park, MD.
- Bravo, D., Contreras, D., Sanhueza, C., 1999. Rendimiento escolar, desigualdad y brecha de desempeño privado/público: Chile 1982–1997. Documento de Trabajo, Departamento de Economía, U de Chile, 1999.
- Bravo, D., Mukhopadhyay, S., Todd, P.E., 2009. How Universal School Vouchers Affect Educational and Labor Market Outcomes: Evidence from Chile. University of Pennsylvania, Mimeo.
- Bunar, N., 2010. Choosing for Quality or Inequality: Current Perspectives on the Implementation of School Choice Policy in Sweden. J. Educ. Policy 25 (1), 1–18.

- Calderon, A., 1996. Voucher Programs for Secondary Schools: The Colombian Experience. Human Capital Development and Operations Policy Working Paper. The World Bank, Washington, DC. http://www.worldbank.org/education/economicsed/finance/demand/related/wp_00066.html.
- Carnoy, M., 1998. National Voucher Plans in Chile and Sweden: Did Privatization Reforms Make for Better Education? Comp. Educ. Rev. 42 (3), 309–337.
- Fiske, E.B., Ladd, H.F., 2000. When Schools Compete: A Cautionary Tale. Brookings Institution Press, Washington, DC.
- Friedman, M., 1962. Capitalism and Freedom. University of Chicago Press, Chicago, IL.
- Gallego, F.A., 2005. Voucher-School Competition, Incentives, and Outcomes: Evidence from Chile. MIT mimeo.
- Gibbons, S., Machin, S., Silva, O., 2008. Choice, Competition, and Pupil Accountability. J. Eur. Econ. Assoc. 6 (4), 912–947.
- Hoxby, C., Murarka, S., Kang, J., 2009. The New York City Charter Schools Evaluation Project: How New York City's Charter Schools Affect Achievement. http://www.nber.org/~schools/charterschoolseval/ how_NYC_charter_schools_affect_achievement_sept2009.pdf (accessed 20.05.10.).
- Hsieh, C.-T., Urquiola, M., 2003. When Schools Compete, How Do They Compete? An Assessment of Chile's Nationwide School Voucher Program. NBER Working Papers 10008.
- King, E., Rawlings, L., Gutierrez, M., Pardo, C., Torres, C., 1997. Colombia's Targeted Education Voucher Program: Features, Coverage and Participation, Working Paper No. 3, Series on Impact Evaluation of Education Reforms, Development Economics Research Group, The World Bank (September).
- King, E., Orazem, P., Wolgemuth, D., 1998. Central Mandates and Local Incentives: The Colombia Education Voucher Program. Working Paper No. 6, Series on Impact Evaluation of Education Reforms, Development Economics Research Group, The World Bank (February).
- Klitgaard, M.B., 2007. School Voucher and Political Institutions: A Comparative Analysis of the United States and Sweden. NCSPE Working Paper No. 153.
- Klitgaard, M.B., 2008. School Vouchers and the New Politics of the Welfare State. Governance Int. J. Policy Administration Inst. 21 (4), 479–498.
- Kremer, M., Muralidharan, K., 2010. The General Equilibrium Effect of School Vouchers: Evidence from a Two-Stage Experiment. Ongoing Research.
- Lavy, V., 2006. From Forced Busing to Free Choice in Public Schools: Quasi-Experimental Evidence of Individual and General Effects. NBER Working Paper Number 11969.
- McEwan, P.J., 2001. The Effectiveness of Public, Catholic, and Non-Religious Private Schools in Chile's Voucher System. Educ. Econ. 9 (2), 103–128.
- McEwan, P.J., Carnoy, M., 2000. The Effectiveness and Efficiency of Private Schools in Chile's Voucher System. Educ. Eval. Policy Anal. 22 (3), 213–239.
- Morales-Cobo, P., 1993. Demand Subsidies: A Case Study of the PACES Voucher Program. Universidad de los Andes, Economics Department, Bogota, mimeo.
- Neal, D., 1998. What Have We Learned About the Benefits of Private Schooling? Econ. Policy Rev. 4, 79–86.
- Psacharapoulos, G., 1987. Public vs. Private Schools in Developing Countries: Evidence from Colombia and Tanzania. Int. J. Educ. Dev. 7 (1), 59–67.
- Ribero, R., Tenjo, J., 1997. University de los Andes. Department of Economics Working Paper.
- Rounds, P.T., 1996. Will Pursuit of Higher Quality Sacrifice Equal Opportunity in Education? An Analysis of the Educational Voucher System in Santiago. Soc. Sci. Q. 77 (4), 821–841.
- Rouse, C.E., 1998. Private School Vouchers and Student Achievement: An Evaluation of the Milwaukee Parental Choice Program. Q. J. Econ. 13 (2), 553–602.
- Sandstrom, M., Bergstrom, F., 2002. School Vouchers in Practice: Competition Won't Hurt You. The Research Institute of Industrial Economics. Working Paper No. 578.
- Sappelli, C., Vial, B., 2002. The Performance of Private and Public Schools in the Chilean Voucher System. Cuad. Econ. 39 (118), 423–454.
- Soderstrom, M., Uusitalo, R., 2010. School Choice and Segregation: Evidence from An Admission Reform. Scand. J. Econ. 112 (1), 55–76.
- Tiebout, C.M., 1956. A Pure Theory of Local Expenditures. J. Polit. Econ. 6, 416-424.

- West, E.G., 1996. Educational Vouchers in Practice and Principle: A World Survey. http://www.world bank.org/html/extdr/hnp/hddflash/workp/wp_00064.html#TofC19.
- Wößmann, L., 2007. International Evidence on School Competition, Autonomy, and Accountability: A Review. Peabody J. Educ. 82 (2), 473.
- Zimmer, R., Bettinger, E., 2007. Getting Beyond the Rhetoric: Surveying the Evidence of Vouchers and Tax Credits. In: Ladd, H., Fiske, E. (Ed.), Handbook of Research in Education Finance and Policy. Routledge, New York, pp. 447–466.
- Zimmer, R., Bettinger, E., 2008. Getting Beyond the Rhetoric: Surveying the Evidence of Vouchers and Tax Credits. In: McGaw, B., Peterson, P., Baker, E. (Eds.), The International Encyclopedia of Education. third ed. Elsevier, Amsterdam.



Dropouts and Diplomas: The Divergence in Collegiate Outcomes John Bound^{*} and Sarah Turner^{**}

*Department of Economics, University of Michigan and NBER, Ann Arbor, MI 48109-1220

**Department of Economics and Curry School of Education, University of Virginia and NBER, Charlottesville, VA 22904-4182

Contents

1.	Introduction and Motivating Data	574
	1.1. Baseline Evidence on College Degree Attainment and Enrollment	575
	1.2. Motivating Economic Analysis of College Degree Attainment	580
2.	Models of Years of Schooling and Application to Degree Attainment	580
	2.1. Empirical Validity: Years of College versus Degrees?	583
3.	Specification of College Degree Attainment	584
	3.1. Where Students Attend and Complete College	585
	3.2. Field of Study	586
	3.3. Modeling College Degree Attainment	587
	3.3.1 Demand-Side Determinants of College Persistence	587
	3.3.2 Supply-Side Determinants of Degree Attainment	588
4.	Evidence on Determinants of Degree Attainment	590
	4.1. Demand-Side Determinants of Attainment	590
	4.1.1 Student Preparation	590
	4.1.2 Student Background and Family Circumstances	593
	4.2. Supply-Side Determinants of Attainment	596
	4.2.1 Changes in Resources and Stratification over Time	598
	4.3. College Matching as a Determinant of Attainment	601
5.	Unanswered Questions in the Study of Degree Attainment	603
	5.1. College Choice	604
	5.2. In-College Experiences	605
	5.3. Market Structure: Competition and Stratification	606
	5.4. Political Economy and Public Subsidies	607
	5.5. International Comparisons	608
	5.6. Economics of College Degree Completion and Public Policy	608
Re	609	

Abstract

Although collegiate attainment rates have risen in many developed and developing countries over the last three decades, they have remained essentially flat in the United States over the same period. In this chapter, we distinguish various models of degree attainment in the general context of theoretical and empirical specifications of educational attainment. To explain collegiate degree attainment, we consider the roles of student demand, the supply side of the postsecondary education market, and

the role of public support in determining outcomes. Although the study of college degree attainment has traditionally focused on demand-side determinants of attainment, including how students finance college attainment and academic preparation, we present here the evidence that supply-side determinants including the level of public subsidies and the associated stratification among colleges and universities are also important determinants of degree attainment. Review of this evidence and research suggests a number of unexplored areas for economic research related to college choice, in-college attainment, and the supply-side determinants of stratification and resources per student.

Keywords

College Degree Attainment Higher Education Colleges and Universities

1. INTRODUCTION AND MOTIVATING DATA

Collegiate attainment rates—the fraction of a population with a college degree—have not kept pace with increases in the demand for skilled workers in the United States, resulting in persistently high returns to a college degree. In the United States, the number of high school graduates receiving a BA degree has increased only modestly over the last three decades, while the earnings premium to a college degree relative to a high school degree has nearly doubled (Goldin and Katz (2008)). While the supply of college graduates relative to workers with only a high school degree among young workers increased rapidly through the 1960s and much of the 1970s, the subsequent years brought a sharp deceleration in the growth in the relative supply of college graduates. Even as the rising college wage premium was met with an increase in the number of students attending college, degree attainment has not increased commensurately. Indeed, the share of college entrants who complete a college degree is little higher today than in the 1970s. Because the supply of graduates affects both returns to education and overall economic growth, economic analysis of degree attainment is an important area of inquiry.

Plausibly, the inadequate preparation of many US high school students has contributed to this stagnation in collegiate attainment (Carneiro and Heckman (2003)). However, the combination of the declining share of college costs covered by state appropriations, the erosion in the real value of federal financial aid from programs like the Pell grant, and the dramatic increase in the stratification of resources among colleges have left many college students with lower collegiate resources than were available three decades ago (Kane (2006); McPherson and Schapiro (2006); Hoxby (2009)). As a result, the stagnation in collegiate attainment not only can be attributed solely to changes in preparation of entering college students but also reflects declines in the availability of collegiate resources.

This analysis begins with a review of the main trends in collegiate attainment and degree receipt in recent decades in the United States and other countries. To explain these outcomes, we consider the roles of student demand, the supply side of the postsecondary education market, and the role of public support in determining outcomes. The concluding section outlines the many important—and unanswered—questions in the economics of higher education.

1.1. Baseline Evidence on College Degree Attainment and Enrollment

Although collegiate attainment rates have risen in many developed and developing countries over the last three decades, they have remained essentially flat in the United States over the same period. Figure 8.1 presents the share of college graduates in the age groups of 25–34 and 45–54, representing individuals born in the years 1974–1983 and the years 1963–1954, for a range of countries in Europe, North America, and Asia for 2008. Most countries show large gains in college attainment for the more recent cohorts. For example, in South Korea, the proportion of the population with a college degree rises from only about 18% for those ages 45–54 to 51% for those ages 25–34. In contrast, the population in the age range of 25–34 in the United States has an identical college attainment rate at 39% to the population





Source: Data are from Table A1.3 in the report "Education at a Glance 2010," released by the Organisation for Economic Co-Operation and Development (OECD). Degrees included for each country include degrees or certificates obtained from 2- to 3-year vocational programs, 3- to 5-year academic programs, and doctoral programs. For the United States, degrees included are bachelor's, master's, vocational associate's, and doctorate degrees. Details on degrees included for other countries can be found in "Education at a Glance 2007: Annex 3," pages 20–40.

in the age range of 45–54. As a result of this stagnation, the United States has slipped from ranking third in college degree attainment among those ages 45–54 to a rank of tenth among those ages 25–34. To be sure, these data suffer from nontrivial problems in alignment in degree types across countries; nevertheless, they are the source of the widely acknowl-edged decline in United States leadership in college degree attainment.

The rate of degree completion—the fraction of college entrants who obtain a degree—is a clear point of difference between the United States and other countries and a factor in the recent divergence in attainment. Among those who start postsecond-ary education, rates of degree completion in the United States are lower than that in most other developed countries. Figure 8.2 presents a comparison of completion rates in the United States relative to those in other countries. Although in the United States, only about half of those who begin first-level degree programs actually obtain their degree, rates of completion exceed 70% in many other countries. The United States' considerably lower college completion rates, combined with its eroding rank in collegiate



Figure 8.2 College Completion Rates, by Country.

Source: OECD. Table A4.1. See Annex 3 for notes (www.oecd.org/edu/eag2010). Chart shows proportion of students entering tertiary education and receiving a degree; some of the students who have not graduated may be still enrolled, or may have finished their education at a different institution than the one they started.

attainment, have produced pointed policy questions in the United States. At issue in the allocation of public resources to improve collegiate attainment are questions of how well US students are prepared for college, how well students are able to match their needs with collegiate opportunities, and how well colleges and universities are furthering degree completion among college entrants. Although the United States has been a historical leader in "mass higher education" (Goldin and Katz (2008)), the institutional and policy emphasis on postsecondary educational access (initial enrollment) has not produced sustained growth in college attainment.¹ Given the significance of collegiate attainment, long-term economic growth, and the substantial public investments made in postsecondary education, it is important to understand how policy goals align with models of collegiate attainment. Although our focus in this paper is on collegiate outcomes in the context of US colleges and universities, the basic issues for analysis are represented in other developed countries, leading to opportunities for further comparative analysis of how differences in institutions, funding mechanisms, and other policies impact degree attainment.

A closer look at the US data reveals dramatic gains in educational attainment in the first three quarters of the twentieth century. Between 1900 and 1970 birth cohorts, the fraction of the population graduating from high school rose from around 6 to 76% (Fig. 8.3a). Not only did high school graduation increase, but the fraction of the population participating in college (panel B) and attaining a BA degree (panel C) also increased markedly, as shown in Fig. 8.3. Although only about 5.4% of men and 4.2% of women born at the turn of the century received a BA degree, about 30% of men and 34% of women born in 1970 could be expected to receive a BA. Not surprisingly, the rise in secondary school attainment preceded the growth in college enrollment and degree completion. For cohorts born from the 1920s to the 1940s, college enrollment and degree completion advanced hand in hand, rising particularly rapidly for men due in some part to the generous educational benefits provided through GI Bills (Bound and Turner (2002)). These increases in educational attainment account for a considerable fraction of the economic growth the United States experienced over this interval (Jorgenson and Griliches (1967); Griliches (1970); DeLong, Goldin, and Katz (2003)).

For US cohorts born after 1945, we see a dramatic reversal in the overall rate of college attainment and the rate of graduation among those students who enrolled in college. This shift has been particularly magnified for men. For example, the proportion of males who completed a BA degree was actually lower for the 1970 birth cohort than the 1950 cohort. More recently, we see a decided increase in both college enrollment and completion, with this change concentrated among women.

Quoted in the *Christian Science Monitor* (December 3, 2008), Pat Callan—president of the National Center for Public Policy and Higher Education—notes "Historically, our strength has been access [to higher education] and our weakness has been completion. We've always said the reason we can't be expected to do so well on completion is because we're generous on access. But now, we see countries catching up to us and surpassing us on access and completion."



Figure 8.3 Educational Attainment by Birth Cohort, 1900–1980. (a) Years of Schooling. (b) College Enrollment. (c) College Degree Attainment.

Source: Data are from Goldin and Katz (2008) and tabulated from 1940 to 2000 Census of Population Integrated Public Use Microdata Samples (IPUMS). Observations are for US native-born individuals adjusted to 35 years of age. Figure 8.3a shows the fraction of each birth cohort with at least a high school degree, Fig. 8.3b shows the fraction of each cohort with some college attendance, and Fig. 8.3c shows the fraction of each cohort with a college degree. For additional details, see DeLong, Goldin, and Katz (2003).

Comparing across the panels shown in Fig. 8.3, it is clear that changes in college degree attainment have not followed changes in college enrollment consistently over the course of the last 25 years. While college enrollment rates have increased fairly consistently, college degree attainment declined before increasing among more recent cohorts. Figure 8.4 presents the trend by birth cohort in the share of enrolled college students who complete a BA degree—essentially the trend shown in Fig. 8.3c divided by the trend in Fig. 8.3b. For both men and women, the rate of college completion has been below 50% for nearly a half century, with this level appreciably below the rate of completion achieved by men in the early part of the century.

A component of this stagnation has been a growing disparity in college completion rates by parental circumstances. For example, for high school students from the top quartile of the family income distribution, completion rates rose slightly from 67.4 to 71% between those starting college in the early 1980s and those starting in the early 1990s, while the college completion rates fell for students from other income groups (Bowen, Chingos, and McPherson (2009)). Indeed, for 1992 high school seniors who enrolled in college, the difference in college completion rates between the students



Figure 8.4 Share of College Entrants Receiving BA Degree.

Notes: The completion rate presented in this figure represents the ratio of the number of college degree recipients (Fig. 8.3c) to the number of individuals with at least some college (Fig. 8.3b). See Fig. 8.3 for additional notes on the data.

from high-income families and students from the other income groups was about 32.6 percentage points. Thus, low college completion rate among lower income students is one factor contributing to the stagnation in supply of college-educated workers in recent decades (Bowen, Chingos, and McPherson (2009)).

1.2. Motivating Economic Analysis of College Degree Attainment

In the United States, limited growth in collegiate attainment and flat college completion rates combined with a declining position in international collegiate attainment have led federal policy makers and policy organizations from across the political spectrum to make strong statements about the desirability of increasing college degree attainment. In introducing the 2010 budget, President Obama declared: "By 2020, America will once again have the highest proportion of college graduates in the world."² Yet how do policy actions such as increasing funding for student aid or for colleges and universities impact attainment? The projected efficacy of these and other policy actions rests on economic models of how students with different levels of preparation navigate the path to college degree attainment and how the supply-side dimensions of higher education affect college completion.

Despite the economic importance of college completion, much of the economics literature subsumes degree receipt in general discussions of years of educational attainment. In this paper, we distinguish theoretically and empirically the economic analysis of degree completion. In Section 2, we consider the empirical validity of the canonical model of educational returns, which uses years of educational attainment. In Section 3, we identify the salient dimensions of college degree completion, including college and program choices, persistence, and the role of the supply side of the market. Section 4 presents evidence on the determinants of degree completion in the US market. The final section identifies the central theoretical and empirical questions for future analysis.

2. MODELS OF YEARS OF SCHOOLING AND APPLICATION TO DEGREE ATTAINMENT

Although the existence of degree effects (or diploma effects) in returns to education may be among the most persistent empirical findings in labor economics (Lange and Topel (2006)), the conceptualization of educational investment as linear in years has endured in both theoretical and empirical work. The convenience of the linearity assumption is

² President Barack Obama, February 24, 2009. Similarly, the Gates Foundation's initiative on postsecondary attainment is explicit in its objective: "We have set an ambitious goal for ourselves and the nation: double the number of young people who earn a postsecondary degree or certificate with value in the marketplace by the time they reach age 26. To accomplish this goal, we must connect the millions of young Americans who have the will to get the education they need with a way to get there: helping them get further, faster—and at far less cost in terms of time and money."

hard to dispute, though the limitations are substantial and there is a practical divergence between this specification and available evidence.

The canonical model of educational investment specifies earnings (γ) as a function of years of educational attainment (S). This model dates to Becker (1967) and Mincer (1974). Nearly every labor economics class taught at the undergraduate or graduate level relies on some formulations of this model. Following Card (1995, 1999), consider the individual choice of S to maximize the utility function $U(S, \gamma)$. Consider a straightforward utility function in terms of log earnings and educational costs as $U(S, \gamma) = \log \gamma - h(s)$, where utility is a function of log earnings and the discounted, presented value of earnings takes the form $\int_{S}^{\infty} \gamma(S) e^{rt} dt = \frac{\gamma(S)e^{-rS}}{r}$. Optimal schooling choices are defined by the equality between marginal costs and marginal benefits $h'(S) = \frac{\gamma'(S)}{\gamma(S)}$. In turn, it is natural to think of heterogeneity in optimal educational attainment as following from differences in the costs of schooling (an individual-specific component in h(S)) and differences in individual returns to schooling $\left(\frac{\gamma'(S)}{\gamma(s)}\right)$.

As it is usually presented, this model sets aside distinctions between different institutional transitions: investing in another year of high school is treated as a choice perfectly analogous to the decision to transition from high school to college. The model misses the richness of the choice set prevailing at the postsecondary level that is largely absent at the elementary and secondary levels. Postsecondary students choose where to attend (from over 3000 colleges and universities in the United States), how many courses to take per term (the rate of attainment), and which combination of classes and subjects to pursue. In short, because collegiate attainment is multidimensional in institutional characteristics and field of study, this simple setup falls short because it only affords student choice in a single dimension: time in school.

A more general exposition of the human capital earnings function applicable to college attainment is the model of heterogeneous human capital. Linking occupation choice with collegiate attainment, Willis (1986) sets forth a framework that distinguishes college degree choices leading to certain occupations, such as accounting and chemical engineering. Individuals are assumed to make collegiate choices consistent with their underlying occupational aptitudes in the spirit of the model of Roy (1951). Degree choice and labor-market outcomes are then determined by the interaction of the underlying distributions of individual aptitude (precollegiate skills) and labor-market demand. A notable feature of this framework is that occupation choice is determined by comparative advantage, and the discrete choice framework does not require a unidimensional index of skill. One might expect that individuals who would be excellent English teachers might make poor accountants, whereas those who would earn high wages as accountants might not excel as English teachers. The resulting problem of self-selection in modeling the returns to education has been addressed at some length, beginning with Roy (1951) and continuing with Willis and Rosen (1979).

Although the self-selection problem has received considerable attention in the labor economics literature examining educational attainment and associated labor-market returns, the supply side of the education market has been largely ignored. In modeling degree attainment, researchers have implicitly assumed a perfectly elastic supply of collegiate options at given levels of resources per student ("quality") and tuition.³ As such, it is straightforward to model educational attainment as a response to opportunities in the labor market. For models of attainment at the elementary and secondary levels, such a simplification may be relatively innocuous given the structure of public provision of K-12 education in the Unites States since the early twentieth century when public education through secondary school has been universally available (Goldin and Katz (2008)). In the higher education market, the presence of selective admissions at some colleges and universities implies that the supply side of the market is not perfectly elastic, particularly at the top of the market. More generally, substantial public and private subsidies to colleges and considerable stratification across institutions in resources per student point to a dynamic on the supply side of the market in which changes in demand will be accommodated differentially across institutions.

Finally, two types of market failures are likely to be more severe as students negotiate the collegiate attainment process than in the elementary and secondary years of schooling. First, capital market failures or credit constraints that increase marginal costs of attainment are likely to become sharply binding for some students at the postsecondary level. Although economists have long written about the threat of credit constraints to optimal educational investments, an implication in the context of the standard model described above is that these constraints may increase sharply and discontinuously at college entry as students must meet tuition costs. The policy expectation that parents will help children to finance college education in the presence of incomplete intergenerational contracts is likely to lead to binding credit constraints and underinvestment in college, with these problems becoming more severe as direct tuition charges increase (Brown, Scholz, and Seshadri (2010)). Second, information constraints are likely to have a much larger impact at the postsecondary stage of educational attainment-affecting both the matching of students to colleges and degree completion. Individuals must solve a complicated nonsequential choice problem in selecting schools to which they will apply, and then they must negotiate what is most often an administratively cumbersome application process. What has received less attention in the literature—but is also surely important—is that colleges and universities have very imperfect processes for assessing the capacities of their

³ Alternatively, one can interpret models that focus solely on student choice as being partial equilibrium. A partial equilibrium analysis is appropriate for understanding student choice but not for understanding shifts in degree attainment among the population.

applicants, potentially leading to too few offers of admission to students for whom the uncertainty about collegiate performance is the greatest. These students may be most likely to come from families or high schools where applying to selective colleges and universities is not the norm, potentially generating some inefficiencies in the allocation of students to colleges where public and institutional subsidies are the largest.

Although much of the literature emphasizes the role of credit constraints and imperfect information at the point of initial college enrollment, these also may be factors throughout the college years, affecting the likelihood of graduation. Limited access to credit markets and limited precautionary savings may not only lead to inefficient tradeoffs for students between work and course enrollment but also mean that small shocks, such as unanticipated health expenses or parental job loss, trigger interruptions in college attendance. Limited information and guidance may also impede students' progress during the college years. For instance, students may experience educational setbacks if they are not well-informed about the prerequisite skills needed for success in different courses of study or if they lack information about course transfer policies.

2.1. Empirical Validity: Years of College versus Degrees?

Although models using years of education hold appeal for the simplicity of representation and the availability of historical data,⁴ the empirical support for this class of models is limited. The linear returns specification implied by the Mincer model has not held up to empirical scrutiny, particularly in more recent cohorts.⁵ Available estimates demonstrate additional return to BA degree receipt beyond measured years of educational attainment. Using years of education measured in one CPS rotation and degree attainment in a following rotation, Jaeger and Page (1996) estimate effects for the BA degree in the range of 25–28% for white men and white women, with somewhat larger BA degree effects in the range of 36–48% for black men and women, respectively. It is noteworthy that the estimated BA degree effects are much larger than degree effects associated with high school completion (6–13%), suggesting a distinct economic phenomenon at the postsecondary level.

Diploma (or "sheepskin") effects have been the subject of much empirical inquiry in the context of screening theories of education. These theories posit that the added return to a degree is the credential effect—that is, a signal to the market of abilities inherent in obtaining a degree that go beyond the skills associated with years of educational attainment (Hungerford and Solon (1987)). Although early work by Layard and

⁴ For the 1940–1980 Decennial Census enumerations and in the Current Population Surveys from the 1940s to 1991, education was measured in terms of years of completed schooling.

⁵ While Card (1999) suggests that the linear functional form provides a quite good fit to the data in his Handbook of Labor Economics chapter, a number of researchers have tested the linearity in years specification directly. A series of articles by Heckman, Lochner, and Todd (2003, 2006, 2008) leads to the clear rejection of the linearity in years specification, particularly at the level of college completion.

Psacharopoulous (1974) found relatively small effects associated with degree completion alone, further empirical analysis provided substantial evidence of degree effects, particularly among women and minorities (Belman and Heywood (1991); Jaeger and Page (1996); Belman and Heywood (1997)). One critique of many of these estimates is that they tended to suffer from substantial measurement error. Degree effects were often inferred from the discontinuity at 12 and 16 years of education, whereas years of attainment aligned imperfectly with units of collegiate attainment. The microdata analysis of Kane and Rouse (1995) answered much of this critique by using accrued credit hours as the measure of the quantity of educational attainment and adding other measures for degree receipt. Their estimates of the BA degree effect on log hourly wages are on the order of 25.7% for men and 32.8% for women, with the effects on annual earnings of 32.4 and 47.6% for men and women, respectively.

Much of this literature on degree effects presents discussion in terms of screening models versus human capital models. But this dichotomy overlooks distinctions between degree recipients and those with similar educational attainment who did not receive degrees. For instance, Lange and Topel (2006) theorize that degree effects demonstrate that in the face of uncertainty about individual returns to education during their college years, students who persist are those most capable of profiting from degree completion. This interpretation is consistent with the option value of schooling models discussed below. A further, perhaps more transparent, explanation rests on differences in knowledge acquisition and course selection between those completing degrees and those who do not.

Here, we focus on BA degree completion, while noting the empirical importance of other degree types including Associate Degrees and post-BA degrees in law (JD) and medicine (MD). As such, analysis of collegiate attainment must be motivated by economic models that specify degree attainment as distinct from the quantity of post-secondary enrollment (credit hours or years of enrollment). In Section 3, we begin by considering the salient characteristics of degree attainment and then turn to the specification of models that address college choice, college persistence, and field of study.

3. SPECIFICATION OF COLLEGE DEGREE ATTAINMENT

Economists have been slow in distinguishing the practical import of a college degree, as well as developing models to capture this educational outcome and the associated measure of labor-market returns. College degree requirements are usually defined by curricular requirements and a specific number of credit hours, typically 120.⁶ Both time

⁶ In fact, there is some variation across institutions and states in the credits required for graduation. [Moreover, colleges may differ appreciably in GPA requirements and distributional requirements.] For example, Boston College requires 114 credits for graduation while Kansas State requires 124 credit hours. While there have been some initiatives aimed at offering a three-year baccalaureate option at colleges and universities, it is the elongation of time to degree, not its acceleration, that is the decisive aggregate trend over the last three decades (Bound, Lovenheim, and Turner (2010b)).

enrolled in college and years of educational attainment are likely to be error-ridden measures of credit attainment, as many students attempt credits that they do not complete (Bound, Lovenheim, and Turner (2010b)) and many students are required to take remedial courses in the college years that do not count toward degree credit (Bettinger and Long (2009)).

As a matter of practical significance, it is worth articulating how a college degree is conceptually distinct from a fixed quantity of educational attainment.⁷ As seen by employers on resumes, college degrees are distinguished by which college or university awards the degree and the field of specialization. Both the college attended and type of major are associated with economically meaningful differences in the returns to a college degree, though given the degree of selection at both margins there is a long-standing challenge in the measurement of causal effects of these dimensions of collegiate attainment.

3.1. Where Students Attend and Complete College

In the United Stated, there are more than 3000 colleges and universities from which students can choose. These institutions differ quite markedly in resources, tuition costs, fields of study offered, modes of instruction and scale; the impacts of these differences in resources will be discussed in more detail in the next section. What is immediately important for modeling collegiate attainment is that the "match" between student attributes and the characteristics of a college or university may have a substantial impact on collegiate attainment. This process of matching depends on both student choice (application, matriculation) and institutional choice (admissions) and represents an important economic problem of nonsequential choice (e.g., of these college choice models, see Manski and Wise (1983); Long (2004); Howell (2004); Pallais (2009)).

For the purposes of modeling the economic returns to degree attainment, a central question concerns whether there are discernable economic benefits to attending a more selective college or a college with greater resources per student. Simple comparisons between students who attend high- and low-resource schools show enormous differences in the distribution of earnings between these groups. Although such simple comparisons almost certainly exaggerate the causal effects, studies that control both family

⁷ In this analysis we concentrate on the BA degree, though the discussion applies more generally to other subbaccalaureate degrees. As a point of historical reference, Oxford began granting the Bachelor of Arts degree at the end of a four-year course of study in the fifteenth century. The term "bachelor" can be traced to degree conferral in the medieval European universities. The University of Paris and its younger cousin, Oxford University, each of which granted licentia docendi—or the license to teach—to its students beginning in the twelfth or thirteenth century (Green (1946); van Scoyoc (1962)).

background and scholastic achievement in high school typically still find significant economic returns to school quality (e.g., Brewer, Eide, and Ehrenberg (1999); Monks (2000); Black and Smith (2004, 2006)). These returns, if anything, have increased over time (Hoxby and Long (1999)). Although such studies have limitations, Hoekstra has recently confirmed large returns to attending a state flagship school with a sharp test-based admission cutoff relative to an alternative using a regression discontinuity design. His estimates suggest that for those at the margin, attending the flagship school increases earnings by 20%.⁸

3.2. Field of Study

Beyond credit or course requirements, contemporary college degree receipt is defined in terms of the completion of a specific course of study that is distinguished by a major, or area of concentration. Degree receipt at most colleges and universities in the United States requires the completion of a specified set of general requirements (often called distributional requirements, which may include foreign language competency and literature and science requirements) as well as completion of a major course of study or specialization. The range of college majors spans arts and science disciplines such as English, economics, and physics to professional fields such as business or education. In the United States, the choice of major is often made after at least a year of undergraduate study (with some continued option for switching), though some institutions have professional schools in areas like engineering or nursing in which students follow a specialized course of study from admission to graduation. For students in many European countries and England, choice of major is concurrent to initial college enrollment—that is, students are admitted to specific programs within institutions.⁹

In the United States, there have been substantial changes over time in the fields that students have chosen as undergraduate majors, reflecting broad trends in student interests, the institutions that students attend, and the skills and aspirations that students bring

⁸ It is important to note some dissension in this literature. Dale and Krueger (2002) do not find significant effects of the return to attending a more selective college based on estimates that control for the selectivity of the schools to which a student applied. While this result has, justifiably, received a good deal of attention, when Dale and Krueger compare the earnings of those who went to more relative to less expensive schools, they find statistically significant differences in earnings between the groups. In addition, it is possible to question the validity of Dale and Krueger's specification (Hoxby (2009)).

⁹ In Scotland, students are typically admitted to the broader unit of a faculty or a university at large. As such, Scottish students tend to choose their area of specialization later in their undergraduate coursework than their counterparts in England and are often required to take a broader range of courses. Malamud (2007) exploits the difference between the English and Scottish arrangements and finds that relatively late specialization is tied to relatively fewer early career occupational switches, which is consistent with the idea that the learning occurring in the system allowing later specialization may lead to improved match quality.

to undergraduate education.¹⁰ Significant individual selection into different majors (Bowen and Turner (1999); Paglin and Ruffolo (1990)) and variation in the economic returns to different majors (Grogger and Eide (1995)) unquestionably complicate the economic analysis of degree completion. The degree to which choice of major is determined by student's preferences and abilities, as well as institutional constraints, is an important dimension of college degree attainment. To this end, we note the presence of ambitious models linking majors and earnings in a dynamic context (Arcidiacono (2004); Arcidiacono, Hotz, and Kang (2010)), while limiting our own ambitions in the remainder of this paper to the consideration of the path from college choice to credit attainment and degree completion.

3.3. Modeling College Degree Attainment

The demand for college degree attainment, like educational attainment more generally, is expected to increase with labor-market returns and decrease with direct costs. In addition, student aptitude is expected to reduce costs and increase net rewards to degree attainment. On these demand-side dimensions, college degree attainment follows the same theoretical outline as educational attainment more generally. Decisions on what college to attend, what curriculum to pursue, and whether to complete the degree distinguish the economic analysis of degree attainment. Students' choices and degree outcomes are, in turn, highly dependent on the supply side of the market for higher education. Although explicit characterization of a collegiate production function is difficult, basic economic theory suggests a positive link between collegiate resources and degree attainment. Although the presentation of a complete model of college degree completion is beyond the scope of this review, our objective is to highlight the key dimensions of the higher education market that frames the economic analysis of degree attainment.

3.3.1 Demand-Side Determinants of College Persistence

We suspect that college is essentially an "experience good"—an event for which benefits are difficult to gauge in advance. Students will discover how well suited they are to postsecondary pursuits or how well-matched they are to a particular institution only

¹⁰ Aggregate trends in choice of major show a substantial "flight" from arts and sciences disciplines between the late 1960s and the mid-1980s, which can be decomposed into two factors: first, shifts in the concentration of higher education enrollment from liberal arts colleges and research universities toward institutions concentrating in pre-professional preparation; and second, changes in student demands within institutions (Turner and Bowen (1990)). Perhaps the most notable change in the distribution of concentrations is seen in the shift from majors in education to majors in business. In 1966, 34 percent of BAs at comprehensive public institutions were awarded in education, while only 13 percent were in business. However, the late 1960s represented a peak in the demand for new teachers, with the progression of the baby-boom generation through elementary and secondary school. By academic year 2005, the proportions had flipped and only 11 percent of majors were in education while 27 percent were in business fields.

after they enroll and take some classes (Smith (2008)). Yet standard models of educational investment assume full ex-ante information about the benefits and costs of college completion (Card (1999)). As Stange (2008) notes, these models are not consistent with considerable attrition during the college years, when the education–earnings profile is relatively flat. A fuller characterization of the college completion decision process is found in option value models, which explicitly incorporate the post-entrance process of learning about personal abilities, and making sequential decisions throughout the college years about persistence toward the degree and pursuit of a specific course of study. Indeed, there is growing empirical evidence of the importance of the value that students place on the ability to make persistence decisions sequentially (Stange (2008); Stinebrickner and Stinebrickner (2010)).¹¹

Although option value models of college persistence provide a theoretical framework for modeling "efficient" attrition from the college attainment pipeline, other demand- and supply-side barriers may act to impede academic progress. In particular, credit constraints and financial shocks are widely referenced as a deterrent to degree completion.¹² Although the theoretical basis for credit constraints in higher education should be quite clear, the empirical evidence is mixed, as discussed below.

3.3.2 Supply-Side Determinants of Degree Attainment

The supply side of the higher education market effectively determines the collegiate choices open to students. The market is highly stratified—more so in the United States than in many other developed countries¹³—with some institutions offering considerably greater resources per student than others (Winston (1999)). Within states in the United States, there is a recognized hierarchy of public institutions, with differentiated levels of support and collegiate missions. The most obvious of such examples is in the state of California, where the Master Plan for Higher Education defines a three-tier structure consisting of the research universities in the University of California system, the California State University, and the California Community Colleges (Sallee, Resch, and Courant (2008)). In general, most states have a flagship public university that is selective and resource-intensive, a range of four-year colleges that are more modest in resources and often minimally selective, and two-year community colleges that are

¹¹ Notably, option values of collegiate attainment are not particularly new, dating at least to Weisbrod (1962).

¹² When students are asked why they leave school, many cite financial problems as the deciding factor though such responses do not provide a causal link between attrition and credit constraints (Barr-Telford, Cartwright, Prasil, and Shimmons (2003); Cabrera, Nora, and Castañeda (1992)).

¹³ US institutions are unambiguously overrepresented at the very top of the distribution of institutional quality, accounting for 54 of the top 100 universities in relation to 4.5% of world population or 24.8% (http://www.arwu.org/ARWUAnalysis2010.jsp). It is, of course, more difficult to prove overrepresentation in the bottom tail of the distribution of institutional quality though one British commentator notes: "The U.S., with 4000 institutions of higher education probably has 50 of the best universities in the world and undoubtedly has 500 of the worst" (Stevens as cited in Bowen, Kurzweil, and Tobin (2005)).

open access in mission. In addition, a number of the most resource-intensive colleges and universities in the United States are private, nonprofit universities that benefit from substantial endowment funding (Winston (1999)).

Stratification on the supply side of the higher education market can lead to a fully efficient matching of students and institutional resources. In their classic 1995 paper, Rothschild and White (1995) present a model in which students ("customers") are inputs into the production process and universities charge differentiated prices in relation to the impact of peer inputs in the educational production process. Rothschild and White's key result is that a frictionless decentralized market in which colleges charge zero-profit prices produces an efficient allocation of students to schools.

The Rothschild-White model abstracts from key features of the US system. According to their model, students and their families pay the full cost of college, and the allocation of students to schools is achieved through the price mechanism. In fact, however, students in the United States receive large institutional subsidies, and access is rationed through selective admissions. Rothschild, in unpublished work, extends the Rothschild–White model to the case where resources exogenously vary across universities, while Sallee, Resch, and Courant (2008) explore the optimal allocation of resources across students. In these cases, if one assumes some kind of complementarity between student's ability and college resources, either because better prepared and better motivated students are in a better position to make the most of generous resources or because more capable students benefit disproportionately from having like-minded peers, then efficient allocation will match well-prepared students with highly resourced schools. Although these models seem to suggest that some degree of stratification is efficient, they do not imply that the degree of stratification we observe in US colleges and universities is efficient. Within public postsecondary education systems, there is little reason to believe that the political process of allocating resources among colleges and universities results in the efficient allocation of resources across schools within a state. In fact, dramatic differences across states suggest that not all of these patterns can be efficient. Within the private sector, endowments significantly subsidize current students. In discussing a model in which colleges are partly funded through alumni gifts, Hoxby (2009) notes: "it is harder to claim an efficiency result in an intergenerational model with endowments than in the static Rothschild-White model where student tuition covers the cost of inputs."

As we noted, the Rothschild–White model assumes a frictionless market with full information available to both sides—students and institutions. In reality, of course, colleges have very imperfect knowledge of student applicants and applicants have a very imperfect understanding of what college will be like and their likelihood for success in a particular collegiate environment. We suspect that this imperfect information has important implications for efficiency in the college market. In particular, it seems inevitable that the returns associated with attending highly resourced institutions will induce inefficient, rent-seeking behavior on the part of applicants and their families (Bound, Hershbein, and Long (2009); Abrams (2005)). The more difficult question is whether any gains associated with stratification are sufficient to counterbalance such costs; see Bound, Hershbein, and Long (2009) for a fuller discussion.

4. EVIDENCE ON DETERMINANTS OF DEGREE ATTAINMENT

Conceptually, college degree attainment is determined by the interaction of student attributes (the demand side) and institutional characteristics (the supply side). Academic preparation, finances, and expected returns shape the benefits of attending and completing college for potential students on the demand side of the market. On the supply side of the market, the level and distribution of resources among colleges and universities determine the opportunities available for postsecondary attainment and how students are sorted among colleges and universities will likely have a substantial effect on overall collegiate attainment. In this section, we consider how student characteristics and institutional resources impact college completion.

4.1. Demand-Side Determinants of Attainment

4.1.1 Student Preparation

The preparation that students bring to college is one of the most important determinants of college completion (Bound, Lovenheim, and Turner (2010a); Bowen, Chingos, and McPherson (2009)). Students who have succeeded—or excelled—in the secondary education experiences will require less effort to complete the requirements of a BA degree. Indeed, potential college students' lack of preparation is often cited as one of the most significant barriers to increasing collegiate attainment and degree completion (Carneiro and Heckman (2003)), and it is striking that the countries with higher attainment and completion rates than the United States are also countries in which the achievement of students on standardized tests in the primary and secondary grades is well above that observed for students in the United States (Hanushek, Jamison, Jamison, and Woessmann (2008)). Therefore, efforts to increase collegiate attainment that act to increase the proportion of poorly prepared students who choose to attend college induce a compositional shift in the pool of enrollees. Since the 1970s, enrollment changes in response to increases in the returns to degree completion likely induced more students with weaker preparation to attend college, with lower completion rates one potential result.¹⁴ At the same time, as we discuss in more detail below, there is

¹⁴ In fact the expected change in the college completion rate (degree recipients relative to entrants) in response to an increase in the return to education is ambiguous. Among inframarginal students, the likelihood of completion will increase while those newly induced to enroll may have lower completion rates.

evidence that stagnating completion rates are also partially attributable to declines in the resources available to students attending public college and universities within the United States.

In a paper focused on the explanation for the observed change in completion rates over time, Bound, Lovenheim, and Turner (2010a) measure degree attainment using data from two longitudinal surveys, the National Longitudinal Survey of the High School Class of 1972 (NLS72) and the National Educational Longitudinal Study (NELS:88), which include detailed data on student preparation, information about which college each student attended, and the timing of attendance and graduation outcomes. Students with relatively low academic achievement are unlikely to complete the BA degree in both cohorts of observation. Focusing on differences in collegiate outcomes by a measure of achievement (math test scores in high school), Fig. 8.5 shows the likelihood of college attendance and college completion for high school graduates in both cohorts, as well as completion rates conditional on college attendance. In the bottom quartile of the test score distribution, the likelihood of attending college increases from 21.7 to 44.0%, which is consistent with a larger percentage of lessprepared students attending college in the later cohort in order to take advantage of the rising returns to education. However, among this group, only 5.6% in the initial period of observation receive a BA, and this percent falls yet further to 5.0% for the later cohort. Focusing on college attendees, the likelihood of completing a BA declined from 25.8 to 11.4% across cohorts for those in the bottom quartile of math test scores, whereas it actually increased from 66.8 to 73% across cohorts for those in the top quartile.

There is no question that the cross-cohort change in the representation of college entrants by preparation affected the aggregate of college completion rate. Although college enrollment rates increased across the board, the changes were somewhat larger among those who were relatively unprepared. The proportion of college entrants with test scores below the median rose from 32 to 39% between the two cohorts. Using a logit decomposition analysis, Bound, Lovenheim, and Turner (2010a) estimate that about one-third of the cross-cohort decline in completion rates (from 50.5 to 45.9%) can be explained by the change in student preparedness. (The supply-side determinants of this change are discussed in the next section.)

The tabulations presented in Fig. 8.5 suggest two basic points that frame discussions about increasing the number of college graduates, recalling that the number of college graduates is mechanically equal to high school graduates \times enrollment rate \times completion rate. First, with college enrollment rates well over 90% for high-achieving students, efforts to increase their enrollment in college are likely to yield limited gains. Second, given the low likelihood of degree completion among students with poor secondary achievement, efforts to further increase their enrollment are unlikely to produce gains in the overall number of college graduates.





Source: Bound, Lovenheim, and Turner (2010a). Authors' calculations from the NLS72 and NELS:88 surveys. NLS72 calculations were made using the fifth follow-up weights included in the survey. Fourth follow-up weights were used for the NELS:88 survey calculations. Only those participating in these follow-ups are included in the regression. School-type samples refer to first institution attended.

4.1.2 Student Background and Family Circumstances

Although parental education has increased for cohorts reaching college age in the last three decades (Kane (1994)), family economic circumstances combined with the rising direct cost of college may increase the challenge of paying for many college students. Indeed, increasing income inequality has likely placed students whose families are in the bottom quartile of the income distribution in worse economic circumstances today than three decades ago.

Increases in college costs relative to family income may limit progression through degree programs if credit constraints lead students to increase employment at the cost of reducing their rate of credit attainment. College costs have increased dramatically in recent years and family incomes have not kept pace. Increases in tuition have been sizeable at all types of institutions, with real tuition costs rising by about 247% between academic years 1976–1977 and 2006–2007 at four-year private institutions and by about 266% at four-year public institutions. Net costs faced by aid-eligible families likely increased at a greater rate as the real value of the Pell grant fell from \$4952 in 1976 to \$4050 in 2006, and family incomes—particularly below the median—increased much less rapidly than the rise in tuition. Because tuition is only a fraction of the total cost of full-time attendance (including room and board), it would seem plausible that an increasing fraction of students may be credit constrained.¹⁵ Figure 8.6 shows the expected costs of college relative to family income across the distribution for families of college-age students.

For college students whose families face significant liquidity constraints, we would expect real tuition increases to lead to a decrease in their persistence in college, an increase in the fraction of college they pay relative to their parents, and an increase in their employment. Indeed, as one would predict if a significant number of the families of college-going students were liquidity constrained, Belley and Lochner (2007) show evidence that, conditional on academic preparedness, the association between family income and college enrollment/attainment has increased, along with student employment among high-achieving, low-income, and moderate-income youth. [Notably, the Belley and Lochner analysis, using the data from 1979 to 1997 cohorts of the National Longitudinal Survey of Youth (NLSY), is limited by data constraints in the measure of the link between family circumstances and college completion.]

¹⁵ Researchers have found evidence suggesting that credit constraints are relevant for a subset of the student population and that the resources available to students have important effects on behavior. For example, Kane (1996) finds high tuition at public colleges tends to induce students to postpone college, while Christian (2007) finds evidence that college enrollment tends to be procyclical for students from low-income families. The question of the size of the population constrained by access to credit in college enrollment and choice of college has been debated extensively in the literature (Carneiro and Heckman (2002)). The important analysis of Carneiro and Heckman (2002) identifies substantial differences by family circumstance in choice of college, with low-income students particularly unlikely to attend four-year institutions, though this analysis does not continue to the college years to assess whether family circumstances affect attainment and degree completion.



Figure 8.6 Public Four-Year College Cost as a Function of Family Income.

Source: Four-year tuition, room and board fees data come from table 5a in the report "Trends in College Prices: 2009," released by the College Board. Prices are based on data reported to the College Board by colleges and universities in the Annual Survey of Colleges and are weighted according to the size of full-time enrollment in each institution. Data on family income came from U.S. Census Bureau, Income Surveys Branch. The table used is table F-3, "Mean Income Received by Each Fifth and Top 5 Percent of Families, All Races: 1966 to 2006." The data includes both families with and without dependents aged 18–24 enrolled full time in a four-year college or university.

In practice, it is quite likely that students face some limits in access to capital markets (Becker, 1994). With relatively modest availability of federal aid and limited institutional financial aid funds outside the most affluent colleges and universities, it is plausible that an increasing number of students attend college part time—thus extending time to degree and reducing completion rates—because they are credit constrained and unable to borrow to finance full-time attendance.¹⁶ If students are limited in their capacity to borrow

¹⁶ The maximum Pell grant was \$2400 in 1992 and the borrowing under the Stafford program was limited to \$2625 for first-year dependent students, meaning a low-income student hoping to attend a residential college full time would face substantial unmet need. See Fitzpatrick and Turner (2008) for further discussion.



Figure 8.7 Hours Worked by College Students.

Source: Data are from authors' tabulations using the October CPS. Individual weights are employed.

or reluctant to do so, the rising college costs may increase the incidence of employment while in school, as well.¹⁷

The number of hours worked by college students has increased in recent decades. As shown in Fig. 8.7, Current Population Survey data show the average (unconditional) weekly hours worked among those enrolled in college increased from 9.5 to 12.4 between 1972 and 1992 among 18- to 21-year-old college students, with a further increase to 13.2 hours per week evident in 2005. Consistent with these observations from the CPS, comparison of the NLS72 and NELS:88 cohorts show average (unconditional) weekly hours worked that increased from 7.1 to 14.9 and increased from 22.4 to 28.6 on the intensive margin for the high school graduating classes of 1972 and 1992, respectively.

Estimating the effect of working while in school on collegiate attainment is difficult because the decision to work and the choice of hours of employment are endogenous. Indeed, to our knowledge, only one study has credibly estimated the effect of time worked on any form of academic achievement in college. Stinebrickner and Stinebrickner (2003)

¹⁷ Keane and Wolpin (2001) show that, in a forward-looking dynamic model with limited access to credit, increases in employment while enrolled in school are the expected response to tuition increases.

use data from Berea College in Kentucky, where all students are expected to work and all admitted students have financial need. Estimates of the effect of hours worked on academic performance using plausibly exogenous variation in job assignments point to substantial negative effects of hours worked on credit attainment, which is consistent with a negative effect of hours worked on rate of credit attainment.

Certainly, the available evidence is suggestive that credit constraints and rising college costs act to increase time to degree and to decrease college completion rates, as well as overall collegiate attainment (Bound, Lovenheim, and Turner (2010a,b)). However, we regard this issue as open to further analysis, particularly on the margins of how time spent working for pay affects collegiate attainment and the extent to which students fully access available credit to finance college. Indeed, students' self-reported assessments of their capacity to finance college do not show an unambiguous link between credit constraints and attrition from college.¹⁸

4.2. Supply-Side Determinants of Attainment

Any model of collegiate attainment is naturally associated with a production function including student's achievement and institutional characteristics and resources as primary arguments. Greater resources per student can lead to the provision of a variety of services and educational offerings that may increase completion rates, including greater financial aid, better advising about course selection and finances, better trained faculty, well-equipped laboratories and other facilities, and a variety of cocurricular opportunities such as research experiences and internships that enrich the educational experience and promote completion. In the United States, there is considerable (and increasing) stratification in the level of resources afforded by colleges and universities, and many of these resources come from public and private subsidies beyond tuition (Winston (1999); Bound and Turner (2007)).

There are substantial differences in college completion rates among colleges and universities in the United States. Table 8.1, based on our work with Michael Lovenheim (Bound, Lovenheim, and Turner (2010a)), presents college completion rates for students from the high school graduating classes of 1972 and 1992, initially enrolling at different types of colleges and universities. Students from the 1992 high school cohort starting at selective private colleges and universities graduated at a rate over 90%, whereas those starting at public four-year institutions outside the top 50 completed at a rate less than 57%, and those starting at community colleges completed at a rate of only 17.6%. To be sure, these observed differences also incorporate differences in student attributes, though the differences in completion rates by institutional type

¹⁸ In the context of Berea College, Stinebrickner and Stinebrickner (2008) observe student responses to queries about the effect of financing and need to borrow on persistence decisions and find that the majority of attrition of students from low-income families is attributable to factors other than credit constraints.

	BA Completion Rate		S/F Ratios (Median)		Instructional Exp. (Median)	
	NLS72	NELS:88	NLS72	NELS:88	NLS72	NELS:88
Full sample	50.5	45.9	25.2	30.4	\$4716	\$4339
Public colleges and universities						
Nontop-50 public	61.8	56.9	24.6	27.3	\$5331	\$5102
Top-50 public	73.5	82.5	22.7	22.2	\$7871	\$9663
Private colleges and universities						
Less-selective private	58.2	70.5	20.8	21.9	\$4732	\$5269
Highly selective private	80.1	90.3	19.8	17.6	\$7646	\$13782
Total community college	20.2	17.6	38.5	57.8	\$3068	\$2610

Table 8.1 Completion Rates and Resources per Student by Cohort and Type of Institution

Source: Bound, Lovenheim, and Turner (2010a). Data on faculty, enrollment, and expenditures are from the HEGIS/ IPEDS surveys from the Department of Education. Expenditures per student are for instructional expenditures only. All financial figures are in real 2007 \$ and are deflated by the Higher Education Price Index (HEPI). Tabulations are weighted by the fifth follow-up weights in NLS72 and are weighted by the fourth follow-up weights in NELS:88.

and resources per student remain quantitatively large after being adjusted for student achievement. When we control for high school achievement based on a standardized assessment, we predict a completion rate advantage (relative to attending a public college outside the top 50) of 35 percentage points for attending a highly selective private institution and a completion rate advantage of about 24 percentage points for attending a top-50 public university, whereas the penalty in the likelihood of completion associated with attending a community college is about 32 percentage points. Notably, the distinctions in completion rates by type of initial institution increased between the high school class of 1972 and the high school class of 1992, with completion rates declining at the less selective public institutions.

There are clear distinctions in the level of resources per student among the institutional categories presented in Table 8.1. Highly selective institutions in the private sector spend appreciably more per student than institutions in the public sector and these differences are, in turn, strongly associated with completion rates. Not only the cross-sector differences in resources per student are substantial but also they have increased markedly over time.

The economic rewards associated with attending a more resource-intensive college operate through two channels. First, as noted above, students at colleges and universities with relatively generous resources have relatively high college completion rates, with degree completion receiving a substantial return in the labor market. In addition to the evidence noted above from Bound, Lovenehim, and Turner (2010a), a range of other empirical evidence substantiates this link (Light and Strayer (2000); Roderick, Nagaoka, Coca, and Moeller (2008)). Beyond the benefits of degree receipt per se, there is a substantial empirical link between college quality and earnings for college graduates (e.g., Hoxby (2001); Black, Daniel, and Smith (2005); Hoekstra (2009)).

Yet, despite the strong evidence of the association between institutional resources and college completion rates, the research evidence on the postsecondary production function or the causal link between inputs and degree attainment is decidedly underdeveloped. There is little evidence, for example, about how the organization of instruction, such as small classes or residential programs, affects student outcomes and degree completion.

4.2.1 Changes in Resources and Stratification over Time

The last three decades have seen an unambiguous upward trend in expenditures per student in higher education. To illustrate, constant dollar current expenditures per student at public colleges and universities have risen from \$14,610 in 1970–1971 to \$17,606 in 1990–1991 to \$22,559 in 2000–2001 (Snyder, Tan, and Hoffman (2006), Table 339). Yet, this aggregate change in resources per student has not been followed with increases in the college completion rate. One point of note is that changes in spending per student combine changes in the price of educational inputs with changes in quantities. Higher education is a very labor-intensive sector, with (arguably) few opportunities to substitute capital for labor as faculty costs increased. Still, in the aggregate, rising faculty–student measures suggest overall increases in resources per student at colleges and universities.

What the aggregate trends obscure is the growing stratification in resources per student and the associated widening in completion rate outcomes across institutions. Focusing first on the period from the early 1970s to the mid-1990s, Bound, Lovenheim, and Turner (2010a) compute completion rates by institution of first entry and resources per student across different institutional types, from highly selective private universities to open-access public four colleges and community colleges. These comparisons, which are presented in Table 8.1, illustrate the striking differences in the changes among institutions over this period. First, while the median college entrant experiences a decline in resources between the high school class of 1972 and 1992, students at private colleges and universities were likely to experience a notable increase in instructional expenditures per student. Second, while the college completion rate reduced overall, this aggregate result combines the rise in completion rates at relatively resource-intensive institutions (private colleges and top public universities) and the decline in completion rates for students starting at less-selective public four-year colleges and community colleges. Finally, the distribution of students among institutions shifted dramatically over this interval with a relative increase in the share of students beginning at community colleges and a decrease in the share of students beginning at the more-selective four-year institutions.¹⁹ Although it is hard to credibly estimate

¹⁹ Between college entrants from the high school class of 1972 and the high school class of 1992, the share starting at community colleges increased from 31.2% to 43.7%, the share beginning at public four-year institutions declined from 46.7% to 37.6%, and the share beginning at private four-year institutions declined from 22.1% to 18.7% (Bound, Lovenheim, and Turner (2009).

the causal effect of this shift in resources on college completion rates, our estimates suggest that these changes on the supply side of the market can explain the observed decline in the completion rates.²⁰

It is important to emphasize that the demand- and supply-side explanations described above are not mutually exclusive: less-prepared students sort into the most elastic sectors of higher education that tend to have the fewest resources. In essence, increased demand for college crowds more students (and more less-prepared students) into community colleges and nontop-50 public universities. Therefore, increases in demand not only lower the resources per student at these institutions but also cause higher dispersion in resources across the sectors of higher education. In effect, changes in demand are not evenly distributed among the hierarchy of institutions (Bound and Turner (2007)). Although those institutions with the greatest resources are unlikely to expand enrollment with increases in student demand, open-access public institutions and community colleges are relatively elastic in supply. Such changes increase stratification in resources because public subsidies are spread more thinly at the least selective institutions, in turn placing downward pressure on completion rates.

Without more recent longitudinal data following students from college enrollment to BA receipt, we are unable to extend the Bound, Lovenheim, and Turner (2010a) analysis to more recent cohorts. Still, we are able to examine the continuing change in the distribution of resources across institutions and student's enrollment decisions. Table 8.2 builds on the careful work of the Delta College Cost Project and shows enrollment and expenditures by type of institution from the late 1980s to 2008. The key message from this table is the continued increase in resource stratification across different types of institutions during this period. Although instructional expenditures per student increased by about 43% between 1987 and 2007 at the most selective private colleges and universities, real instructional expenditures were essentially flat at public colleges and universities, evidencing a modest rise between 1987 and 1997 followed by a slight decline between 1997 and 2007. This same trend is also apparent when we examine the broader classification of educational expenses that includes academic support services in the middle columns of Table 8.2. Yet, the revenue sources shifted appreciably over this period as state and local appropriations to public colleges and universities fell and were replaced by increased reliance on tuition in public sector institutions. The result, shown in the last columns of Table 8.2, is the appreciable decline in the higher education subsidy for students attending public colleges and universities. This decline in subsidy associated with attendance at a public colleges contrasts sharply with the case of the most selective privates where spending per student has increased while

²⁰ When one accounts for declines in student preparation as well, Bound, Lovenheim, and Turner (2009) "over-explain" the total observed decline, which is consistent with the other student background characteristics shifting in ways that would suggest an increase in completion.

	Instructional Expenditures per Student (2007\$), Median		Education and General Exp. per Student (2007\$), Median			Net Subsidy per Student (2007\$), Median			
	1987	1997	2007	1987	1997	2007	1987	1997	2007
Public colleges and universities	3								
Nontop-50 public	4993	4980	4871	11,544	12,467	11,735	6517	5847	4915
Top-50 public	9039	10,230	9752	24,381	28,596	27,654	9407	8722	7491
Private colleges and universitie	s								
Less-selective private	4592	4926	5408	15,145	14,782	14,475	3086	2344	2563
Highly selective private	11,037	14,105	15,797	32,396	35,858	40,538	6273	10,462	11,872
Total community colleges	2834	2951	2783	6416	6999	6468	4359	4250	3880

Table 8.2 Expenditures and Subsidies by Type of Institution

Source: Microdata from "Trends in College Spending: 1998-2008" A Report of the Delta Cost Project," http://www.deltacostproject.org/analyses/delta_reports.asp. "Instructional Expenditures" includes activities directly related to instruction, including faculty salaries and benefits, office supplies, and administration of academic departments. "Education and General" expenditures includes spending for instruction and student services, plus spending on academic and institutional support and for operations and maintenance of buildings. The average subsidy measure reflects the difference between "Education & General" expenditures and tuition revenues exclusive of institutional financial aid awards. Note that tabulations in Tables 8.1 and 8.2 are not strictly comparable; Table 8.1 figures are weighted by the enrollment of firsttime students in the high school graduating classes of 1972 and 1992.

the subsidy provided through nontuition sources such as endowment income has also increased.

4.3. College Matching as a Determinant of Attainment

The substantial effects of type of institution and institutional resources noted above raise significant questions about how students are matched with colleges and universities initially. Indeed, the effects of initial college choice are so large that the analysis of how students match to colleges should be included in the consideration of the determinants of completion rates. What is more, family circumstances interact substantially with initial college choice as high-achieving students from relatively affluent families are much more likely to attend the most highly ranked institutions (and less likely to attend community colleges) than their less-affluent peers (Table 8.3). To this end, how students from different circumstances negotiate the college choice process may have substantial impacts on degree attainment.

"Matching" of students to colleges can be quantified in terms of the alignment between student's achievement and the achievement of the overall population of students at a given college (Roderick, Nagaoka, Coca, and Moeller (2008); Dillon and Smith (2009)). To be sure, optimal matching goes well beyond maximizing college selectivity to a more general definition: "... finding the right college means more than gaining acceptance at the most competitive college possible. It is about finding a place that is a good 'fit:' a college that meets a student's educational and social needs, as well as one that will best support his or her intellectual and social development. Match is just one consideration of the larger process of engaging in an effective college search, but it is also an important indicator of whether students are engaged more broadly in a search that incorporates the larger question of fit" (Roderick, Nagaoka, Coca, and Moeller (2008)). College matching is, in many ways, similar to other ways that market participants engage in search—either for jobs or housing—where high costs of search produce low match quality for particular demographic groups.

Recent empirical work highlights the magnitude of potential "undermatch." Bowen, Chingos, and McPherson (2009) link secondary academic records and collegiate outcomes for the state of North Carolina and find that completely 40% of the students with SAT scores and high school grades in the range needed to enroll at a very selective university failed to do so, with this finding appreciably more pronounced among students in the bottom quartile of family income (59% of SAT-taking students) than among students from the top quartile (27%). Such evidence of "undermatching" is well documented in other areas including Chicago (Roderick, Nagaoka, Coca, and Moeller (2008)) and Virginia (Avery and Turner (2010)).

Low-income students apply to fewer colleges and colleges that are less aspirational than their more-affluent peers. Pallais and Turner (2006) find that low-income students

Table 8.3 Distribution of College Enrollment and Completion by Type of Institution Enrollment by Family Income and Test Scores, NELS:88, 1992 High School Graduates

		Distribution of Enrolment by Type of institution						
	Overall Enrollment Rate (%)	Public 4-Year Nontop 50	Public 4-Year Top 50	Private 4-Year Less Selective	Private 4-Year Highly Selective	Community College		
Family Incom	ne < \$50,000							
Lowest Math	45.1	29.3	0.7	6.0	0.0	64.1		
Second Math	62.0	33.4	2.2	9.6	0.1	54.8		
Third Math	73.9	41.6	4.4	12.4	1.3	40.4		
Top Math Quartile	89.2	41.3	13.9	15.8	6.5	22.4		
Family Income $> $ \$50 000								
Lowest Math Quartile	53.3	33.0	1.9	6.6	0.0	58.5		
Second Math	80.7	31.4	8.4	13.5	2.2	44.5		
Quartile Third Math Quartile	90.6	34.9	9.5	14.6	4.7	36.3		
Top Math Quartile	97.3	32.3	21.2	17.9	16.5	12.1		

Distribution of Enrollment by Type of Institution

Source: Authors' calculations from the NLS72 and NELS:88 surveys. NLS72 calculations were made using the fifth follow-up weights included in the survey. Fourth follow-up weights were used for the NELS:88 survey calculations. Only those participating in these follow-ups are included in the regression. School-type samples refer to first institution attended.

are less likely than their more-affluent peers to send standardized test scores to top private universities and top liberal arts colleges; in addition, these students are generally less likely to send scores to public flagship universities. More than 70% of students from low-income families who are very high achieving and would likely be strong candidates for admission at many top colleges and universities, as indicated by SAT scores over 1400, do not send scores to a single institution in the "most competitive" category, with only 8% following the guidance of the College Board and professional admission counselors of establishing an application portfolio that includes five to eight schools. Relatively high-income students apply to more colleges and include somewhat higherquality institutions (as measured by test scores of matriculating students) in their choice sets than their low-income counterparts do. To the extent that more applications generate more choices after admission decisions, wide differences in college application behavior by family circumstances are a significant determinant of the income gap in collegiate outcomes conditional on student's achievement.

The question of how students make decisions about the set of schools to which to submit applications can be motivated by a utility maximizing model in which a student considers the benefits that would follow from each potential collegiate option, the likelihood of admission, and the costs (both pecuniary and non-pecuniary) of application.²¹ As has been widely noted, the informational requirements necessary for a student to choose the optimal set of colleges are actually quite onerous, requiring the evaluation of many combinations of colleges while assessing the benefits of even one additional application may be difficult if information about the benefits of a particular college choice or the likelihood of admission may be difficult to ascertain (Pallais (2009)). The information requirements and data necessary for individuals to assess this problem may impose a particularly large burden for first-generation college students, low-income students, and students without extensive networks of peers making similar choices. Indeed, there is a well-established market for professional private admission counselors to help students develop application strategies and navigate the application process. Although it is possible that differences in application behavior by socioeconomic circumstances follow from broad searches that fully consider preferences and family finances, the weight of existing evidence suggests that just the opposite is true as the absence of information and guidance contributes to quite-limited application strategies among low-income students (Roderick, Nagaoka, Coca, and Moeller (2008)).

5. UNANSWERED QUESTIONS IN THE STUDY OF DEGREE ATTAINMENT

That college completion is a central outcome of higher education and a critical input for labor-market success and economic growth is not in dispute. More complicated is the question of whether reasoned investments at the postsecondary level can substantially change the number of college graduates entering the labor force. Our first conclusion

²¹ As noted by Chade, Lewis, and Smith (2006), the problem is trivial in the case of certainty in admissions (a student simply applies to the best college to which he or she will be admitted) or in the case of no admissions costs (a student would apply to all colleges). Howell (2004) and Pallais (2009) present similar models characterizing the expected utility associated with any application strategy chosen by an individual as a function of the expected utility associated with attendance at any particular institution, the probability of each of these institutions, and the cost of application. Note that there are 2ⁿ potential combinations of application sets and with 4314 postsecondary institutions and 2629 four-year universities this number is potentially quite large, though practically somewhat smaller if one were to distinguish national and regional institutions.

is that substantial further investments to increase college *enrollment* are unlikely to have an appreciable effect on the number of college graduates. Although the numbers are not zero, we find little evidence of large numbers of students well prepared to complete college, who are not already enrolling.

Improving the quality of K–12 education has potentially large impacts on collegiate attainment. However, the evidence suggests that there are also potential opportunities to improve graduation rates that involve the tertiary sector. The evidence we have discussed earlier in this chapter suggests that increased public investments to raise resources available to those students who are presently attending colleges outside the public flagship universities and the most selective private institutions would lead to gains in persistence to degree attainment. Moreover, there may also be opportunities to increase the efficiency of the resources currently being used at the collegiate level through the adoption of technology or improved student guidance, though the mechanisms to achieve such gains are not well understood. Further, we observe many students—particularly those from the least-advantaged circumstances—who appear well prepared to benefit from resource-intensive college experiences but instead attend colleges and universities with low-funding levels and poor-graduation prospects. For this reason, we hypothesize that policies that operate at the margin of college choice may improve persistence to degree completion.

The study of college degree attainment has traditionally focused on demand-side determinants of attainment, including how students finance college attainment and academic preparation. Yet, basic models of collegiate attainment that assume an elastic supply-side determinant of the market likely miss an important determinant of degree attainment. The incomplete adjustment of public subsidies in responses to changes in population leads to clear evidence of differences in college degree attainment and the rate of college completion among students who enroll (Bound and Turner (2007); Bound, Lovenheim, and Turner (2010a)). Moreover, as the stratification in collegiate resources has increased in recent decades, so has the difference among institutions in degree outcomes.

Review of this evidence and research suggests a number of unexplored areas for economic research related to college choice, in-college attainment, and the supply-side determinants of stratification and resources per student. Indeed, there has been much thoughtful new research in the United States and abroad in the last decade in the area of the economics of education and, more recently, a wave of attention to postsecondary attainment.

5.1. College Choice

It is well documented that many students do not apply to or attend the institutions with greatest resources or likelihood of successful attainment for which they are qualified. Although it is widely suggested that there is a "market failure" in the college choice
process, the barriers to optimal choice are poorly understood. Among the unanswered questions, some are as follows:

- 1. To what extent do students and their parents understand differences among colleges in expected graduation rates, future incomes, and the associated academic requirements of different courses of study?
- 2. To what extent do students and their parents understand the distinction between the posted tuition price and the "net" cost of college given the availability of financial aid? Do students fully understand the availability and functioning of different types of aid including grants, loans, and college work study?
- **3.** To what extent do students from different academic and economic backgrounds receive appropriate guidance from school counselors?
- **4.** Do students (as well as parents and guidance counselors) make mistakes on process including missing deadlines as failing to complete the FAFSA that limits college choice and enrollment?
- 5. What sources of information do students and their families rely on when making decisions about both preparing for college and applying to various schools. When students make sensible choices, is this because they understand their options, or is this because they are guided in their choices?

In short, there is much that is not well known about the nature of information problems potentially impeding college choice that needs to be ascertained or identified before policy makers will be able to use the lever of improving consumer choice as a successful lever to improve the allocation of students among institutions.

5.2. In-College Experiences

How the organization of the college production function and the choices made by students during college affect attainment remain questions that are not well addressed in the current research literature. Indeed, there is some risk that the research focus on college choice will crowd-out the study of behaviors and resources during college. In hypothesizing about why students leave college without receiving a degree, the research literature has posited many ideas ranging from learning about own ability to clear "mistakes" in the utilization of financial aid or the navigation of complicated collegiate requirements. On the supply side of the market, it is unambiguously clear that there is variation in completion rates that cannot be explained completely by student's characteristics and is associated with institutional resources. Yet, it is far from clear "how" and "why" resources impact collegiate attainment. In short, the production function for higher education is not well understood, and as such, it is hard to infer how changes in the organization of college education such as the introduction of more on-line instruction, changes in class size, and so forth would affect attainment. To understand why students do not complete college, we suspect that existing data, whether it would be the NCES longitudinal data sets or the administrative data sets increasingly used by researchers, are not well suited for the task. The in-depth longitudinal evidence on the experiences that students have as they navigate college outcomes collected by Stinebrickner and Stinebrickner for a sample of students enrolled at Berea College represents an example of the kind of innovative approach to measuring the determinants of student persistence that we have in mind (Stinebrickner and Stinebrickner (2008, 2010)). The work by Stinebrickner and Stinebrickner and Stinebrickner is clearly limited by their focus on a single, very atypical college. However, we believe that a great deal can be learned from their work, both substantively and methodologically.

In addition, carefully designed experiments to understand how program design, advising, course structure, and other academic resources impact student's degree attainment are an important tool to understand "what works" and the higher education production function. For example, to test the hypothesis that low ratios of academic counselors per student at community colleges contribute to suboptimal course selection and degree progress, the MDRC Opening Doors counseling experiment in Ohio community colleges randomly assigned students to receive extra guidance and provided a modest stipend to low-income participants with the expectation that students would meet their counselor at least two times per semester for two semesters to discuss academic progress and resolve any issues that might affect their schooling (Scrivener and Weiss (2009)). Although the treatment group did earn slightly more credits than the control group, the program did not have a substantial impact on credits earned after the termination of treatment. In the context of a randomized control trial, Angrist, Lang, and Oreopoulous (2009) find that the provision of modest incentives for academic performance combined with study assistance had a modest impact on grades and credit attainment at a Canadian university. Despite the relatively modest results from these experimental innovations, there is much to be gained from incorporating well-designed strategies for evaluation in university-level program innovations.

5.3. Market Structure: Competition and Stratification

The unique mixed-market institutional structure of higher education in the United States—with a combination of nonprofit, for-profit, and public providers and a mix of funding from student, philanthropic, state, and federal sources—presents many challenges for textbook models of organizational behavior and industrial organization. The contrast with postsecondary systems in other countries with greater public control and public support is striking. How institutional control and sources of revenue affect the nature of competition in the market, the level of stratification among institutional offerings and student outcomes remain open questions for those working in applied I/O and applied microeconomics more generally. It is reasonable to conjecture

that greater endowment support (and the growth thereof) and research support lead to more market stratification, though the proof of this proposition is not established. As emphasized here and other work (Hoxby (1997); Hoxby (2009); Bound, Hershbein, and Long (2009)), the stratification of resources in higher education has increased dramatically in the last three decades among US institutions.²² These substantial changes in the distribution of resources likely have important implications for degree receipt and future returns, given that a substantial share of enrollment expansion has occurred at community colleges and open-access public institutions in recent decades.

Indeed, given the importance of market structure in determining the distribution of resources among students, there are rich opportunities for applied theoretical work that builds on the framework set forth in papers like Rothschild and White (1995) to incorporate market imperfections. Limited access to credit markets and information barriers are likely to be primary factors that impeded efficient allocation of students among institutions. Further work on socially efficient pricing strategies and allocation of public subsidies in the presence of such constraints may have a high return.

Moreover, empirical evidence on how the nature of competition in local and regional higher education markets impacts student choice and degree attainment is decidedly limited. One hypothesis following from evidence in the K–12 arena is that in markets where students have greater choice between public and private institutions, productive efficiency will be enhanced. It should be straightforward to investigate whether those regions in which students have more choices owing to historical differences in the presence of private colleges or the location of public universities have lower educational costs and higher rates of degree completion.

5.4. Political Economy and Public Subsidies

The evidence on the stagnant rate of collegiate attainment in the United States together with the decline in public subsidies for tertiary schooling leads to reconsideration of long-standing questions about "who pays?" and "who benefits?" from public investments in higher education. Although these questions were first asked in higher education nearly four decades ago (Hansen and Weisbrod (1969)), the answers remain somewhat elusive. Johnson (2006) re-examines the issue using recent data and has argued that public subsidies for higher education are, if anything mildly progressive, though Johnson did not try to estimate the general equilibrium effects of such subsidies.

²² To what extent these shifts have work to increase the efficiency of higher education in the United States as Hoxby has sometimes argued or done the reverse as Bound, Hersbhein, and Long (2009) suggest is possible, remains an open question. Credible evidence on this question will be hard to come by. Bound, Hersbhein, and Long try to shed some light on this question by comparing outcomes across states. Alternatively, one might imagine cross-country comparisons, though such comparisons are always open to alternative interpretations.

Related questions involve the incentives states have to invest in higher education. One reason might be that such investments benefit the state taxpayers who send their children to state universities. Beyond this, it has often been argued that states benefit from having have educated workforces (Goldin and Katz (1999)). However, given the mobility of college-educated labor, investing in educating college-age individuals may not be that effective in increasing the number of college graduates in a state's workforce (Bound et al. (2004); Kennan (2010)). More generally, the local effects of having vibrant colleges and universities in a state are not well understood.

State spending on higher education relative to per capita income has unambiguously declined over the last quarter century (College Board (2010)). One explanation for this decline involves crowd-out from mandated state budget items such as Medicaid which serves to magnify contractions in higher education spending during cyclical downturns (Kane, Orszag, and Gunter (2003)). An alternative hypothesis for declines in state support for higher education is that the magnitude of local spillovers from colleges and universities has declined as the mobility of workers has increased. Research efforts to understand the relationships among local spillovers to public investments in colleges and universities, worker mobility, and state support for higher education may serve to identify the determinants of state subsidies.

5.5. International Comparisons

While much of the evidence on college degree completion in this chapter comes from the US experience, we emphasize that there is a potentially high return from international comparative assessments. For example, while the mechanism for college choice in the U.S. is highly decentralized, relying on individual application and institutional evaluations, other countries maintain more centralized systems of college admission (such as the UCAS system in the United Kingdom). One question to consider is how alternative application and admission regimes affect the distribution of students across colleges and universities with different levels of resources. In addition, differences across countries in the level and distribution of public support for higher education and funding mechanisms, including tuition and financial aid, may have substantial effects on the sorting of students and college completion. Although direct assessments using cross-country variation may be limited by data constraints, cross-country comparisons are one of the few avenues to provide evidence on the efficacy of different mechanisms for matching students to colleges and the extent to which greater stratification in resources affect collegiate attainment.

5.6. Economics of College Degree Completion and Public Policy

The focus of this chapter is squarely on the importance of distinguishing the collegiate degree attainment process from very general models of educational attainment. Efforts to more fully model the features of degree attainment which take into account the salient features of student college choice and persistence decisions and the important dimensions of the supply side of the postsecondary market will yield substantial returns both in the academic study of the economics of education and the application of this analysis to key questions for public policy.

Because there are many important unanswered questions about the nature of student choice and college and university behavior in the market, considerable caution is warranted in targeting degree outcomes as a policy objective. For example, adding rewards (or sanctions) to institutional degree completion rates is a policy tool that likely has many consequences unintended by the proponents of increased collegiate attainment as one clear incentive is the dilution of standards and requirements for completion. Generating many low-quality degrees would likely be the outcome, ultimately imposing costs on those individuals with nominal degrees and few labor-market rewards. To this end, it is important to recognize that the degree is not the end objective of the collegiate experience but rather the production of skills that are recognized and rewarded by the labor market.

A further cautionary note follows from the recognition that individuals bring diverse skills and interests to postsecondary education and career selection. To this end, it is naïve to assert that everyone will benefit from collegiate attainment ("That leaves only one path out of poverty: education—a college education," *Postsecondary Success: focusing on completion*, Bill & Melinda Gates Foundation (2009)). Beyond the examples of individuals of incredible wealth who did not complete college, there are many others who achieve personal satisfaction and success in specialized crafts without college degrees. The optimal college completion rate is not 100%, even as it may be higher than the rate currently observed, particularly for students from modest economic circumstances.

REFERENCES

Abrams, S.J., 2005. Unflagged SATs. Educ. Next 5 (3).

- Angrist, J., Lang, D., Oreopoulos, P., 2009. Incentives and Services for College Achievement: Evidence from a Randomized Trial. Am. Econ. J. Appl. Econ. 1 (1), 136–163.
- Arcidiacono, P., 2004. Ability Sorting and the Returns to College Major. J. Econom. 121 (1-2), 343-375.
- Arcidiacono, P., Hotz, V.J., Kang, S., 2010. Modeling College Major Choices Using Elicited Measures of Expectations and Counterfactuals. NBER Working Paper No. 15729.

Avery, C., Turner, S., 2010. Aid and Application Awareness. Unpublished Manuscript.

- Barr-Telford, L., Cartwright, F., Prasil, S., Shimmons, K., 2003. Access, Persistence and Financing: First Results from the Postsecondary Education Participation Survey (PEPS). 81-595-MIE2003007. Statistics Canada, Ottawa, Canada.
- Becker, G. (Ed.), 1994. Human Capital Revisited: Human Capital, third ed. University of Chicago Press, Chicago, pp. 15–28.
- Becker, G., 1967. Human Capital and the Personal Distribution of Income: An Analytical Approach. University of Michigan, Ann Arbor, MI.
- Belley, P., Lochner, L., 2007. The Changing Role of Family Income and Ability in Determining Educational Achievement. J. Hum. Capital 1 (1), 37–89.

- Belman, D., Heywood, J., 1991. Sheepskin Effects in the Returns to Education: An Examination of Women and Minorities. Rev. Econ. Stat. 73 (4), 720–724.
- Belman, D., Heywood, J., 1997. Sheepskin Effects by Cohort: Implications of Job Matching in a Signaling Model. Oxf. Econ. Pap. 49 (4), 623–637.
- Bettinger, E., Long, B., 2009. Addressing the Needs of Under-Prepared College Students: Does College Remediation Work? J. Hum. Resour. 44 (3), 736–771.
- Bill & Melinda Gates Foundation, 2009. Postsecondary Success: Focusing on Completion. Bill & Melinda Gates Foundation, Seattle, WA. http://www.gatesfoundation.org/learning/Documents/postsecondaryeducation-success-plan-executive-summary.pdf (accessed 04.05.11.).
- Black, D., Daniel, K., Smith, J., 2005. College Quality and Wages in the United States. Ger. Econ. Rev. 6 (3), 415–443.
- Black, D., Smith, J., 2004. How Robust Is the Evidence on the Effects of College Quality? Evidence from Matching. J. Econom. 121(1–2), 99–124.
- Black, D., Smith, J., 2006. Evaluating the Returns to College Quality with Multiple Proxies for Quality. J. Lab. Econ 24 (3), 701–728.
- Bound, J., Groen, J., Kezdi, G., Turner, S., 2004. Trade in university training: Cross state variation in the production and use of college educated labor. J. Econom. 121 (1-2), 143–173.
- Bound, J., Hershbein, B., Long, B.T., 2009. Playing the Admissions Game: Student Reactions to Increasing College Competition. J. Econ. Perspect. 23 (4), 119–146.
- Bound, J., Lovenheim, M., Turner, S., 2010a. Why Have College Completion Rates Declined? An Analysis of Changing Student Preparation and Collegiate Resources. Am. Econ. J. Appl. Econ. 2 (3), 129–157.
- Bound, J., Lovenheim, M., Turner, S., 2010b. Increasing Time to Baccalaureate Degree in the United States. NBER Working Paper 15892.
- Bound, J., Turner, S., 2002. Going to War and Going to College: Did the G.I. Bill Increase Educational Attainment. J. Lab. Econ. 20 (4), 784–815.
- Bound, J., Turner, S., 2007. Cohort Crowding: How Resources Affect Collegiate Attainment. J. Public Econ. 91(5–6), 877–899.
- Bowen, W., Chingos, M., McPherson, M., 2009. Crossing the Finish Line: Completing College at America's Public Universities. Princeton University Press, Princeton.
- Bowen, W., Kurzweil, M., Tobin, E., 2005. Equity and Excellence in American Higher Education. University of Virginia Press, Charlottesville, VA.
- Brewer, D., Eide, E., Ehrenberg, R., 1999. Does it Pay to Attend an Elite Private College? Cross-Cohort Evidence on the Effects of College Type on Earnings. J. Hum. Resour. 34 (1), 104–123.
- Brown, M., Scholz, J., Seshadri, A., 2010. A New Test of Borrowing Constraints for Education. Mimeo, At http://www.ssc.wisc.edu/~scholz/Research/Ed_Borrowing_Constraints.pdf.
- Cabrera, A., Nora, A., Castañeda, M., 1992. The Role of Finances in the Persistence Process: A Structural Model. Res. High. Educ. 33 (5), 571–594.
- Card, D., 1999. The causal effect of education on earnings. In: Ashenfelter, O., Card, D. (Eds.), Handbook of Labor Economics, vol. 3A. Elsevier, Amsterdam, pp. 1801–1863.
- Card, D., 1995. Earnings, Schooling and Ability Revisited. In: Polochek, S. (Ed.), Research in Labor Economics. vol. 14. JAI Press, Greenwich, CT, pp. 23–48.
- Carneiro, P., Heckman, J., 2002. The Evidence on Credit Constraints in Post-Secondary Schooling. Econ. J. 112 (482), 705–734.
- Carneiro, P., Heckman, J., 2003. Human Capital Policy. In: Heckman, J.J., Krueger, A.B. (Eds.), Inequality in America. MIT Press, Cambridge, MA, pp. 77–240.
- Chade, H., Lewis, G., Smith, L., 2006. The College Admissions Problem under Uncertainty. Society for Economic Dynamics Meeting Papers, no. 125.
- Christian, M., 2007. Liquidity Constraints and the Cyclicality of College Enrollment in the United States. Oxf. Econ. Pap. 59 (1), 141–169.
- College Board, 2010. Trends in College Pricing. The College Board, Washington, DC. http://trends .collegeboard.org/downloads/College_Pricing_2010.pdf (accessed 20.06.11.).
- Dale, S.B., Krueger, A., 2002. Estimating the Payoff to Attending a More Selective College: An Application of Selection on Observables and Unobservables. Q. J. Econ. 117 (4), 1491–1527.

- DeLong, B., Goldin, C., Katz, L., 2003. Sustaining U.S. Economic Growth. In: Aaron, H., Lindsay, J., Nivola, P. (Eds.), Agenda for the Nation. Brookings Institution, Washington, DC, pp. 17–60.
- Dillon, E., Smith, J., 2009. Determinants of Mismatch between Student Ability and College Quality. Mimeo. Population Association of America, Detroit, MI. http://paa2009.princeton.edu/download .aspx?submissionId=91536 (accessed 20.06.11.).
- Fitzpatrick, M., Turner, S., 2008. Blurring the Boundary: Changes in the Transition from College Participation to Adulthood. In: Danziger, S., Rouse, C. (Eds.), The Economics of the Transition to Adulthood. Russell Sage Foundation, New York.
- Goldin, C., Katz, L., 1999. The Shaping of Higher Education: The Formative Years in the United States, 1890 to 1940. J. Econ. Perspect. 13, 37–62.
- Goldin, C., Katz, L., 2008. The Race between Education and Technology. Harvard University Press, Cambridge, MA.
- Green, H., 1946. The Origin of the A. B. Degree. J. High. Educ. 17 (4), 205–10. At http://www.jstor .org/stable/1974966.
- Griliches, Z., 1970. Notes on the role of education in production functions and growth accounting. In: Lee Hansen, W. (Ed.), Education Income and Human Capital, NBER Studies in Income and Wealth, vol. 35. Columbia University Press, New York.
- Grogger, J., Eide, E., 1995. Changes in College Skills and the Rise in the College Wage Premium. J. Hum. Resour. 30 (2), 280–310.
- Hansen, L., Weisbrod, B., 1969. The Distribution of Costs and Direct Benefits of Public Higher Education: The Case of California. J. Hum. Resour. 4 (2), 176–191.
- Hanushek, E., Jamison, D., Jamison, E., Woessmann, L., 2008. Education and Economic Growth. Educ. Next 8 (2), 62–70.
- Heckman, J., Lochner, L., Todd, P., 2003. Fifty Years of Mincer Earnings Regressions. NBER Working Paper No. 9732. At www.nber.org/papers/w9732.
- Heckman, J., Lochner, L., Todd, P., 2006. Earnings Functions, Rates of Return and Treatment Effects: The Mincer Equation and Beyond. In: Hanushek, E., Welch, F. (Eds.), Handbook of the Economics of Education. Elsevier, Amsterdam.
- Heckman, J., Lochner, L., Todd, P., 2008. Earnings Functions and Rates of Return. NBER Working Paper No. 13780. At www.nber.org/papers/w13780.
- Hoekstra, M., 2009. The Effect of Attending the Flagship State University on Earnings: A Discontinuity-Based Approach. Rev. Econ. Stat. 91 (4), 717–724.
- Howell, J., 2004. A Structural Equilibrium Model of the Market for Higher Education: Assessing the Impact of Eliminating Affirmative Action. PhD dissertation, University of Virginia, Charlottesville, VA.
- Hoxby, C.M., 1997. How the changing market structure of U.S. higher education explains college tuition. National Bureau of Economic Research Working Paper No. 6323. At http://www.nber.org/papers/ w6323 (accessed 20.06.11.).
- Hoxby, C., 2001. The return to attending a more selective college: 1960 to present. In: Devlin, M., Meyerson, J. (Eds.), Forum Futures Exploring the Future of Higher Education, 2000 Papers. Jossey-Bass, San Francisco, pp. 13–42.
- Hoxby, C., 2009. The Changing Selectivity of American Colleges and Universities: Its Implications for Students, Resources, and Tuition. J. Econ. Perspect. 23 (4), 95–118.
- Hoxby, C., Terry, B., 1999. Explaining rising income and wage inequality among the college educated. NBER Work. Pap. 6873.
- Hungerford, T., Solon, G., 1987. Sheepskin Effects in the Returns to Education. Rev. Econ. Stat. 69 (1), 175–177.
- Jaeger, D., Page, M., 1996. Degrees Matter: New Evidence on Sheepskin Effects in the Returns to Education. Rev. Econ. Stat. 78 (4), 733–740.
- Johnson, W., 2006. Are public subsidies to higher education regressive? Educ. Finance Policy 1 (3), 288-315.
- Jorgenson, D.W., Griliches, Z., 1967. The Explanation of Productivity Change. Rev. Econ. Stud. 34 (3), 249–283.

- Kane, T.J., 1994. College entry by blacks since 1970: The role of college costs, family background, and the returns to education. J. Polit. Econ. 102 (5), 878–911.
- Kane, T., 1996. College Cost, Borrowing Constraints and the Timing of College Entry. East. Econ. J. 22 (2), 181–194.
- Kane, T., 2006. Public intervention in post-secondary education. In: Hanushek, E., Welch, F. (Eds.), Handbook of the Economics of Education. Elsevier, Amsterdam, pp. 1369–1402.
- Kane, T., Orszag, P., Gunter, D., 2003. State Fiscal Constraints and Higher Education Spending. In: Urban-Brookings Tax Policy Center Discussion Paper No. 11, Brookings Institution, Washington, DC.
- Kane, T., Rouse, C., 1995. Labor Market Returns to Two- and Four-year College. Am. Econ. Rev. 85 (3), 600–614.
- Keane, M., Wolpin, K., 2001. The Effect of Parental Transfers and Borrowing Constraints on Educational Attainment. Int. Econ. Rev. 42 (4), 1051–1103.
- Kennan, J., 2010. Higher Education Subsidies and Human Capital Mobility. http://www.ssc.wisc.edu/ ~jkennan/research/HumanCapitalMobility.pdf.
- Lange, F., Topel, R., 2006. The Social Value of Education and Human Capital. In: Hanushek, E., Welch, F. (Eds.), Handbook of the Economics of Education. Elsevier, Amsterdam, pp. 459–509 (Chapter 8).
- Layard, R., Psacharopoulos, G., 1974. The Screening Hypothesis and the Returns to Education. J. Polit. Econ. 82 (5), 985–998.
- Light, A., Strayer, W., 2000. Determinants of College Completion: School Quality or Student Ability? J. Hum. Resour. 35 (2), 299–332.
- Long, M., 2004. College Applications and the Effect of Affirmative Action. J. Econ. 121 (1-2), 319-342.
- Malamud, O., 2007. Breadth vs. Depth: The Timing of Specialization in Higher Education. Harris School Working Paper No. 08.09. At http://harrisschool.uchicago.edu/About/publications/working-papers/ abstract.asp?paper_no=08.09.
- Manski, C., Wise, D., 1983. College Choice in America. Harvard University Press, Cambridge, MA.
- McPherson, M., Schapiro, M., 2006. Higher Education Finance. In: Handbook of the Economics of Education. Elsevier, Amsterdam.
- Mincer, J., 1974. Schooling, Experience and Earnings. Columbia University Press, New York.
- Monks, J., 2000. The Returns to Individual and College Characteristics: Evidence from the National Longitudinal Survey of Youth. Econ. Educ. Rev. 19 (3), 279–289.
- Paglin, M., Rufolo, A., 1990. Heterogenous Human Capital, Occupational Choice, and Male-Female Earnings Differences. J. Labor Econ. 8 (1), 123–144.
- Pallais, A., Turner, S., 2006. Opportunities for Low Income Students at Top Colleges and Universities: Policy Initiatives and the Distribution of Students. Natl. Tax. J. 59 (2), 357–386.
- Pallais, A., 2009. Small Differences that Matter: Mistakes in Applying to College. MIT Manuscript. At http://econ-www.mit.edu/files/4030.
- Roderick, M., Nagaoka, J., Coca, V., Moeller, E., 2008. From High School to the Future: Potholes on the Road to College. Consortium on Chicago Schools Research.
- Rothschild, M., White, L., 1995. The Analytics of the Pricing of Higher Education and Other Services in Which the Customers Are Inputs. J. Polit. Econ. 103 (3), 573–586.
- Roy, A.D., 1951. Some Thoughts on the Distribution of Earnings. Oxf. Econ. Pap. 3, 135-146.
- Sallee, J., Resch, A., Courant, P., 2008. On the Optimal Allocation of Students and Resources in a System of Higher Education. B.E. J. Econ. Anal. Policy 8 (1), Berkeley Electronic Press.
- Scrivener, S., Weiss, M.J., 2009. More Guidance, Better Results? Three-Year Effects of an Enhanced Student Services Program at Two Community Colleges. MDRC Report, August.
- Smith, J., 2008. Heterogeneity and higher education. In: McPherson, M., Schapiro, M. (Eds.), College Success: What it Means and How to Make it Happen. College Board, New York, pp. 131–144.
- Snyder, T., Tan, A., Hoffman, C., 2006. Digest of Education Statistics, 2005. National Center for Education Statistics, Washington, DC.
- Stange, K., 2008. An Empirical Examination of the Option Value of College Enrollment. Am. Econ. J. Appl. Econ. (forthcoming).

- Stinebrickner, R., Stinebrickner, T., 2003. Working during School and Academic Performance. J. Labor. Econ. 21 (2), 449–472.
- Stinebrickner, R., Stinebrickner, T., 2008. The Effect of Credit Constraints on the College Drop-Out Decision: A Direct Approach Using a New Panel Study. Am. Econ. Rev. 98 (5), 2163–2184.
- Stinebrickner, R., Stinebrickner, T., 2010. Learning About Academic Ability and the College Drop-Out Decision. At http://economics.uwo.ca/faculty/stinebrickner/recentpapers/learningabout.pdf.
- Turner, S., Bowen, W., 1990. The 'Flight from the Arts and Sciences': Trends in Degrees Conferred. Science 250.
- Turner, S.E., Bowen, W.G., 1999. Choice of major: The changing (unchanging) gender gap. Ind. Labor Relat. Rev. 52 (2), 289–314.
- van Scoyoc, M., 1962. Origin and Development of the University. Peabody J. Educ. 39 (6), 322–333. At http://www.jstor.org/stable/1490086.
- Weisbrod, B., 1962. Education and Investment in Human Capital. J. Polit. Econ. 70 (5), 106-123.
- Willis, R., 1986. Wage determinants: A survey and reinterpretation of human capital earnings functions. In: Ashenfelter, O., Layard, R. (Eds.), Handbook of Labor Economics. vol. 1. Elsevier, Amsterdam, pp. 525–602.
- Willis, R.J., Rosen, S., 1979. Education and Self-Selection. J. Polit. Econ. 87 (5), S7–36. University of Chicago Press, Chicago.
- Winston, G., 1999. Subsidies, Hierarchy, and Peers: The Awkward Economics of Higher Education. J. Econ. Perspect. 13 (1), 13–36.



The Political Economy of Education Funding¹

Gerhard Glomm^{*}, B. Ravikumar^{**} and Ioana C. Schiopu^{***}

*Department of Economics, Indiana University, Bloomington, IN 47405

**Research Division, Federal Reserve Bank of St. Louis, P.O. Box 442, St. Louis, MO 63166-0442

****Department of Economics, ESADE, Universitat Ramon Llull, Barcelona, Spain

Contents

1.	Introduction	617
2.	Public Education Spending with a Private Option	618
	2.1. A Benchmark Model	618
	2.1.1 Universal Public Education	623
	2.1.2 Child Labor and Opting Out of Education	624
	2.2. Public Education Spending in the Presence of Transfers	624
	2.3. Private Supplements to Publicly Provided Education	627
	2.4. Stages of Education	629
	2.5. Choice of Systems	631
3.	Vouchers	632
	3.1. Uniform Vouchers	632
	3.2. Private School Vouchers	635
	3.3. Vouchers Targeting the Poor	637
4.	Heterogeneity	638
	4.1. Ability Distribution	638
	4.2. Preference Distribution over Types of Education	639
	4.3. The Young and Old	642
5.	Education Funding in Multi-community Models	644
	5.1. A Simple Model of Residential Segregation	646
	5.2. Local Human Capital Externalities	647
	5.3. Housing Markets	648
	5.4. Commuting Costs	651
	5.5. Interjurisdictional Spillovers	653
	5.6. Education Funding Systems	654

¹ We thank the editors and the participants at the *Handbook of the Economics of Education* conference at CESifo in Munich for their comments and suggestions. Part of this work was developed while Schiopu was visiting the Institute for Economic Analysis (IAE-CSIC), Barcelona, during the academic year 2008–2009. The hospitality and financial support of this institution are gratefully acknowledged. She also thanks Banc Sabadell for partial financial support.

	5.7. Results from Calibrated Models	656
	5.7.1 Local versus Centralized Funding with No Private Option	656
	5.7.2 Tiebout Models with Property Taxes, Peer Effects, and Private Education	657
	5.7.3 Tiebout Models with Commuting Costs	660
6.	Dynamics	662
	6.1. Private versus Public Education: Growth and Inequality	662
	6.2. Dynamic Tiebout Sorting	666
	6.2.1 Residential Sorting over the Life Cycle	668
7.	Empirical Issues	669
	7.1. Educational Expenditure	669
	7.2. Competition and School Productivity	670
	7.3. Substitution between Education Quality and Other Goods	671
	7.4. Voting Patterns and Political Support for Public Education	671
8.	Future Directions	673
Re	eferences	675

Abstract

In most countries, the government is the main provider of education services. Even when a private education sector exists, it is often subsidized. Given the substantial involvement of governments in the education sector and the importance of skill acquisition for individual and national welfare, understanding how societies allocate public resources for education is a crucial issue. The purpose of this chapter is to review positive models of public funding for education. Models reviewed in this chapter consist of a private layer and a political economy layer. In the private layer, firms and households make their decisions taking as given the public policies. In the political economy layer, voters or groups with conflicting interests determine the public policy, taking into account the private sector response to the policy. The questions addressed by the models in this chapter include: What is the majority preferred level of funding for public education when private options are available? How do various dimensions of household heterogeneity (e.g., income, age, ability, tastes) alter the political equilibrium? What is the level of public funding in each community when households can sort themselves into multiple communities? Why are large-scale vouchers in education so rare across the world? Why are public education expenditures as a fraction of GDP rising along the development path? The focus of this chapter is theory, but calibrated versions of the theory that rely on empirical work are also included. We also review the empirical evidence that has bearing on the theoretical models in this chapter.

Keywords

Education Funding Political Economy Majority Voting Opting Out Tiebout Sorting Human Capital Accumulation

1. INTRODUCTION

The governments play a substantial role in the education sector everywhere. Public spending on education is one of the largest government outlays: during 2000–2007, 13% of total government expenditure was spending on education in developed countries and 17% in developing countries. During the same period, public investment in education accounted for 5.15% of GDP in developed economies and 4.5% in developing countries (UNESCO (2010)). In most countries, the government is the main provider of education services, even though these services can, in principle, be provided privately. For instance, in 2007 the average enrollment in public elementary and secondary education was 88% in OECD countries, ranging from 44% in Belgium to over 95% in Ireland, Turkey, Norway, and Poland (OECD (2010)). Even when the provision of education is private, public money could represent a significant source of funding for private schools. For example, private schools are publicly funded in Belgium and partially supported by the government in France and New Zealand.

From a normative point of view, the government intervention in education can be justified by externalities or market failures. From a positive point of view, the level of funding for public education and the quality of public education is the result of the interaction of market forces and political decisions involving groups. For instance, if the provision of public education is a form of redistribution between groups (e.g., between rich and poor, between old and young), the political decisions have to aggregate conflicting preferences regarding taxation, redistribution, or income inequality.

The main purpose of this chapter is to review positive models of public funding for education. Models of public funding typically consist of two layers: a private layer in which firms and households make their decisions taking as given the public policies and a political economy layer in which voters or groups determine the public policy taking into account the private sector response to the policy. In the models reviewed in this chapter, the policy in the political economy layer is endogenized through majority voting. Results from Black (1948) on existence of majority voting equilibrium and on identity of the decisive voter are assumed to be familiar to the readers of this chapter.

The questions addressed by these models include: What is the majority preferred level of funding for public education when private options are available? What are the trade-offs for the allocation of the public education budget across different stages of education, that is, primary and secondary versus tertiary education? What determines the organization of the education systems (e.g., purely public or a public–private mix in which the private institutions may or may not be subsidized)? How do various dimensions of household heterogeneity (e.g., income, age, ability, tastes) alter the political equilibrium? Why are public education expenditures as a fraction of GDP higher in rich countries? We present the benchmark model in Section 2 and use its extensions in Sections 2 and 4 to address these questions.

Why are large-scale vouchers in education so rare across the world? Section 3 covers the political economy issues concerning vouchers and models presented there can shed some light on this question. What is the majority preferred public education spending in a decentralized system, where the share of local funding in total resources is substantial? What degree of centralization garners the most political support? These questions are addressed in Section 5 that covers issues related to the (de)centralization of education funding and the interplay of residential and education choices.

Why are public education expenditures as a fraction of GDP rising along the development path? How does inequality affect the public education funding? What is the most preferred system (public or private) in the short and long run? The dynamic frameworks covered in Section 6 deal with such questions.

The central focus of this chapter is theory, but calibrated versions of the theory that rely on empirical work are also included in this chapter. Toward the end of the chapter, we raise some empirical issues in Section 7. Suggestions for fruitful future work are presented in Section 8.

2. PUBLIC EDUCATION SPENDING WITH A PRIVATE OPTION

In most countries, the public and private education sectors live side by side. The government typically collects taxes and uses the tax revenue to provide education services to all school-age children, at a zero or almost zero price. Each family can use these services or it can opt out and choose a private school that charges tuition. The fundamental political economy problem considered here is: What is the funding level for public education that a majority of society would prefer when the possibility to opt out into private education exists? The difficulty that arises in these types of models is that the private school option generates non-single-peaked preferences over funding levels for voters (see Stiglitz (1974)). Consequently, Black's (1948) median voter theorem cannot be applied. This problem has been recently studied by Bearse, Glomm, and Janeba (2001), Bearse, Glomm, and Patterson (2005), Epple and Romano (1996), Glomm and Ravikumar (1998), Gutierrez and Tanaka (2009), and Luelfesmann and Myers (2010) in a static setting with one jurisdiction and income heterogeneity. This literature on the static one jurisdiction model has been extended to allow for preference heterogeneity by Alesina, Glaeser, and Sacerdote (2001), Cohen-Zada and Justman (2003), Levy (2005), Preston (2003), and others.

2.1. A Benchmark Model

In this section, we describe a benchmark model in which both private and public education coexist. Consider an economy populated by a large number of households who are differentiated only by income. Normalize the population size to 1. Let y_i denote the income of household *i*. Assume that income is distributed according to some c.d.f.

F, with support \mathbb{R}_+ . Also, assume that mean income exceeds median income, which is typical for measured income distributions. In situations where income is the only source of heterogeneity, we will refer to the household with income *y* as household *y*.

We assume that all households have the same utility function over the quality of education (e) and a numeraire consumption good (c). This utility function is given by

$$u(c_i, e_i) = \frac{1}{1 - \sigma} \{ c_i^{1 - \sigma} + \delta e_i^{1 - \sigma} \}.$$
(9.1)

The parameter σ is assumed to be greater than 0 and not equal to 1 in Eq. (9.1); when $\sigma = 1$, the utility function is assumed to be logarithmic. Some of the available analytical results can be established under more general conditions without assuming a specific functional form. (Since much of the literature uses calibrated versions of such models to quantitatively evaluate the effects of various policies, it relies, by necessity, on particular functional forms.)

The government collects taxes on income at a uniform rate τ . Since income is exogenous, these taxes are non-distortionary. All tax revenue is used to finance public education. Public education is made available free of charge to all children, but households are free to opt out of public education and choose private schooling. The private school choice entails paying full tuition costs. Letting N denote the number of families choosing public schools and Y the total tax base, we obtain public school spending per student as

$$E = \frac{\tau Y}{N}.$$
(9.2)

Equating public school quality with expenditure per student, we obtain the allocations for households who choose public education as $c_i = (1 - \tau)y_i$ and $e_i = E$. (See Hanushek (2006) for empirical evidence regarding education production functions.) Denote the associated indirect utility function as $V^U(y_i, \tau, E)$.

$$V^{U}(\gamma_{i},\tau,E) = \frac{1}{1-\sigma} \left\{ \left(1-\tau\right)^{1-\sigma} \gamma_{i}^{1-\sigma} + \delta \left(\frac{\tau Y}{N}\right)^{1-\sigma} \right\}.$$

Any household who chooses private education maximizes the utility function in (9.1) subject to the constraint $c_i + e_i = (1 - \tau)\gamma_i$. Just as in public education, the production technology is linear in the private sector, turning one dollar into one unit of quality of education. There is no markup of price over marginal cost of quality, for any level of quality. That is, there is a competitive market at each quality level. Each household's maximization problem yields the optimal choices specific to the household:

$$c_i = \frac{1}{1+\delta^{\frac{1}{\sigma}}}(1-\tau)\gamma_i, \quad e_i = \frac{\delta^{\frac{1}{\sigma}}}{1+\delta^{\frac{1}{\sigma}}}(1-\tau)\gamma_i$$

and generates an indirect utility function $V^{R}(\gamma_{i}, \tau)$.

$$V^{R}(\gamma_{i},\tau) = \frac{\left\{1 + \delta^{\frac{1}{\sigma}}\right\}^{\sigma}}{1 - \sigma} (1 - \tau)^{1 - \sigma} \gamma_{i}^{1 - \sigma}.$$

The economic choice in these models is easy as there are only two discrete options: public school and private school. Each household chooses public or private education in order to maximize:

$$V(\gamma_i, \tau, E) = \max\{V^U(\gamma_i, \tau, E), V^R(\gamma_i, \tau)\}.$$
(9.3)

Glomm and Ravikumar (1998) show that there exists a critical income \hat{y} such that all households with incomes below (above) \hat{y} choose public education (private education). Furthermore, \hat{y} is a continuous function of τ and N, and there exists a unique $N^* \in (0,1)$ that solves the consistency condition $N = F(\hat{y})$ for all $\tau \in (0,1)$. Denote the fixed point as $N(\tau)$. Substituting the government budget constraint (9.2) into (9.3) generates household y_i 's indirect utility function over funding levels or tax rates: $V\left(y_i, \tau, \frac{\tau Y}{N(\tau)}\right)$.

Among the many possible political mechanisms to determine the equilibrium funding level for public schools, the lion's share of the literature has zeroed in on simple majority rule. (Other mechanisms include probabilistic voting, citizen candidate, etc.) At least since Stiglitz (1974), it has been recognized that, in general, the preferences over funding levels are not single peaked.

In Fig. 9.1, at a low tax rate, the level of public school funding and hence public school quality is rather low, so a typical household will opt for private school. Increasing the tax rate marginally from that level only decreases the after-tax income of the household with no substantial increase in the quality of public education, so the household still chooses private school. However, if the tax rate is increased further, a point is reached where the public school quality is large enough to make the household just



Figure 9.1 Indirect Utility over Tax Rates for Different Households.

indifferent between public and private schools. Before that tax rate is reached, the utility over tax rates is declining. Beyond that tax rate, the utility is rising until the household reaches its most preferred mix of consumption and quality of education. For tax rates above the one corresponding to the preferred mix, the utility declines. Thus, the preferences over tax rates are not single peaked, so we cannot invoke the majority voting equilibrium existence theorem in Black (1948).

However, our benchmark model has more structure than the problem studied by Stiglitz (1974). For $\sigma \leq 1$, Glomm and Ravikumar (1998) established sufficient conditions for the existence of a majority voting equilibrium directly using the properties of the utility functions $V\left(\gamma_i, \tau, \frac{\tau Y}{N(\tau)}\right)$. (This is also the approach followed by Luelfesmann and Myers (2010).) Let γ_m denote the median income, and γ_d denote the income of the decisive household. Glomm and Ravikumar (1998) show that the decisive voter γ_d chooses public education, and the most preferred tax rate of household γ_d is the unique solution to

$$\max u\left((1-\tau)\gamma_{\rm d}, \frac{\tau Y}{N(\tau)}\right). \tag{9.4}$$

The decisive voter in this case is household y_m since the most preferred tax rate is a monotonically declining function of household income.

For $\sigma > 1$, there are no sufficient conditions guaranteeing the existence of a majority voting equilibrium. The decisive voter is implicitly defined by $F(\hat{y}) - F(y_d) = 0.5$, where \hat{y} is the income of the household that is just indifferent between public and private education at the equilibrium tax rate. The equilibrium has to be computed numerically.

An alternative approach to determine the majority voting equilibrium in the presence of non–single peakedness is due to Epple and Romano (1996). It is perhaps easiest to follow their approach by considering Fig. 9.2, which illustrates preferences for a typical



Figure 9.2 Majority Voting Equilibrium.

household or voter over public policies in (E, τ) space. Voter *i* is indifferent between public and private schools on all points on the line segment *AB*. To the left of that line segment, public school quality *E* is too low, given the tax rates, and hence, the household chooses private school. To the right of the line segment, the household chooses public school. Since taxes are a "bad" and public school quality is a "good," the indifference curves are upward sloping and utility is increasing along the southeast direction as indicated by the arrow.

It is easy to show that, for a household $\gamma' > \gamma$, the indifference locus A'B' is to the right of the indifference locus AB of the household γ .

In general, there is no reason for the slope of the indifference curves to vary in a particular way with income. Epple and Romano (1996) consider two cases: *SDI*, where the slope of the indifference curves is declining in income, and *SRI*, where the slope of the indifference curves is rising in income. The SDI condition is satisfied when $0 < \sigma < 1$. In this case, the indifference curves (of different households) can cross only once in the entire (E, τ) space, including the flat part of the indifference curves where private education is chosen. This is the *single-crossing property*. When the single-crossing property prevails, it is easy to establish that a majority voting equilibrium exists and that the household with median income is the decisive voter. In Fig. 9.3, let point G denote the policy (E^*, τ^*) that is most preferred by the household with median income and that is feasible given the government budget constraint in Eq. (9.2). In order to verify that (E^*, τ^*) is a majority voting equilibrium, it suffices to check that (E^*, τ^*) beats points such as A and B. Points like C need not be considered since they are not affordable. In a direct comparison of A and G, G emerges as the winner since the richest 50% of the population prefer G over A. Similarly, the poorest 50% of the population prefer G over B.

For the case $\sigma > 1$, the single-crossing property does not hold and the existence of a majority voting equilibrium is not guaranteed. Epple and Romano (1996) argue that $\sigma > 1$ is the empirically relevant case and proceed to solve the model numerically. In this



Figure 9.3 Non-single Crossing Indifference Curves.

case, an "ends against the middle" equilibrium arises, where the poor and the rich form a coalition in favor of low taxes against the middle-income group that prefers high taxes.

Remark 1. Our benchmark model assumes that a large menu of quality is available at marginal cost to the households if they want to opt out of public education and choose a private school. Such a large menu may indeed be available in metropolitan areas such as Boston, Chicago, or Detroit. In smaller cities and in rural areas, competitive markets might be absent at each quality level. Models of monopoly may be more appropriate.

Remark 2. The benchmark model completely abstracts from any subsidy to the private sector. Although this is consistent with the US experience where most of the private schools are operated by religious institutions and a strict separation of church and state prohibits subsidies, in many other countries, subsidies to private schools are common and can be large (see, e.g., James (1993)).

Remark 3. Note that in problem (9.4), voters internalize the effect of tax rate on public school enrollment. If the voting were myopic, they would instead maximize $u((1-\tau)\gamma_i, \frac{\tau Y}{N})$ taking N as given. Non-single-peakedness does not arise in this case. In fact, the most preferred tax rate is a monotonic function of income:

$$\tau_i = \left[1 + \delta^{-\frac{1}{\sigma}} \left(\frac{\gamma_i}{Y} N\right)^{\frac{1-\sigma}{\sigma}}\right]^{-1}$$

Thus, household $\gamma_{\rm m}$ is the decisive voter for all σ .

2.1.1 Universal Public Education

It is useful to contrast the benchmark model with an environment where the private school option is not available to any household. This is sometimes called "universal public education." Under universal public education, N=1 and $e_i = E = \tau Y$ for all *i*. Then, each household's most preferred tax rate solves

$$\max_{\tau \in [0,1]} \frac{1}{1-\sigma} \left\{ \left[(1-\tau) \gamma_i \right]^{1-\sigma} + \delta(\tau Y)^{1-\sigma} \right\}.$$

These preferences are single peaked in τ for all σ , and the majority voting equilibrium existence theorem in Black (1948) applies. Furthermore, the most preferred tax rate of household γ_i is given by

$$\tau_i = \left[1 + \delta^{-\frac{1}{\sigma}} \left(\frac{\gamma_i}{Y}\right)^{\frac{1-\sigma}{\sigma}}\right]^{-1}$$

Clearly, the most preferred tax rate is monotonic in income, so the median-income household is the decisive voter for all σ .

2.1.2 Child Labor and Opting Out of Education

Unlike the society described in the benchmark model, in poor countries a substantial fraction of the population opts out of education altogether. Instead of attending schools, the children work and contribute to the household income. Gutierrez and Tanaka (2009) extend the benchmark model by adding an option not to attend school at all and to choose child labor instead. Unlike the public or private school options that augment human capital, the child labor option generates a fixed human capital level. They show that sufficiently low-income households choose the child labor option. Similar to the benchmark model, the household with median income does not choose private school, and the equilibrium tax rate is always lower than the one preferred by the median-income household. However, the decisive voter is different from the one in the benchmark model. The difference is precisely determined by the fraction of the population opting for child labor; this set, like the private school choosers, also prefers zero public education funding.

2.2. Public Education Spending in the Presence of Transfers

Although funding public education is one of the means of redistribution available to the government, it is not the only one. Public pensions and other transfers can play similar roles. In fact, public pensions often dwarf public education in terms of expenditures. This is especially true in rich countries. It is thus natural and important to question whether the results from Section 2.1 survive in the presence of transfers.

Following Bearse, Glomm, and Janeba (2001), we consider a model that is identical to the benchmark model in Section 2.1 in terms of preferences and in terms of the income distribution. The only difference is the government budget constraint, which now contains public education expenditures and transfer payments. Letting τY stand for total tax revenue and Δ for the fraction of public revenue allocated to public education, the public education quality is given by

$$E = \frac{\Delta \tau Y}{N}$$

and the household's post tax and transfer income is

$$(1-\tau)\gamma_i + (1-\Delta)\tau Y,$$

where $(1 - \Delta)\tau Y$ represents the amount of the lump-sum transfers received. In order to simplify the voting problem and avoid well-known issues that arise when voting is multidimensional (see Ordeshook (1986), for example), we keep the overall government size, that is, the tax rate τ , fixed exogenously and only consider voting on Δ , the fraction of the government budget allocated to public education.

In this case, for a household γ , preferences over (E, Δ) can be represented by indifference curves that are analogous to the ones in Fig. 9.2. The crucial difference between the model here and the benchmark model is how the slope of the indifference curves in (E, Δ) space changes with income. In the part of the (E, Δ) space where the household chooses public schools, the slope of the indifference curves is given by

$$\frac{\partial \Delta}{\partial E} = \frac{\delta}{\tau Y E^{\sigma}} \left[(1 - \tau) \gamma_i + (1 - \Delta) \tau Y \right]^{\sigma}.$$
(9.5)

It is then easy to check that condition SDI is not satisfied, even if $0 < \sigma < 1$. In fact, taking the derivative of the expression in (9.5) with respect to γ_i yields,

$$\frac{\mathrm{d}\frac{\partial\Delta}{\partial E} \mid V^{U} = \mathrm{constant}}{\mathrm{d}\gamma^{i}} = \frac{\delta\sigma}{\tau Y E^{\sigma}} \left[(1-\tau)\gamma_{i} + (1-\Delta)\tau Y \right]^{\sigma-1} (1-\tau).$$
(9.6)

Thus, for any $\sigma > 0$, the slope of the indifference curve is increasing in income. In other words, the sufficient condition used in the benchmark model to prove the existence of a majority voting equilibrium never applies. This is illustrated in Fig. 9.4 below.

It is useful to investigate the underlying reasons for this difference in the two models. In the benchmark model, public funding for education is inextricably linked to the redistributive role of the government since it is the only policy instrument, and its financing operates like a linear tax. The voters' attitude toward redistribution is determined by both their income and the curvature of the utility function.

Figure 9.5 illustrates the budget constraints for households y_1 and y_2 , with $y_1 < y_2$. As we vary the tax rate from 0 to 1, the entire budget line is traced out from the intercept on the horizontal axis to the intercept on the vertical axis. Increasing income from y_1 to y_2 is analogous to a decrease in the relative price of consumption good. Such a "price change" is always associated with an income and a substitution effect. When 0 $< \sigma < 1$, the substitution effect dominates, and hence, the household with the higher income is less willing to trade private consumption *c* for higher public education funding levels. The indifference curves in (E, τ) space are flatter, and the household prefers lower



Figure 9.4 Non-single Crossing Indifference Curves.

funding levels of education. When $\sigma > 1$, the income effect dominates and this ordering is reversed. The slope of the indifference curves is rising in income, and for those house-holds who choose public education, the most preferred tax rate is rising in income. Of course, for private school choosers, the most preferred tax rate is zero.

When there are two redistributive programs in the government budget, public education funding is no longer inextricably linked to redistribution. When voting on the share of the government budget allocated to public education, the primary determinant of the slopes of the indifference curves in the (E, τ) space is the marginal rate of substitution between E and c. An increase in Δ raises E and lowers c through a decrease in transfers, at least for those households choosing public school. The marginal rate of substitution between E and c and, hence, the slope of the indifference curves in the (E, τ) space depend on own income γ_i only in so far as c_i depends on γ_i through the budget constraint. It is obvious that an increase in γ_i increases c_i , which increases the marginal rate of substitution and, hence, SRI obtains.

In the case of SRI, the majority voting equilibrium (if one exists) needs to be computed numerically. Table 9.1 below, which is taken from Bearse, Glomm, and Janeba (2001),



Figure 9.5 Budget Constraints in the Space (c, e) for Different Households.

(Inverse of) Elasticity of Substitution in Preferences (σ)	Share Parameter in Preferences (δ)	Fraction of Tax Rates Where Cycles Appear (%)	
0.9	0.1	20	
0.9	0.5	13	
1.1	0.1	23	
1.1	0.5	17	
1.5	0.1	18	
1.5	0.5	13	

Tuble Still Found Cycle	Tabl	e 9.1	Voting	Cyc	les
-------------------------	------	-------	--------	-----	-----

reveals that the issue of non-existence of a voting equilibrium (or existence of cycles) is pervasive when there are transfers in addition to public education expenditures.

The set of tax rates in which voting cycles appear, that is, where there does not exist a majority voting equilibrium, are "randomly" distributed over the [0, 1] interval. These cycles also appear for many parameters of the lognormal income distribution. It is thus difficult to see any hope for an existence theorem. These negative results suggest that other political mechanisms, for example, citizen candidate models or probabilistic voting models, could facilitate the study of public education spending in the presence of transfers as the existence problem is more easily solved in the alternative mechanisms. Recent work has indeed shifted more toward these mechanisms (see, e.g., de la Croix and Doepke (2009) and Arcalean and Schiopu (2010)).

2.3. Private Supplements to Publicly Provided Education

In the benchmark model, all households attending public schools are assumed to get the same quality of education. For those households attending private schools, school quality is an increasing function of income (see Fig. 9.6).

The dashed line in Fig. 9.6 indicates that in a purely private school system where each household purchases its own education, the quality of education is a linear function of income since the preferences are homothetic. For households with incomes between $\underline{\gamma}$ and $\hat{\mathbf{y}}$, the publicly provided education quality is below their desired quality in a (hypothetical) purely private regime. Therefore, there might be an incentive to supplement public school quality with private expenditure.



Figure 9.6 Quality of Education by Income Level.

Bearse, Glomm, and Patterson (2005) introduce a supplementation technology of the form

$$e_i^p = (E^\alpha + \psi s_i^\alpha)^{\frac{1}{\alpha}}$$

where *E* stands for public education quality, s_i stands for privately financed education supplements, $\psi > 0$ is the productivity of the supplementation technology, and $\alpha \le 1$. In this version, public inputs and private supplements are combined via the CES aggregation. The simple case in which $\alpha = 1$ (the supplementation technology is linear) is theoretically compelling. A teacher who instructs a child in algebra on a Wednesday morning in the public school is a component of *E*. The same teacher can be hired as a private tutor on a Friday afternoon and is a component of s_i . Both of these types of instructions ought to be perfect substitutes.

If $\alpha = 1$, a household who chooses public schooling solves the problem:

$$\max_{\{c_i, s_i\}} \frac{1}{1 - \sigma} \left\{ c_i^{1 - \sigma} + \delta (E + \psi s_i)^{1 - \sigma} \right\}$$

s.t. $c_i + s_i = (1 - \tau) \gamma_i, s_i \ge 0, c_i \ge 0.$

The solution to this problem is given by

$$s_{i} = \begin{cases} \frac{(\delta \psi)^{\frac{1}{\sigma}}}{\psi + (\delta \psi)^{\frac{1}{\sigma}}} (1 - \tau) \gamma_{i} - \frac{1}{\psi + (\delta \psi)^{\frac{1}{\sigma}}} E, & \text{if } (\delta \psi)^{\frac{1}{\sigma}} (1 - \tau) \gamma_{i} \ge E\\ 0, & \text{otherwise.} \end{cases}$$

It is clear that a household who chooses public school is indifferent between positive and zero supplements if

$$\tau = 1 - \frac{1}{\left(\delta\psi\right)^{\frac{1}{\sigma}}\gamma_i}E.$$
(9.7)

Equation (9.7) describes a locus in the (E, τ) space for each income level y_i . The absolute value of the slope of this locus is decreasing in y_i .

The problem of a household who chooses private education is exactly the same as before. Equalizing indirect utility of a public school chooser (with positive supplement) and a private school chooser yields the following indifference locus:

$$\tau = 1 - \frac{1}{(\kappa - \psi)\gamma_i} E, \tag{9.8}$$

where $\kappa = \left(\psi^{\frac{1}{\sigma}} + (\delta\psi)^{\frac{1}{\sigma}}\right)^{\frac{\sigma}{1-\sigma}} \left(\psi + (\delta\psi)^{\frac{1}{\sigma}}\right)^{-\frac{\sigma}{1-\sigma}}$. The absolute value of the slope of this locus is also decreasing in income.

The relative location of these two loci determines how the (E, τ) space is divided into three mutually exclusive regions. These regions are illustrated in Fig. 9.7.



Figure 9.7 Individual Choice of Private School and Public School with or without Supplements. (a) High Values of ψ . (b) Intermediate Values of ψ . (c) Low Values of ψ .

In Fig. 9.7a, where the supplementation technology is very productive (high values of ψ), no one chooses a private school. In this case, the locus from Eq. (9.8) separates those households who will supplement the publicly provided quality, which occurs in region B, from those households who will not supplement public education in region C. Going from Fig. 9.7a to 9.7c, the productivity of the supplementation technology decreases, and the private schooling option becomes more desirable, especially when the public school quality is low. Thus, households in region A of Figs 9.7b and 9.7c choose private education. When ψ is very low, no household will use the supplementation technology.

We can verify that a majority voting equilibrium exists when $\psi \ge 1$ (no household chooses a private school). When ψ is sufficiently low to allow for a private sector, SDI is not satisfied and hence the existence of a majority voting equilibrium cannot be guaranteed. The failure of SDI is illustrated in the Fig. 9.8.

2.4. Stages of Education

Su (2004) documents very large differences in public expenditures across three stages of education (primary, secondary, and tertiary). Gradstein (2003) documents that in an income breakdown by quintiles, the highest quintile in poor countries receives over four times the educational resources as households in the lowest quintile and that the inequality in educational resources is associated with unequal spending across elementary, secondary, and tertiary education. Although there is some research on the political economy of spending on a particular stage such as higher education (e.g., Haupt (2005)), there are not many models that take into consideration the technological link in the learning process between the stages, the individual decision to drop out of school after a particular stage, the explicit financial trade-offs between the stages and the voting



Figure 9.8 Non-single Crossing.

or, more generally, the political decisions across the socioeconomic distribution to support public education at the various stages.

Cardak and Givon (2004) explain the low college enrollment of students from poor socioeconomic backgrounds by emphasizing the importance of human capital buildup across multiple education stages: The first stage, including primary and secondary education, is mandatory, and the second stage, tertiary education, is optional. The level of attainment during the first-stage rations participation in the second stage, generating an endogenous participation constraint (in addition to the credit constraints discussed in the literature). The public education in the first stage is financed by a separate tax rate that is determined through majority voting. In the second stage, the government guarantees a certain quality of education that is exogenous. The tax rate that funds the tertiary education is the one that balances the budget.

In this environment, Cardak and Givon (2004) analyze how various education finance policies modify the participation of low-income students. Universal first-stage public education addresses the attainment constraint but leaves poor students credit constrained. Universal public higher education alleviates the credit constraint but not the attainment constraint. A mixed school system at the first stage, where households can opt out of public schools and university is universally public, addresses both constraints but can magnify the relative differences in college attendance between high- and low-income households.

Su (2006) analyzes the policy preferences of the top class in a model with public education in both stages. In less-developed economies, where the rich might have substantially more political power than the poor, the distribution of resources greatly favors the rich. The resulting political equilibrium implies higher spending on advanced education at the expense of basic education funding. The developed countries are characterized by a more balanced allocation of spending. Zhang (2008) connects income inequality and allocation of public education spending for different stages of education. She documents that more unequal countries tend to spend less on secondary education and more on tertiary education. She rationalizes these findings in a model in which different socio-economic groups can influence the education policies through lobbying in a model in which the formation of lobbies is endogenous. Thus, different societies have different lobby configurations and, hence, different allocations for secondary and tertiary education.

Romero (2009) allows for the possibility of opting out into the private education at each stage. Under special assumptions regarding the voting behavior of the rich, the "ends against the middle" result survives in this setting, and majority voting delivers higher funding for the first stage and lower funding for the second stage relative to the allocations preferred by the median-income household. The decisive voter has lower than median income, is unlikely to choose tertiary education, and thus favors relatively higher funding for primary and secondary education.

2.5. Choice of Systems

As mentioned in Section 1, the organization of the education process varies greatly across countries. Some political systems provide education publicly and disallow a private alternative while other systems allow, and in some cases even subsidize, a private alternative. Luelfesmann and Myers (2010) call the former a one-tier regime. The latter, without subsidies, in their terminology is a two-tier regime. The question whether simple majority voting favors the one-tier or the two-tier system arises naturally. Luelfesmann and Myers (2010) compare three systems—one-tier, two-tier, and a purely private education system with no government involvement—and examine majority preferences for each system.

In their environment, there are two cases to distinguish. In the first case in which education and the numeraire good are substitutes, they show the following:

- 1. The median-income voter is decisive in both systems.
- 2. Those voters who opt out of public education are better off in the two-tier system since they have an option that is not available in the one-tier system.
- **3.** Those voters who remain in public education are better off in the two-tier system than in the one-tier system since at the same tax rate, they receive higher education quality.

Combining (1)–(3) allows us to reach the conclusion that the two-tier system is majority preferred to the one-tier system. In fact, this decision is unanimous. Furthermore, the two-tier system beats the pure private system as well since rich households who prefer a zero tax are in minority.

The case in which the education and the numeraire good are complements is slightly more complicated. The most preferred tax rate is increasing in income in the one-tier system; it is increasing in income only for the public school choosers in the two-tier system. The top layer of the income distribution (i.e., those households who choose private education) switch from a higher most preferred tax rates in the one-tier system to a most preferred tax rate of zero in the two-tier system. As a consequence, the tax rate falls from that preferred by the median-income household to a lower tax rate that is favored by a "coalition" of the rich and the poor. This creates a conflict in the choice of regimes between the middle class that prefers relatively higher taxes and the ends of the income distribution that prefer lower tax rates. Luelfesmann and Myers (2010) show that the coalition of the rich and the poor beats the middle class and implements the two-tier system. When comparing the two-tier system and the private system, the middle class prefers the two-tier regime with a positive tax, in contrast to the rich and the poor who prefer the private system. In this case, however, the middle class can succeed in garnering majority support in favor of the twotier regime.

3. VOUCHERS

An alternative to the mixed private and public education regime studied in the benchmark model is a voucher regime. A voucher is simply a piece of paper issued by a government that allows the bearer to purchase education services in the amount specified on the paper. In principle, there are many different types of vouchers that a government can design. It is then possible to tailor the voucher amount to each household. For instance, if all households receive the same voucher amount and they are precluded from privately spending more than the voucher amount, the voucher scheme is identical to the universal public education scheme analyzed in Section 2.1.1. The literature on the political economy of voucher funding reflects a variety of voucher designs.

3.1. Uniform Vouchers

In this subsection, we examine a voucher design recommended by Friedman (1962). All households receive the same (uniform) voucher amount v so that the government budget constraint is $v = \tau Y$. The voucher imposes only a lower bound on each household's quality of education (or equivalently, an upper bound on consumption). Households can top up the voucher amount and obtain a higher quality subject to an overall resource constraint. Thus, the two constraints faced by household γ are

$$c + e \le (1 - \tau)\gamma + \tau Y, \quad c \le (1 - \tau)\gamma.$$
 (9.9)

Here, vouchers are merely instruments used by the government to finance education; the government does not directly provide education. Thus, in our model of uniform vouchers, there are no "public schools."

Each household chooses the pair (c, e) so as to maximize u(c, e) subject to (9.9). For households who are constrained by the lower bound on quality ($e \ge \tau Y$ binds), the optimal choice is

$$c = (1 - \tau)\gamma, \quad e = \tau Y.$$

This is illustrated in Fig. 9.9. For the household in Fig. 9.9, at the constrained bundle, the marginal rate of substitution of education for consumption is less than the marginal rate of transformation. Even though the household's overall resources are large enough to move to a higher indifference curve, it cannot since the voucher amount is restricted to educational expenditures.

For the households who are not constrained by the lower bound on quality, the optimal choice is

$$c = \frac{(1-\tau)\gamma + \tau Y}{1+\delta^{\frac{1}{\sigma}}}$$
$$e = \delta^{\frac{1}{\sigma}} \frac{(1-\tau)\gamma + \tau Y}{1+\delta^{\frac{1}{\sigma}}}$$



Figure 9.9 Optimal Choice for Households Constrained by the Lower Bound on Quality of Education.



Figure 9.10 Optimal Choice for Unconstrained Households.

For such households, the marginal rate of substitution of education for consumption is greater than the marginal rate of transformation as illustrated in Fig. 9.10. They can move to a higher indifference by reducing their consumption and topping up the voucher amount. It is thus easy to show that there exist low-income households who do not supplement their voucher and rich households who do supplement their voucher.

Bearse, Cardak, Glomm, and Ravikumar (2009) show that the households' preferences over τ are single peaked, so there exists a majority voting equilibrium tax rate. They also show that the decisive household is constrained, and its most preferred tax rate solves

$$\max_{\tau \in [0,1]} \frac{1}{1-\sigma} \left\{ \left[(1-\tau) \gamma_{\mathrm{d}}^U \right]^{1-\sigma} + \delta(\tau Y)^{1-\sigma} \right\},\$$

where y_d^U is the income of the decisive household. If $\sigma \leq 1$, then the decisive voter is household y_m , and the majority preferred tax rate is given by

$$\frac{\delta(\tau Y)^{-\sigma}}{\left((1-\tau)\gamma_{\rm m}\right)^{-\sigma}} = \frac{\gamma_{\rm m}}{Y},$$

and if $\sigma > 1$, then the decisive voter is implicitly determined by $1 - F(Y) + F(\gamma_d^U) = 0.5$, and the majority preferred tax rate is given by

$$\frac{\delta(\tau Y)^{-\sigma}}{\left((1-\tau)\gamma_{\rm d}^U\right)^{-\sigma}} = \frac{\gamma_{\rm d}^U}{Y}.$$

Chen and West (2000) and Bearse, Cardak, Glomm, and Ravikumar (2009) study whether a move from the benchmark model to a uniform voucher system would garner sufficient political support. In their study, voters anticipate the political outcome for public funding in each of the two regimes. The distinction between SDI and SRI again is crucial. Under SDI, the median-income voter is decisive in both regimes. When SRI prevails, the decisive voter in the universal voucher regime has lower income than in the benchmark model. The reason is simple: In the "ends against the middle" configuration of the voucher regime, more voters prefer zero taxes—in the benchmark model, only the households who opt out for private education prefer zero taxes, whereas in the voucher regime, all households with incomes above average prefer zero taxes. Both papers show that the voucher regime will be defeated in a plebiscite against the benchmark regime. One of the reasons is that public revenue in the voucher regime will be shared among all school-age children, whereas in the benchmark regime, it is shared only among those who choose public education. This feature clearly imparts welfare losses on all those households who do not supplement the vouchers.

3.2. Private School Vouchers

Rangazas (1995a), Hoyt and Lee (1998), Nechyba (2000), and Piolatto (2009), for example, consider vouchers that are given only to private school choosers. The basic setup of their model is the same as our benchmark model. The families who choose public education receive no voucher, and their indirect utility is $u((1 - \tau)\gamma, E)$. Any household choosing private school receives a voucher in amount v and obtains indirect utility $u((1 - \tau)\gamma + v - e^*, e^*)$, where e^* is the quality of private school chosen by the household.

There are now three groups of the population to consider as illustrated in Fig. 9.11 below.

Households with incomes below $\gamma^*(E, \nu)$ always choose public schooling regardless of the voucher amount. Households with incomes between $\gamma^*(E, \nu)$ and $\gamma^*(E, 0)$ choose public education when there are no vouchers and private education when there are vouchers. Families with incomes above $\gamma^*(E, 0)$ always choose private schooling.

To understand voting behavior toward vouchers, we have to examine how changes in the voucher amount v influences welfare of families in each one of the categories in



Figure 9.11 School Choice and Income Level with Private School Vouchers.

Fig. 9.11. One crucial element in this analysis is whether the introduction of the voucher lowers or raises taxes (relative to the benchmark model).

For an exogenously fixed public education quality, Hoyt and Lee (1998) establish the following result: For a fixed public education quality E, the tax rate is lower in the private school voucher regime, that is, $\tau(E, \nu) < \tau(E, 0)$, if

$$[E-\nu] [N(\gamma^*(E,0)) - N(\gamma^*(E,\nu))] > [1 - N(\gamma^*(E,0))]\nu.$$
(9.10)

This condition has a straightforward interpretation. On the left-hand side, E - v is the reduction in public expenditure for each student who is induced by the vouchers to leave public school, and $N(\gamma^*(E, 0)) - N(\gamma^*(E, v))$ is the total number of students leaving public schools due to the vouchers so that the left-hand side is the total education expenditure reduction. The right-hand side is the number of private school students (under the no-voucher regime) multiplied by the voucher amount. Whether inequality (9.10) is satisfied depends on the income distribution and preferences.

Ruling out the case where the public education sector vanishes (E = 0), it is easy to see that if condition (9.10) holds, majority voting will result in a voucher $v^* \in (0, E)$. No-voucher amount larger than E will be approved by majority because such an amount will increase the tax rate relative to the benchmark model.

When inequality (9.10) is violated, there is no guarantee that introducing private school vouchers would decrease the tax rate. If the tax rate rises due to the voucher, families whose children remain in the public school are made worse off since their after-tax consumption of the numeraire good falls, and by assumption, the quality of public education is unaffected. For those households who choose private school only after the introduction of the vouchers, the welfare effect of the vouchers is ambiguous. They receive better education but lower consumption of the other good. For those families who choose private education even without the vouchers, the welfare losses or gains depend on the size of the tax increase. If the increase in taxation exceeds the voucher amount, the family is made worse off. The very rich, that is, those whose incomes satisfy $\gamma(\tau(v^*) - \tau(0)) > v^*$ will experience the largest welfare losses and thus most vigorously oppose the vouchers.

The preceding analysis has been carried out under the assumption that the quality of public education or the public funding level for education is independent of the size of the voucher contemplated. The literature often determines one aspect of fiscal policy endogenously while fixing another dimension of policy exogenously. For example, Piolatto (2009) studies the determination of public education funding while fixing the voucher policy exogenously. This approach fails to recognize the interaction of voucher and public education funding. Rangazas (1995a) and Hoyt and Lee (1998) attempt to solve this joint determination of voucher and public school funding. Their analysis highlights three possible effects: (1) increasing the voucher amount decreases

public school enrollment, thereby decreasing the cost of improving public school quality, which in turn may increase support for public education funding; (2) increasing the voucher amount decreases demand for public education and this shifts the identity of the decisive voter to the left (to lower income), which in turn decreases support for public education; (3) increasing the voucher amount will in general change the after-tax income that will change the political equilibrium as well. The relative sizes of these effects are unknown; Hoyt and Lee (1998) conclude that "... the impact of a voucher on the quality of public education cannot be predicted on theoretical grounds alone."

Rangazas (1995a) appeals to empirical evidence to pin down the relative sizes of the above three effects. He finds that the first effect, decreasing the cost of improving public school quality, is so large as to dominate the other effects, implying that the introduction of vouchers causes an increase in public education expenditure per pupil and that this effect is large. On the other hand, the numerical results in Nechyba (2000), which are obtained in a very rich model with several jurisdictions, housing markets, and peer-group effects, show that an increase in the voucher level tends to decrease spending per student in the public schools, indicating that the second effect mentioned above may be dominant. The many possible general equilibrium effects are discussed in Section 5.7.2.

3.3. Vouchers Targeting the Poor

Uniform vouchers in Section 3.1 have obvious distributional consequences: richer families will supplement the vouchers, which generates higher inequality in access to educational resources. In turn, to the extent that access to educational resources determines future human capital and income, vouchers can generate higher (future) income inequality. The remedy for such inequality might be to make vouchers means-tested.

Chen and West (2000) consider a selective voucher to the poor: all household below a certain level of income receive a fixed-amount voucher. Above that level of income, households receive no voucher. Modeling vouchers as a step function of income is very stark and at odds with most in-kind subsidies in the US. Means-tested vouchers that are smooth declining functions of income are studied by Bearse, Glomm, and Janeba (2001) and Bearse, Cardak, Glomm, and Ravikumar (2009). For instance, a linear function given by

$$\nu(\gamma_i, \alpha, \beta) = \max(\alpha - \beta \gamma_i, 0), \qquad (9.11)$$

where $\alpha > 0$ and $\beta > 0$, phases out the voucher amount as income increases. Households with incomes above $\frac{\alpha}{\beta}$ receive no vouchers at all. (Uniform vouchers is a special case of $(9.11) - \text{set } \beta = 0$ and the voucher amount to α for each household.) To endogenize the funding for vouchers under such a means-testing scheme, we have to determine the tax rate τ (i.e., the total revenue) and the slope, β , of the means-testing function. (Given any pair (β, τ) , the government budget constraint uniquely determines α .) The voting problem is thus two dimensional. It is well known (see Ordeshook 1986, ch. 4.7) that simultaneous multidimensional majority voting equilibria may not exist. One way to avoid this non-existence issue is to vote sequentially on the two policy variables. For instance, if voting takes place first over τ and then over β , preferences over policies in each variable turn out to be single peaked so that a majority voting equilibrium exists.

4. HETEROGENEITY

There are at least five different types of heterogeneity that could potentially influence public education via any given political process. These are (a) income heterogeneity, (b) heterogeneity in student ability, (c) heterogeneity in preferences for the type of education, (d) heterogeneity in the intensity of preferences for education or the child's human capital (e.g., differences in preferences for education between the young and the old), and (e) heterogeneity in the number of children. The previous two sections of this chapter were devoted to models in which income was the only source of heterogeneity. In this section, we review models with other types of heterogeneity.

4.1. Ability Distribution

Epple and Romano (1998, 2008) incorporate income and ability heterogeneity into models of education provision by public and private schools. They study the equilibrium distribution of student types across public and private schools and the effects of different voucher schemes on student sorting and school productivity.

Households are heterogeneous in income, and the demand for education quality is increasing in income. A student is characterized by ability and income, which are drawn from a continuous bivariate distribution. Public-sector schools offer free admission to all students and are homogenous in equilibrium. There is no competition among schools within the public sector while there is free entry and exit in the private sector. Private schools maximize profits and price discriminate based on ability and income. In these models, schooling outcomes depend on own ability, as well as the peer effect, the average ability of the school's student body.

Epple and Romano (1998) show that in such an economy, the equilibrium is characterized by a strict hierarchy of school qualities where the public sector has the lowest ability peer group. In the private sector, the students sort themselves along the ability and income dimensions. Thus, the model produces "diagonal stratification" by ability and income among private schools as relatively richer but lower ability students crosssubsidize relatively low-income, high-ability peers. Consequently, private schools "cream skim" the richer and higher ability students. Moreover, universal vouchers magnify the cream skimming. They impart (1) losses to low-ability and low-income students who remain in the public schools and (2) gains to low-income and high-ability students. Epple and Romano (2008) extend the previous analysis to conditional vouchers. In addition to peer effects, here schooling quality is also a function of per-student spending. The tax rate and, implicitly, the per-student education expenditures in public schools are chosen by majority voting. When voting, households take as given the distribution of the other students in public and private schools. However, the voucher system is assumed to be exogenous and financed out of tax revenues while maintaining per-student expenditure on educational inputs at the majority chosen level with no voucher. They show that combining a tuition restriction with ability-linked vouchers induces equal-quality, competing, and technically efficient schools, even under voluntary school participation in the voucher system. Thus, appropriately conditioned, vouchers can reap the benefits of school competition without necessarily generating higher stratification.

4.2. Preference Distribution over Types of Education

Alesina, Baqir, and Easterly (1999; ABE henceforth) studied the implications of preference heterogeneity for different types of public education. A slightly modified version of their model is presented in this section. All households have the same income, the same preferences over the amount or quality of education, but different preferences over the types of education. In this context, education types can be thought of as bilingual or monolingual instruction, inclusion or exclusion of religious instruction, academic or vocational training, etc.

The per-student spending in public schools is τY , where τ is the tax rate and Y is the average (and aggregate) income. The after-tax income, $(1 - \tau)y$, of the representative household is spent entirely on the private consumption good: $c = (1 - \tau)y$. (Since all households have the same income, y must equal Y.)

Preferences over the private consumption and public education of household *i* are given by

$$\frac{1}{1-\sigma} \bigg\{ c^{1-\sigma} + \delta \big[e(1-d_i) \big]^{1-\sigma} \bigg\},\,$$

where d_i is a household utility parameter that captures the heterogeneity of preferences for schooling.

Using the government budget constraint ($e = \tau Y = \tau \gamma$), the indirect utility can be written as

$$\frac{1}{1-\sigma}\Big\{\big[(1-\tau)\gamma\big]^{1-\sigma}+\delta\big[\tau\gamma(1-d_i)\big]^{1-\sigma}\Big\}.$$

In ABE, the political process is carried out in two stages. In stage 1, majority voting determines the tax rate τ^* . In this stage, all voters are forward looking and form (correct) expectations over the type of education determined by voting in the second stage.

ABE are silent on the precise nature of the political process in stage 2 and simply assume that the median type will be chosen. This political choice determines the distance from each voter's ideal point to the chosen type. We denote this distance by d_i^m or the "median distance from the median."

The political problem in the first stage is now easy to solve. Each voter's ideal tax rate is now given by

$$\tau_{i} = \left[1 + \delta^{-\frac{1}{\sigma}} (1 - d_{i}^{m})^{-\frac{(1-\sigma)}{\sigma}}\right]^{-1}.$$
(9.12)

Inspection of Eq. (9.12) reveals the following results:

- 1. If $0 < \sigma < 1$, an increase in median distance of the median type decreases each voter's support for public education.
- 2. If $\sigma > 1$, an increase in the median distance of the median type increases each voter's support for public education.

ABE only consider the case of perfect substitutability between public education and private consumption. Part (1) above corresponds to ABE's main theoretical result.

The ABE model can be easily extended to allow for income heterogeneity. We let y_i be distributed according to some c.d.f. *F*. The public expenditure per student remains $e = \tau Y$, but the consumption of the numeraire good is $c_i = (1 - \tau)y_i$.

The indirect utility of voter *i* over funding levels becomes

$$\frac{1}{1-\sigma}\Big\{\big[(1-\tau)\gamma_i\big]^{1-\sigma}+\delta\big[\tau Y(1-d_i)\big]^{1-\sigma}\Big\}.$$

The household's most preferred tax rate (in stage 1) is given by

$$\tau_i = \left[1 + \frac{\delta^{-\frac{1}{\sigma}}}{(1 - d_i^m)^{\frac{1-\sigma}{\sigma}}} \left(\frac{\gamma_i}{Y}\right)^{\frac{1-\sigma}{\sigma}} \right]^{-1}.$$
(9.13)

It is clear from expression (9.13) that the household's most preferred tax rate is increasing in income if $\sigma > 1$ and decreasing in income if $0 < \sigma < 1$. The monotonic relationship between most preferred tax rate and income implies that the equilibrium tax rate is

$$\tau_m = \left[1 + \frac{\delta^{-\frac{1}{\sigma}}}{\left(1 - d_i^m\right)^{\frac{1-\sigma}{\sigma}}} \left(\frac{\gamma_m}{Y}\right)^{\frac{1-\sigma}{\sigma}} \right]^{-1}.$$
(9.14)

Evidence from Bergstrom, Rubinfeld, and Shapiro (1982), for example, suggests that σ is greater than 1. This evidence is inconsistent with the case emphasized by ABE, $\sigma < 1$. More recent evidence collected by Cohen-Zada and Justman (2003) seems to be closer to the case emphasized by ABE.

Equation (9.14) suggests that the relationship between income inequality and public funding for education is dependent on the elasticity of substitution in the utility function. The ratio of $\frac{\gamma_m}{Y}$ leads to an increase (decrease) in τ_m if $\sigma > 1$ ($\sigma < 1$). The evidence on the relationship between income inequality and support for public education funding is mixed. In a voting model calibrated to Colombia, Gutierrez, and Tanaka (2009) find that increasing income inequality decreases the political support for public education. A number of papers have found that support for redistribution and public goods provision is weaker in more unequal or more heterogenous societies (Perotti (1996), Goldin and Katz (1997), Alesina, Baqir, and Easterly (1999), Luttmer (2001)). However, Sylwester (2000) finds in a cross-section of countries that higher income inequality is associated with higher public education expenditures. More recent papers by Corcoran and Evans (2009) and Boustan, Ferreira, Winkler, and Zolt (2010) using district level data from the United States find that higher inequality is associated with higher public spending.

When there is a substantial degree of preference heterogeneity, different groups may form their own preferred types of schools if this is permitted by the legal system. There is ample evidence for this type of school segregation. For example, some charter schools explicitly cater to African-American students. Lankford and Wyckoff (2001) and Fairlie and Resch (2002) find that larger minority populations in public schools are associated with larger fractions of white students enrolled in private schools. Campbell, West, and Peterson (2005) and Buddin, Cordes, and Kirby (1998) find no evidence of such "white flight," and Brunner, Imazeki, and Ross (2010) attribute white flight not to race per se, but to associate concerns with limited academic performance. Nevertheless, increased sorting into diverse and more private schools as "taste" heterogeneity increases is a distinct possibility, even within a given jurisdiction.

As heterogeneity rises, more private schools may start up and attract students with these given preferences. As a consequence, preference heterogeneity among those students remaining in public schools may actually decrease. In turn, the support for public schools among these families may rise. Preliminary simulations in a model with such private school options by Van Alstine (2009) indicate that the theoretical relationship between preference heterogeneity and public education funding may be non-monotonic. For some cities in Texas, complementary empirical results by Van Alstine (2009) also point toward non-monotonicities in this relationship.

How an increase in preference heterogeneity influences the prevalence of private school choices, the variety of private school options, or even the adoption of the curriculum in the public schools are largely unexplained questions. There is little empirical work on this issue, and there is even less theoretical work on this question. In suitable theoretical models, the precise nature of preference heterogeneity ought to determine the number and types of private schools entering the market. This might induce a
reaction by the public schools regarding the curriculum and quality choices, which in turn will influence the political support for public school funding.

Cohen-Zada and Justman (2003) explicitly introduce heterogeneity in religious values in addition to income heterogeneity. The two types of heterogeneity on income and on religious values are assumed to be independent of each other. The religious preference distribution simply implies that households with high (low) enough religious preferences will choose parochial (non-sectarian) private schools. Apart from the religious dimension, their model is similar to the benchmark model. Their model is calibrated to US aggregate data and to data on the 50 states. The calibration allows them to back out the fraction of the population that would rather send their child to a religious school instead of a nonreligious school if both charge the same tuition. Of course, the fraction of households choosing religious private schools depends on the private school subsidies. For plausible values of religious school subsidies, religious school preferences loom large. They then use their model to study the effects of private school vouchers. As the voucher amount rises from 60% of public school budgets to 100%, non-sectarian private school enrollment rises by over 400%, while religious enrollment rises by about 130%, lending credence to the claim that private school enrollment is heavily influenced not only by income but also by preference heterogeneity.

Cardak (1999) and Preston (2003) have explored the implications of heterogeneity in the intensity of preferences for children's human capital. Preston (2003), for example, finds that funding levels under a pure public education regime are always higher than under a voucher system and then traces out the implications of this result for the evolution of the entire human capital and income distribution over successive generations. These models are too stylized at this stage to be used for quantitative exercises.

4.3. The Young and Old

Levy (2005) provides a simple stylized model of preference heterogeneity that can be interpreted as a model of the young and old. In this context, it is natural to think of the young as deriving utility from education and from a numeraire good. In contrast, the old derive little utility from education (or no utility as actually posited in the model). The two types of preferences are complemented by two income types, the rich and the poor, so that there are four types of voters: the rich young, the poor young, the rich old, and the poor old.

The government collects a tax at the uniform rate τ and uses all tax revenue for two items in the government budget, public provision of education (to the young only) and a lump-sum income distribution. All young individuals optimally choose a private supplement to the publicly provided education.

The political mechanism differs from the one considered in the previous sections. Each type of voter is allowed to field one "candidate." All "candidates" are free to form parties. Parties can be singletons, that is, no candidate has to join a party or coalition. Each party can offer policies only in the Pareto set of its members. Majority voting among all households determines the winning policy.

The results of this paper include the following:

- **1.** Per capita provision of public education is higher when the young are in a minority than when they are in a majority.
- **2.** The lump-sum transfer is lower when the young are in a minority than when they are in a majority.
- **3.** The poor and the rich are more likely to be equally educated when the young are in a minority than when they are in a majority.
- **4.** The effect of income inequality on the public policies depends on whether the old or the young form a majority.

These are novel, and perhaps counterintuitive results. The first result that public education funding declines as the young, who value education, become a majority is especially counterintuitive. When the young are a minority, a coalition of the rich old and poor young forms and its policies beat the bliss point of the poor old, who prefer a 100% tax rate and no public education provision. However, when the young are a majority, the group that needs to be defeated is the poor young. Public education is relatively expensive, and the rich can form a coalition with the old poor and focus redistribution away from education to transfers.

To what extent these predictions are consistent with empirical evidence is an open question. The first two predictions seem at odds with the evidence: poor countries that have high-fertility rates (and thus potentially a majority of the young) allocate smaller fractions of GDP to public education than rich countries with lower fertility rates. Also, social security and welfare budgets in high-fertility countries are dwarfed by social security and welfare budgets in low-fertility countries.

The fourth prediction is useful as it might contribute to the large but inconclusive literature on the relationship between inequality and public policies since it suggests inclusion of an interaction effect between inequality and demography that has previously been ignored.

Gradstein and Kaganovich (2004) develop an overlapping generations model in which the rise in the fraction of the elderly is due to rising life expectancy. In their economy, rising life expectancy generates higher savings and capital accumulation. Complementarities between physical capital and human capital result in higher public education expenditures. Rising longevity thus generates two effects: first, higher public education expenditures for each young individual, and second, within each period, a larger fraction of the old who attach little or no value to public education expenditures. Gradstein and Kaganovich show that in their specification, the first effect always dominates, implying that population aging leads to higher public education funding. They also use a two-region version of their model in which exogenous migration generates different ratios of the old relative to the young. The cross-sectional implications of their model are that a larger fraction of the old population (induced by immigration) is associated with lower public education funding levels; however, their earlier time series implication for a positive relation between aging and public education funding survives.

5. EDUCATION FUNDING IN MULTI-COMMUNITY MODELS

As observed by Tiebout (1956), households can segregate in communities that offer their preferred level of public goods and taxation that is, they can "vote with their feet." The Tiebout hypothesis connects location choices and *local* public good provision. When local resources represent a significant share of public schools budgets, sorting across communities is relevant for the determination of public education funding.

The share of local spending in total government expenditures for education varies dramatically across countries. Local spending on elementary and secondary education in 2006 accounted for less than 6% of total public education spending in Belgium, Portugal, or Spain, but for more than 70% in Norway, Poland, and the United Kingdom. Among the emerging economies, local spending is sizeable in Russia (69%) and Brazil (37%), but less than 6% in Chile, India, and Mexico (OECD (2010) and authors' calculations.)

Historically, in the United States, public education has been the responsibility of local governments, and the main source of revenue to fund public schools has been property taxes. Starting with Serano vs. Priest ruling of California Supreme Court in 1971, the last four decades were characterized by a trend of increased centralization of education funding. The aim of Serano vs. Priest decision was to reduce the inequality of education spending across different school districts. Such reforms of education finance entailed using redistribution schemes to equalize spending across districts and setting limits on property taxes that can be levied. Consequently, the share of local funding decreased from 52.5% in 1970–1971 to 43.9% in 2006–2007, while the state funding share increased from 39.1% in 1970-1971 to 47.6% in 2006-2007 (Snyder and Dillow (2010), Table 172). Although there is a lot of variation across states in the United States in the importance of local funding, 75% of students enrolled in public schools are enrolled in states where the local financing share is at least 30%. Figure 9.12 illustrates the cumulative public education enrollment in the US states as a function of how much of the state's public education spending comes from local finance (i.e., school districts within the state).

If the level of public goods provision were the only concern behind household location decisions, then the fall in the transportation costs during the last century is likely to have increased the sorting between communities. However, Cutler, Elmendorf, and Zeckhauser (1993), Cutler, Glaeser, and Vigdor (1999), Kremer (1997), and Rhode and Strumpf (2003) found no evidence of increased sorting. These results suggest there



Figure 9.12 Percent Local Finance and Cumulative Enrollment in the US Public Schools, 2006–2007. *Source: Snyder and Dillow (2010; Tables 35 and 173).*

are other factors such as housing quality, local amenities (air quality, traffic congestion, etc.), or peer effects in a community that might be important in explaining the observed sorting patterns observed. There is ample evidence that sorting is driven both by differences in public good provision and by demographic characteristics (Black (1999), Bayer, Ferreira, and McMillan (2007), Calabrese, Epple, Romer, and Sieg (2006)). Thus, households' mobility across jurisdictions can lead to significant heterogeneity in education quality, not only due to different spending levels but also because of the various neighborhood characteristics. This heterogeneity can have important consequences in terms of welfare and inequality, which in turn alter the political equilibrium and the resulting policies.

In this section, we consider the literature that studies the endogenous determination of education funding in connection with residential segregation of households between different communities. The essential features of this class of models are (1) heterogeneous communities (e.g., school districts); (2) heterogeneous households with respect to income and other characteristics (e.g., age, preferences for education); (3) mobility across communities; (4) an education production function that connects the financial and non-financial inputs to the future earnings of households in that community; (5) a political economy process that translates the household preferences over the quality of education into a collective choice of the tax rate and hence the level of local spending on education.

In Section 5.1, we present a simple model of Tiebout sorting by income in order to illustrate some basic theoretical results. We then discuss the implications of peer effects on the quality of schooling in this model, the role of housing markets, commuting costs, and the effect of interjurisdictional spillovers. Next, we move on to discuss results from the recent work on residential segregation and education funding when private schooling is available. This body of literature uses more complex, calibrated models with several dimensions of household heterogeneity as well as institutional details on school

funding. These models have been used to shed light on the mechanisms that generate sorting patterns across communities that are consistent with data. On the normative side, they are useful in analyzing policy issues such as the choice of the education funding system (centralized funding, local funding, or vouchers).

5.1. A Simple Model of Residential Segregation

The model in this subsection was developed by Gradstein, Justman, Meier (2005). The population consists of two income groups, the poor and the rich, indexed by i = p, r. Assume that the poor are more numerous than the rich. Each household has one child. Households choose to reside in one of two communities, indexed by j = 1, 2. For simplicity, we abstract from the housing markets and assume that public education is financed entirely by a local income tax, chosen by majority voting within the community. There is no private provision of education. Children must attend the public school in their community of residence. (See Epple and Romano (2003) for a model in which households can choose a public school located in another community.)

Households derive utility from consumption of a numeraire good and the quality of public schooling in their community, which is determined by the amount of public spending per child. Denoting by τ_j , the tax rate in the community *j*, the quality of public education (or the spending per student) is $E_j = \tau_j Y_j$, where Y_j is the mean income level in community *j*. Using the CRRA preferences in (9.1), the utility of a household from income group *i* residing in community *j* is given by

$$u(c_{ij}, E_j) = \frac{1}{1 - \sigma} \left\{ \left[(1 - \tau_j) \gamma_i \right]^{1 - \sigma} + \delta(\tau_j Y_j)^{1 - \sigma} \right\},\$$

where consumption c_{ij} equals the after-tax income of household *i* in community *j*.

The tax rate τ_{ij} preferred by household γ_j in community j is given by

$$\tau_{ij} = \left[1 + \delta^{-\frac{1}{\sigma}} \left(\frac{\gamma_i}{Y_j}\right)^{\frac{1-\sigma}{\sigma}}\right]^{-1}$$

An equilibrium is characterized by an allocation of households across communities such that no household has an incentive to move and the tax rates are supported by the majority of the residents in each community. As the preferences are single peaked, the desired tax rate in each community will be the one preferred by the median-income household in that community. When the preferred tax rate is increasing in income ($\sigma > 1$), two types of equilibria obtain. In one case, there is perfect segregation, with poor and rich households residing in different communities, with low and high taxes, respectively. The other type of equilibrium is non-segregating: one community is populated by the poor and the other is mixed, with the rich being the majority. In the mixed community, the taxes and the quality of schooling are higher. The poor are indifferent between the two communities, as the disutility from higher taxes is compensated by the higher quality of public education.

The mixed equilibrium is inefficient compared with the segregating one because of the externality generated by the choice of residence. Households do not take into account the effect of their location decision on the tax base and implicitly, on the quality of education in the community. Consequently, a transfer from the mixed community to the poor one is Pareto improving as it reduces the fraction of poor in the mixed community and raises the spending on schooling in both communities without changing the tax rates. The key insight is that policies that have a positive effect on the poorest community (such as income redistribution or in-kind transfers toward educational expenditures) could produce overall welfare improvements. Thus, there is no equityefficiency trade-off in this case.

Fernandez and Rogerson (1996) study these issues in a model of residential segregation with two communities and three income groups. Under the standard assumption of single-peaked preferences, a stratified equilibrium obtains in which the quality of public education differs across the two communities. Households from the high- and the low-income types reside exclusively in the high- and the low-quality communities, respectively, while the middle-income households live in both communities. Subsidizing some middle-income households to move from the rich to the poor community is welfare improving.

5.2. Local Human Capital Externalities

As discussed above, in addition to local public goods provision, there are other segregating forces that include different forms of local "social capital": peer effects such as human capital spillovers, role models, norms of behavior, neighborhood effects such as crime levels (Bénabou (1994)). Peer effects are perhaps the most relevant form of social capital in the education process.

Peer-group effects can be modeled by assuming that education quality in community j is a function of both spending per child and the average household human capital in the community. For simplicity, we assume the household's human capital is proxied by its income. Thus, the quality of public education is

$$E_j = f(s_j, Y_j),$$

where $f(\cdot, \cdot)$ has the usual properties of concavity and monotonicity, and $s_j = \tau_j Y_j$ is the spending per student.

The same qualitative results obtain as in the simple model in Section 5.1. In the non-segregating equilibrium, one community is populated by the poor and the other is mixed. The poor residing in the mixed community pay higher taxes but are compensated by higher spending per student and stronger peer effects. In the case in which there are more than two types of income, households with income above a certain

threshold will reside in one community and those with below the threshold income will reside in the other one. Subsidizing some marginal households to move to the poor community is welfare improving. Similar results are obtained by Bénabou (1993, 1994) in dynamic models that connect the local human capital spillovers with segregation and income inequality.

5.3. Housing Markets

As mentioned before, in the United States, an important fraction of the public school spending is financed locally, through property taxes. In this context, local education funding can be regarded as a particular case of public good provision under residential mobility. Epple, Filimon, and Romer (1984, 1993), Epple and Romer (1991), and Epple and Sieg (1999) study the properties of multi-district models, where taxation of a location specific asset (e.g., housing) is used to finance a generic local public good.

In the following, we incorporate competitive housing markets into the basic twodistrict model of residential segregation in Section 5.1, drawing on Epple, Filimon, and Romer (1993). Public schooling in jurisdiction j (j = 1, 2) is financed by a local property tax, τ_j , which is determined by majority voting. Households choose their location and the consumption of the numeraire and housing services. Denoting with p_j^h , the net-of-tax price of a unit of housing services in the district j, the gross-of-tax price in district j is $p_j = (1 + \tau_j)p_j^h$. We augment the preferences in (9.1) with housing services. The household with income γ_j residing in district j solves the following maximization problem:

$$u(c_{ij}, E_j, h_{ij}) = \max_{\{c_{ij}, h_{ij}\}} \frac{1}{1 - \sigma} \left\{ c_{ij}^{1 - \sigma} + \delta E_j^{1 - \sigma} + \beta h_{ij}^{1 - \sigma} \right\}$$

s.t. $c_{ij} + p_j h_{ij} \le \gamma_i$,

where c_{ij} and h_{ij} are the consumption of the numeraire and housing services, respectively, and E_j is the quality of public education in district *j*. The local government budget is balanced:

$$E_j = \tau_j p_i^h H_j, \tag{9.15}$$

where H_i is the average housing services in district *j*.

The housing demand of household y_i residing in district *j* is

$$h(p_j, y_i) = y_i \frac{p_j^{-\frac{1}{\sigma}}}{\beta^{-\frac{1}{\sigma}} + p_j^{1-\frac{1}{\sigma}}}.$$

An equilibrium for the decentralized regime with no private schooling is an allocation of households across communities, an allocation of consumption good and housing across households $\{c_{ij}, h_{ij}\}_{i,j=1}^{2}$, a vector of house prices $\{p_{j}^{h}\}_{j=1}^{2}$, and a vector of public policies $\{E_{j}, \tau_{j}\}_{j=1}^{2}$, satisfying the following conditions:

- 1. Given the housing prices and the public policies of districts, each household maximizes its utility.
- 2. No household wants to move from its district.
- 3. Given the policy choices in the other district, the local public policies $\{E_j, \tau_j\}_{j=1}^2$, are preferred by a majority of residents in district *j*.
- 4. Housing markets clear in each district.
- 5. Local governments balance their budgets.

Next, we move on to the equilibrium analysis. The indirect utility of household i in district j is

$$V(p_j, E_j, \gamma_i) = u(\gamma_i - p_j h(p_j, \gamma_i), E_j, h(p_j, \gamma_i)).$$

Denote by

$$M(p, E, \gamma) = -\frac{\partial V(p, E, \gamma)}{\partial E} \Big/ \frac{\partial V(p, E, \gamma)}{\partial p}$$

the marginal rate of substitution between E and p of a household with income y. The absolute value of M(p, E, y) is the slope of a household's indifference curve in the (E, p) space.

Under the assumption that $M(p, E, \gamma)$ is increasing in income, the indifference curves of households with different incomes cross only once in the (E, p) space (single-crossing property). The indifference curve of the poorer household crosses that of the richer household from above. Epple, Filimon, and Romer (1993) show that given the single-crossing property and given a set of residents, a majority voting equilibrium over tax rates exists in each district, and the equilibrium tax rate is the most preferred tax rate of the median-income resident. (If the distribution of housing types is exogenous, then the single-crossing property is not necessary for the existence of an equilibrium (see Dunz (1989) and Nechyba (1997)).

Assume that the voters in district j are myopic when they choose the tax rate, that is, they do not anticipate the effects of the tax rate on the aggregate housing demand, the pre-tax housing price and the location choices of other households. (Epple and Romer (1991) consider non-myopic voting.) However, the voters anticipate how changes in the tax rate affect the local budget. Using the government budget constraint (9.15) in the expression of the gross-of-tax price, we obtain

$$pj = p_j^h + \frac{E_j}{H_j}.$$
 (9.16)

Consequently, the most preferred per-student spending of household γ_i residing in district *j*, E_{ij} , is a solution to the following maximization problem:

$$\max_{E_{ij} \ge 0} V(p(E_{ij}), E_{ij}, \gamma_i) \text{ s. t.}$$
(9.17)

Equivalently, E_{ij} solves

$$M(p(E_{ij}), E_{ij}, \gamma_i) = p'(E_{ij}),$$
(9.18)

that is, the marginal rate of substitution between E and p must equal the marginal change in the price of housing. The education quality (and implicitly the tax rate) most preferred by the median-income household voter is the solution to (9.18), where $y_i = y_m$.

Incorporating competitive housing markets into the basic model does not alter the main insights derived in Section 5.1. Housing is just another normal good in the utility function and its demand is increasing with income. Households face a trade-off between higher quality of education and higher house prices. Higher income households prefer higher quality of education and thus choose to live in districts with higher taxes and house prices. In equilibrium, the marginal household is indifferent between the two districts as the higher housing price is exactly compensated by the higher quality of schooling. In an economy with J jurisdictions in which household income is distributed uniformly over the interval $[\gamma_L, \gamma_H]$, Epple, Filimon, and Romer (1993) prove that the equilibrium is stratified: each district is formed by a single-income interval and the equilibrium set of communities fully partitions $[\gamma_L, \gamma_H]$ into J intervals.

Recall that in the basic model in Section 5.1, as well as in the model described in this section, there is no private education. When the option of private schooling is available, the resulting equilibrium is not stratified. To get some intuition, consider the simple case with two income groups and two school districts with the possibility of opting out of public education into private education. Private education is not specific to district and households in both districts can choose their desired quality of private school. Suppose that the economy is in the segregating equilibrium in which the rich live in one district and the poor live in the other district. For simplicity, assume the supply of housing is fixed in both districts. The possibility of opting out of public education breaks the link between residential location and school choice and can reduce the residential segregation. Rich households can be better off moving to the poor district where they can benefit from lower taxes and house prices and enjoy private schooling. This happens if the difference in the house prices between the two districts (or the difference in incomes between the rich and the poor) is sufficiently large. If the rich move to the poor district, in the new equilibrium the housing prices in the poor district rise while they fall in the rich district. Moreover, the poor district gets additional

tax revenues, and hence, the quality of the public schooling rises in the poor district while it remains unchanged in the rich one. The change in the welfare of the poor is ambiguous. It depends on how the gain from better quality of the public schools compares with the loss from higher house prices.

Mora (2006) studies an economy with multiple communities and multiple-income groups when opting out of public education is allowed. His main result is that some intermediate-income households choose private schooling in poorer neighborhoods since that choice is affordable while a better public school in a richer neighborhood with higher housing prices is not affordable.

Another extension of this class of models considers zoning regulations (e.g., restrictions on building height, lot size, etc.). These restrictions impose a minimum required housing consumption for district residents. Fernandez and Rogerson (1997) study the effects of both exogenous and endogenous zoning regulation in a two-district model with local public schooling. The zoning regulations affect the location decision of the marginal individuals and thus change the income distribution in each district. If the relatively poor households are pushed out from the rich district, the average income increases in each district. When zoning is endogenized, the two policies are chosen sequentially: households vote on zoning first and then on the tax rate. Fernandez and Rogerson show via numerical simulations that endogenous zoning always generates higher tax rates and quality of schooling in both districts, whereas in the case of exogenous zoning, this result obtains only over a particular range of zoning levels. In Calabrese, Epple, Romano (2007), households vote simultaneously on the zoning restrictions and the local tax rate that finances the public good. They find that zoning improves the aggregate welfare, although with significant welfare losses for poorer households.

5.4. Commuting Costs

Recent literature attempts to merge the Tiebout framework with elements from urban location theory (Alonso (1964), followed by Mills (1967) and Muth (1969)), where the sorting of households across different communities is driven by the presence of spatial amenities (e.g., de Bartolome and Ross (2003) and Epple, Gordon, and Sieg (2010)). Incorporating the spatial component into the model gives rise to an additional trade-off between land use and accessibility of employment when choosing the residence. On the one hand, commuting to the workplace, which might be located in the Central Business District (CBD), is costly for households, and the opportunity cost of commuting is increasing in income. On the other hand, if land is a normal good, the rich households are likely to move away from the city. Such models yield predictions not only with respect to the sorting of different households across districts but also regarding their distance from CBD.

Hanushek, Sarpça, Yilmaz (2007), Hanushek and Yilmaz (2007, 2010), merge both public good provision and spatial constraints as main determinants of households' location in multi-community models of education choice. Consider a simple setup based on Hanushek, Sarpça, Yilmaz (2007), with two districts and two income groups (poor, P, and rich, R). The households also differ in their valuation of education quality (high, H, and low, L). Thus, there are four types of households: RH, RL, PH, PL. The two districts, West and East, are separated by a straight line passing through the CBD, where the employment opportunities are located. Households can live anywhere on this line. They care about consumption of the numeraire good, quality of housing, and quality of public education. As in the basic setup in Section 5.1, there are no private schools.

Housing quality is proxied here by the lot size. The plots are owned by absentee landlords and are sold by auction to the highest bidders. The rent per unit of land in district *j*, $R_j(d)$, depends on the distance *d* from CBD. In both jurisdictions, the public schools are financed by property taxes on the value of land owned. The tax rates are determined through majority voting. Households incur a commuting cost that is proportional to the distance. The income net of transportation costs of a household with earnings w_i is

$$\gamma_i(d) = w_i - (a + bw_i)d,$$

where a, b > 0. Households choose the distance to their residence taking into account the price of land and the transportation cost.

For simplicity, suppose that preferences are logarithmic. The household with income $y_i(d)$ located in district *j* solves the following problem:

$$\max_{\{c_{ij}, h_{ij}\}} \alpha \ln c_{ij} + \delta_i \ln E_j + \beta_i \ln h_{ij}$$

s.t. $c_{ij} + (1 + \tau_j) R_j(d) h_{ij} \le \gamma_i(d),$

where $\alpha + \delta_i + \beta_i = 1$. The utility maximizing budget shares for consumption and lot size are $\alpha/(\alpha + \beta_i)$ and $\beta_i/[(\alpha + \beta_i)(1 + \tau_j)R_j(d)]$, respectively. Thus, the indirect utility function can be written as

$$V(.) = (\alpha + \beta_i) \ln \gamma_i - \beta_i \ln \left[(1 + \tau_j) R_j(d) \right] + \delta_i \ln E_j + \ln k_i,$$

where $k_i = (\alpha / (\alpha + \beta_i))^{\alpha} (\beta_i / (\alpha + \beta_i))^{\beta_i}$.

The maximum rent a household i is willing to pay per unit of land in district j (located at distance d from CBD) for a given utility level u is the solution to the following problem:

$$\Psi(d, E_j, \tau_j, u) = \max_{c_{ij}, h_{ij}} \left\{ \frac{\gamma_i(d) - c_{ij}}{(1 + \tau_j)h_{ij}} \mid U(c_{ij}, h_{ij}, E_j) = u \right\}.$$

The left-hand side is the bid rent function, and it can be expressed analytically as

$$\Psi(d, E_j, \tau_j, u) = \frac{(\gamma_i(d))^{\frac{\alpha+\beta_i}{\beta_i}}}{1+\tau_j} (k_i E_j^{\delta_i})^{\frac{1}{\beta_i}} e^{-\frac{u}{\beta_i}}.$$

The slope of the bid rent curve is negative:

$$\frac{\partial \Psi}{\partial d} = -\frac{\alpha + \beta_i}{\beta_i} \Psi(.) \frac{a + bw_i}{\gamma_i(d)} < 0,$$

meaning that all types of households are willing to pay lower rents at farther distances from CBD. The equilibrium rent curve $R_j(d)$ is the upper envelope of the bid rent curves of all household types and is decreasing in d.

Households with steeper bid curves locate closer to CBD. For example, if the bid rent curves of two different households intersect at distance d^* from CBD, $\Psi(d^*, u_1) = \Psi(d^*, u_2)$, then household 1 has a steeper curve if the following condition is satisfied:

$$\frac{(\alpha+\beta_1)\beta_2}{(\alpha+\beta_2)\beta_1} \frac{a+bw_1}{a+bw_2} \frac{\gamma_2(d)}{\gamma_1(d)} > 1.$$

If the households live in the same district and have the same income, then the household with a high valuation of education (high δ , low β) will live closer to CBD. When comparing households who have the same preferences over education but different incomes, the outcome depends on how the gain from cheaper rents compares with the increase in transportation costs as *d* increases. Building on this analysis, we discuss results from numerical simulations in Section 5.7.3.

5.5. Interjurisdictional Spillovers

So far, we have considered multi-community models of public education provision in which parents move across jurisdictions such as to maximize their utility from private consumption (including housing) and the education quality available for their children. However, this setup ignores the production side and the long-run consequences that labor mobility can have on the provision of local public goods.

Although jurisdictions spend tax revenues to provide public education, the beneficiaries may well choose to reside and work in a different community if the benefits of relocating exceed the costs. Thus, labor mobility induces an interjurisdictional spillover in addition to the previously considered intrajurisdictional spillovers such as peer effects.

In this context, under decentralization, public education becomes an instrument of fiscal competition as each jurisdiction takes into account the policies implemented by its neighbors. Equilibrium policies depend on the balance of political forces in the economy. Justman and Thisse (2000) consider a simple model that looks at public provision of education when mobility increases with human capital. They show that decentralization results in underprovision of education, unless the high mobility types are relatively influential in the political process. These results on strategic education spending and their welfare implications mirror findings from the literature on fiscal competition in productive public inputs (Bayindir-Upmann (1998), Dhillon, Wooders, and Zissimos (2007), Arcalean, Glomm, Schiopu, and Suedekum (2010)).

5.6. Education Funding Systems

As noted earlier, in multi-community models, an important policy question is what type of education funding system should be adopted – pure local, centralized, or a mixture of local and centralized. Under a local or decentralized funding regime examined in the previous subsections, public education is financed entirely through local taxes as in the benchmark model in Section 5.1. The pure centralized regime corresponds to a state-funded system with voter-determined taxes under which all public schools receive the same funding per pupil, and districts are not allowed to supplement spending. The spending per student is thus $E = \tau Y$, where τ is the statewide income tax rate, and Y is the economy-wide income as in Section 2.1.1.

The mixed regimes of education financing combine local and central funding sources in various ways and aim to equalize the disparities in per-student spending. One example of a mixed scheme is the state foundation grant (in the next section, we discuss other mixed modes of funding). Under the state foundation grants regime, the state funds a minimal level of education spending using tax revenues raised centrally. Each community has to spend at least the amount of the foundation grant and can use local funds raised through majority-backed property taxes to supplement the state funds. Such a system can be found in some European countries, such as Germany. The states in the United States have increasingly adopted such schemes in order to reduce the inequality in education spending across districts generated under local funding.

Gradstein, Justman, Meier (2005) compare the political equilibria under different education funding regimes, but with exogenous residential segregation. For simplicity, assume that households' income is identical within a district, that is, the index of the household income is also the index of a district. Thus, the support for a certain funding regime depends on the tax prices of public education associated with each type of funding. Districts with above average income favor local relative to state funding as the tax price of providing the same level of education quality is lower under the former regime. In general, districts with above average income prefer the centralized regime. However, an "ends against the middle" equilibrium can arise in which the very poor districts prefer a lower tax than that under the centralized system and thus support the local regime, together with the rich communities.

Consider now the state foundation grant regime. First, households vote on the central tax that funds the foundation grants and then on the community tax that is used to fund the local supplementary spending. Thus, a pure local funding regime can be considered a special case of the mixed regime in which the central tax and the foundation grant are set to zero. As local and state spending are perfect substitutes, the preferred taxes at various income levels reflect the tax prices associated with each type of funding. As discussed above, for the households residing in districts with income below the mean, the tax price of centralized funding is lower than that of local funding. These households prefer a zero local tax and thus support full centralized funding. By the same token, the households in districts with income above the mean prefer a zero central tax and support full local funding. An income distribution in which the median is lower than the mean implies that the decisive voter prefers the full centralized funding. A majority voting equilibrium exists for both $\sigma < 1$ and $\sigma > 1$.

For $\sigma < 1$, the preferred tax rate is decreasing in income and the decisive voter of the central tax lives in the district with median income and prefers pure central funding with no local supplementation (zero local tax). Denoting by $\tau_{\rm m}$ the preferred central tax, the spending on foundation grant is $f = \tau_{\rm m} Y$, where Y is the economy-wide income. When there is no local supplementation, the public spending on education equals the amount of the foundation grant, E = f. Assuming the preferences in (9.1), the preferred central tax rate of the median-income household is

$$\tau_{\rm m} = \left[1 + \delta^{-\frac{1}{\sigma}} \left(\frac{\gamma_{\rm m}}{Y}\right)^{\frac{1-\sigma}{\sigma}}\right]^{-1}.$$

As discussed above, richer districts prefer a zero central tax and full local funding. Thus, when the foundation grants are positive, there is an income threshold above which districts choose to supplement the foundation grants. These districts spend less on public education under a foundation grants regime than under a pure local regime. This is because the central funding is more expensive than local funding for these communities. Along the same lines, they will spend more under the foundation grants regime than under pure centralization. Conversely, the relatively poorer districts that do not supplement the central funding spend more under the foundation grants regime than under pure local funding but less than under pure central financing.

For $\sigma > 1$, the preferred tax rate is increasing in income, and there is an "ends against the middle" equilibrium. Districts with income higher than the average, which prefer zero central tax and zero foundation grants, form a coalition with low-income districts, which prefer smaller foundation grants than the middle-income districts.

5.7. Results from Calibrated Models

Models that incorporate multiple dimensions of household heterogeneity and institutional details of public and private education are too complex to yield closed form analytical solutions. Consequently, researchers have resorted to simulations in order to deliver implications of the model for observed patterns and to perform policy experiments. The goal is to specify a theoretical framework that is rich enough so that its parameters can be calibrated or estimated to reproduce essential features of the data, such as the distribution of house prices or stratification patterns across communities and type of schools, given the current institutional setting in the education market. Such models have been used to study counterfactual policy experiments, such as the introduction of state-funded vouchers or the switch from a local to a pure centralized or mixed-funding system.

5.7.1 Local versus Centralized Funding with No Private Option

Fernandez and Rogerson (2003) use a calibrated model to study the implications for various education funding systems: the local funding regime, the pure centralized funding regime, and some mixed regimes that combine local and central funding. Examples of mixed schemes they consider are (1) state foundation grants and (2) district power equalization (DPE) with and without recapture.

In the district power equalization scheme with recapture (PER), each district chooses its own tax rate τ_j . The state guarantees a common tax base *S*, determined endogenously: Thus, the spending in district *j* is obtained by applying its tax rate τ_j to the statewide tax base:

$$s_j = \tau_j S.$$

The actual tax revenues raised equal $\tau_j Y_j$, where Y_j is the tax base in community *j*. The difference between the guaranteed tax base and the actual one is negative in the rich districts $(S < Y_j)$, where the government collects the balance to finance the transfers to poor districts $(S > Y_j)$. Finally, the state levies a tax rate τ_s in order to balance the budget:

$$\tau_S Y = \Sigma_j \tau_j \left(S - Y_j \right),$$

where *Y* is the average state income.

In the power equalization scheme with no recapture (PEN), the tax revenue in districts with mean income greater than the statewide average income is determined by taxing their own tax base, Y_i :

$$s_j = \begin{cases} \tau_j S, & \text{if } Y_j \leq S; \\ \tau_j Y_j, & \text{if } Y_j > S. \end{cases}$$

As in the case of PER scheme, a statewide tax rate is used in order to balance the budget.

In a series of papers, Fernandez and Rogerson (1998, 1999, 2003) use calibrated models with pure public education (no opting out) to compare different funding regimes in terms of welfare, total education spending and inequality of spending across districts. A key result is that the mixed-funding regimes generate the highest spending on education. Numerical simulations reveal that for a large range of parameter values, the PER system is welfare superior to other funding regimes, and it is the regime preferred by the majority of voters. When the elasticity of substitution is close to 1, the foundation grant regime beats the PER system.

5.7.2 Tiebout Models with Property Taxes, Peer Effects, and Private Education

Nechyba (1999) develops a multi-community model with housing markets and local public schools where the option of private schooling is available. The private schools are modeled as clubs of parents that share the cost equally. An essential ingredient of Nechyba's framework is the education production, which is a function of both per-pupil spending and average peer quality within the school. The peer quality is assumed to be positively correlated with the household's income. The weight attached to peer effects in the education production function is essential to reproduce the observed levels of private school attendance. Simulations reveal, for example, that if the peer effects are not sufficiently important for the quality of education, the households with high-quality peer effect do not have an incentive to opt out of the public school system. Thus, an equilibrium with no private schooling obtains. (In a more recent paper, Calabrese, Epple, Romer, and Sieg (2006) find that peer effects, defined to include both consumption and education externalities are essential in generating the pattern of tax rates in the data.) When the peer effects are strong enough, the private school markets exist and the degree of the residential segregation by income and peer quality is lower since the richer households have an incentive to choose a private school in a poorer district in order to benefit from lower house prices. At the same time, the school segregation by income increases.²

This framework has been used in two major policy applications: the introduction of private school vouchers and changing the degree of centralization of the public education funding.

Policy Experiments on Private School Vouchers

Nechyba (1999) illustrates the policy implications of private education vouchers in a variety of institutional settings:

- 1. A system of pure local funding of public schools.
- **2.** A state-funded system with voter-determined income taxes, under which all public schools receive the same funding per pupil.

 $^{^2}$ For a detailed discussion on how the results obtained in this framework compare with other models of student sorting, see Nechyba (2006).

- **3.** A local funding system with state foundation grants funded through a state income tax.
- **4.** A local funding system with district power equalization (DPE), where the state guarantees a per-pupil tax base in each district.

The voucher program is similar to the one in Section 3.2. Each household is entitled to a voucher that can be used to pay the private school tuition. (See also Rangazas (1995b), Nechyba (2000), and Ferreyra (2007) for the effects of vouchers in multi-community environments.)

The introduction of vouchers in a local funding system exacerbates the desegregating effects induced by the availability of the private option, as the marginal households who used the public schools in the non-voucher regime now opt out and migrate to the poor communities. Thus, the private schools appear first in the poorest community and attract the rich newcomers, whereas the middle- to high-income residents leave for other communities. Simulations imply that the quality of the public schools in the rich communities decreases, whereas the public schools in the poorer communities may improve under the voucher regime.

The peer quality in the poor community's public schools falls as the high-peer-quality residents who were previously using the public schools leave the community and cluster in the private school. At the same time, enrollment in public schools decreases, and the support for public spending declines. However, as the rich families using the private school live in relatively higher priced houses, they pay higher property taxes. This effect counteracts the reduced support for public spending, and as a result, perpupil funding in the poor communities rises in equilibrium with the voucher level. Nonetheless, as more families choose the private schools, at high voucher values there is no longer a majority support for public education and the public school system collapses.

Recall that the quality of schooling is a function of both peer quality and spending per student. If the decrease in peer quality is more than offset by the positive effect on per-pupil spending, the quality of public schools in the poor communities can improve after the introduction of vouchers. This happens if the weight of resources spent per pupil in the education production is sufficiently high relative to that of peer effects. In contrast, in the rich communities, the quality of public schools unambiguously decreases. As high-income and peer-quality households choose private schools in other communities, the peer quality and the public support for high per-pupil spending diminish.

The negative effects of vouchers on the stratification by income described above also arise under regimes (2), (3), and (4). However, the response of per-student public spending in various communities will differ, depending on the pre-voucher funding regime. When vouchers are introduced in a centralized state-funded system (2), for instance, the positive change in the per-pupil public spending in poor communities is smaller than under pure local funding. The response of private school enrollment to the voucher level is weaker. Under centralized funding, the benefit of high-income households from migrating and choosing private schools in the relatively poorer communities is lower than under the decentralized regime, as they still have to pay the state tax, regardless of their residence. In contrast to case (1), vouchers improve per-pupil spending in public schools located in rich communities, as some households opt out of the public system.

Under regimes (3) and (4), the quality of public schools is decreasing in the voucher level in all communities. In the foundation grants regime, the driving force behind the result is the response of private education markets to an increase in the level of grants. For instance, at high levels of foundation grants, the private school market disappears in the poor community. The change in the voucher level does not impact the per-pupil spending in poor community's public school. High-peer-quality households migrate to other districts and they are not replaced by high-income newcomers who use the private school as in the case of previous experiments. With no counteracting effect for the decrease in peer quality, the quality of public education in the poor community falls.

Unlike the foundation grants, the DPE system generates high disparities in the price of public schools across districts. The price of public schooling increases substantially in the rich district, generating a rapid switch to private schools. Thus, high voucher levels can produce a collapse of the public education in the wealthiest community as the voted property tax is driven to zero. This undermines the transfers from rich to poor communities that take place under DPE. The per-pupil spending in the poor community still rises but the effect is not as strong as in the case of local and state funding.

Policy Experiments on Public School Funding

Nechyba (2003) extends the framework in Nechyba (1999) to investigate the general equilibrium effects on mobility and quality and enrollment in private and public schools as a result of a change in the public school finance policies, such as a switch from a pure local funding system to a centralized (state-funded) regime. Under the assumption that state and locally financed schools are equally productive, the switch to a centralized funding system leads to higher resources and higher quality of public schools in the poor communities. Consequently, the enrollment in private schools decreases. The opposite occurs in the rich communities. However, the existence of housing markets introduces a general equilibrium effect that generates a further decrease in the private school enrollment after the switch to a state-funded regime. Under the centralized system, private school users who move to the poor community in order to benefit from lower house prices and property taxes, have fewer incentives to do so. This is because the state funding reduces the public school quality disparity between the poor and rich communities and the tax rate is uniform across districts. Simulation results suggest that the

aggregate private school enrollment is slightly lower under state funding. The same result holds for the aggregate per-student spending in public schools.

The change to a mixed-funding regime generates the same effects discussed above. However, the decline in the private school enrollment is larger than in the previous experiment because under the mixed regimes, the various communities are still able to tailor the amount of public spending to their preferences by changing local property taxes.

5.7.3 Tiebout Models with Commuting Costs

Hanushek, Sarpça, and Yilmaz (2007) calibrate and solve numerically a model with commuting costs, heterogeneity in income, and heterogeneity in preferences for education quality. They augment the model presented in Section 5.4 with peer effects and private school markets.³ Within each income group, there are two types of households, according to the value they attach to education (high or low). There is one private school that admits students from both communities. As in Nechyba (1999), the voters are myopic, that is, they take the school and residence choice of the other households as given when voting on local taxes.

The equilibria with and without the private option exhibit income heterogeneity in both communities. This is in line with theoretical results (Epple and Platt (1998)) and empirical studies (Davidoff (2005)). Similar to the theoretical results in Section 5.4, the land rents decrease with the distance from the CBD. At the same income level, the households with a lower valuation of education (higher valuation of land) will choose to reside farther from the CBD. When comparing households with the same preferences for education, higher income generates two opposite effects: higher demand for land and thus an incentive to locate farther from the CBD, and an increase in the opportunity cost of commuting time. Simulation results suggest that the first effect dominates. Thus, the PH (poor with high valuation of education) households reside closest to the CBD, followed by PL, RH, and RL communities. In equilibrium, one of the districts has higher taxes and is inhabited only by households who use public schools. Different communities partition each jurisdiction into semirings with the CBD at the center.

When the private option is not available, the rich households segregate by valuation type—RH in the high-tax district and RL in the low-tax district—and locate farthest from the CBD. The PH types live in the high-tax district, close to the CBD, and PL types live in both communities. the PL types in the low-tax district live close to the CBD, whereas in the high-tax district, they locate between PH and RH, accepting higher taxes in exchange of a higher lot size.

³ Hanushek and Yilmaz (2007) and Hanushek and Yilmaz (2010) also consider a similar Tiebout model with commuting costs. However, their policy experiments are focused on the case where private schooling is not available.



Figure 9.13 Spatial Distribution of Communities in a Two-District, One-Center Model.

When private schools are available, the equilibrium exhibits less residential segregation, with each type of household present in both districts (see Fig. 9.13).

The private school has only RH types, who choose to reside in the lower-tax district. The per-student spending in the public school is higher in the low-tax district as rich households who use the private schools contribute a significant portion to the district's education budget. However, the lower peer-quality effect turns out to be stronger, generating a lower quality of public education in the low-tax district.

Relative to the equilibrium with only public education, when private schools are available, both the quality of education and the per-student spending are lower in the low-tax district and higher in the high-tax one. Most of the differences in quality between communities are due to peer effects rather than differences in spending.

A switch from a decentralized funding regime to a centralized system with private schooling results in lower quality of education and lower per-student spending in both districts. The difference in quality of public schools between districts is larger than in the decentralized funding regime. As the low-valuation households are the majority, the common tax rate in the centralized regime is lower than in the decentralized regime. The decrease in residential segregation has a negative effect on rents and thus on the tax base, so the overall public education spending decreases. The low quality of public schools makes the private option more attractive. Although the peer quality in the private school is unaffected (only the RH types send their children there), the education spending declines. As a result, the quality of private schooling decreases as well.

6. DYNAMICS

The models in the previous sections are purely static. Including dynamic components in the model would provide implications for the evolution of income distribution and hence evolution of per capita income, growth rate and inequality. The dynamic components would also help us analyze the short- and long-run effects of changes in education policies. The next subsection presents a simple model developed by Glomm and Ravikumar (1992). It extends the benchmark framework presented in Section 2.1 by connecting education spending choices in the current period to income in the next period. To keep the dynamics analytically tractable, it contrasts two extreme cases: pure public education and pure private education.

6.1. Private versus Public Education: Growth and Inequality

Consider a two-period overlapping generation model with constant population of size 1. All agents born at t = 0, 1, 2, ... have identical preferences represented by

$$\ln n_t + \ln c_{t+1} + \ln e_{t+1},$$

where n_t is leisure in the first period of their life and c_{t+1} and e_{t+1} are consumption and education spending in the second period of their life. Here education spending is a bequest to the offspring of the agent. In a pure public education regime, we will denote the education spending by E to indicate that it is a collective choice. In a pure private education regime, e is chosen by each agent.

Agents are endowed with 1 unit of time in the first period of their lives. They can use their non-leisure time to accumulate human capital according to

$$h_{t+1} = \theta (1 - n_t)^\beta e_t^\gamma h_t^\delta,$$

where h_t is the human capital of the corresponding parent, $\theta > 0$, β , γ , $\delta \in (0, 1)$. At time t + 1 the agent's income is h_{t+1} .

At t = 0, the initial generation of old agents is endowed with human capital h_0 that is distributed according to the lognormal distribution (denoted by G_0) with parameters μ_0 and σ_0^2 . The distributions of human capital in subsequent periods (G_t) are endogenous.

In the public education regime, the education spending is financed by a flat tax on income and the government maintains a balanced budget: $E_{t+1} = \tau_{t+1}H_{t+1}$ where $H_{t+1} = \int h \, dG_{t+1}(h)$. The agent's problem at time *t* is to choose n_t in order to maximize

$$\ln n_t + \ln \left\{ (1 - \tau_{t+1}) \theta (1 - n_t)^\beta E_t^\gamma h_t^\delta \right\} + \ln E_{t+1},$$

given E_t , h_t , E_{t+1} , and τ_{t+1} . Although the tax rate and education spending are the same across all agents, the parent's human capital is different across agents. In order to endogenize

the public education spending, we assume that only the old agents vote in each period. Each old agent's most preferred tax rate maximizes

$$\ln\left\{(1-\tau_{t+1})h_{t+1}\right\} + \ln\left\{\tau_{t+1}H_{t+1}\right\},\$$

where h_{t+1} is his own income, and H_{t+1} is the mean income in period t + 1. The equilibrium tax rate in period t + 1 is the one preferred by a majority of old agents in that period. It is easy to see that (1) the optimal time allocated to human capital accumulation is given by $1 - n_t = \frac{\beta}{1+\beta}$ and $h_{t+1} = \theta \left(\frac{\beta}{1+\beta}\right)^{\beta} E_t^{\gamma} h_t^{\delta}$ and (2) the majority preferred tax rate is $\tau_{t+1} = \frac{1}{2}$. The distributions of human capital in periods t = 1, 2, ... are also lognormal, characterized recursively by the parameters

$$\mu_{t+1} = \ln A + (\gamma + \delta)\mu_t + \gamma \frac{\sigma_t^2}{2}$$
$$\sigma_{t+1}^2 = \delta^2 \sigma_t^2,$$

where $A = \theta \left(\frac{\beta}{1+\beta}\right)^{\beta} \left(\frac{1}{2}\right)^{\gamma}$.

In the private education regime, each agent chooses n_t and e_{t+1} to

$$\max \ln n_{t} + \ln c_{t+1} + \ln e_{t+1}$$

s.t. $h_{t+1} = \theta (1 - n_{t})^{\beta} e_{t}^{\gamma} h_{t}^{\delta},$
 $c_{t+1} + e_{t+1} = h_{t+1}.$

The optimal split of future income between own consumption and bequest to the offspring is given by $c_{t+1} = e_{t+1} = \frac{1}{2}h_{t+1}$ and the optimal allocation of time to human capital accumulation is $1 - n_t = \frac{\beta}{\frac{1}{2} + \beta}$, so $h_{t+1} = \theta \left(\frac{\beta}{\frac{1}{2} + \beta}\right)^{\beta} \left(\frac{1}{2}\right)^{\gamma} h_t^{\gamma + \delta}$. The sequence of distributions of human capital is characterized by lognormals with parameters

$$\mu_{t+1} = \ln B + (\gamma + \delta)\mu_t$$

$$\sigma_{t+1}^2 = (\gamma + \delta)^2 \sigma_t^2,$$

where $B = \theta \left(\frac{\beta}{\frac{1}{2} + \beta}\right)^{\beta} \left(\frac{1}{2}\right)^{\gamma}$.

For illustrating the implications for growth and inequality, consider the case $\gamma + \delta = 1$, so there are constant returns to augmentable factors. Both regimes will then exhibit positive endogenous growth assuming θ is sufficiently large. Growth rate of per capita income at any point in time is given by $\ln H_{t+1} - \ln H_t$. Since the distributions are lognormal, we can express the mean income in terms of parameters of the lognormal and write the growth rate as $\mu_{t+1} - \mu_t + \frac{1}{2}(\sigma_{t+1}^2 - \sigma_t^2)$.

In the private education regime, the growth rate is $\ln B$ and in the public education regime the growth rate is $\ln A + \frac{1}{2}\sigma_t^2(\gamma + \delta^2 - 1)$. It is easy to see that the growth rate in

the public education regime is lower. The incentive to accumulate human capital under public education is lower because each agent ignores the effect of his income on the education spending that is, 1 - n is less in the public education regime or, $\ln A < \ln B$. This has a negative effect on growth.

The income inequality (measured by σ) will decline over time in the public education regime, but it will remain constant in the private education regime. With heterogeneity only in initial human capital, the rank of each household in the income distribution does not change over time—if a household was richer than another household at t = 0 it will continue to be richer forever. Finally, Glomm and Ravikumar (1992) show that if choice of regime was put to vote each period, then a majority would prefer public education, provided the income inequality has not completely disappeared. This stems from the fact that the decisive vote in such an election is the agent with median income, and with lognormal distribution, median income is less than the mean income.

Subsequent papers that consider the growth-inequality relationship in political economy models include Saint-Paul and Verdier (1993), Perotti (1993), and Gradstein and Justman (1997). Gradstein and Justman (1997), for example, incorporate both human capital externalities and distortions from taxes that finance public schooling. They show that periodic swings between public and private spending are possible. There is no conflict between the popular support of a system and economic growth: the system that is the choice of the majority also generates the highest growth.

Bénabou (1996b) shows that including random ability among children in addition to the initial human capital heterogeneity could potentially change the growth-inequality trade-off. In an infinite horizon model with local and economy-wide spillovers, he shows that integration might slow down growth in the short run, but raise it in the long run. Soares (2003) uses a dynamic general equilibrium model to assess whether capitallabor complementarities are sufficient to generate large support for public education expenditures without appealing to any form of altruism. He finds that these complementarities are so strong that savers benefit substantially from high investment in human capital. His model delivers a human capital investment ratio of close to 5% of GDP, a result close to the observed public expenditure ratio in the United States. In a companion piece, Soares (2006) finds that the influence of changing social security funding on the level of support for public education is modest.

Cardak (2004b) merges the static benchmark model in Section 2.1 with the dynamic model of Glomm and Ravikumar (1992). He solves for the majority voting equilibrium when public and private education coexist and then feeds the outcome into a human capital production function. This process allows him to trace out the evolution of the income distribution. The human capital production function exhibits diminishing returns to the augmentable factors ensuring that incomes converge to one of two steady-state levels, one for those dynasties attending public school and another, a higher one, for those attending private education. The equilibrium thus exhibits a bimodal

distribution with modes at the two steady states. This separation of the population into two classes is a natural consequence of the dichotomous school choice; but from an empirical point of view, this separation seems stark. In a similar setup, Cardak (2004a) allows for constant returns to scale in the augmentable factors and finds that per capita output along the equilibrium trajectory is highest in the pure private regime, lowest in the pure public regime, with the mixed regime in the middle. Cardak (2005) introduces selective vouchers à la Rangazas (1995b) and Hoyt and Lee (1998) into this model. He finds that increases in the voucher amount increase economic growth through a direct effect—more households choose private education—and an indirect effect through an increase in the tax base that raises public school quality.

Although Cardak's assumption of Cobb-Douglas preferences generate unanimity over funding levels for those households choosing public education, Beauchemin (2001) allows for isoelastic preferences that generate heterogeneity in preferences over funding levels. The production function for human capital here is linear in current human capital and public education expenditure. The overlapping generations are structured so that voters in each generation have to form expectations over future policies chosen. He shows that the dependence of current policies on expected future policies depends crucially on the elasticity of substitution in preferences. A rapidly growing population, for instance, can lead to stagnation in output per capita as the education spending level preferred by the majority (i.e., the young) goes to zero.

Cardak (2004a,b) and Lee and Roemer (1998) aggregate public and private education expenditures into final output via a neoclassical production function. Preferences over own consumption and the bequest left for the offspring are Cobb-Douglas. The Stone–Geary functional form for the production technology ensures that only households with sufficiently high wealth will invest in private education. For all other households, private education expenditures are zero. Each generation votes for the level of public education funding. There are three main channels through which inequality influences the evolution of the income distribution:

- 1. Higher inequality induces higher taxation that diminishes private education investment.
- 2. Higher income inequality can shift households from positive private investment below the threshold to zero private investment.
- 3. Higher inequality induces higher public expenditures.

Their theory does not pin down the relative sizes of these three effects. However, simulations reveal that the inequality-growth relationship can be non-monotonic.

In most of the above models, public education is universal and uniform, that is, all school-age children have access to equal amounts or qualities of education services. Under this assumption, the limiting human capital distribution is often degenerate as tomorrow's human capital is a concave function of today's human capital. The evidence from Gradstein (2003) contradicts this assumption. Roemer (2005) allows for the political process to deliver unequal amounts of education to households. In particular, he allows richer families to receive more education. He finds that under some conditions on the political bargaining within political parties, the income distribution does not necessarily become degenerate. In Bénabou (1996b, 2002) and Seshadri and Yuki (2004), individuals face uninsurable ability shocks every period, so the income distribution does not become degenerate.

6.2. Dynamic Tiebout Sorting

As noted in Section 5.2, there is a body of literature that emphasizes the role of spillovers in explaining the socioeconomic segregation and income inequality across generations. See, for instance, de Bartolome (1990), Durlauf (1996), Bénabou (1993, 1994, 1996a,b), Fernandez and Rogerson (1996), and Fernandez and Rogerson (1998). In this section, we consider dynamic Tiebout models.

As an illustration of a simple dynamic Tiebout economy, consider the setup in Fernandez and Rogerson (1996) with two districts. Individuals live for two periods: youth and old age. When young, agents receive education, while during the old age, they consume and earn an income proportional to their education and produce an offspring. Thus, each household consists of an old agent (parent) and one young agent (child). An old agent's income depends only on the quality of the education received when young and an idiosyncratic random shock. There are no peer effects. The old agents derive utility from consumption of a private good and housing services and the expected income of their child. The preferences of parent *i* residing in district *j* are given by:

$$u(c_{ij}, h_{ij}) + E\gamma_i^c$$

where c_{ij} and h_{ij} are the old agent's consumption of the private good and housing services, respectively. The term $E\gamma_j^c$ is the next period expected income of the young agent in the household and is a function of the quality of public education provided in the district j, q_j .

Assuming that the income of a household can take values in the interval $[y_L, y_H]$ and denoting by $f_z(q_j)$ the probability that an agent has income y_z when old given that the quality of education she received when young is q_j , the expected income when young is

$$E\gamma^{c} = \sum f_{i}(q_{j})\gamma i = \nu(q_{j}).$$

After the realization of the income shock, the old agents decide in which community to reside. Then each community chooses the local tax on house expenditures through majority voting in which each voter takes as given the composition of the community. The tax proceeds are used to finance public education. The quality of public education determines the income distribution for the next generation of adults. Thus, the model can be solved period-by-period in manner analogous to the static model described in Section 5.1.

Fernandez and Rogerson (1998) calibrate their model with two jurisdictions such that in the steady-state equilibrium under local financing, the model matches some relevant features of the US economy. They then examine long run and transitional effects of switching from a local to a centralized funding regime. Their calibrated steady state is stratified, with households above a certain threshold income living in the rich district and the poorer households in the other. In the centralized regime, the housing price is the same across all jurisdictions since the property taxes, and hence, the quality of education is equal in all districts. The spending per student in the steady state of the centralized regime is between the corresponding values for the two communities under the local funding system. The same result holds for the tax rate and the price of housing. However, the regime change increases steady-state average income and the fraction of income allocated to public education.

In order to evaluate the welfare changes, the expected utility at time t is computed as follows:

$$EU_t = \sum_i \lambda_{it} V_{it} \tag{9.19}$$

where λ_{it} is the fraction of the population with income γ_i in period *t* and V_{it} is the indirect utility of a household with income γ_i . As we can see from (9.19), the steady-state-expected utility in the new regime is the result of changes in the income distribution (λ_{it}) and changes in individual utilities (V_{it}). In Fernandez and Rogerson's calibration, all household utilities V_i are higher under the local financing system, as this allows different households to choose the community that suits their preferences for public good provision. On the other hand, the centralized system reduces the dispersion in education expenditures, producing improvements in the income distribution and a higher mean income. In Fernandez and Rogerson's benchmark calibration, the second effect more than offsets the changes in the household utilities. Consequently, they find substantial gains in terms of expected utility from switching to the centralized financing regime—in order to achieve indifference between the steady states of the two systems—the total income in the centralized financing regime would have to be decreased by 3.2% every period.

Bearse, Glomm, and Ravikumar (2001) study the same education financing reform in a calibrated dynamic two-community model in which the private school option is available. Their framework differs from Fernandez and Rogerson's in two other dimensions:

1. There are no housing markets and the public education is financed through income taxes.

2. The two generations are connected by a learning technology of the form:

$$\gamma_{i,t+1} = \beta_{0i} e_{it}^{\beta_1} \gamma_{it}^{\beta_2},$$

where $\gamma_{i,t+1}$ is the income of the household head *i* in period t+1, $\gamma_{i,t}$ denotes the income of the household head *i* in period t, β_{0i} is the dynasty-specific total factor productivity, $\beta_1 > 0$, $\beta_2 > 0$, $\beta_1 + \beta_2 < 1$.

The equilibrium in this model does not exhibit stratification by income. The low education quality jurisdiction is inhabited by the rich (who opt out of the public system) and the poor. The middle class live in the high-tax district. A switch to centralized funding regime generates less public funding and higher income inequality than the decentralized regime. This result holds both in the steady state and on the transition path.

Although the results are not directly comparable with those in Fernandez and Rogerson (1998) because of the different framework used, they show that the availability of the private school option is an important margin to be taken into account when comparing the education financing regimes. The private school option generates large differences in the income of the decisive voter across the two regimes. In the centralized regime with the possibility of opting out, the private school provides a sorting mechanism. The equalization of public spending per student generates a larger demand for private schools and some of the middle-income households opt out of the public system. These households prefer a zero tax rate. Consequently, the decisive voter is poorer than in the decentralized regime and the equilibrium tax rate is lower. The introduction of vouchers generates a lower aggregate public spending than in the decentralized regime but higher than in the centralized funding system both in the short and long run. This is because under the centralized regime, the fraction of the population choosing the private school is the largest. Thus, the tax rate, and hence the public education budget, is the lowest under the centralized system.

6.2.1 Residential Sorting over the Life Cycle

In the multi-community models presented above, the households keep the residence they chose for their entire life time. However, households' preferences for the public good change over the life cycle: families with school-aged children locate in districts with high-quality public schools and move out when the children graduate.

Epple, Romano, and Sieg (2009) is the first attempt to incorporate the life-cycle component in a multi-community framework. Households live for three periods: childhood, youth, and old age. Households face a mobility cost if they want to relocate to a new community when they enter the old age and their children leave the house. The mobility costs are heterogeneous and entail both monetary and psychic costs. When young households decide their initial location, they also choose their old-age community. When voting, the young households take as given the house

prices and the variables in their old-age community, which simplifies the problem. The authors examine the steady states of the model.

The parameters of the moving cost distribution are chosen such that they match the ratios of old to young cohorts in four stylized communities in the Boston Metropolitan Area. The model implies mobility patterns and political decisions consistent with those in the data. The communities with high levels of spending in public education consist of a majority of young households, while the old residents are concentrated in low-spending communities.

A switch from pure local public funding to a centralized regime yields lower average per-pupil spending and hence lower mean income in steady state, in line with results in Bearse, Glomm, and Ravikumar (2001). However, the private schooling is not available in the Epple, Romano, and Sieg (2009) framework. Extending the model in this direction could provide additional insights regarding the mechanisms behind community formation and collective choice of public education spending.

7. EMPIRICAL ISSUES

In this section, we review the empirical evidence that have a bearing on the quantitative implications of the models in this chapter. The empirical issues can be divided into two areas: the economic realm and the political realm.

In the economic realm, quantitative answers to the following questions are crucial: What are the inputs to the production function for human capital? What role do public expenditures play in the production function? What is the degree of substitutability between public and private inputs? What is the relationship between parental educational background and human capital of the child? What is the extent of peer-group effects in human capital production?

7.1. Educational Expenditure

In the theoretical models reviewed earlier, educational expenditure is one of the inputs to the human capital production function. Higher expenditure, all else equal, is typically assumed to deliver higher quality of education. The empirical literature on education production functions following Hanushek (1986) presents strong evidence that the effects of increasing educational resources (e.g., decreasing class size, extending the school day or school year, hiring a larger fraction of teachers with master's degrees, etc.) on learning outcomes are modest. Perhaps the best evidence on this issue comes from quasi-experimental settings such as Project STAR in Tennessee.

Schanzenbach Whitmore (2006) review the experimental design of Project STAR and concludes that "... reducing class size is a reasonable economic investment: the benefits are sizable and long lasting, especially for black students, and the overall benefits outweigh the costs" (p. 224). A comprehensive review of the evidence on this question

is contained in Hanushek (2006), a chapter in volume 2 of this handbook series. He concludes that "... the evidence – whether from aggregate school outcomes, econometric investigations, or a variety of experimental or quasi-experimental approaches suggests that pure resource policies that do not change incentives are unlikely to be effective. Importantly, the results appear similar across both developed and developing countries." More recent contributions such as Ferreyra (2007) seem to imply much larger effects of education spending.

Some of the arguments for a small effect of public spending on education are based on large and persistent increases in public education expenditures that are not matched by large increases in education outcomes. It is sometimes argued that bureaucratic obstacles stand in the way of translating extra resources into better outcomes. However, in the United States, the student-teacher ratio in the private education sector has fallen from 23 in 1970 to 13.5 in 2005, a decrease which is even larger than the one observed in the public sector. It is a puzzle why the private sector, which surely is under the influence of market forces, would follow a pattern of behavior in the public sector that is argued to be inefficient.

An estimate of the elasticity of substitution between public and private education resources would help point the theory research toward the essential ingredients in a model. Schmidt (1992), Downes and Schoeman (1998), and Goldhaber (1999) find that higher public school funding is correlated with lower private school enrollment. Houtenville and Conway (2008) find that parental time and public school inputs are substitutable and suggest that this can potentially explain the low productivity of public school inputs puzzle. Substitutability between a public school student's time and private materials purchased by parents remains an open question.

7.2. Competition and School Productivity

Much of the policy debate in education focuses on school choice or vouchers, and the typical discussions hinge on the effect of increased competition on productivity in schools and/or on learning outcomes. The report by Witte (1991) on school choice has been extended in the empirical literature in several directions. In a series of papers, Hoxby (1998, 2000, 2001, 2002) studies the competition effects (broadly defined) on schooling inputs and schooling outcomes. Barrow and Rouse (2009) evaluate the voucher experiment in Milwaukee. The literature on the role of vouchers in education or school choice more broadly defined is summarized in Neal (2002), as well as Bettinger (2011) in this volume. The latter surveys the international school competition outcomes.

It is questionable whether any of the empirical findings from this literature based on small-scale experiments, that is, the lower tail of the socioeconomic distribution in a small geographic area such as Milwaukee, are applicable to large-scale voucher reforms that were placed on the ballot in California for statewide implementation. For any evidence on productivity effects of large-scale vouchers, we have to look outside the United States to countries like Czech Republic, Chile, or New Zealand. Filer and Munich (2000) report that in the Czech Republic, new gymnasia open up in metro-politan areas where public preexisting gymnasia were of poor quality and that gymnasia in areas exposed to substantial new competition move up in the quality ranking. (Here quality is defined as the fraction of gymnasia students admitted to a university.) Ladd (2002) and Hsieh and Urquiola (2006), on the other hand, find that reading scores and math scores are negatively correlated with new private school penetration after the voucher reform in Chile and that grade repetition rates and drop out rates are positively correlated with private school penetration. Moreover, these adverse effects seem to get magnified over time. Why would large-scale vouchers have a positive effect in one case and negative effect in others is an open question. One can conjecture that the discrepancy in results has its roots in the specifics of the institutional design.

7.3. Substitution between Education Quality and Other Goods

When Epple and Romano (1996) calibrated their model the evidence available at that time pointed to a low elasticity of substitution between education quality and the numeraire good ($\sigma > 1$ in the utility function (9.1)). In this case, the majority voting equilibrium, if one exists, is computed numerically. In that equilibrium, Director's Law holds, where the ends of the income distribution have common interests against the middle class. This kind of "ends against the middle" case is based on estimates of price elasticities from Rubinfeld and Shapiro (1989) that fall in the range approximately between -0.4 and -0.7. More recently, Cohen-Zada and Justman (2003) exploit the substantial variation within the United States in private school enrollments across communities and estimate a high elasticity of substitution. With a high elasticity, the majority voting equilibrium exists, and the decisive voter is the median-income household, so the equilibrium can be computed easily. Using Colombian data, Gutierrez and Tanaka (2009), however, conclude that the elasticity is low. There is no systematic evidence across countries that points in the direction of one range of elasticities versus another.

7.4. Voting Patterns and Political Support for Public Education

The second realm of empirical issues refers to political aspects. To determine policy endogenously, a political mechanism is used to aggregate households' preferences over policy dimensions, taking the economic environment as given. Among the many political mechanisms, the models reviewed in this chapter have focused on simple majority voting. Few political decisions concerning public education are actually made through majority voting directly, so this approach must maintain that whatever the concrete political mechanisms in place are, the outcomes are determined "as if" through majority voting in a direct democracy. Empirically the question is: Are the actual policies in place the ones that are most preferred by the median household? In simple rational voter models (e.g., Meltzer and Richard (1981)), the decisive voter is the median-income household. The empirical literature following Meltzer and Richard (1983) has typically assumed that the decisive voter is the median-income household and proceeded to examine the policy. Brunner and Ross (2010) study two referenda in California and find that the "jurisdictional median income accurately captures the expected outcomes of majority votes on public services spending and that voters understand the impact of small changes in the identity of the decisive voter." However, in several models reviewed earlier, the decisive voter's income is less than the income of the median household. Moreover, the difference between the median-income household and the identity of the decisive voter is magnified by the fact that political participation (voting) rises with income.

In the United States far fewer than 50% of households have children. If the models in this chapter are interpreted literally, then households without children and households who will send their child to private school would be able to block any spending on public education. Clearly, such an implication is contradicted by the data. There are many reasons why people in other demographic groups may vote for positive education expenditures, such as grandparents, aunts and uncles, those expecting to have children, etc. Perhaps individuals have stronger preferences for education of own children than of grandchildren or of nieces and nephews. Our empirical knowledge of altruism from various demographic groups and voting patterns is inadequate.

There might be other indirect factors that influence one's voting behavior on education issues. For instance, public school quality is often capitalized into housing prices (see Bogart and Cromwell (1997, 2000), Black (1999), Dee (2000), Weimer and Wolkoff (2001), Cheshire and Sheppard (2004), Barrow and Rouse (2004), Figlio and Lucas (2004), Reback (2005), Epple and Ferreyra (2008)). The capitalization effect then connects the age/demographic structure of a community to the popular support for public education. For example, Harris, Evans, Schwab (2001), Brunner and Balsdon (2004), and Bergstrom, Rubinfeld, and Shapiro (1982) document substantial support of public education among the elderly. In a recent paper Hilber and Mayer (2009) find evidence that this house price capitalization is the largest in areas with little undeveloped land and that public education spending is strongly negatively correlated with the fraction of undeveloped land in the jurisdiction. Brunner and Sonstelie (2003) document that homeowners in districts with good public schools were opposed to vouchers in the California referendum.

In some political economy models of public education funding, externalities take center stage. In Bénabou (1993, 1994, 1996a,b), for instance, a high skilled worker who sends his or her child to a private school may support public education expenditures because his or her own productivity can depend on the productivity of low-skilled workers who attended public schools in the same or a neighboring district. Early childhood intervention programs, such as the Perry Preschool program, have the potential for large rates of return for children at the bottom end of the cognitive skill distribution (see Heckman and Masterov (2007)). To the extent that complementarities among skilled and unskilled labor are large, these high potential rates of return might generate substantial political support for such programs. The magnitude of the effect of externalities and complementarities on voting behavior is an open empirical question.

Other potential more indirect determinant of voting behavior on education issues is ideology. Merzyn and Ursprung (2005) find that after controlling for socioeconomic determinants, ideology as reflected by party affiliation can have strong effects on voting behavior.

8. FUTURE DIRECTIONS

Much of the political economy literature on education funding ignores industrial organization features of the education sector. (An exception is Piolatto (2009).) The theoretical and empirical literature in industrial organization analyzes vertical and horizontal differentiation and their interplay, and how heterogeneity of consumer preferences and incomes impacts the strategic interaction between firms in price and quality (horizontal and vertical) competition. Few of the tools and insights from that literature have been systematically applied to the education sector. Below are a few examples where the industrial organization literature might be helpful. (1) There is a clear ranking of quality of institutions in higher education in the United States. When quality is clearly ranked, models of vertical competition might help to provide a better understanding of the education sector. (2) The significant scale effects in education documented by Kenny (1982) and O'Shaughnessy (2007) are typically bypassed by calibrating the models to US- or state-level data, where scale effects are most likely irrelevant. In small cities or rural areas, however, it is reasonable to expect that voters might take into account the impact of private school entry on public school quality. Downes and Greenstein (1996), in the context of post Serano vs. Priest California, find that public school characteristics and local community characteristics matter for the private school entry decisions and that secular and religious private schools respond differently to these characteristics. (3) Schools might compete with each other via horizontal product differentiation. Models with such theoretical underpinnings might explain the positive impact on school quality of the introduction of vouchers in the Czech Republic, as documented by Filer and Munich (2000) and the adverse or neutral effects on learning outcomes documented for the Chilean case by Hsieh and Urquiola (2006). Industrial organization-type models of the education sector, including vertical and horizontal differentiation, may be fruitful for understanding a variety of political economy issues in education

such as why vouchers or other school choice mechanisms are adopted under certain circumstances but not in others.

In models with income heterogeneity, many researchers use the annual household income distribution in the calibrated versions of their models. Funding (private) education in the United States is a long-run phenomenon as is evidenced by the existence and popularity of the 529 savings plans. As soon as a child is born, parents can start saving for future college expenditures or for private school expenditures at the elementary or secondary level. If parental savings are not sufficiently high, students can borrow against future income. In the United States, for instance, the fraction of full-time undergraduate students in 2007–2008, who had taken out a loan, including federal loans, was 53% (Aud, Hussar, Planty, Snyder, and Bianco (2010)). Using dynamic models with a measure of wealth distribution would be an interesting approach to better understand education finance. Whether these models can be calibrated both with wealth and annual income distributions and whether the substantive results on the political support for vouchers or for (de)centralization survive are open questions.

Most of the education finance literature uses majority voting as the political mechanism. In the United States, for example, education funding decisions are outcomes of complicated bargaining and voting procedures in the state legislature or the local school district. Whether other mechanisms such as probabilistic voting or legislative bargaining are better descriptions of the aggregation of preferences is an open question. Consider, for instance, the typical majority voting approach in the models reviewed earlier in this chapter. It assumes that every household participates in elections and that each household has one vote. The evidence, however, indicates that voting participation is increasing in income. Put differently, higher income groups might have more "political power" in the sense that their votes are weighted more than those of the lower-income groups. In probabilistic voting models political power can be introduced by incorporating weights into the social welfare function, as is done in de la Croix and Doepke (2009) or Arcalean and Schiopu (2010). These weights typically are functions of income and not of particular occupation groups. However, those who work in the education sector might have more weight in education politics and they are also more likely to show up at the polls. Sandy (1992) presents early evidence that support for vouchers is decreasing among public school employees. Brunner, Sonstelie, and Thayer (2001) also report that support for the voucher decreased in precincts with larger percentages of the workforce employed in education services. Gokcekus, Phillips, and Tower (2004) find that legislators receiving contributions from the political action committees (PACs) of the teachers' unions were less likely to vote in favor of vouchers. Of course, actual political support for any proposal will depend on the details of the proposal and how they are perceived by the voters. Kenny (2005) presents a summary of how some of the details of voucher proposals influenced the voting outcomes. The large uncertainty surrounding the cost of Michigan's proposal H in 1978, for example, is said to have been responsible for the failure of the proposal. Incorporating alternative mechanisms into the political economy models of education finance might be a fruitful topic for future research.

REFERENCES

- Alesina, A., Baqir, R., Easterly, W., 1999. Public Goods and Ethnic Divisions. Q. J. Econ. 114, 1243–1284. Alesina, A., Glaeser, E., Sacerdote, B., 2001. Why Doesn't the United States Have a European-Style Wel-
- fare State? Brookings Pap. Econ. Act. 32 (2), 187–278.
- Alonso, W., 1964. Location and Land Use. Harvard University Press, Cambridge, MA.
- Arcalean, C., Glomm, G., Schiopu, I., Suedekum, J., 2010. Public Budget Composition, Fiscal (de)centralization, and Welfare. Can. J. Econ. 43 (3), 832–859.
- Arcalean, C., Schiopu, I., 2010. Inequality and Education Funding: Theory and Evidence from the US School Districts. CAEPR Working Paper No. 009-2010, Indiana University.
- Aud, S., Hussar, W., Planty, M., Snyder, T., Bianco, K., et al., 2010. The Condition of Education 2010 (NCES 2010-028). National Center for Education Statistics, Institute of Education Sciences. U.S. Department of Education, Washington, DC.
- Barrow, L., Rouse, C.E., 2004. Using Market Valuation to Assess Public School Spending. J. Public Econ. 88 (9–10), 1747–1769.
- Barrow, L., Rouse, C.E., 2009. School Vouchers and Student Achievement: Recent Evidence, Remaining Questions. Annu. Rev. Econ. 1 (1), 17–42.
- Bayer, P., Ferreira, F., McMillan, R., 2007. A Unified Framework for Measuring Preferences for Schools and Neighborhoods. J. Polit. Economy 115 (4), 588–638.
- Bayindir-Upmann, T., 1998. Two Games of Interjurisdictional Competition when Local Governments Provide Industrial Public Goods. Int. Tax Public Finance 5, 471–487.
- Bearse, P., Cardak, B., Glomm, G., Ravikumar, B., 2009. Why do Education Vouchers Fail? CAEPR Working Paper No. 014-2009, Indiana University.
- Bearse, P., Glomm, G., Janeba, E., 2001. Composition of Government Budget, Non-Single Peakedness and Majority Voting. J. Public Econ. Theory 3 (4), 471–481.
- Bearse, P., Glomm, G., Patterson, D.M., 2005. Endogenous Public Expenditures on Education. J. Public Econ. Theory 7 (4), 561–577.
- Bearse, P., Glomm, G., Ravikumar, B., 2001. Education Finance in a Dynamic Tiebout Economy. Indiana University, Manuscript.
- Beauchemin, K.R., 2001. Growth or Stagnation? The Role of Public Education. J. Dev. Econ. 64 (2), 389-416.
- Bénabou, R., 1993. Workings of a City: Location, Education, and Production. Q. J. Econ. 108 (3), 619-652.
- Bénabou, R., 1994. Human Capital, Inequality, and Growth: A Local Perspective. Eur. Econ. Rev. 38 (3–4), 817–826.
- Bénabou, R., 1996a. Equity and Efficiency in Human Capital Investment: The Local Connection. Rev. Econ. Stud. 63 (2), 237–264.
- Bénabou, R., 1996b. Heterogeneity, Stratification, and Growth: Macroeconomic Implications of Community Structure and School Finance. Am. Econ. Rev. 86 (3), 584–609.
- Bénabou, R., 2002. Tax and Education Policy in a Heterogeneous-Agent Economy: What Levels of Redistribution Maximize Growth and Efficiency? Econometrica 70 (2), 481–517.
- Bergstrom, T., Rubinfeld, D., Shapiro, P., 1982. Micro-Based Estimates of Demand Functions for Local Public Goods. Econometrica 50, 1183–1205.
- Bettinger, E., 2011. International Experience with Choice. In: Hanushek, E., Machin, S., Woessmann, L. (Eds.), Handbook of the Economics of Education. vol. IV. North Holland/Elsevier, Amsterdam.
- Black, D., 1948. On the Rationale of Group Decision Making. J. Polit. Economy 56, 23-34.

- Black, S.E., 1999. Do Better Schools Matter? Parental Valuation of Elementary Education. Q. J. Econ. 114, 577–599.
- Bogart, W.T., Cromwell, B.A., 1997. How Much More Is a Good School District Worth? Natl. Tax J. 50, 215–232.

Bogart, W.T., Cromwell, B.A., 2000. How Much is a Neighborhood School Worth? J. Urban Econ. 47 (2).

- Boustan, L., Ferreira, F., Winkler, H., Zolt, E.M., 2010. Inequality and Local Government: Evidence from U.S. Cities and School Districts, 1970–2000. Law and Economics Workshop, Berkeley Program in Law and Economics, UC Berkeley, 280–305.
- Brunner, E., Balsdon, E., 2004. Intergenerational Conflict and the Political Economy of School Spending. J. Urban Econ. 56 (2), 369–388.
- Brunner, E., Imazeki, J., Ross, S., 2010. Universal Vouchers and Racial and Ethnic Segregation. Rev. Econ. Stat. 92 (4), 912–927.
- Brunner, E., Ross, S., 2010. Is the median voter decisive? Evidence from referenda voting patterns. J. Public Econ. 94(11-12), 898–910.
- Brunner, E., Sonstelie, J., 2003. Homeowners, Property Values, and the Political Economy of the School Voucher. J. Urban Econ. 54 (2), 239–257.
- Brunner, E., Sonstelie, J., Thayer, M., 2001. Capitalization and the Voucher: An Analysis of Precinct Returns from California's Proposition 174. J. Urban Econ. 50 (3), 517–536.
- Buddin, R., Cordes, J., Kirby, S., 1998. School Choice in California: Who Chooses Private Schools? J. Urban Econ. 44, 110–174.
- Calabrese, S., Epple, D., Romano, R., 2007. On the Political Economy of Zoning. J. Public Econ. 91 (1–2), 25–49.
- Calabrese, S., Epple, D., Romer, T., Sieg, H., 2006. Local Public Good Provision: Voting, Peer Effects, and Mobility. J. Public Econ. 90 (6–7), 959–981.
- Campbell, D.E., West, M.R., Peterson, P.E., 2005. Participation in a National, Means-Tested School Voucher Program. J. Policy Anal. Manage. 24 (2), 523–541.
- Cardak, B.A., 1999. Heterogeneous Preferences, Education Expenditures and Income Distribution. Econ. Rec 75 (228), 63–76.
- Cardak, B.A., 2004a. Education Choice, Endogenous Growth and Income Distribution. Economica 71, 57-81.
- Cardak, B.A., 2004b. Education Choice, Neoclassical Growth, and Class Structure. Oxf. Econ. Pap. 56 (4), 643–666.
- Cardak, B.A., 2005. Education Vouchers, Growth, and Income Inequality. Macroecon. Dyn. 9 (01), 98–121.
- Cardak, B.A., Givon, D., 2004. Why the Poor Don't Go to University: Attainment Constraints and Two-Staged Education. La Trobe University, Manuscript.
- Chen, Z., West, E.G., 2000. Selective versus Universal Vouchers: Modeling Median Voter Preferences in Education. Am. Econ. Rev. 90, 1520–1534.
- Cheshire, P., Sheppard, S., 2004. Capitalising the Value of Free Schools: The Impact of Supply Characteristics and Uncertainty. Econ. J. 114 (499), 397–424.
- Cohen-Zada, D., Justman, M., 2003. The Political Economy of School Choice: Linking Theory and Evidence. J. Urban Econ. 54 (2), 277–308.
- Corcoran, S., Evans, W., 2009. Income Inequality, The Median Voter, and the Support for Public Education. Technical Report, New York University.
- Cutler, D.M., Elmendorf, D.W., Zeckhauser, R.J., 1993. Demographic Characteristics and the Public Bundle. Public Finance 48, 178–198.
- Cutler, D.M., Glaeser, E.L., Vigdor, J.L., 1999. The Rise and Decline of the American Ghetto. J. Polit. Economy 107, 455–506.
- Davidoff, T., 2005. Income Sorting: Measurement and Decomposition. J. Urban Econ. 58 (2), 289-303.
- de Bartolome, C., 1990. Equilibrium and Inefficiency in a Community Model with Peer Group Effects. J. Polit. Economy 98 (1), 110–133.
- de Bartolome, C., Ross, S.L., 2003. Equilibria with Local Governments and Commuting: Income Sorting vs Income Mixing. J. Urban Econ. 54 (1), 1–20.

- de la Croix, D., Doepke, M., 2009. To Segregate or to Integrate: Education Politics and Democracy. Rev. Econ. Stud. 76 (2), 597–628.
- Dee, T.S., 2000. The Capitalization of Education Finance Reforms. J. Law Econ. 43 (1), 185-214.
- Dhillon, A., Wooders, M., Zissimos, B., 2007. Tax Competition Reconsidered. J. Public Econ. Theory 9 (3), 391-423.
- Downes, T.A., Greenstein, S.M., 1996. Understanding the Supply Decisions of Nonprofits: Modelling the Location of Private Schools. Rand J. Econ. 27 (2), 365–390.
- Downes, T.A., Schoeman, D., 1998. School Finance Reform and Private School Enrollment: Evidence from California. J. Urban Econ. 43 (3), 418–443.
- Dunz, K., 1989. Some Comments on Majority Rule Equilibria in Local Public Good Economies. J. Econ. Theory 47, 228–234.
- Durlauf, S.N., 1996. A Theory of Persistent Income Inequality. J. Econ. Growth 1 (1), 75-93.
- Epple, D., Ferreyra, M.M., 2008. School Finance Reform: Assessing General Equilibrium Effects. J. Public Econ. 92 (5–6), 1326–1351.
- Epple, D., Filimon, R., Romer, T., 1984. Equilibrium among Local Jurisdictions: Toward an Integrated Treatment of Voting and Residential Choice. J. Public Econ. 24 (3), 281–308.
- Epple, D., Filimon, R., Romer, T., 1993. Existence of Voting and Housing Equilibrium in a System of Communities with Property Taxes. Reg. Sci. Urban Econ. 23 (5), 585–610.
- Epple, D., Gordon, B., Sieg, H., 2010. DRs. Muth and Mills Meet, Dr. Tiebout: Integrating Location-Specific Amenities into Multi-Community Equilibrium Models. J. Reg. Sci. 50 (1), 381–400.
- Epple, D., Platt, G.J., 1998. Equilibrium and Local Redistribution in an Urban Economy When Households Differ in both Preferences and Incomes. J. Urban Econ. 43 (1), 23–51.
- Epple, D., Romano, R., 1996. Ends against the Middle: Determining Public Service Provision When There Are Private Alternatives. J. Public Econ. 62, 297–325.
- Epple, D., Romano, R., 1998. Competition between Private and Public Schools, Vouchers, and Peer-Group Effects. Am. Econ. Rev. 88 (1), 33–62.
- Epple, D., Romano, R., 2003. Neighborhood Schools, Choice and the Distribution of Educational Benefits. In: Hoxby, C.M. (Ed.), The Economics of School Choice. University of Chicago Press, Chicago, IL, 227–286.
- Epple, D., Romano, R., 2008. Educational Vouchers and Cream Skimming. Int. Econ. Rev. 49 (4), 1395–1435.
- Epple, D., Romano, R., Sieg, H., 2009. Life Cycle Dynamics within Metropolitan Communities. Tepper School of Business. Paper 444. http://repository.cmu.edu/tepper/444.
- Epple, D., Romer, T., 1991. Mobility and Redistribution. J. Polit. Economy 99 (4), 828-858.
- Epple, D., Sieg, H., 1999. Estimating Equilibrium Models of Local Jurisdictions. J. Polit. Economy 107 (4), 828–858.
- Fairlie, R.W., Resch, A.M., 2002. Is There "White Flight" into Private Schools? Evidence from the National Educational Longitudinal Survey. Rev. Econ. Stat. 84 (1), 21–33.
- Fernandez, R., Rogerson, R., 1996. Income Distribution, Communities, and the Quality of Public Education. Q. J. Econ. 111 (1), 135–164.
- Fernandez, R., Rogerson, R., 1997. Keeping People Out: Income Distribution, Zoning, and the Quality of Public Education. Int. Econ. Rev. 38 (1), 23–42.
- Fernandez, R., Rogerson, R., 1998. Public Education and Income Distribution: A Dynamic Quantitative Evaluation of Education-Finance Reform. Am. Econ. Rev. 88 (4), 813–833.
- Fernandez, R., Rogerson, R., 1999. Education Finance Reform and Investment in Human Capital: Lessons from California. J. Public Econ. 74 (3), 327–350.
- Fernandez, R., Rogerson, R., 2003. Equity and Resources: An Analysis of Education Finance Systems. J. Polit. Economy 111 (4), 858–897.
- Ferreyra, M., 2007. Estimating the Effects of Private School Vouchers in Multi-District Economies. Am. Econ. Rev. 18 (3), 789–817.
- Figlio, D.N., Lucas, M.E., 2004. Do High Grading Standards Affect Student Performance? J. Public Econ. 88 (9–10), 1815–1834.
- Filer, R.K., Munich, D., 2000. Responses of Private and Public Schools to Voucher Funding: The Czech and Hungarian Experience. CERGE-EI Working Papers wp160, The Center for Economic Research and Graduate Education-Economic Institute, Prague.
- Friedman, M., 1962. The Role of Government in Education. In: Capitalism and Freedom. Chicago University Press, Chicago, 85–108.
- Glomm, G., Ravikumar, B., 1992. Public versus Private Investment in Human Capital: Endogenous Growth and Income Inequality. J. Polit. Economy 100 (4), 818–834.
- Glomm, G., Ravikumar, B., 1998. Opting Out of Publicly Provided Services: A Majority Voting Result. Social Choice and Welfare. Soc. Choice Welfare 15 (2), 187–199.
- Gokcekus, O., Phillips, J.A., Tower, E., 2004. School Choice: Money, Race, and Congressional Voting on Vouchers. Public Choice 119 (1–2), 241–254.
- Goldhaber, D., 1999. An Endogenous Model of Public School Expenditures and Private School Enrollment. J. Urban Econ. 46 (1), 106–128.
- Goldin, C., Katz, L.F., 1997. Why the United States Led in Education: Lessons from Secondary School Expansion, 1910 to 1940. NBER Working Papers 6144, National Bureau of Economic Research, Inc., Cambridge, MA.
- Gradstein, M., 2003. The Political Economy of Public Spending on Education, Inequality, and Growth. Policy Research Working Paper 3162, The World Bank.
- Gradstein, M., Justman, M., 1997. Democratic Choice of an Education System: Implications for Growth and Income Distribution. J. Econ. Growth 2 (2), 169–183.
- Gradstein, M., Justman, M., Meier, V., 2005. The Political Economy of Education: Implications for Growth and Inequality, CESifo Book Series. The MIT Press, Cambridge, MA.
- Gradstein, M., Kaganovich, M., 2004. Aging Population and Education Finance. J. Public Econ. 88 (12), 2469–2485.
- Gutierrez, C., Tanaka, R., 2009. Inequality and Education Decisions in Developing Countries. J. Econ. Inequality 7 (1), 55-81.
- Hanushek, E., 2006. School Resources. In: Hanushek, E., Welch, F. (Eds.), Handbook of Economics of Education. Vol. 2, North Holland/Elsevier, Amsterdam, 865–909.
- Hanushek, E., Sarpça, S., Yilmaz, K., 2007. School Choice: Traditional Mechanisms and Extending the Poor's Ability to Choose. TÜSIAD-Koç University, Working Paper 0709.
- Hanushek, E., Yilmaz, K., 2007. The Complementarity of Tiebout and Alonso. J. Hous. Econ. 16 (2), 243–261.
- Hanushek, E., Yilmaz, K., 2010. Household Location and Schools in Metropolitan Areas with Heterogeneous Suburbs; Tiebout, Alonso, and Government Policy. NBER Working Paper 15915.
- Hanushek, E.A., 1986. The Economics of Schooling: Production and Efficiency in Public Schools. J. Econ. Lit. 24 (3), 1141–1177.
- Harris, A.R., Evans, W.N., Schwab, R.M., 2001. Education Spending in an Aging America. J. Public Econ. 81 (3), 449–472.
- Haupt, A., 2005. The Evolution of Public Spending on Higher Education in a Democracy. CESifo Working Paper Series (1631).
- Heckman, J., Masterov, D., 2007. Education Spending in an Aging America. Rev. Agric. Econ. 29 (3), 446-493.
- Hilber, C.A., Mayer, C., 2009. Why Do Households without Children Support Local Public Schools? Linking House Price Capitalization to School Spending. J. Urban Econ. 65 (1), 74–90.
- Houtenville, A.J., Conway, K.S., 2008. Parental Effort, School Resources, and Student Achievement. J. Hum. Resour. 43 (2), 437–453.
- Hoxby, C.M., 1998. What do America's "Traditional" Forms of School Choice Teach Us about School Choice Reforms? Econ. Policy Rev. 4 (1), 47–59.
- Hoxby, C.M., 2000. Does Competition among Public Schools Benefit Students and Taxpayers? Am. Econ. Rev. 90 (5), 1209–1238.
- Hoxby, C.M., 2001. Rising Tide: New Evidence on Competition and the Public Schools. Educ. Next 1 (4), 68–74.
- Hoxby, C.M., 2002. How School Choice Affects the Achievement of Public School Students. In: Hill, P. (Ed.), Choice with Equity. Hoover Press, Stanford, 141–177.
- Hoyt, W., Lee, K., 1998. Educational Vouchers, Welfare Effects and Voting. J. Public Econ. 69, 211-228.
- Hsieh, C.-T., Urquiola, M., 2006. The Effects of Generalized School Choice on Achievement and Stratification: Evidence from Chile's Voucher Program. J. Public Econ. 90 (8–9), 1477–1503.

- James, E., 1993. Why Do Different Countries Choose a Different Public-Private Mix of Educational Services? J. Hum. Resour. 28, 242–572.
- Justman, M., Thisse, J.-F., 2000. Local Public Funding of Higher Education When Skilled Labor is Imperfectly Mobile. Int. Tax Public Finance 7 (3), 247–258.
- Kenny, L., 1982. Economies of Scale in Schooling. Econ. Educ. Rev. 2 (1), 1-24.
- Kenny, L., 2005. The Public Choice of Educational Choice. Public Choice 124 (1), 205–222.
- Kremer, M., 1997. How Much Does Sorting Increase Inequality? Q. J. Econ. 112, 115–139.
- Ladd, H.F., 2002. School Vouchers: A Critical View. J. Econ. Perspect. 16 (4), 3-24.
- Lankford, H., Wyckoff, J., 2001. Who Would Be Left Behind by Enhanced Private School Choice? J. Urban Econ. 50 (2), 288–312.
- Lee, W., Roemer, J.E., 1998. Income Distribution, Redistributive Politics, and Economic Growth. J. Econ. Growth 3 (3), 217–240.
- Levy, G., 2005. The Politics of Public Provision of Education. Q. J. Econ. 120 (4), 1507-1534.
- Luelfesmann, C., Myers, G., 2010. Two-tier Public Provision: Comparing Public Systems. J. Public Econ. Article in Press.
- Luttmer, E.F.P., 2001. Group Loyalty and the Taste for Redistribution. J. Polit. Economy 109 (3), 500-528.
- Meltzer, A.H., Richard, S.F., 1981. A Rational Theory of the Size of Government. J. Polit. Economy 89 (5), 914–927.
- Meltzer, A.H., Richard, S.F., 1983. Tests of a Rational Theory of the Size of Government. Public Choice 41 (5), 403–418.
- Merzyn, W., Ursprung, H.W., 2005. Voter Support for Privatizing Education: Evidence on Self-Interest and Ideology. Eur. J. Polit. Economy 21 (1), 33–58.
- Mills, E.S., 1967. An Aggregate Model of Resource Allocation in a Metropolitan Area. Am. Econ. Rev. 57, 197–210.
- Mora, F.M., 2006. The Existence of Non-Elite Private Schools, J. Public Econ. 90 (8-9), 1505-1518.
- Muth, R.F., 1969. Cities and Housing: The Spatial Pattern of Urban Residential Land Use. University of Chicago Press, Chicago, IL.
- Neal, D., 2002. How Vouchers Could Change the Market for Education. J. Econ. Perspect. 16 (4), 25-44.
- Nechyba, T.J., 1997. Existence of Equilibrium and Stratification in Local and Hierarchical Tiebout Economies with Property Taxes and Voting. Econ. Theory 10, 277–304.
- Nechyba, T.J., 1999. School Finance Induced Migration and Stratification Patterns: The Impact of Private School Vouchers. J. Public Econ. Theory 1 (1), 5–50.
- Nechyba, T.J., 2000. Mobility, Targeting, and Private-School Vouchers. Am. Econ. Rev. 90 (1), 130-146.
- Nechyba, T.J., 2003. Centralization, Fiscal Federalism, and Private School Attendance. Int. Econ. Rev. 44 (1), 179–204.
- Nechyba, T.J., 2006. Income and Peer Quality Sorting in Public and Private Schools. In: Hanushek, E., Welch, F. (Eds.), Handbook of Economics of Education. Vol. 2. North Holland/Elsevier, Amsterdam.
- OECD, 2010. Online Education Database, Organisation for Economic Co-operation and Development.
- Ordeshook, P.C., 1986. Game Theory and Political Theory: An Introduction. Cambridge University Press, New York.
- O'Shaughnessy, T., 2007. Parental Choice and School Quality When Peer and Scale Effects Matter. Econ. Educ. Rev. 26 (4), 501–515.
- Perotti, R., 1993. Political Equilibrium, Income Distribution, and Growth. Rev. Econ. Stud. 60 (4), 755–776.
- Perotti, R., 1996. Growth, Income Distribution, and Democracy: What the Data Say. J. Econ. Growth 1 (2), 149–187.
- Piolatto, A., 2009. Education and Selective Vouchers. MPRA Paper 8934, University Library of Munich, Germany.
- Preston, R., 2003. Public Education or Vouchers? The Importance of Heterogeneous Preferences. Econ. Rec. 79, 74–84.
- Rangazas, P., 1995a. Vouchers and Voting: An Initial Estimate Based on the Median Voter Model. Public Choice 82 (3–4), 261–279.

- Rangazas, P., 1995b. Vouchers in a Community Choice Model with Zoning. Q. Rev. Econ. Finance 35, 15–40.
- Reback, R., 2005. House Prices and the Provision of Local Public Services: Capitalization under School Choice Programs. J. Urban Econ. 57 (2), 275–301.
- Rhode, P.W., Strumpf, K.S., 2003. Assessing the Importance of Tiebout Sorting: Local Heterogeneity from 1850 to 1990. Am. Econ. Rev. 93, 1648–1677.
- Roemer, J.E., 2005. Will Democracy Engender Equality? Econ. Theory 25, 217-234.
- Romero, G., 2009. Does the Possibility of Opting Out of Public Education Favor Expenditure on Basic Education? Manuscript, University of Santiago, Chile.
- Rubinfeld, D.L., Shapiro, P., 1989. Micro-estimation of the Demand for Schooling : Evidence from Michigan and Massachusetts. Reg. Sci. Urban Econ. 19 (3), 381–398.
- Saint-Paul, G., Verdier, T., 1993. Education, Democracy and Growth. J. Dev. Econ. 42 (2), 399-407.
- Sandy, J., 1992. Evaluating the Public Support for Educational Vouchers: A Case Study. Econ. Educ. Rev. 11 (3), 249–256.
- Schanzenbach Whitmore, D., 2006, 2007. What Have Researchers Learned from Project STAR? Brookings Pap. Educ. Policy 205–228.
- Schmidt, A.B., 1992. Private School Enrollments in Metropolitan Areas. Public Finance Rev. 20 (3), 298–320.
- Seshadri, A., Yuki, K., 2004. Equity and Efficiency Effects of Redistributive Policies. J. Monet. Econ. 51 (7), 1415–1447.
- Snyder, T.D., Dillow, S.A., 2010. Digest of Education Statistics 2009 (NCES 2010-013). National Center for Education Statistics, Institute of Education Sciences. U.S. Department of Education, Washington, DC.
- Soares, J., 2003. Self-interest and Public Funding of Education. J. Public Econ. 87 (3-4), 703-727.
- Soares, J., 2006. A Dynamic General Equilibrium Analysis of the Political Economy of Public Education. J. Popul. Econ. 19 (2), 367–389.
- Stiglitz, J.E., 1974. The Demand for Education in Public and Private School Systems. J. Public Econ. 3, 349–385.
- Su, X., 2004. The Allocation of Public Funds in a Hierarchical Educational System. J. Econ. Dyn. Control 28 (12), 2485–2510.
- Su, X., 2006. Endogenous Determination of Public Budget Allocation Across Education Stages. J. Dev. Econ. 81 (2), 438–456.
- Sylwester, K., 2000. Income Inequality, Education Expenditures, and Growth. J. Dev. Econ. 63 (2), 379-398.
- Tiebout, C.M., 1956. A Pure Theory of Local Expenditures. J. Polit. Economy 64, 416-424.
- UNESCO, 2010. UNESCO Institute for Statistics Data Centre. Montreal, http://stats.uis.unesco.org.
- Van Alstine, J., 2009. Diversity and Education Choices. Ph. D. thesis, Indiana University.
- Weimer, D., Wolkoff, M., 2001. School Performance and Housing Values: Using Non-Contiguous District and Incorporation Boundaries to Identify School Effects. Natl. Tax J. 54, 231–253.
- Witte, J.F., 1991. Milwaukee Parental Choice Program. First Year Report. Technical Report, University of Wisconsin-Madison.
- Zhang, L., 2008. Political Economy of Income Distribution Dynamics. J. Dev. Econ. 87 (1), 119-139.

INDEX

Note: Page numbers followed by f, t and n indicate figures, tables and notes, respectively.

Α

Abecedarian program, 210 Achievement test scores, IQ and, 54-63 AFQT, see Armed Forces Qualifying Test Aggregate incidence, 296-298 Agrarian economy, 468 Agricultural Land Reform Amendment Act, 484 Allocative efficiency of health inputs, 224, 229-230 Ambiguity aversion, 78 American College Test (ACT), 40 Andhra Pradesh Randomized Evaluation Study (APRES), 520 Anterior cingulate cortex (ACC), 46 Arbitrariness, 89n Armed Forces Qualifying Test (AFQT), 40, 61, 62f, 63n, 101, 132 Armed Services Vocational Battery (ASVAB), 132 Assessment-based accountability systems, 504 Axis II personality disorders, 54

В

Balfour Education Act, 477 Behavioral economics, 14, 75 Bellman equation, 501 Benchmark model, 618-624, 620-622f Bequest dynamics, 455-457 factor prices, 449 fundamental assumptions, 451-452 Galor-Zeira model, 447 individuals, 449-451 inequality and economic development, 460-462 persistence of, 458-459 parental transfers and, 453-455 robustness dynasties, 465 human capital and alternative utility functions, concave production of, 466 labor-augmenting technological progress, 463-465 random shocks, 465

skilled labor-intensive sector, 448, 463 skilled worker, income of, 452-453 skills composition income distribution and, 457-458 and income per capita, 459-460 unskilled labor-intensive sector, 448, 463 unskilled worker, income of, 452 Bequest dynamics, 455-457 threshold level of, 454f, 457, 460 BES, see British Election Surveys Big Five personality traits, 11, 11t, 44-45t, 48 educational attainment and, 92t emotional stability, 52-53 intelligence and, 93f, 99f juvenile delinquency and, 116f Birth cohort, 577, 578f Brain lesion studies, 126–127 British Election Surveys (BES), 264 Budget constraint, 451 preferences and, 450

С

Calibrated models, 656-661 California Achievement Test (CAT), 134 Campbell analyses, 505 Campbell's Law, 522 Canonical model, 581 Career mobility, 313-314 Cattell's investment theory, 61 Centralized vs. local fund, with no private option, 656-657 Chemical and laboratory interventions, 127 Chicago Child Parent Center (CPC), 207, 209 Chicago public schools, 518 Child labor and opting out of education, 624 Childhood interventions, 207, 208t Children of the National Longitudinal Survey of Youth (CNLSY), 211 Chilean system, 559 Chilean voucher program, 558-559, 564

Classical hypothesis, 443-444 Classical measurement error, 370 CNLSY, see Children of the National Longitudinal Survey of Youth Cognition, 37-43 predictive validity of tests, 40-43 Cognitive ability, see Intelligence Cognitive control, 43 Cognitive skills distribution of, 95f vs. noncognitive skills, 205n Cognitive traits, personality and, 54-63 College completion rates, 576f, 576 College degree attainment baseline evidence on, 575-580 determinants evidence of college matching, 601-603 demand-side determinants, 590-596 supply-side determinants, 596-601 economic analysis of, 580 schooling and application, 580-584 specification of field of study, 586-587 modeling, 587-590 supply-side determinants of, 588-590, 596-601 unanswered questions in college choice, 604-605 college degree completion and public policy, economics, 608-609 international comparisons, 608 market structure, 606-607 political economy and public subsidies, 607-608 College degree completion, and public policy, 608-609 College production function, 605 Collegiate attainment rates, 574 Colombia's voucher program, 557-558 Commuting costs, 651-653 Tiebout models with, 660-661, 661f Compulsory schooling laws, 206 Conduct Problems Prevention Research Group (CPPRG), 213 Conscientiousness, 89-91, 99, 103, 106, 112 benefits of, 104 predictive power of, 105t Construct validity, 18

Consumer Price Index for All Urban Consumers (CPI-U), 216, 217t Convergent validity, 31-33 Core self-evaluations, 53 CPC, see Chicago Child Parent Center CPPRG, see Conduct Problems Prevention Research Group CPS, see Current Population Survey Credit market imperfections channel, 445-446 Crime, 116f, 116-117 contemporaneous, 192 contemporaneous schooling and, 205-207 economics of education and, 186-193 educational attainment and, 193-202 human capital-based model of, 220-221 school attendance, 192 school quality and, 202-205 Criminal ability, 186, 188-189 Crystallized intelligence, 39-40, 40n predictive validities of various tests of, 41-42t Culture-free IQ test, 48 Current Population Survey (CPS), 264, 365, 595

D

Delta College Cost Project, 599 Demand-side determinants of college persistence, 587-588 of degree attainment, 590-596 Determinants ability, 299-300 age, 298 ethnicity, 298-299 gender, 298 Diagnostic and Statistical Manual (DSM), 54 Dickens-Flynn model, 124 Differential Aptitudes Test (DAT), 61, 63f, 63n Diploma effects on education, 583 Discriminant validity, 31-33, 35 District power equalization (DPE), 656, 659 Dominant theories of personality, 48 Downgrading and ethnic networks, 374-375 DPE, see District power equalization Dynamic complementarity, 129 Dynamic Tiebout sorting, 666-669 Dynamic time allocation model, crime, 186 Dynamics model of education, 662-669

Ε

Earnings findings, 309-312 immigrants assimilation, 414-424t different generations of, 408 growth of, 362 impact on, 302-312 measurement error, 306-309 omitted variable bias, 304-306 over/underschooling, 321-326t EB, see Eurobarometer Economic development, inequality and, 444, 460-462 Economic growth, 471–473 Economic model definition of personality within, 23-27 integrating traits into, 75-84 personality traits within, 14-30 Economic preference parameters, stability of, 139–150 Economics of college degree completion and public policy, 608-609 Education, 329 acquisition of, 339-343 child labor and opting out of, 624 and democracy, macro evidence on, 262-264 effect on health, 271-274 evidence on causal effects of, 128-139, 130f expenditure, 669-670 fund in multicommunity models, 644-661, 645f and health, future research, suggestions, 260-261 in home country, 354 of immigrant children, 404-407 industrial demand for, 475-480 international, 339n and mortality, simple model of, 225-230 over-investment in, 287-290 preference distribution over types of, 639-642 private vs. public, 662-666 quality and goods, substitution between, 671 reforms, 475, 485-487 land reforms and, 482-485 political reforms and, 485-487 stages of, 629-631, 630f tertiary, 339 Education Act, 477, 486 Education effects on citizenship, 261-269 estimation, 264-269, 265t, 266-267t

on crime, 184-223 on democracy, 261-269 health and mortality, 224-261, 238-239t evidence on, 230-258 on political participation, 261-269 Education funding systems, 654-655 Educational attainment, 90-99, 93f, 95f, 96t, 97-98f achievement and, 90-106 Big Five traits and, 92t course grades, 99-101, 99-100f, 103-106, 105t of emigrant population, 337t foreign-born (FB) population, 331, 332t, 335t, 336 of migrants, 331-339 native-born population (NB), 331 standardized achievement test scores, 101-103, 102f diverge, 103-106, 105t Educational Code, 483 Educational Maintenance Allowances (EMA), 201 Effective discount rate, 226 Elmira Nurse Home Visitation Program, 209n EMA, see Educational Maintenance Allowances Emigration and human capital investment, 392-395 Empirical issues, 669-673 competition and school productivity, 670-671 educational expenditure, 669-670 Empirical validity, 583-584 Enduring traits, 26 Ethnic capital, 398, 402 Ethnic networks, downgrading and, 374-375 Eurobarometer (EB), 264 Executive function, 43, 46

F

Fast Track program, 213
Finance, 557
First-generation immigrants, 348f, 348
Five Factor Theory, 29
Five-factor model, 48, 50–54
alternative systems, 50–54
Fluid intelligence, 39–40, 40n
predictive validities of tests of, 41–42t
Flynn effect, 39
Foreign graduates, 341
Foreign-born (FB) population
educational attainment of, 331, 332t, 335t, 336
share of, 345t
Freeman's analysis, 288

G

Galor–Zeira model, 445, 447, 465–466
GED, see General Educational Development
General Aptitude Test Battery (GATB), 40
General Educational Development (GED), 212n, 242, 292
General intelligence, hierarchical scheme of, 38f
General mental ability, see Intelligence
General Social Surveys (GSS), 264
German Socio–Economic Panel (GSOEP), 80, 369
Graduates

foreign, 341
retention rates of, 341

Gross Domestic Product, 519
GSS, see General Social Surveys

Н

Health and mortality education effect, 247-258 evidence on effects of education, 230-258 Health inputs, 225, 227-288 Health policy, implications for education and, 258-260 Heritability of preferences and personality, 151-155t Heterogeneity, 638-644 ability distribution, 638-639 norm, 503 preference distribution over types of education, 639-642 talent, 503 young and old, 642-644 High School and Beyond (HS&B), 264, 268 Higher Education and Graduate Employment in Europe, 299 High-risk youth, 203 High-School Graduates, 285f Home country, skill investment in, 354 Housing markets, 648-651 HS&B, see High School and Beyond Human capital, 312-313 based model of crime, 220-221 channel, 472 externalities, local, 647-648 investment, 349-355 choosing optimal migration period, 358-360 emigration and, 392-395

immigration and, 395-397 optimal in, 356 Human capital accumulation, 467 nonfinancial hurdles for agrarian economy, 468 gender inequality, 471 social-political transitions, 470-471 Human capital formation, 471-473 industrialization and, 467-468, 473-487 education, industrial demand for, 475-480 Industrial Revolution, 473-474 land concentration, 481-482 land reforms and education reforms, 482-485 political and education reforms, 485-487 technological progress, 474 Unified Growth Theory, 473-474

I

IHDP, see Infant Health and Development Program Immigrant children intergenerational transmission and ethnic networks, 411-412 of human capital, 407-410 and language, 410-411 PISA outcomes of, 406t schooling outcomes of, 404-407 Immigrants, 329n assimilation and adaptation, 362-365 duration shares, 345t earnings assimilation, 414-424t different generations of, 408 growth of, 362 educational attainment of, 331, 332t first-generation, 348, 348f and intergenerational mobility, 399-401 literature on, 425-428t model of, 401-404 native students and, 347t negative selection scenarios of, 381 next generation, 346-349 positive selection scenarios of, 381 well-educated, 356 Immigration, and human capital investment, 395-397 Immigration Act of 1924, 363 Immigration and Nationality Act of 1965, 363

Incapacitation effect, 206 Incidence aggregate incidence, 296-298 determinants ability, 299-300 age, 298 ethnicity, 298-299 gender, 298 dynamics, 300-302 overschooling and underschooling, 297t Income distribution, 454, 459 and skills composition, 457-458 Industrial Revolution, 473-474 Industrialization, 467-468, 473-487 and human capital formation, 471-473 education, industrial demand for, 475-480 Industrial Revolution, 473-474 land concentration, 481-482 land reforms and education reforms, 482-485 political and education reforms, 485-487 technological progress, 474 Unified Growth Theory, 473-474 Inequality, 80, 471-473 and economic development, 444, 460-462 persistence of, 458-459 Infant Health and Development Program (IHDP), 207, 209, 211 Instrumental variable (IV), 224 Integrating traits as constraints, 76-77 into economic model, 75-84 as preferences, 77-84 multidimensionality, 77-81 specifications and consequences, 81-83 Intelligence, 37n fluid vs. crystallized, 39-40 hierarchical scheme of general, 38f Intelligence tests, 47 incentives and performance on, 57-60t Intensive production function, 448 Intergenerational transmission and ethnic networks, 411-412 of human capital, 407-410 and language, 410-411 Interjurisdictional spillovers, 653-654 International Adult Literacy Survey, 300 International child support (ICS), 519 International education, 339n

Interventions chemical and laboratory, 127 evidence from, 134–139, 136f evidence on causal effects of, 128–139, 130f on personality, effect of, 140–149t IQ tests, 47–48 and achievement test scores, 54–63 culture-free, 48 measures of, 7 IV, *see* Instrumental variable

J

Job analysis, 292 Job competition, 314 Job requirements job analysis, 292 realized matches, 292–294 self-assessment, 291–292 Jovanovic model, 501

Κ

KIRIS system, 517

L

Labor market theories career mobility, 313-314 human capital, 312-313 job competition, 314 preferences, 315 search and frictions, 315-316 signaling and screening, 314 Labor-augmenting technological progress, 463-465 Land reforms and education reforms, 482-485 Language, 370-374 Latent variables, causal status of, 36 Learning ability, 186 Learning in first period, choosing optimal investment in, 360-362 Least Squares estimator, 399 Life cycle broad evidence on changes in traits over, 118-122 dynamics of traits, 27-28 framework, for organizing studies and integrating evidence, 130f residential sorting over, 668-669 Linear factor models, 30-31 Local average treatment effect, 231n Local human capital externalities, 647-648

Local vs. centralized fund, with no private option, 656–657 Locus of control, 52–53 Log-linear utility function, 450

Μ

Malleability of preferences and personality, 151–155t MAP system, 529 Markets, 543-545 Measurement issues job requirements job analysis, 292 realized matches, 292-294 self-assessment, 291-292 overskilled, 294-296 Measurement systems, implementing, 36-64 cognition, 37-43 operationalizing the concepts, 47–48 personality traits, 43-46 Metamorphosis, 443 Meta-regressions, 312t Mexican-US migration, 386 Migrants, 349-376 educational attainment of, 331-339 empirical studies, 362-376 implications of model, 355-362 return, 387-388 learning and multidimensional skills, 388-391 selection of, 387-388 multiple skill model of, 378-382 selection patterns, 384-387 skill prices, productive capacity and selection, 382 - 384skill enhancement and return decisions of, 353-354 Migration, 330 and acquisition of education, 339-343 decisions, 349-351, 355 optimal investment in skills, 355-356 permanent, 356-357 return, 343-346, 351 temporary, 357-358 time dimension of, 344 Mincerian wage equation, 302 Montreal Longitudinal Experimental Study, 214 Moral hazard assessment-based accountability systems, 504 common design flaws, 527-538

generalizability, 515-521 hidden actions, 521-523 human capital production technology, 506 incentives, 508-515 KIRIS system, 517 optimal bonus rate, 506 targets and prizes, 523-527 Mortality effect, 228-229 Mortality-weighted average effect, 250 Motivation, 51 Moving-to-Opportunity (MTO), 204 MTO, see Moving-to-Opportunity Multicommunity models education fund in, 644-661, 645f features of. 645 Multidimensional education system, 386 Multidimensional model, 388

Ν

NAEP, see National Assessment of Educational Progress National Assessment of Educational Progress (NAEP), 397 National Childhood Development Study (NCDS), 246, 250 National Educational Longitudinal Study (NELS), 591 National Election Surveys (NES), 264 National Longitudinal Survey of Youth 1979 (NLSY79), 132 National Longitudinal Survey of Youth (NLSY), 194, 250, 593 Native students, immigrant and, 347t Native-born population (NB), 331 NCDS, see National Childhood Development Study Negative reciprocity, 80 NELS, see National Educational Longitudinal Study Neoclassical paradigm, 443, 487 NES, see National Election Surveys Neuroticism, 90, 94, 99, 112 New entrants, 284, 286 NLSY, see National Longitudinal Survey of Youth No Child Left Behind Act (NCLB), 518 Noncognitive skills, distribution of, 95f Nonhierarchical sorting, 378 Nonlinear migration costs, 385f, 385 Norm heterogeneity, 503

0

Oblique factor analysis, 33 OLS, see Ordinary least squares Openness, 90-91, 99 Operationalization and construct validation, 19 in psychology, 18 Optimal investment choosing, 360-362 in skills, 355-356 Optimal migration duration, 355-356 choosing, 358-360 Optimal migration plan, 377 Optimization, 451 Ordinary least squares (OLS), 198 Overeducation literature, 288 Overschooling, 297t, 311t Overskilling, 294-296

Ρ

PACES, see Programa de Ampliación de Cobertura de la Educación Secundaria Panel Survey of Income and Dynamics (PSID), 211 Parental education and child health, 245-247 Parental investment evidence of change in traits from, 131-132 evidence on causal effects of, 128-139, 130f Partial transferability, 389-390 Peer effects, 657-660 Peer-group effects, 647 PER, see Power equalization scheme with recapture Performance contests, common design flaws of aggregation, 536-538 high standards, 527-529 incorrect handicapping, 529 relative performance systems, 534-536 steps forward, 538-543 assessment-based incentives, 541 heterogeneity, 541-542 screening, 543 team competition, 539-541 subjective targets, 532-534 wrong sample, 530-531 Performance pay systems, 510-514t Performance related pay (PRP), 532 Permanent migration, 356-357 Perry Preschool Program, 4, 5f Personality, 7–14 competing taxonomies of, 51f

conceptualizing, 14-30 constructs, 48-50 definition of, 23-27 disorders, 54 dominant theories of, 48 and health, 112-115, 113-115f interventions on, effect of, 140-149t mapping preferences into, 73-74 measuring, 30-36, 65-73, 90 causal status of latent variables, 36 discriminant and convergent validity, 31-33 faking, 35-36 linear factor models, 30-31 predictive validity, 33-35 and preference parameters, 65-84 evidence on, 65-73 measures of, 66-69t predicting real world behavior, 74-75 Roberts's model of, 8, 9f strategy functions, 5 Personality predict main findings, 89-90 measures of, 90 Personality psychology definitions, 7-14, 8n history of, 7-14 models in, 28-30 pioneers of, 10-12 subsequent flourishing of, 12-14 Personality Research Form (PRF), 51 Personality traits, 6-14 adding learning and uncertainty, 22-23 allowing for multiple tasking, 17-18 approach based on comparative advantage, 15 - 17Big Five, 11, 11t, 44-45t, 48 broad evidence on changes in, 118-122 cognitive, 54-63 within economic models, 14-30 evidence ontogenic and sociogenic change, 122-125 from psychotherapy, 139 external changes to biology, 126-127 identifying, 18-20 implementing measurement systems, 43-46 life cycle dynamics of, 27-28 measuring, 34 predictive power of, 84-117

Personality traits (Cont.) and preferences goals, 21-22 link between, 85-88t Roberts' definition of, 8 schooling on cognitive and, effects of, 132-134, 133f, 135f stability and change in, 117-150 Person-situation debate, 12-14 evidence on. 64 PISA, see Programme for International Student Assessment POINT program, 527 Policy experiments on private school vouchers, 657-659 on public school funding, 659-660 Policy lessons, 216-220 additional, 219-220 Political economy channel, 446-447 and public subsidies, 607-608 Political reforms and education reforms, 485-487 Positive core self-evaluations, 53 Positive reciprocity, 80 Postmigration schooling investment, 375–376 Power equalization scheme with recapture (PER), 656-657 Predictive power of conscientiousness, 105t of personality traits, 84-117 crime, 116f, 116-117 educational attainment and achievement, 90-106 labor market outcomes, 106-112, 107f, 108t, 109f main findings, 89-90 personality and health, 112-115, 113-115f Predictive validity, 33-35 of tests of cognition, 40-43, 41-42t Preference parameters, 65-84 evidence on, 65-73 measures of, 66-69t predicting real world behavior, 74-75 Prefrontal cortex (PFC), 46 Private education, 650 growth and inequality, 662-666 vs. public education, 662-666 Private school vouchers, 635f, 635-637 policy experiments on, 657-659

Private sectors public sectors and, 556 utilization of, 556 Private supplements, to publicly provided education, 627f, 627-629, 629f Productive efficiency, 224 Programa de Ampliación de Cobertura de la Educación Secundaria (PACES), 557 Programme for International Student Assessment (PISA), 346 outcomes of immigrant children, 406t test scores, raw difference in, 347t Promoting Alternative Thinking Strategies (PATHS), 137 PSID, see Panel Survey of Income and Dynamics Psychology literature, analogous measures in, 66-69t Psychopathology, relating Big Five to measures of, 53 - 54Psychotherapy, evidence from, 139 Public education, 630 choice of systems, 631-632 growth and inequality, 662-666 spending in presence of transfers, 624-627, 625-626f, 626t with private option, 618-632 universal, 623, 630 voting patterns and political support for, 671-673 vs. private education, 662-666 Public policy, college degree completion and, 608-609 Public school funding, policy experiments on, 659-660 Public sectors, 559 and private sectors, 556 Public subsidies, political economy and, 607-608 Publicly provided education, private supplements to, 627f, 627-629, 629f

Q

Quantum Opportunity Program, 214 Quasi-cognitive traits, 43

R

Randomization, 560 RBI, *see* Reducing Burglary Initiative RD, *see* Regression discontinuity Realized matches method, 292–294

Reciprocity, 79 Reducing Burglary Initiative (RBI), 201 Reform Act, 486 Refugee sorting, 381 Region-specific cohort trends, 240 Regression discontinuity (RD) design, 231 estimation, 231 strategies, 231 techniques, 224 Residential segregation, simple model of, 646-647 Residential sorting over life cycle, 668-669 Retention rates of graduates, 341 Return migration, 343-346, 351, 368-370 immediate and postponed, 359f model of, 351f, 351-355 skill enhancement and decisions of, 353-354 Roberts's model of personality, 8, 9f Robustness dynasties, 465 human capital and utility functions, concave production of, 466 labor-augmenting technological progress, 463-465 random shocks, 465 Rockoff and Staiger retention policy, 500 Rosenberg Self-Esteem Scale, 52, 132 Rothschild-White model, 589 Rotter Locus of Control Scale, 52, 132 Roy model, 377, 387 extensions of, 20-21

S

Schooling on health European instrumental variable studies, 236-237t US-based instrumental variable studies, 234-235t on mortality and survival, 233t stage, 186-187, 189 Schooling age laws, 198n Screening mechanisms, 504 Screening teachers Bellman equation, 501 Jovanovic model, 501 norm heterogeneity, 503 Rockoff and Staiger retention policy, 500 screening mechanisms, 504

steady-state analysis, 500 talent heterogeneity, 503 Seattle Social Development Project, 214 Second-generation immigrants, 348f, 348 Selection scenarios, 383f Self-assessment, 291-292 Self-control, 43, 116 Self-esteem, 52-53 Self-Sufficiency Project (SSP), 138 SES, see Socioeconomic status SIE, see Survey of Income and Education Signaling and screening, 314 SIMCE, see Sistema de Medición de Calidad de la Educación Single-crossing property, 622 Single-input model, 229 SIPP, see Survey of Income and Program Participation Sistema de Medición de Calidad de la Educación (SIMCE), 564 Situational specificity hypothesis, evidence on, 63 - 64Skilled labor-intensive sector, 448, 463 Skills composition income distribution and, 457-458 and income per capita, 459-460 Skin conductance responses (SCRs), 127 Social and emotional learning (SEL) programs, 137 Social cognitive theories, 29-30 Social Democratic Party, 559 Social multiplier effects, 249 Social-political transitions, 470-471 Socioeconomic status (SES), 207 Stability of economic preference parameters, 139-150 of preferences and personality, 151-155t Standardized achievement test scores, 101-103, 102f diverge, 103-106, 105t Stanford-Binet test, 47 STAR system, 529 State-specific cohort trends, 240 Static complementarity, 129 Steady-state analysis, 500 Strong transferability, 390 Students with immigrant background, 346, 347t migrations, 329n

Super transferability, 389
Survey of Income and Education (SIE), 375
Survey of Income and Program Participation (SIPP), 241
Sweden's voucher program, 559–560
Synthetic panel methodology, 363
System GMM estimator, 263–264, 264n
Systematic measurement error, 370

Т

Talent heterogeneity, 503 TAP program, 530 Technical Instruction Act. 477 Temperament, 49, 49n measuring, 49n Temporary migration, 357-358 Tennessee Comprehensive Assessment Program (TCAP), 527 Tertiary education, 339 foreign students in number and share of, 340t origin of, 343t, 344f Texas Assessment of Academic Skills (TASS), 518 "The Overeducated American" (Freeman), 284 Tiebout hypothesis, 644 Tiebout models with commuting costs, 660-661 with property taxes, peer effects, and private education, 657-660 Timeline return migration, 389f TIMSS test data, 565 Two-dimensional skill framework, 391f

U

Underschooling, 297t, 311t Unified Growth Theory, 473–474 Unified theory of inequality and growth, 466–468 Uniform vouchers, 632–635, 633–634f Unions, teacher, 556 Unique mixed-market institutional structure, 606 Universal public education, 623, 630 Unskilled crimes, 190–191, 191n Unskilled labor-intensive sector, 448, 463

V

Vaginal birth after a caesarean (VBAC), 254 Value of a statistical life (VSL), 228, 229n, 259 Value-added models (VAM), 529 VBAC, see Vaginal birth after a caesarean Ventromedial prefrontal cortex (VMPC), 126-127 VMPC, see Ventromedial prefrontal cortex Vouchers, 632-638 Chilean program, 564 effects on school system, 565-568 on voucher users, 560-565 negative effects of, 658 policy experiments on private school, 657-659 private school, 635f, 635-637 programs, 557-560 Sweden program, 559-560 targeting the poor, 637-638 uniform, 632-635, 633-634f VSL, see Value of a statistical life

W

Wages, 90, 108t
Wechsler Adult Intelligence Scale (WAIS), 47
Wechsler Intelligence Scale for Children (WISC), 47
White-collar crimes, 190–191, 191n
Work stage of life, 186–187, 189n