

Essays in Applied Microeconomics

*Inauguraldissertation zur Erlangung des akademischen Grades eines
Doktors der Wirtschaftswissenschaften der Universität Mannheim*

Xi Lin

Winter Semester 2024

Head of the Department of Economics: Prof. Dr. Thomas Tröger

Advisor: Prof. Dr. Antonio Ciccone

Co-advisor: Prof. Katja Maria Kaufmann, Ph.D.

Co-advisor: Prof. Philipp Ager, Ph.D.

Defense: 13 November 2024

Acknowledgments

I would like to express my deepest gratitude to my supervisors, Antonio Ciccone and Katja Kaufmann, who have also been my coauthors in Chapters 2 and 3 of this dissertation. Antonio has been an extraordinary mentor, and his guidance has profoundly shaped how I approach research as an applied economist. From him, I learned the importance of rigorous attention to detail and critical thinking in every aspect of research. His ability to derive insights from data and connect them with theory has taught me how to bridge empirical evidence with economic models. Antonio's passion for exploring unfamiliar areas and his determination to uncover the true story behind the data have been invaluable to my growth as a researcher. I am especially grateful for the immense effort Antonio has invested in me—his unwavering support, patience, and understanding of both my work and personal circumstances, particularly as a foreigner and father navigating the challenges of the pandemic.

Katja has played a pivotal role in my development from the beginning of my PhD. She introduced me to the exciting field of applied economics, particularly in the area of child development, which became my primary research focus. When I was formulating my research agenda, she generously introduced me to the rich Dutch administrative data. This unparalleled research opportunity became the foundation for all the chapters in this dissertation. Her extensive network and insightful feedback on my projects were invaluable in advancing my work. Katja is an incredibly encouraging and understanding mentor, and her support was especially valuable during the ups and downs of my academic journey. Even after she left Mannheim, her support never wavered; for that, I am truly appreciative.

I am also privileged to have collaborated with Pia Pinger on the third chapter of this dissertation, and I am sincerely grateful for her guidance and support. Her expertise in subjective expectations and survey data significantly enhanced the research, and her vision and insightful feedback were crucial in shaping this chapter. Our frequent remote discussions during COVID ensured steady progress, even during challenging times, and I am thankful for her continued support throughout.

I would also like to express my appreciation to Alex Göppert, Yogam Tchokni, and Jakob Yi Bao for their outstanding research assistance and editing help.

Special thanks go to Lukas Mahler for generously sharing his expertise and advice on calibration, which greatly accelerated my familiarity with this approach. I am also grateful to Mark Spils and Felix Rusche for the countless inspiring conversations we had while sharing offices, and to my talented and supportive cohort of Ph.D. students, who greatly enriched my doctoral experience. Furthermore, I extend my gratitude to many other researchers in the economics department at the University of Mannheim, as well as CRC members and participants of various seminars, for their valuable feedback and engaging discussions.

I would like to thank Stefanie Liwall for her many years of dedicated administrative support. I gratefully acknowledge funding for the administrative microdata provided by the German Research Foundation (DFG) through CRC TR 224 (Projects A04 and C01). This research made use of data from Statistics Netherlands (CBS), which was made available through the microdata service team, and I sincerely thank them. I am deeply grateful to Hans Martin von Gaudecker and Britta Altenburg for their efforts in organizing the CRC research group using CBS data and for their dedicated coordination between CRC and Statistics Netherlands.

Finally, I would like to express my deepest gratitude to my parents, Shen, and my grandparents. They have shaped who I am today, and even though they are on the other side of the continent, they have always managed to motivate and support me, pushing me to become a better person. To Chuan and Nomi, thank you for all your love, patience, and companionship all the way along the long journey of my PhD. None of this would have been possible without you.

Contents

Acknowledgments	ii
Contents	iv
List of Figures	vii
List of Tables	xi
Preface	1
1 Neighborhood Exposure Effects in Cognitive Skills and the Role of Primary Schools	3
1.1 Introduction	4
1.2 Data and Institutional Background	8
1.2.1 Data Source and Sample Selection	8
1.2.2 Neighborhoods and Movers	8
1.2.3 Institutional Setting	10
1.2.4 Variable Definitions	11
1.2.5 Summary Statistics	12
1.3 Identifying Childhood Exposure Effects	12
1.3.1 Step 1: Estimating Neighborhood-Level Predicted Outcomes	12
1.3.2 Step 2: Estimating the Impact of Neighborhood Moves on Child Outcomes	13
1.3.3 Identification Assumptions	14
1.3.4 Validation of Identification Assumptions	16
1.3.5 Relevance of School Performance for Long-Term Educational Outcomes	17
1.4 Mechanisms	18
1.4.1 Buurt-Level Exposure Effects	18
1.4.2 Spatial Decay of Neighborhood Effects	19
1.4.3 Separating School and Neighborhood Effects	20
1.5 Conclusion	23
1.6 Tables and Figures	25

Appendices to Chapter 1	32
1.A Mediation Analysis of School Performance and Educational Attainment	33
1.A.1 Effects on Educational Attainment	33
1.A.2 Mediation Analysis	35
1.A.3 Results	36
2 Improving Equality of Opportunity at School Entry? Primary Education Enrollment in the Netherlands	38
2.1 Introduction	39
2.2 Kindergarten and Primary Education in the Netherlands: Background and a First Look at the Data	42
2.2.1 Kindergarten and Primary Education in the Netherlands . .	42
2.2.2 Age at the Start of Primary Education	43
2.2.3 Primary and Secondary Education Outcomes	45
2.3 The Model	46
2.4 Age at Enrollment in Primary Education	50
2.4.1 Skill Endowments, Skill Shocks, and Enrollment Policies . .	50
2.4.2 Ordered Logit Model for Age at Enrollment in Primary Education	51
2.4.3 Evaluation of the Primary Education Enrollment Model . .	52
2.5 Primary and Secondary Education Outcomes	54
2.5.1 Calibrating the Models for Education Outcomes	55
2.5.2 Evaluating the Model for Education Outcomes	57
2.6 Counterfactual Analysis of Enrollment Policies	62
2.6.1 Two Counterfactual Primary Education Enrollment Policies	63
2.6.2 Comparing Outcomes in Primary and Secondary Education	63
2.7 Summary	69
2.8 Tables and Figures	70
Appendices to Chapter 2	86
2.A Data Sources and Summary Statistics	87
2.B Additional Tables and Figures	91
3 Intra-Household Allocation of Parental Investment: Parents' Equity-Efficiency Preferences and Beliefs	107
3.1 Introduction	108
3.2 Theoretical Framework	113
3.2.1 Model Set-Up	113
3.2.2 Interplay of Perceived Returns and Preferences	114

3.3	Survey Design	116
3.3.1	Beliefs about the Human Capital Production Function . . .	117
3.3.2	Equity-Efficiency Preferences	118
3.3.3	Parental Time Investment	119
3.4	Data and Descriptive Statistics	119
3.4.1	Sample Selection	120
3.4.2	Variable Description	121
3.4.3	Descriptive Statistics	122
3.5	Empirical Strategy	123
3.5.1	Parental Beliefs in Returns to Time Investment	123
3.5.2	Analyzing Parental Equity-Efficiency Preferences	124
3.5.3	Equity-Efficiency Preferences and Child Outcome Gap . . .	125
3.5.4	Equity-Efficiency Preferences and Parental Investment Gap .	126
3.5.5	Equity-Efficiency Preferences, Perceived Returns, and Parental Investments	127
3.6	Results	128
3.6.1	Descriptive Evidence on the Perceived Production Function	129
3.6.2	Descriptive Results on Equity-Efficiency Preferences	131
3.6.3	Equity-Efficiency Preferences and Outcome Gap	134
3.6.4	Equity-Efficiency Preferences, Perceived Production Function and Investment Gap	135
3.7	Conclusion	138
3.8	Figures	139
3.9	Tables	144
Appendices to Chapter 3		152
3.A	Survey Questions: Exact Wording	153
3.A.1	Equity-Efficiency Preferences	153
3.A.2	Parental Time Investment	153
3.B	Equality Preferences Types	154
Bibliography		157

List of Figures

1.1	Mean Test Scores for Children of Permanent Residents at the 25th Income Percentile	28
1.2	Childhood Exposure Effects on Test Scores at Age 12	29
1.3	Exposure Effects Estimation Using Within-Municipality Moves	30
1.4	Spatial Decay	31
1.A.1	Reduced-Form Effects on Educational Attainment	34
1.A.2	Correlation Between Parental Income Rank and Educational Transmission	36
2.1A	Age at Enrollment in Primary Education	71
2.1B	Parental Income and Primary Education Enrollment Age 7	71
2.1C	Age at the Start of Primary Education	71
2.2A	Non-Retention in Primary Education	72
2.2B	End-of-Primary-Education Test Score	72
2.2C	Share Attending Academic Secondary School	72
2.3A	Predicted and Actual Age at Enrollment in Primary Education	73
2.3B	Age at Primary Education Enrollment with Homogeneous Learning in Kindergarten	73
2.3C	Primary Education Enrollment with Identical Skill Threshold for October-December Birthdays	73
2.4A	Predicted and Actual Age at the Start of Primary Education	74
2.4B	Predicted and Actual Variance of Age at Start of Primary Education Among Children Born on the Same Day	74
2.4C	Predicted and Actual Age at Start of Primary Education and Income	74
2.5A	Children's Predicted and Actual Age at the Start of Primary Education	75
2.5B	Actual and Predicted Between-Sibling Difference in Age at the Start of Primary Education	75
2.6A	Predicted and Actual End-of-Primary-Education Test Score	76
2.6B	Predicted and Actual Non-Retention in Primary Education	76
2.6C	Predicted and Actual Attendance of Academic Secondary School	76
2.7A	Predicted and Actual Variance in End-of-Primary-Education Test Scores Among Children Born on the Same Day	77

2.7B	Predicted and Actual Variance in Non-Retention Among Children Born on the Same Day	77
2.7C	Predicted and Actual Variance in Academic-Secondary-School Attendance Among Children Born on the Same Day	77
2.8A	Predicted and Actual Test Scores and Income	78
2.8B	Predicted and Actual Non-Retention in Primary Education and Income	78
2.8C	Predicted and Actual Enrollment in Academic Secondary School and Income	78
2.9A	Actual and Predicted Between-Sibling Difference in Test Scores . . .	79
2.9B	Actual and Predicted Between-Sibling Difference in Non-Retention . .	79
2.9C	Actual and Predicted Between-Sibling Difference in Enrollment in Academic Secondary School	79
2.10A	Regression Discontinuities on October 1 in Data Versus Model	80
2.10B	Regression Discontinuities on January 1 in Data Versus Model	80
2.11A	Primary Education Enrollment Policies and Age at the Start of Primary Education	81
2.11B	Standard Deviation in Age at the Start of Primary Education Among Children Born on the Same Day	81
2.11C	Correlation Between Age at the Start of Primary Education and Skill Endowments	81
2.12A	Primary Education Enrollment Policies and Skills at the Start of Primary Education	82
2.12B	Standard Deviation in Skills at the Start of Primary Education Among Children Born on the Same Day	82
2.13A	Primary Education Enrollment Policies and End-of-Primary-Education Test Scores	83
2.13B	Primary Education Enrollment Policies and Non-Retention in Primary Education	83
2.13C	Primary Education Enrollment Policies and Secondary School	83
2.14A	End-of-Primary-Education Test Scores Among Children with Low Skill Endowments	84
2.14B	Non-Retention in Primary Education Among Children with Low Skill Endowments	84
2.14C	Academic Secondary School Among Children with Low Skill Endowments	84
2.15A	End-of-Primary-Education Test Scores and Adverse Skill Shocks . . .	85
2.15B	Non-Retention in Primary Education and Adverse Skill Shocks	85
2.15C	Academic Secondary School and Adverse Skill Shocks	85
2.B.1A	PIRLS	93
2.B.1B	Maternal Schooling and Primary Education Enrollment Age 7	93

2.B.2A	End-of-Primary-Education Test Scores and Age at the Start of Primary Education	94
2.B.2B	Non-Retention in Primary Education and Age at the Start of Primary Education	94
2.B.2C	Attendance of Academic Secondary School and Age at the Start of Primary Education	94
2.B.3	Calendar Year of Enrollment in Primary Education of Children by Birthdates and Skill Endowments	95
2.B.4A	End-of-Primary-Education Test Scores and Age at the Start of Primary Education	96
2.B.4B	Non-Retention in Primary Education and Age at the Start of Primary Education	96
2.B.4C	Attendance of Academic Secondary School and Age at the Start of Primary Education	96
2.B.5A	End-of-Primary-Education Test Scores and Age at the Start of Primary Education	97
2.B.5B	Non-Retention in Primary Education and Age at the Start of Primary Education	97
2.B.5C	Attendance of Academic Secondary School and Age at the Start of Primary Education	97
2.B.6A	Actual and Predicted End-of-Primary-Education Test Scores at the Individual Level	98
2.B.6B	Actual and Predicted Non-Retention at the Individual Level	98
2.B.6C	Actual and Predicted Enrollment in Academic Secondary School at the Individual Level	98
2.B.7A	Actual and Predicted End-of-Primary-Education Test Scores at the Individual Level—Excluding Implausible Values for Income and Education	99
2.B.7B	Actual and Predicted Non-Retention at the Individual Level—Excluding Implausible Values for Income and Education	99
2.B.7C	Actual and Predicted Attendance of Academic Secondary School at the Individual Level—Excluding Implausible Values for Income and Education	99
2.B.8A	Age at the Start of Primary Education—Girls and Boys	100
2.B.8B	Test Scores at the End of Primary Education—Girls and Boys	100
2.B.8C	Non-Retention Primary Education—Girls and Boys	101
2.B.8D	Academic Secondary School—Girls and Boys	101
2.B.9A	End-of-Primary Education Test Scores and Age at the Start of Primary Education—Instrumental-Variables Results and Data	102

2.B.9B	End-of-Primary Education Test Scores and Age at the Start of Primary Education—Instrumental-Variables and Calibration Results	102
2.B.10A	Primary Education Enrollment Policies and Skills at the Start of Primary Education—Income Only Sample	103
2.B.10B	Standard Deviation in Skills at the Start of Primary Education Among Children Born on the Same Day—Income Only Sample	103
2.B.11A	Primary Education Enrollment Policies and End-of-Primary-Education Test Scores—Income Only Sample	104
2.B.11B	Primary Education Enrollment Policies and Non-Retention in Primary Education—Income Only Sample	104
2.B.11C	Primary Education Enrollment Policies and Academic Secondary School—Income Only Sample	104
2.B.12A	End-of-Primary-Education Test Scores Among Children with Low Skill Endowments—Income Only Sample	105
2.B.12B	Non-Retention in Primary Education Among Children with Low Skill Endowments—Income Only Sample	105
2.B.12C	Academic Secondary School Among Children with Low Skill Endowments—Income Only Sample	105
2.B.13A	End-of-Primary-Education Test Scores and Adverse Skill Shocks—Income Only Sample	106
2.B.13B	Non-Retention in Primary Education and Adverse Skill Shocks—Income Only Sample	106
2.B.13C	Academic Secondary School and Adverse Skill Shocks—Income Only Sample	106
3.1	Within-Family Correlation between Parental Investment and Child Ability	139
3.2	Production Function: Perceived Return to Parental Investment (Probability University Degree)	140
3.3	Production Function: Perceived Return to Parental Investment (Expected Income)	141
3.4	Equity-Efficiency Preferences in Choice Probabilities	142
3.5	Change in Likelihood of Response	143
3.B.1	Equality Preferences Type: by Parental Education	154
3.B.2	Equality Preferences Type: by Number of Children	155
3.B.3	Equality Preferences Type: by Family Origin	156

List of Tables

1.1	Summary Statistics	25
1.2	Estimates of Test Score Exposure Effects	26
1.3	School Shares	27
2.1	Ordered Logit Model for Age at Enrollment in Primary Education	70
2.A.1	Data Sources	89
2.A.2	Summary Statistics	90
2.B.1	Calibration Results for Education Outcome Models	91
2.B.2	RDD Estimates for Primary and Secondary Education Outcomes	91
2.B.3	Instrumental-Variables Estimation	92
3.1	Model Predictions for Different Parent Types: Investment and Outcome Gaps	144
3.2	Summary Statistics of Parents	144
3.3	Perceived Return to Parental Investment	145
3.4	Preference Types and Parental Characteristics	146
3.5	Preferences for Outcome Inequality	147
3.6	Response to Additional Efficiency Cost	148
3.7	Preferences and Actual Outcome Inequality	149
3.8	Preference Types and Actual Investment Gap	150
3.9	The Interaction of Preference and Beliefs on Actual Investment	151

Preface

This dissertation explores three fundamental factors in child development—neighborhood, school, and family—focusing on how these factors influence equality of opportunity. Across three self-contained chapters, it combines rich administrative data from the Netherlands with innovative survey insights to investigate how the environments children grow up in shape their educational outcomes. From the quality of the neighborhoods children live in and the age at which they enroll in school to how parents allocate resources among siblings, this research uncovers how these elements impact both academic success and the opportunities available to children. The dissertation contributes to a deeper understanding of the forces shaping child development and social mobility.

The first chapter, which I authored alone, investigates how childhood residential location affects cognitive skills, focusing on the roles of neighborhood and primary school quality in shaping children’s school performance. Using administrative data from the Netherlands, I estimate the causal effect of neighborhood exposure—defined as the impact of time spent in a neighborhood—on children’s test scores at the end of their primary education. By comparing children who move at different ages, I separate the effects of exposure from those of sorting into neighborhoods. The results show that for each additional year, a child spends in a neighborhood with higher expected test scores, their test scores improve by approximately 2.5% relative to the total gap between the lower- and higher-performing neighborhoods. As families can choose primary schools without geographical restrictions in the Netherlands, I can further isolate improvements attributable to school quality. Approximately 40% of the observed test score improvements can be explained by differences in primary school quality. These findings highlight the critical roles of neighborhood environments and school quality in reducing spatial educational inequalities.

The second chapter, co-authored with Antonio Ciccone, examines the Dutch policy of enrolling children in primary education based on developmental skills rather than date of birth, as is common in other countries. We show that compared to enrollment by date of birth, the Dutch enrollment policy reduces differences in outcomes in primary and secondary education across birthdays and raises

outcomes of children with skill endowments at the bottom of the distribution. We also find that an enrollment policy solely based on children's skills would lead to further improvements in equality of opportunity in primary and secondary education.

The third chapter, co-authored with Katja Kaufmann and Pia Pinger, investigates parents' intra-household investment decisions, focusing on how parents' beliefs about the productivity of their time investment (conditional on child ability) and their equity-efficiency preferences influence their investment into the human capital of their children. We design and implement a survey to elicit parents' beliefs and preferences using innovative survey instruments and analyze responses from parents in the Netherlands. We uncover a negative correlation between a child's academic potential and the parents' investment into this child relative to its siblings. To explain this finding, we show that Dutch parents perceive higher marginal returns from learning-related investments in less academically able children. Moreover, on average, parents exhibit equality-focused preferences in the treatment of siblings, i.e., they prefer to invest more into the less able child. We introduce a unified framework demonstrating how the interplay of equity-efficiency concerns and diminishing returns to parental involvement can moderate disparities in intra-family investments. By linking our survey data to administrative data from the Dutch Statistics Bureau (CBS), we show that parents' equality-focused preferences lead them to invest in a way that reduces the gap in academic outcomes among their offspring. Data on parents' beliefs and preferences also help in predicting differential investments between siblings across families. Actual investment differences are smallest for parents with equality-in-investment preferences. Parents who prioritize equality in outcomes are more likely to invest additional resources in the less academically able child, while parents who focus on efficiency and believe that higher ability leads to greater returns on investment tend to direct more resources toward the more academically able child.

Chapter 1

Neighborhood Exposure Effects in Cognitive Skills and the Role of Primary Schools

1.1 Introduction

Recent experimental and quasi-experimental studies provide evidence that the neighborhood in which a child grows up is associated with long-term differences in educational and labor market outcomes (Chetty and Hendren, 2018a; Deutscher, 2020; Chetty et al., 2020; Laliberté, 2021). However, the mechanisms through which neighborhoods influence these outcomes remain less well understood. Emerging research suggests that neighborhoods exert influence through both contemporaneous (situational) effects shaped by the current environment and developmental (exposure) effects that accumulate over time¹. Further investigation into these mechanisms is essential for informing policies aimed at improving childhood environments and expanding opportunities for disadvantaged populations.

In this study, I investigate whether neighborhood exposure—defined as the effect of time spent in a specific neighborhood—affects children’s cognitive skills development, particularly their school performance, as measured by standardized test scores at the end of primary education. Additionally, I explore how much of the observed neighborhood exposure effects can be attributed to differences in school quality. By isolating the role of school quality, I aim to determine whether improvements in test scores are primarily driven by the neighborhood environment itself or by the quality of the schools children attend within those neighborhoods.

My empirical analysis draws on detailed administrative data from the Netherlands that offer several advantages for studying these issues. First, the data include standardized measures of academic performance at the end of primary school and long-term outcomes, enabling consistent comparisons across regions. Second, the dataset tracks children’s residential histories from birth, allowing for precise measurement of neighborhood exposure over time. Third, the Dutch system of free school choice decouples residential location from school attendance, facilitating a clear separation of neighborhood and school quality effects.

In my main analysis, I define neighborhood at the municipality level, the lowest tier of government in the Netherlands. I begin by documenting the variation in school performance across municipalities. The results show substantial differences in standardized test scores between municipalities, even after accounting for family background. For instance, among children with parental income at the 25th percentile of the national income distribution, the difference in expected test scores between the highest- and lowest-performing municipalities can be as large as 15 percentage points in the national test score ranking.

To further explore the relationship between neighborhood exposure and school performance, I apply a *mover design*, following the framework developed by Chetty

¹See recent review by Chyn and Katz (2021).

and Hendren (2018a). This approach compares children who move between neighborhoods at different ages. It allows me to estimate how much their academic outcomes converge toward those of children who have always lived in the destination neighborhood (i.e., permanent residents). The results suggest that moving to a neighborhood with higher expected school performance is associated with gradual improvements in own school performance over time: for each additional year of exposure, children close the gap in end-of-primary school performance by approximately 2.5%, relative to the difference between lower- and higher-performing neighborhoods. This identification strategy assumes that selection effects related to moving into neighborhoods with different school performance levels do not vary systematically with the child's age at the time of the move. I control for family fixed effects and examine subject-specific convergence patterns to test this assumption. The results of these robustness checks are consistent with my main findings.

I then examine the granularity of neighborhood dynamics. I replicate the baseline analysis, restricting the sample to moves at a finer geographic level—specifically, *buurten*, the smallest geographical unit used by the Statistical Bureau Netherlands with populations ranging from 1,000 to 5,000 residents. This allows for a more granular analysis of neighborhood effects. The results indicate that neighborhood effects are primarily localized. I also find that only the closest *buurten* have a significant impact on children's outcomes, with effects weakening as distance from the child's home increases. My findings suggest that policies aiming to improve educational outcomes may need to target very local areas to be most effective.

Finally, the institutional context of the Dutch education system, which has a free school choice policy (Patrinos, 2011), allows me to decompose the total neighborhood exposure effects into components attributable to school quality, neighborhood quality, and family quality. The results suggest that school quality accounts for around 40% of the observed variation in test scores due to neighborhood exposure. While these findings highlight the potential importance of school quality, other neighborhood characteristics also appear to play a role in shaping children's academic outcomes.

This paper contributes to two strands of literature: the impact of neighborhoods on child development and the role of schools in mediating this impact. First, this study adds to the literature on neighborhood effects on children's outcomes. Chetty and Hendren (2018a) and Chetty and Hendren (2018b) document the long-term impacts of neighborhood environments on children's economic and educational trajectories, showing that the duration of exposure to better neighborhoods—measured using outcomes of the permanent residents—is an important determinant of future success. Chyn and Katz (2021) review this literature, emphasizing the importance

of understanding how neighborhood effects translate into different outcomes. Moreover, studies such as those by Guryan et al. (2021) and Wodtke (2018) show that neighborhood characteristics, such as access to resources and exposure to violence, influence cognitive development. At the same time, Jackson (2020) and Rossin-Slater (2018) highlight the role of early childhood environments in shaping long-term achievement. Building on this literature, my paper provides new evidence from the Netherlands by analyzing how neighborhood exposure affects school performance, using data from a national standardized test administered to nearly all pupils. This earlier measure of academic outcomes helps bridge the gap between childhood exposure and longer-term educational outcomes, while the use of large-scale administrative data allows for a comprehensive analysis. Parallel to my study, Webbink, ter Weel, and Odding (2023) examines neighborhood exposure effects on income at age 30 in the same Dutch context, noting that defining movers based on parental addresses can introduce measurement errors, especially for older children (ages 16-24). In my study, which focuses on children up to age 16, this issue is less pronounced, as most children still live and move with their parents in the Netherlands. However, I exclude cases where children do not move with their parents to minimize potential bias.

Second, this paper contributes to the literature on the role of schools in mediating neighborhood effects. While extensive literature has explored the effects of school and neighborhood separately, studies combining these two contexts are still underdeveloped. Sykes and Musterd (2011) uses Dutch data to examine the correlation between the characteristics of neighborhoods and schools and educational outcomes. Unlike their approach, I use outcome-based measures of neighborhood and school quality, which allows me to avoid the issue of selecting specific observable characteristics to proxy for quality. Card, Domnisoru, and Taylor (2018) uses state- and county-level data from the early 20th century, showing that variations in school quality were key drivers of regional differences in upward mobility. Similarly, Rothstein (2019) examines the impact of K-12 school quality on intergenerational mobility using aggregate data at the commuting zone (CZ) level across U.S. cities. His findings suggest that school quality explains only a small portion of the variation in intergenerational income mobility across regions, with broader factors like neighborhood characteristics and economic conditions playing a more substantial role. In comparison, my use of micro-level data allows for a more detailed analysis by linking childhood environments directly to later educational outcomes, offering more precision in identifying the specific effects of school and neighborhood quality. Consistent with the findings of Gibbons and Silva (2018) and Aizer and Currie (2019), I find that school quality is an important determinant of later education attainment. My study also provides evidence that variation in test scores due to neighborhood

differences can be attributed to differences in school quality.

A directly related study is Laliberté (2021), which developed an innovative approach to decomposing neighborhood exposure effects into the portion attributable to school quality and the portion due to non-school neighborhood factors. Set in Montreal, Canada, the study leverages the special institutional context, where the primary school catchment area boundaries for French and English schools do not perfectly overlap, to disentangle the long-term effects of school quality and neighborhood characteristics on educational outcomes. The findings reveal that 50 to 70 percent of the total effect of living in a better neighborhood on educational attainment is attributable to school quality, with the remaining portion tied to non-school neighborhood characteristics.

While my study build on the decomposition approach developed in Laliberté (2021), my study features several differences. First, the primary school system in the Netherlands has no formal school catchment areas, meaning parents are free to send their children to the school of their choice (Patrinos, 2011). This unique feature disentangles school choice from residential location choice, allowing for a clearer separation of school and neighborhood effects. Second, unlike in Laliberté (2021), where children with different mother tongues might face cultural barriers, children attending different schools in the Netherlands typically encounter no language-based barriers with their neighbors. This situation is more common globally and may be more representative. This difference could explain why, in my study, the neighborhood share of the effect is higher than in Laliberté (2021)’s findings. Third, given that the government equally funds both private and public primary schools in the Netherlands, and all school enrollments are well documented, I can examine the entire population of the Netherlands, capturing a more diverse cross-section of the population compared to the sample of public schools from Montreal Island used in Laliberté (2021). Lastly, although I could not employ a boundary discontinuity design, observing family income and schooling from the administrative data allows me to better proxy for family environment and separate those effects from school and neighborhood quality.

The remainder of this paper proceeds as follows. Section 2 provides an overview of the Dutch education system and details the data used in the study. Section 3 outlines the empirical strategy and presents the main results. Section 4 conducts a mediation analysis to disentangle the effects of school quality from other neighborhood factors. Section 5 concludes with policy implications.

1.2 Data and Institutional Background

1.2.1 Data Source and Sample Selection

I use a comprehensive administrative database from the Netherlands, managed by Statistics Netherlands (Centraal Bureau voor de Statistiek, CBS).² This database covers the country's entire population and provides detailed information that enables tracking individuals both geographically and over time. Different datasets in the database can be linked through individual random identification numbers.

The selection of birth cohorts is based on the availability of three primary datasets: the municipal registry (1995–2022), the CITO test scores at the end of primary education (2006–2019), and education enrollment records (2003–2022). I restrict the sample to individuals whose moving history is fully observable from age 1, including birth cohorts from 1994 onward. Due to the cancellation of the CITO test in 2019 during the pandemic, children born after 2007 are excluded, as some of them did not take the test as they otherwise would have.

I assess educational attainment at age 24, by which time nearly all individuals in the Netherlands have typically completed or are concluding their education. This age selection aligns with previous studies that evaluate educational outcomes at entry into the labor market. Given that 2022 is the last available year for education enrollment records, the latest birth cohort for which I can observe educational attainment at age 24 is 1998.

The main sample comprises all children who meet the following criteria: (i) inclusion in the municipal population register from birth, (ii) birth between 1994 and 2007, (iii) at least one parent is identifiable, and (iv) their CITO test scores are available for the years 2006 to 2019. Every individual who has resided in the Netherlands since 1995 is included in the municipal population register. Undocumented immigrants and asylum seekers, who are typically not registered, are the primary groups excluded from the sample. Children are linked to parents using data on legal parent(s) from the municipal population registers, and only those with at least one identifiable legal parent are included in the sample.

1.2.2 Neighborhoods and Movers

I determine the neighborhoods of parents and children using their home addresses registered in the municipal population register, which are continuously updated based on administrative data from schools and social security agencies. The municipal population register is the government's primary means for communicating

²The database is accessible via a remote-access computer after a confidentiality statement has been signed.

with citizens on various issues, including taxes, income, and social security matters. The Dutch data uniquely enable identifying the addresses of children and those of their parents separately, allowing for precise analysis. I code relocations based on the child's address in the baseline analysis. When individuals relocate, they must notify the municipal administration of their old and new addresses, along with the exact date of the move. Using birth date data, I can calculate the age at the time of the move with day-level precision.

The geographic information available for this study is exceptionally detailed, allowing observations of an individual's address down to the specific building and enabling precise calculations of moving distances. To define relevant neighborhoods, I use two levels of granularity. My baseline analysis focuses on municipalities. As of 2011, the Netherlands has 419 municipalities with an average population of approximately 40,000. For comparison, the study of Chetty and Hendren (2018a) is at the level of 741 commuting zones, averaging approximately 380,000 inhabitants. Municipalities are crucial in administering various local services, including education, as they implement national education laws and regulations at the local level. This administration includes overseeing the establishment and maintenance of schools, ensuring compliance with educational standards, and supporting special education needs.

For a more detailed analysis of local interactions, I define neighborhoods at a smaller scale, specifically, *buurten* in Dutch. A *buurt* is the smallest geographical unit used by the Statistical Bureau Netherlands, typically corresponding to a well-defined area within a city or town. There are approximately 11,000 *buurten* across the country with populations ranging from 1,000 to 5,000 residents. The number of schools per *buurt* varies, with urban areas having more schools, while many rural *buurten* possibly having none. The number of primary school-aged children per *buurt* ranges from 100 to 300. This level of granularity allows for a more precise examination of local dynamics and social interactions, which is crucial for studying the impact of localized factors on educational outcomes.

I divide the sample into permanent residents (or stayers) and movers. Permanent residents in each neighborhood are defined as the subset of children born and residing in a single neighborhood up to age 15. Movers are non-permanent residents, with one-time movers defined as those who move exactly once with their parents across neighborhoods before the age of 15³. The main sample has approximately 1.7 million children for whom I observe CITO scores, including about 1.3 million permanent residents and 295,000 one-time movers.

³In a recent study, Webbink, ter Weel, and Odding (2023) argue that including parental moves without children could introduce measurement error in this type of design. Therefore, I only consider cases where the addresses of both children and parents change simultaneously when defining movers.

1.2.3 Institutional Setting

Education in the Netherlands is compulsory from age 5 to 16, with many children beginning at age 4. The system is split into primary, secondary, and tertiary education.

Primary education in the Netherlands is characterized by a unique system of free school choice: families are not limited by residential catchment areas when selecting schools. Parents can enroll their children in any primary school—public, private, religious, or special education—regardless of their geographic location. All schools must meet the same national educational standards and receive government subsidies to ensure equitable funding across different types of schools. The central government provides the majority of financial resources for schools, including funding for teacher salaries, instructional materials, and student support services. This centralized funding model is designed to ensure that all schools, irrespective of their type or location, can deliver high-quality education and adhere to the national curriculum guidelines. It highlights the freedom of choice, the financial equality across school types, and the role of government subsidies in ensuring consistent educational standards across all schools in the Netherlands.

This structure contrasts with the U.S. context studied by Chetty and Hendren (2018a), where educational opportunities, especially those provided by the states, are more directly tied to the neighborhood in which a child resides. In the U.S., school catchment areas typically restrict families to public schools within their residential zone (e.g. Eppele and Romano, 2003), making it difficult to disentangle place effects from school effects. In contrast, the Dutch system’s freedom of choice allows parents to bypass geographic limitations, offering a unique opportunity to identify the separate effects of neighborhoods and schools on educational outcomes.

A structured and diverse secondary education system complements the freedom of primary school choice in the Dutch primary system. After primary school, students are tracked into one of three main paths—VMBO (pre-vocational), HAVO (senior general), or VWO (pre-university)—based on their performance. These tracks cater to students’ varying abilities and aspirations, preparing them for tertiary education in vocational training (MBO) or at universities of applied sciences (HBO) or research universities (WO).

MBO offers vocational education and training at four levels, preparing students for the labor market or further studies. Students are typically directed into one of these four pathways after completing secondary education, following the same track as in their secondary schooling. HBO institutions focus on professionally oriented programs, while WO institutions emphasize academic and research-based

education. Both HBO and WO offer bachelor's and master's degrees, with the possibility of advancing to doctoral programs at research universities.

The entry into tertiary education is highly structured, with most students beginning their bachelor's programs immediately after secondary school. Most students complete their education efficiently, with many reaching the final stage of their studies or having already finished their degrees by 24 (OECD, 2023; European-Commission, 2019).

1.2.4 Variable Definitions

In this section, I define the key variables used in the analysis: *parental income*, *CITO test scores*, and *educational attainment*.

Parental Income. Parental income is defined as the sum of the disposable incomes of the father and mother when the child is between 9 and 12. If parents are separated, parental income is calculated as the average of the mother's disposable income and that of her spouse. If this information is unavailable, the income is based on the mother's disposable income alone. The income variable is top-coded at €1 million. Following Chetty and Hendren (2018a), I convert incomes into percentile ranks relative to the national distribution for the child's birth cohort. This method improves comparability across individuals and reduces the influence of outliers and lifecycle income variability.

CITO Test Scores. Children's academic performance at the end of their primary education is measured using the so-called CITO test. In the Netherlands, primary education consists of six grade levels, and children typically complete their primary education at age 12. Schools can choose the provider of the end-of-primary education test they administer. Approximately 85% schools opt for the CITO test, renowned for its comprehensive assessment of key subjects such as mathematics, Dutch language, and study skills. Participation is mandatory for all enrolled students once a school decides to administer the CITO test. End-of-primary education test scores are high-stakes as teachers consider them in making recommendations for each student's secondary school track. The CITO scores range from 500 to 550 and are standardized across years to ensure consistency. As with parental income, I convert the raw test scores into percentile ranks within the child's national birth cohort. This approach allows for a more nuanced comparison of relative academic performance across regions and periods. Percentile ranks are calculated both overall and by subject area.

Educational Attainment. Educational attainment is measured use the highest degree or academic qualification obtained by the age of 24. Degrees are then converted into years of schooling, following the International Standard

Classification of Education (ISCED) system. This conversion ensures a consistent and comparable measure of educational achievement across individuals, facilitating a robust analysis of how early-life factors affect long-term educational outcomes.

1.2.5 Summary Statistics

Table 2.1 lists the variables included in the study and provides summary statistics for the primary analysis samples—which consist of birth cohorts from 1994 to 2007—and presents these statistics for the entire population and by moving status.

Consistent with previous studies, I find that movers and permanent residents have similar characteristics. For instance, the median family disposable income is nearly €43,000 for both groups. The same pattern applies to children’s educational attainment and test scores, which are comparable between both groups.

1.3 Identifying Childhood Exposure Effects

I use the empirical framework of Chetty and Hendren (2018a) to estimate how childhood exposure affects school performance. This method involves two key steps. First, I use the outcomes of permanent residents to predict the expected outcomes for children growing up in various neighborhoods. In the second step, I focus on children who moved once during childhood. I analyze how moving one year earlier influences children’s outcomes by estimating how the expected outcomes from those in the origin neighborhood converge to those in the destination neighborhood. Specifically, the model captures the rate at which a child’s outcomes converge toward the outcomes of permanent residents in the destination area, with each additional year spent in the destination.⁴

1.3.1 Step 1: Estimating Neighborhood-Level Predicted Outcomes

In the first step, I generate predicted outcomes for children growing up in different neighborhoods using data from permanent residents. To do this, I estimate the relationship between parental income ranks and children’s CITO scores within each municipality. This relationship is captured by the following linear regression model:

$$CITO_i = \alpha_c + \pi_c p_i + \epsilon_i, \quad (1.1)$$

⁴For a comprehensive and formal introduction to this identification strategy, see Chetty and Hendren (2018a).

where $CITO_i$ denotes the child's CITO score, p_i is the parental income rank, α_c represents the neighborhood (municipality) fixed effect, and π_c measures the effect of parental income rank p_i on child outcomes within each neighborhood. The error term ϵ_i accounts for unobserved factors that may influence a child's outcomes.

In my baseline specification, I focus on the effects of parental income rank, omitting variation by birth cohorts, unlike Chetty and Hendren (2018a). This is because, unlike other educational systems that group students by birth year, school entry in the Netherlands is determined by developmental readiness, making cohort distinctions less relevant. In my robustness analysis, however, I account for neighborhood predicted outcomes within birth calendar years to ensure the robustness of my findings. Consistent with Chetty and Hendren (2018a), I use rank-based measures to avoid issues related to attenuation and life-cycle bias.

After estimating the model, I calculate two key predicted outcomes for each child in the sample of movers: $\bar{y}_{op}$ denotes the predicted outcome if the child grew up entirely in their origin neighborhood. And $\bar{y}_{dp}$ denotes the predicted outcome if the child grew up entirely in their destination neighborhood. These predictions will later be used to compute the “expected gains” from moving between neighborhoods by tracking the change from outcomes expected in the origin neighborhood to those expected in the destination neighborhood.

Figure 1.1 maps children's test scores by municipalities for children whose parents are at the 25th income percentile. Children's outcomes vary significantly across municipalities. For example, among children with parents at the 25th percentile, CITO scores are approximately 15 percentage points higher in municipalities at the top (95th percentile) of the mean test score distribution than those at the bottom (5th percentile).

1.3.2 Step 2: Estimating the Impact of Neighborhood Moves on Child Outcomes

In the second step, I use the sample of movers to assess how the age at which children move between neighborhoods affects their outcomes. Specifically, I test whether children who move at younger ages demonstrate more of the predicted difference between their origin and destination neighborhoods. To quantify this, I estimate a child's eventual CITO score as a function of two key factors: their predicted outcome in the origin neighborhood ($\bar{y}_{op}$) and the predicted change in outcomes resulting from the move to the destination neighborhood ($\Delta dp = \bar{y}_{dp} - \bar{y}_{op}$), interacted with the child's age at the time of the move (m).

Consistent with Chetty and Hendren (2018a), the model is estimated using a

semi-parametric approach, represented by the following specification:

$$CITO_i = \sum_{m=0}^{15} I(m_i = m) [\alpha_m + \phi_m p_i + \zeta_m \bar{y}_{op} + b_m \Delta odp] + \epsilon_i, \quad (1.2)$$

where $CITO_i$ is the eventual CITO score of child i , α_m is an intercept that varies with the age at which the child moves, ϕ_m is the age-specific effect of parental income rank p_i , and ζ_m is the age-specific coefficient on the predicted outcome in the origin neighborhood ($\bar{y}_{op}$).

The key parameters of interest are the b_m coefficients, which measure how much of the predicted change in outcomes from moving to a new neighborhood (Δodp) is realized by children moving at different ages. In other words, these coefficients capture the degree to which children's outcomes shift toward the outcomes predicted for their destination neighborhood, depending on their age at the time of the move. The differences between these coefficients—such as $b_m - b_{m+1}$ —are interpreted as the impact of additional exposure to a neighborhood with higher predicted outcomes.

The rate of convergence to the outcomes of permanent residents in the destination neighborhood is used to measure the effects of neighborhood exposure. In this context, the model identifies how much of the difference between origin and destination outcomes is absorbed by children, depending on when the move occurs.

1.3.3 Identification Assumptions

The coefficient b_m captures both selection and exposure effects. Selection effects arise because families who move to better neighborhoods may differ systematically from those who stay, even after controlling for observed socio-economic factors such as income or education level. These differences may be driven by unobserved characteristics, such as preferences for education or aspirations for their children, which could bias the estimated effects of the neighborhood. In contrast, exposure effects reflect the causal impact of the time spent in a new neighborhood on a child's educational outcomes.

To separate the exposure effect from the selection effect, I make two key identification assumptions following Chetty and Hendren (2018a): **Constant Selection Effects** assumes that selection effects do not systematically vary with the child's age at the time of the move. In other words, the characteristics driving a family's decision to move are assumed to be constant across different ages. This allows me to attribute differences in b_m for children who move at different ages primarily to differences in exposure time rather than unobserved differences in family characteristics. **Linearity of Exposure Effects** suggests

that the effect of neighborhood exposure grows linearly with the time a child spends in the destination neighborhood. Under this assumption, each additional year spent in the new neighborhood contributes equally to the child's outcomes. The improvement in outcomes from moving one year earlier can then be interpreted as a constant exposure effect.

Given the first assumption, the difference between b_m values for children who move at different ages can be interpreted as the causal effect of neighborhood exposure. The difference $b_{m+1} - b_m$ reflects the annual exposure effect, which measures how much a child's outcomes improve with each additional year spent in the new neighborhood. This framework allows me to estimate how outcomes converge toward those of permanent residents in the destination neighborhood over time.

Assuming exposure effects are linear, as posited in the second assumption, I can further parametrize equation (1.2) as follows:

$$\begin{aligned} CITO_i = & \sum_{m=0}^{15} I(m_i = m) [\alpha_m + \zeta_m \hat{y}_{op} + \phi_m p_i] \\ & + K(m_i \leq 12) (\gamma' + \gamma(12 - m_i)) \Delta_{odp} \\ & + C(m_i > 12) (\rho' + \rho(m_i - 12)) \Delta_{odp} + \epsilon_i. \end{aligned} \quad (1.3)$$

In this specification, the coefficient γ represents the annual exposure effect for children who move at age 12 or earlier. It captures the average effect of moving one year earlier to a neighborhood where permanent residents score one percentile higher on educational outcomes. Similarly, the coefficient ρ measures the corresponding slope for children who move after age 12, capturing how exposure to a new neighborhood affects older children's outcomes.

By estimating these coefficients, I can calculate the annual exposure effect for different age groups and infer the overall impact of neighborhood exposure on long-term outcomes. This approach allows me to introduce additional controls such as family fixed effects.

Figure 1.2 displays estimates of b_m from Equation 1.2, revealing two main patterns: selection effects after age 12 and exposure effects before age 12. The positive values of b_m for ages $m > 12$ clearly indicate selection effects, as moves after age 12 cannot causally affect CITO test scores, which are obtained at age 12. This finding suggests that children in families who move to better neighborhoods often have favorable unobservable attributes. Furthermore, the degree of selection remains constant across ages above 12, as evidenced by a statistically insignificant slope of 0.001 when regressing b_m on m . This stability aligns with the assumption that selection does not significantly vary based on the child's age at the time of

moving.

Figure 1.2 also shows a steady decline in b_m estimates with the age at move (m) for $m < 12$. According to the first assumption, this downward trend provides evidence of exposure effects, meaning that moving to a better neighborhood earlier in childhood results in greater benefits. The linear relationship between b_m and the age at move (m) for ages below 12 suggests that the exposure effect remains relatively constant across different ages. A regression of b_m on m for ages below 12 estimates an average annual exposure effect of 0.025, indicating that children’s outcomes improve and align with those of permanent residents at a rate of 2.5% per year of exposure up to age 12.

1.3.4 Validation of Identification Assumptions

One key threat to identification is the possibility that families who move at different times do so for different unobservable reasons. For instance, families moving when their children are young may prioritize long-term educational opportunities, while those moving when their children are older might do so for reasons related to employment or financial constraints. This would violate the first Assumption. The result could bias estimates of the neighborhood effects, as variation in b_m across ages may capture differences in family characteristics and not just differences in exposure time.

To address the concern of family selection effects, I incorporate family fixed effects into the model following the approach of Chetty and Hendren (2018a). This approach allows me to compare siblings who moved at different ages but share the same (observed and unobserved) family background, including the learning environment in the family and the educational preferences and aspirations of the family. Put differently, by examining within-family variation, I can isolate the effects of neighborhood exposure from family-specific selection effects. The key insight is whether the difference in school outcomes between siblings is proportional to their difference in exposure time, once family-level confounders are held constant.

However, time-varying factors, such as a parent’s new job or changes in household financial circumstances coinciding with the move, could still introduce bias. To account for this, I also implement outcome-based placebo tests designed to test whether neighborhood exposure specifically affects the outcome of interest rather than unrelated outcomes. The detailed Dutch data allows me to examine subject-specific test scores, providing a novel test of whether the neighborhood’s advantage in a particular subject, such as math, translates into stronger gains in that same subject for children exposed to the neighborhood. If the causal model holds, a child’s math test scores should correlate more strongly with the neighborhood’s

math performance than unrelated subjects like the Dutch language. This approach ensures that the observed effects are specific to the neighborhood's influence rather than being driven by unobserved time-varying factors or general improvements unrelated to subject-specific advantages.

I implement these robustness tests in Table 2.B.3. Column 1 presents estimates of the average annual exposure effect of 3%, which are robust across various specifications and outcome definitions. Column 2 shows that controlling for maternal education and immigration background yields similar results. Columns 3 and 4 show that even when analyzing the data by subject, the convergence effects persist. Column 5 provides estimates where permanent residents' outcomes are measured based on their birth year. Column 6 shows results using variation in age at moves within families when parents relocate. Across these various robustness tests, the results remain consistent, reinforcing the robustness of the exposure effect estimates.

1.3.5 Relevance of School Performance for Long-Term Educational Outcomes

The findings demonstrate substantial effects of childhood exposure to higher-quality neighborhoods on school performance, as measured by CITO test scores. The remaining question is whether these early improvements in school performance translate into better long-term educational outcomes, such as years of schooling or overall educational attainment. This issue is examined in detail in *Appendix A*, where the relevance of school performance at age 12 for predicting educational attainment at age 24 is analyzed. The methodology builds on Rothstein (2019), who explores how schools mediate the intergenerational transmission of income, focusing on educational outcomes as a mediator for long-term socioeconomic mobility. Compared to Rothstein (2019), who decomposes income variation at the commuting zone level, a key advantage of the Dutch data is the ability to link end-of-primary-school performance to educational attainment at age 24 directly.

The mediation analysis in *Appendix A* decomposes the total effect of parental income on children's long-term educational attainment into direct and indirect effects. The indirect effects focus on primary school performance as a mediator, allowing us to assess whether improvements in CITO test scores—driven by exposure to higher-quality neighborhoods—significantly contribute to educational attainment, measured in years of schooling. This analysis clarifies how school performance is a crucial mechanism linking neighborhood quality to children's long-term educational trajectories.

The results suggest that differences in school performance account for approximately

40% of the variation in educational attainment across neighborhoods. This highlights that improvements in CITO test scores, resulting from better neighborhood environments, have significant implications for long-term outcomes. The role of schools in my context appears larger than in Rothstein (2019) (He estimated skills mediate 11% of the spatial income variation). This may be due to two reasons: First, my analysis focuses on primary school performance and its impact on educational attainment, whereas Rothstein (2019) examines income as the outcome. Second, the school tracking system in the Netherlands makes primary school performance more decisive for future educational success, as early school performance determines secondary school placement, shaping long-term educational trajectories. Therefore, improvements in primary school performance have more immediate and far-reaching consequences within the Dutch educational system.

1.4 Mechanisms

To determine the mechanisms underlying my findings, I identify neighborhood effects more granularly by evaluating the impact of moves across *buurten*, the smallest administrative units in the Netherlands. I then explore the spatial decay of these effects, investigating how geographic proximity influences neighborhood outcomes. Finally, I examine the role of educational institutions in explaining neighborhood effects.

1.4.1 Buurt-Level Exposure Effects

I first replicate the exposure effects analysis from Section 3, focusing on children who moved across different *buurten* while staying within the same municipality. *Buurten* are highly localized areas, with approximately 1,400 residents on average and a range of around 500–2,000 residents. Hence, *buurten* are much smaller than municipalities; in urban areas, they can be as small as a few city blocks.

Each *buurt* typically contains one or two primary schools, depending on the density and size of the population, although some smaller *buurten* may not have any schools within their boundaries. *Buurten* may also differ substantially in their availability of local amenities such as churches, small parks, playgrounds, or community centers.

Given the diversity and the small scale of *buurten*, they provide a granular view of interactions that operate locally compared with the larger, more heterogeneous municipalities. Because I focus on moves across *buurten* within the same municipality, the municipal policies—such as educational reforms, public services, or economic initiatives—should affect both the origin and destination *buurten* in a similar

way. Therefore, any outcome differences after a move are likely driven by local within-*buurten* factors, such as peer interactions, social networks, local norms, or environmental characteristics.

By focusing on such small-scale geographic units, I can better isolate the influence of local dynamics and community interactions that may not be captured when analyzing larger areas such as municipalities. The fact that these moves often occur over distances of only a few kilometers emphasizes the importance of neighborhood-level social and environmental factors, as opposed to broader municipal policies.

The estimation results, shown in Figure 1.3, reveal that moves across *buurten* also exhibit clear exposure effects, with patterns similar to those observed in inter-municipality moves. Even with highly localized moves, often within a very short distance, children experience significant gains in standardized test scores. This suggests that neighborhood effects operate highly granularly, driven by localized social and environmental dynamics within each *buurt*, rather than broader municipal-level interventions.

1.4.2 Spatial Decay of Neighborhood Effects

The available data allow me to investigate to what extent neighborhood effects are within a *buurt* rather than across *buurten* located close to each other. My analysis is similar to that of Chetty and Hendren (2018a) at the census-tract level in the U.S. To capture the spatial decay of neighborhood effects, I estimate a regression in which the predictions for the origin and destination neighborhoods ($\hat{y}_{op}$ and $\hat{y}_{dp}$) are replaced with predictions for the child's specific *buurt*. In addition to these immediate neighborhood effects, I include interactions between the child's age at the move and the average observed outcomes in the 10 closest *buurten* to the origin and destination neighborhoods. The selection of these ten closest *buurten* is based on the distance between their geographic centers.⁵

The decision to include the 10 closest *buurten* is motivated by the fact that, in the Netherlands, most children attend schools and participate in activities within a radius of 1 to 3 kilometers from their homes. Therefore, nearby *buurten* often share key community resources, such as schools and recreational facilities, and children frequently interact across these small boundaries. By incorporating this surrounding *buurten*, I capture potential spillover effects from neighboring environments, such as shared peer groups, social networks, or common access to public goods.

⁵The geographic distance between *buurten* is the straight-line distance between their geographical centers. The 10 closest *buurten* are then ranked based on proximity, from nearest to farthest.

The results, presented in Figure 1.3, reveal a clear pattern of spatial decay. Moving to a higher-performing buurt—where permanent residents achieve test scores one standard deviation higher—has the strongest effect in the child’s own buurt. However, the impact of nearby buurten rapidly diminishes with distance. By approximately 1 to 1.5 kilometers away, the effect becomes negligible. These findings suggest that policy interventions aimed at improving child development in specific areas should be targeted at the micro-level, within the *buurten* of the children. Broader municipal or regional policies may not have any effects.

Figure 1.4 presents the results, showing that relocating to a neighborhood with higher test scores earlier in childhood significantly improves a child’s educational performance. In contrast, moving to a neighborhood where only the surrounding neighborhoods have higher test scores, without a corresponding improvement in the child’s immediate neighborhood, does not significantly affect outcomes. This finding suggests that the beneficial effects of high-performing neighborhoods are hyperlocal and directly tied to the child’s immediate environment.

1.4.3 Separating School and Neighborhood Effects

When children move to a new neighborhood in the Netherlands, they experience changes in school quality and non-school neighborhood factors. To understand the impacts of these changes, I decompose the total effect of moving into two components: a school-related effect, which captures differences in school quality across locations, and a non-school-related effect, reflecting other neighborhood factors such as income composition, public services, and social networks.

The approach I adopt here closely follows the methodology developed by Laliberté (2021), who introduced a framework to decompose the effect of neighborhood moves into separate school and non-school components. This approach is particularly well-suited to the Netherlands due to its unique education system and neighborhood structure. In the Netherlands, students can access schools beyond their immediate residential neighborhood owing to a relatively liberal school choice policy. As a result, students from the same neighborhood may attend different schools, and students attending the same school may come from different neighborhoods. This variation provides the necessary identification to separate school and non-school effects. In many other countries, the primary school children attend is often tightly linked to their neighborhood (and the quality of a school is tightly linked to the neighborhood’s affluence). In the Netherlands, the decoupling of residential location and school attendance allows me to separate the contributions of school quality and neighborhood characteristics to children’s educational outcomes.

I estimate the effects of school and neighborhood quality separately for permanent

residents—those who remain in the same neighborhood throughout their schooling. Using a similar specification to Laliberté (2021), the model for permanent residents is given by:

$$CITO_{it} = \Omega_{s(i)} + \Lambda_{l(i)} + X_i\beta + \epsilon_i, \quad (1.4)$$

where $CITO_{it}$ is the primary school performance at time t of child i , who lives in neighborhood $l(i)$ and attends school $s(i)$. $\Omega_{s(i)}$ captures the quality of the schools attended by children who live in l_i , while $\Lambda_{l(i)}$ captures the separate effect of the neighborhood where children live. X_i includes observable characteristics of children, and ϵ_i accounts for unobserved factors. The model is identified because students from the same neighborhood often attend different schools, and students in the same school often come from different neighborhoods, allowing me to separate the effects of school and neighborhood on student outcomes.

For students who move between neighborhoods, I estimate the total exposure effect of moving on their educational outcomes. The total difference in outcomes between permanent residents of the origin neighborhood o and destination neighborhood d is represented as $\Delta_{od} = \bar{y}_o - \bar{y}_d$.

The primary estimating equation for movers is the following:

$$y_{imod} = \gamma(m_i\Delta_{od}) + \beta X_{imod} + \alpha_{od} + \alpha_m + \epsilon_{imod}, \quad (1.5)$$

In this model, y_{imod} represents the educational outcome of student i , who lived in neighborhood o (origin) at baseline and moved to neighborhood d (destination) at age m . The coefficient of interest, γ , captures the annual rate at which the outcomes of movers converge to those of permanent residents in their destination neighborhood. Unlike equation (3), which includes income-specific effects, equation (5) focuses purely on the spatial exposure effects by incorporating origin-by-destination fixed effects (α_{od}) to control for differences across neighborhoods. Additionally, unobserved differences between children who move at different ages—such as disruption costs from the move—are controlled through age-at-move fixed effects (α_m). The model fundamentally compares children who began in the same origin neighborhood and moved to the same destination neighborhood, but at different ages, to isolate the effect of neighborhood exposure independent of income-specific factors.

Following the decomposition approach in Laliberté (2021), the total effect of moving is separated into contributions from school quality and non-school factors. The difference in outcomes between origin and destination neighborhoods Δ_{od} can be written as:

$$\Delta y_{od} = \Delta\Omega_{od} + \Delta\bar{y}_{od}^{ns}, \quad (1.6)$$

where $\Delta\Omega_{od}$ represents the difference in the expected quality of the schools children attend in neighborhoods o and d . This is calculated as a weighted average of school quality estimated in equation 1.4, with weights corresponding to the share of permanent residents' children attending different schools. This specification implicitly assumes that movers will make similar school choices to those of the permanent residents, allowing us to estimate the Intention-to-Treat (ITT) effects. Meanwhile, $\Delta\bar{y}_{od}^{ns}$ captures the difference in expected non-school (ns) factors, including both neighborhood fixed effects and family composition differences.

Because of the unique structure of the Dutch education system, I can separately estimate the following two models.

$$y_{imod} = \gamma_s(m_i\Delta\Omega_{od}) + \beta X_{imod} + \alpha_{od} + \alpha_m + \epsilon_{imod}, \quad (1.7)$$

and

$$y_{imod} = \gamma_{ns}(m_i\Delta\bar{y}_{od}^{ns}) + \beta X_{imod} + \alpha_{od} + \alpha_m + \epsilon_{imod}. \quad (1.8)$$

The variation in school attendance and neighborhood residency patterns in the Netherlands enables estimation of these two effects separately. This flexibility is key to disentangling school quality from broader neighborhood characteristics.

The total exposure effects can now be written as the sum of two effects:

$$\begin{aligned} \gamma &= \frac{\text{cov}^r(y_{imod}, \Delta y_{od})}{\text{var}^r(\Delta y_{od})} \\ &= \underbrace{\frac{\text{cov}^r(y_{imod}, m_i\Delta\bar{y}_{od}^{ns})}{\text{var}^r(m_i\Delta\bar{y}_{od}^{ns})}}_{\gamma_{ns}} \frac{\text{var}^r(m_i\Delta\bar{y}_{od}^{ns})}{\text{var}^r(m_i\Delta y_{od})} + \underbrace{\frac{\text{cov}^r(y_{imod}, m_i\Delta\Omega_{od})}{\text{var}^r(m_i\Delta\Omega_{od})}}_{\gamma_s} \frac{\text{var}^r(m_i\Delta\Omega_{od})}{\text{var}^r(m_i\Delta y_{od})} \end{aligned} \quad (1.9)$$

where cov^r and var^r refer to the covariance and variance obtained using the residuals of regressions on the controls employed when estimating school and non-school neighborhood effects.

The relative contributions of school and non-school factors to the total effect of moving can now be obtained following as:

$$F^{\text{school}} = \frac{\gamma_s \text{var}^r(m_i\Delta\Omega_{od})}{\gamma \text{var}^r(m_i\Delta y_{od})}, \quad (1.10)$$

and

$$F^{\text{non-school}} = \frac{\gamma_{ns} \text{var}^r(m_i\Delta\bar{y}_{od}^{ns})}{\gamma \text{var}^r(m_i\Delta y_{od})}. \quad (1.11)$$

Table 1.3 presents the results. The upper panel highlights statistically significant

total exposure effects, ranging from 0.023 to 0.033, which closely mirror the estimation results derived from the specifications outlined in Section 3. The breakdown of these effects reveals notable disparities in the contributions of school and non-school factors. In particular, approximately 39% of the total exposure effect (0.023) is attributed to school-related factors, indicating a significant influence of educational institutions on overall test scores. For math performance, γ_s is estimated at 0.288, and around 65% of the observed changes after relocation can be ascribed to differences in the quality of schools attended by children in different neighborhoods. Non-school factors, represented by γ_{ns} , contribute to a lesser extent account for approximately 35% of the total exposure effect. A different pattern is observed in Dutch proficiency. Here, school factors explain roughly 26% of the total change in Dutch proficiency across neighborhoods, while non-school factors contribute to 74% of the total exposure effect.

These findings reveal the role of schools in shaping educational outcomes following moves, particularly in cognitive development. The effect sizes associated with school-related factors emphasize the importance of investing in and improving educational resources and opportunities within neighborhoods, particularly for subjects such as math.

1.5 Conclusion

This study provides a comprehensive analysis of how neighborhoods affect educational opportunities, showing the crucial role of primary school performance in mediating children’s educational attainment. The findings emphasize the significant variation in educational outcomes across neighborhoods and illustrate that moving to a more advantageous neighborhood during childhood can substantially improve test scores and educational attainment.

By examining administrative data from the Netherlands, I obtain several key insights. First, Every additional year a child spends in a neighborhood with higher average test scores improves their test score rank by approximately 2.5% per year of childhood exposure, up to age 12. The longer a child is exposed to a better-performing neighborhood, the closer their test score converges to those of children who have always lived in that neighborhood.

Second, the mediation analysis demonstrates that primary school performance accounts for roughly 40% of the variation in educational attainment across neighborhoods. The decomposition analysis further supports this, revealing a correlation coefficient of 0.785 between primary school performance and educational attainment. This finding emphasizes improving primary school quality to bridge educational disparities.

The analysis at the smallest administrative unit, the *buurt*, yields similar

exposure effects, indicating that neighborhood impacts are highly localized. Even moving within a municipality can, therefore, significantly affect educational outcomes.

Decomposing the total exposure effect on educational attainment into school and non-school factors reveals that school quality contributes to approximately 39% of the improvement in test scores at the end of primary school. In math, 65% of the exposure effect is attributable to school quality, indicating the critical role of schools in reducing educational inequalities. Non-school neighborhood amenities account for the remaining 35%.

These findings imply that improving primary school quality in disadvantaged neighborhoods can help reduce educational disparities by shaping early cognitive development and promoting better educational trajectories. When it comes to non-school neighborhood effects, it is important to recognize that these are highly localized. As a result, it is important that policies narrowly target disadvantaged neighborhoods. More broadly, a better understanding of the impact of neighborhood exposure and the underlying mechanisms should allow policymakers to target resources better to reduce educational disparities and promote socioeconomic mobility across diverse communities.

1.6 Tables and Figures

Table 1.1: Summary Statistics

	Mean	Std. Dev.	Median	Number of Obs.
A. Total				
Native Parents	0.803	0.398	1	1,732,497
Maternal Schooling	16.03	3.173	16	1,107,395
Parental Income	49,228	45,570	42,941	1,731,210
CITO Std. Scores	535.4	9.802	537	1,732,497
Child Schooling	17.27	2.104	17	631,896
B. Permanent Residents				
Native Parents	0.808	0.394	1	1,334,623
Maternal Schooling	15.90	3.166	16	820,807
Parental Income	48,996	40,823	43,253	1,334,623
CITO Std. Scores	535.3	9.792	537	1,334,623
Child Schooling	17.24	2.078	17	485,443
C. Movers				
Native Parents	0.783	0.412	1	397,874
Maternal Schooling	16.41	3.160	17	286,014
Parental Income	50,007	58,754	41,682	397,874
CITO Std. Scores	535.7	9.829	537	397,874
Child Schooling	17.37	2.187	17	146,453

Notes: This table presents summary statistics for the analysis sample used in this study. The sample is restricted to children born between 1994 and 2007, whose moving histories are fully observable from age 1 onward, and who have valid CITO test scores from 2006 to 2019. The sample is divided into three groups: (1) The “Total Sample” includes all children with complete residential and educational records throughout the observation window, encompassing both movers and permanent residents. (2) “Permanent Residents” refers to children who continuously lived in the same neighborhood from age 1 to age 15, providing the baseline for neighborhood-level predicted outcomes as these children experienced constant exposure to a single neighborhood throughout their childhood. (3) “Movers” include children who relocated between neighborhoods at least once before age 15. This group is further analyzed based on their age at the time of the move, differentiating between those who moved at or before age 12 and those who moved after age 12. This differentiation allows for investigating the timing of neighborhood exposure on educational outcomes. Parental income is measured as the average disposable income of the household when the child is between 9 and 12 years old, adjusted for inflation and standardized across years. CITO test scores are standardized nationally and converted into percentile ranks to facilitate comparison across cohorts. All other variables, including parental education and child schooling years, are similarly standardized to account for variations in measurement across different data sources and years.

Table 1.2: Estimates of Test Score Exposure Effects

Dependent Variable	(1)	(2)	(3)	(4)	(5)	(6)
Overall CITO	Overall CITO	Overall CITO	Math	Dutch	Overall CITO	Overall CITO
Exposure Effects γ	0.0304*** (0.00704)	0.0274*** (0.00874)	0.0223*** (0.0103)	0.0348*** (0.00889)	0.0242*** (0.00450)	0.0250* (0.0137)
ρ	-0.0208 (0.0465)	-0.0433 (0.0545)	-0.0169 (0.0655)	-0.0744 (0.0540)	-0.0646* (0.0326)	-0.110 (0.0720)
Observations	250,159	231,112	231,112	231,112	106,102	106,102
Controls	no	yes	yes	yes	no	no
Within Year of Birth	no	no	no	no	yes	yes
Family Fixed Effects	no	no	no	no	no	yes


Note: This table reports estimates from the regression outlined in specification 1.3. Each column shows the results of regressing a child's test score rank at age 12 on the difference between the predicted ranks of permanent residents in the destination and origin neighborhoods, interacted with the child's age at the time of the move (m). The coefficient γ represents the effect of spending an additional year of childhood in a municipality where permanent residents have test scores one percentile higher. The coefficient ρ assesses whether selection effects vary over time. Column 1 presents the baseline results, while column 2 adds controls for parental education, immigration status, and gender. Columns 3 and 4 provide separate regressions by subject. Columns 5 and 6 focus on movers with siblings, and in these columns, the expected outcomes for neighborhoods are predicted using permanent residents within the same birth cohort, with column 6 additionally controlling for family fixed effects. Standard errors are reported in parentheses. *** $P < 0.01$, ** $P < 0.05$, * $P < 0.1$.

Table 1.3: School Shares

Dependent Variables	(1) CITO	(2) Math	(3) Dutch
Total Exposure Effects			
γ	0.0255*** (0.00655)	0.023*** (0.00635)	0.033*** (0.00757)
School and Non-School Components			
γ_s	0.0215** (0.00965)	0.0288*** (0.00885)	0.0191* (0.01101)
γ_{ns}	0.0263*** (0.00826)	0.013 (0.00792)	0.0338*** (0.00852)
School Shares(s^{school})	39%	65%	26%
Observations	43,640	43,640	43,640


Note: This table presents estimates based on the regression specifications outlined in section 1.4.3. The upper panel reports total exposure effects as specified in Equation (1.5). The middle panel provides the estimated coefficients for the school component as specified in Equation (1.7) and for the non-school component as specified in Equation (1.8). The bottom panel reports the school shares as specified in Equation (1.10). Results are provided for CITO test scores, math scores, and Dutch language scores. Standard errors are shown in parentheses. *** $P < 0.01$, ** $P < 0.05$, * $P < 0.1$.

Figure 1.1: Mean Test Scores for Children of Permanent Residents at the 25th Income Percentile




Note: This plot illustrates children's expected percentile ranks in the national distribution of test scores, conditional on having parents at the 25th income percentile. Darker-shaded colors correspond to higher outcomes for children, while gray indicates areas with fewer than 40 children, where data are insufficient to estimate outcomes. The sample includes all children in the analysis sample who are permanent residents (i.e., those families who do not move across municipalities before their children turn 16). To create these estimates, I first regress children's test score ranks on a constant, and their parents' family income ranks separately for each municipality. I then calculate the model-fitted test score rank for children having parents at the 25th income percentile in each municipality.

Figure 1.2: Childhood Exposure Effects on Test Scores at Age 12




Note: This figure shows the relationship between the coefficients b_m and the child's age at the time of moving (m), based on the semi-parametric model described in Equation 1.2. The analysis focuses on children's test scores at age 12, using a sample of children who moved exactly once across municipalities between 1994 and 2007. The b_m coefficients represent the effect of moving to a neighborhood where permanent residents score one percentile higher at a particular age (m). These coefficients are estimated by regressing a child's test score rank on the predicted difference between permanent residents' ranks in the origin and destination neighborhoods, interacted with the child's age at the time of the move. A dashed vertical line is placed at age 12, which corresponds to the age when most children take the CITO test. The best-fit lines, obtained from unweighted OLS regressions of the b_m coefficients, approximate the annual effects of childhood exposure for children who moved before age 12. The best-fit line for children who moved after age 12 tests whether selection effects vary over time. See details in section 1.3.

Figure 1.3: Exposure Effects Estimation Using Within-Municipality Moves



Note: This figure illustrates the relationship between the coefficients b_m and the age at which a child moves (m), using the semi-parametric model in Equation 1.2. The analysis evaluates children's test scores at age 12, and the sample includes children in the primary analysis who moved exactly once across different *buurten* while staying within the same municipality. The estimation approach is identical to the analysis of across-municipality moves. For details, see figure notes in figure 1.2 and in section 1.3.

Figure 1.4: Spatial Decay

Note: The plot visualizes 11 coefficients of interaction terms between the child's age at the time of moving and neighborhood outcomes. This analysis is based on a mover design incorporating the mean observed outcomes of permanent residents from the 10 closest *buurten* to the origin and destination *buurt*.

Appendices to Chapter 1

1.A Mediation Analysis of School Performance and Educational Attainment

This appendix provides a deeper exploration of the role of school performance, measured by CITO test scores, as a mediator in the relationship between neighborhood quality and long-term educational attainment. Specifically, I aim to assess whether improvements in school performance at age 12 can serve as a path by which neighborhood characteristics, such as total years of schooling, affect long-term educational attainment.

1.A.1 Effects on Educational Attainment


A key policy question is whether neighborhoods that enhance school performance at age 12 also positively influence long-term outcomes. In other words, does improving CITO scores translate into higher educational attainment? To investigate this possibility, I replace the outcome variable in the main analysis with the educational attainment of movers, measured in years of schooling, and estimate the following equation:

$$S_i = \sum_{m=0}^{15} (\alpha_m + \phi_m + \zeta_m p_i + b_m \Delta odp) + \epsilon_i \quad (1.A.1)$$

where S_i represents the educational attainment of movers, measured in years of schooling; b_m measures the effect of moving to a neighborhood where test scores are one percentile point higher on total years of schooling; Δodp captures the difference in neighborhood quality (as measured by test scores) between the origin and destination neighborhoods; p_i represents the child's parental income rank; and other terms such as α_m and ϕ_m capture age-specific effects and additional controls.

The results, illustrated in Figure 1.A.1, show a steady decline in b_m estimates with the age at move (m) for children younger than 12, which provides evidence of neighborhood exposure effects on long-term outcomes. The findings indicate that moving to a neighborhood with better school performance earlier in childhood leads to more significant gains in educational attainment. The magnitude of this effect suggests that a one standard deviation increase in CITO test scores leads to an approximate gain of 0.025 years of schooling. If a child was born in such a neighborhood, the cumulative improvement would be around 0.6 years, or roughly 35% of a standard deviation in schooling.

Figure 1.A.1: Reduced-Form Effects on Educational Attainment



Note: This figure shows the relationship between the coefficients b_m and the child's age at the time of moving (m), using the semi-parametric model outlined in Equation 1.A.1. The analysis evaluates children's educational attainment up to age 24, and the sample consists of children who moved exactly once between 1994 and 1998. The b_m coefficients represent the effect of relocating to a neighborhood where permanent residents achieve CITO test scores that are one percentile higher at a given age (m). These coefficients are derived by regressing a child's years of schooling (S_i) on $\Delta odp = T_{op} - T_{dp}$, which is the difference in predicted ranks between permanent residents of origin and destination neighborhoods, interacted with the child's age at the time of the move (m). The best-fit lines are calculated using unweighted OLS regressions of the b_m coefficients on m . The slopes approximate annual childhood exposure effects for children who moved at ages $m \leq 12$.

1.A.2 Mediation Analysis

The mediation analysis focuses on breaking down the total effect of parental income on long-term educational outcomes into direct and indirect effects. In this context, *indirect effects* refer to the role of primary school performance (CITO test scores) as a mediator that links parental income to long-term educational attainment. In contrast, *direct effects* capture the influence of parental income on educational outcomes that bypass school performance, reflecting other pathways such as family resources, social capital, or access to information.

This section aims to quantify the extent to which primary school performance can explain the relationship between parental income and children's educational attainment. This will provide insights into whether school performance serves as an important mechanism or pathway by which socioeconomic status is transmitted across generations.

The following system of equations models the relationships between parental income, school performance, and educational attainment:

1. Reduced-form transmission of parental income to educational attainment

$$S_i = \phi_c + \theta_c p_i + \xi_i \quad (1.A.2)$$

In this equation, S_i represents the child's long-term educational attainment; p_i is the child's parental income rank; θ_c measures the total effect of parental income on children's educational attainment, which includes both direct and indirect effects; and ξ_i captures unobserved factors.

2. Transmission of parental income through school performance

$$S_i = \kappa_c + \lambda_c T_{ic} + \mu_c p_i + u_i \quad (1.A.3)$$

In this equation, T_{ic} represents the child's primary school performance (CITO test scores); λ_c captures the extent to which primary school performance contributes to the child's long-term educational attainment; μ_c measures the direct effect of parental income on educational attainment, independent of school performance; and u_i captures unobserved factors specific to the child.

The key parameter of interest here is λ_c , which quantifies the impact of school performance on long-term outcomes, holding parental income constant. This allows us to estimate how much parental income's total effect on educational attainment operates through school performance.

By substituting Equation 1.A.3 into Equation 1.A.2, the total effect of parental income rank on long-term educational attainment becomes:

$$S_i = (\kappa_c + \lambda_c \alpha_c) + (\lambda_c \pi_c + \mu_c) p_i + (\lambda_c \epsilon_i + u_i) \quad (1.A.4)$$


Thus, the reduced-form transmission of parental income to children's educational attainment can be expressed as a combination of direct and indirect effects:

$$\theta_c = \lambda_c \pi_c + \mu_c \quad (1.A.5)$$

1.A.3 Results

The mediation analysis results suggest that, on average, 40 percent of the variation in long-term educational attainment across neighborhoods can be attributed to differences in primary school performance. This finding indicates that school performance is a meaningful mediator in the transmission of socioeconomic status from parents to children.

Figure 1.A.2: Correlation Between Parental Income Rank and Educational Transmission



Note: This figure illustrates the relationship between coefficients θ (parental income rank to schooling years) and π (parental income rank to test scores). Each point represents a municipality. The fitted line and confidence interval show an unweighted regression of θ on π , and the text indicates the correlation coefficient between the two coefficients across municipalities.

As shown in Figure 1.A.2, there is a strong positive correlation between test score transmission (π) and educational attainment transmission (θ), with a correlation

coefficient of 0.785. This result suggests that improvements in CITO test scores are strongly linked to gains in long-term educational attainment. The decomposition analysis further reveals that primary school performance accounts for approximately 40% of the variation in educational attainment across neighborhoods, emphasizing the critical role of early school performance in shaping long-term outcomes.

These findings underscore the importance of primary school performance as a key mechanism through which neighborhoods influence children's future outcomes. By improving primary education in disadvantaged neighborhoods, policymakers can potentially reduce educational disparities and promote upward mobility. Furthermore, the mediation analysis highlights that while school performance is a critical pathway, other neighborhood factors also play a role in shaping long-term outcomes, suggesting the need for a multifaceted approach to improving childhood environments.

Chapter 2

Improving Equality of Opportunity at School Entry? Primary Education Enrollment in the Netherlands

with Antonio Ciccone

2.1 Introduction

In most countries, the default policy is to enroll children in primary education by date of birth. However, in the Netherlands, elementary schools are explicitly discouraged from applying this policy. Instead, the institution supervising schools recommends that children are enrolled in primary education solely based on their developmental skills (Huizenga and Damstra, 2016a).

We analyze if and how enrolling children in primary education based on their skills affects equality of opportunity in education. Compared to enrollment by date of birth, enrollment policies based on children's skills can in principle improve equality of opportunity along two interlinked dimensions. First, enrolling children by date of birth implies that their skills at the start of primary education end up reflecting randomness in date of birth. That is, children's skills at school entry are partially determined by a birth-date lottery. For example, children born a single day apart may be enrolled one year apart in age-related skills. Enrolling children based on their skills may reduce the effect of randomness in date of birth on skills at school entry and on later education outcomes (e.g. Bedard and Dhuey (2006)). Second, enrollment in primary education by date of birth fully preserves differences in skill endowments at school entry. If enrollment were based on skills, children with low skill endowments would be allowed more time to develop before enrolling in primary education. As a result, children with skill endowments at the bottom of the distribution may achieve better education outcomes.

To understand how enrolling children in primary education based on their skills affects equality of opportunity, we develop and estimate a model of primary education enrollment in the Netherlands and the relationship between skills at the start of primary education and outcomes in primary and secondary education. We then compare the Dutch enrollment policy with two alternatives: enrollment based solely on children's dates of birth and enrollment based solely on children's skills. Our counterfactual policy analysis holds average age at enrollment in primary education constant. The key difference between the alternative primary education enrollment policies is who gets more time to develop before enrolling in primary education.

Our analysis yields two main results. First, compared to enrollment by date of birth, the policy in the Netherlands improves equality of opportunity in education by reducing differences in primary and secondary education outcomes across birthdays *and* by improving education outcomes of children with skill endowments at the bottom of the distribution. Second, notwithstanding the recommendation of their supervisory institution, schools in the Netherlands do not enroll children in primary education solely based on their

developmental skills. We find that doing so would nearly eliminate differences in primary and secondary education outcomes across birthdays and further raise outcomes of children with low skill endowments.

A directly observable feature of primary education enrollment in the Netherlands is that, compared to enrollment by date of birth, there are smaller differences in average age at the start of primary education across birthdays. For example, the largest difference across birthdays in the Netherlands is 6 months, while children born on different days are up to one year apart in age at school entry when enrollment is by date of birth. As a result, the Dutch enrollment policy implies smaller differences in education outcomes across birthdays than enrollment by date of birth. For example, we find that the largest difference across birthdays in average end-of-primary-school test scores is reduced by more than half; the largest difference across birthdays in retention rates in primary education is reduced by 9 percentage points; and the largest difference across birthdays in the share of children attending the most academic track in secondary school is reduced by 6 percentage points.

Another feature of primary education enrollment in the Netherlands, again compared to enrollment by date of birth, is that the smaller differences in age at the start of primary education across birthdays are achieved at the cost of larger differences in age at the start of primary education among children with the same birthday. How this affects differences in children's skills at the start of primary education and differences in later education outcomes depends on which children get more time to develop before starting their primary education. We find that in the Netherlands it is children with low skill endowments who are relatively older when they start their primary education.¹ As a result, compared to enrollment by date of birth, children in the Netherlands with low skill endowments achieve higher test scores at the end of primary education; are less likely to be retained in primary education; and are more likely to attend the most academic track in secondary school. Our counterfactual policy analysis also yields that education outcomes of children with low skill endowments could be further improved by a policy that admits children to primary education solely based on their skills.

To the best of our knowledge, there is no previous research evaluating alternatives to the primary education enrollment policy based on date of birth. There is an extensive literature documenting how enrolling children based on date of birth

¹By contrast, in the US it is children with relatively high skill endowments who are relatively older at the start of primary education, see Bassok and Reardon (2013) for example. This appears to be an exception. In Appendix 2.B, Figure 1A, we analyze the relationship between maternal education and age at the start of primary education for a wide range of countries using data from the Progress in International Reading Literacy Study (Foy, 2013, 2017). In most countries, children with less educated mothers are relatively older at the start of their primary education.

affects outcomes in primary and secondary education, see Bedard and Dhuey (2006); McEwan and Shapiro (2008); Mühlenweg and Puhani (2010); Crawford, Dearden, and Greaves (2013); Schneeweis and Zweimüller (2014); Cascio and Schanzenbach (2016); Dhuey et al. (2019); Oosterbeek, ter Meulen, and van der Klaauw (2021). The main results from our perspective are that children who are younger at the start of primary education are more likely to be retained in a grade later; do worse on cognitive tests; and are less likely to attend the most academic track when there is tracking in secondary school.² Our goal is to analyze how different primary education enrollment policies affect the distribution of these outcomes across children with different birthdays and skill endowments.³

Our analysis examines the consequences of primary education enrollment policies for outcomes in primary and secondary education. Later-life outcomes, like completed years of education and earnings, are not available as the children in our data are still too young. To assess the consequences for later-life outcomes, we must therefore build on the existing literature. Our evaluation of different enrollment policies keeps average age at the start of primary education constant. As a result, studies that find the effect of age at the start of primary education on later outcomes to be linear—and independent of skill endowments—would imply the same aggregate outcomes for the different enrollment policies we consider.⁴ In this case, the increase in equality of opportunity in primary and secondary education resulting from the primary education enrollment policy based on skills would come without a cost in terms of aggregate later outcomes.

Compared to enrollment by date of birth, a key feature of the primary education enrollment policy based on skills is that children with low skill endowments are

²Children who are younger at the start of schooling are also more likely to be diagnosed with a learning disability; have lower self-assessed ability, self-confidence, and mental health; participate less in leadership activities in school even when differences in age and weight are accounted for; have poorer quality relationships with classmates and teachers; are more likely to be victimized; and have higher suicide rates; see Dhuey and Lipscomb (2008); Elder and Lubotsky (2009); Mühlenweg (2010); Crawford, Dearden, and Greaves (2013); Chen, Fortin, and Phipps (2015); Page, Sarkar, and Silva-Goncalves (2019); Yamaguchi, Ito, and Nakamuro (2023); Matsubayashi and Ueda (2015) and references therein.

³An interesting question is how enrollment by date of birth and the alternative enrollment policies we consider compare regarding the non-cognitive outcomes documented in the literature cited in the previous footnote. As the alternative enrollment policies improve skills at the start of primary education in the lower tail of the distribution, they should reduce diagnoses of learning disability. For the other outcomes, a key question—not yet addressed by the literature—is to what extent they are driven by physical maturity and to what extent by cognitive maturity (only the analysis of Dhuey and Lipscomb (2008) can account for physical and cognitive maturity). If both physical and cognitive maturity play a role, enrollment by date of birth implies that the youngest children at the start of primary education are subject to two coincident disadvantages that may aggravate each other. This coincidence would be reduced by an enrollment policy based on skills.

⁴For linear estimates of age at the start of primary education on later-life outcomes in Germany, Japan, Belgium, and Spain, see Dustmann and Schönberg (2016); Kawaguchi (2011); Oosterbeek, ter Meulen, and van der Klaauw (2021); Valdés and Requena (2024).

enrolled in primary education somewhat older than children with high skill endowments. To assess the effect of this policy feature on later-life education and economic outcomes, we require estimates of the effects of age at the start of primary education that account for individual skill endowments. Fredriksson and Öckert (2014) provide such estimates based on data for around 2 million individuals born in Sweden between 1935 and 1955. Their analysis distinguishes between individuals with low parental education (both parents with education levels below a threshold) and high parental education. For completed years of education, they find a positive effect for both groups and a stronger effect for low parental education. For earnings age 25 to 54, Fredriksson and Öckert find a positive effect for low parental education and a negative effect for high parental education. The reason is that the positive effect of more education on earnings outweighs the negative effect of less labor-market experience for low parental education, while the opposite is the case for high parental education. The findings of Fredriksson and Öckert therefore imply that the primary education enrollment policy based on skills would raise aggregate completed years of education and aggregate earnings.⁵

2.2 Kindergarten and Primary Education in the Netherlands: Background and a First Look at the Data

We start with some background on kindergarten and primary education in the Netherlands and then take a first look at the data.

2.2.1 Kindergarten and Primary Education in the Netherlands

Kindergarten education and primary education in the Netherlands are integrated into the same educational institution, the so-called *basisschool*. The premise for the integrated *basisschool*, as formulated in the 1981 Basic Education Act, is that “to promote an uninterrupted development of pupils, it is desirable to combine

⁵Black, Devereux, and Salvanes (2011) estimate heterogeneous effects of age at the start of primary education on completed years of education and on earnings at age 24 and age 35 in Norway. Their heterogeneity analysis groups individuals based on regressions predicting years of education (for the effect on years of education) and earnings at age 35 (when they examine earnings at age 24 and 35) given three family background characteristics (mother’s education, family size, and birth order). They find no evidence for statistically significant heterogeneous effects on education for men, but a statistically significant positive effect of age at the start of primary education on completed years of education for women in the bottom 75 percent of their grouping. For earnings age 35, they do not find any statistically significant heterogeneity for women, but a statistically significant negative effect for men in the bottom 25 percent of their grouping.

the separate forms of pre-school and ordinary primary education into a form of education aimed at a continuous development process.” Children can enter *basisschool*, which we will refer to as elementary school, the day after they turn 4 years old and must enter by age 5. The education program children go through is organized in 8 so-called *groeps*. In *groep 1* and *groep 2*, children receive a typical kindergarten education. There is a focus on learning through play, arts and crafts, social skills, gross and fine motor skills, independence, and getting accustomed to routines. Usually *groep 1* and *groep 2* are combined in the same class, called *kleuterklas*. In *groeps 3* through 8, which can be thought of as grade levels, children receive a typical primary education. There is a focus on basic academic skills like reading, writing, and doing mathematical calculations. After elementary school, children start their secondary education. Secondary education in the Netherlands is structured into 3 school tracks: pre-vocational secondary education; senior general secondary education; and pre-university education (which we will refer to as the most academic secondary school track).

A key decision made in elementary school is how much time children spend in kindergarten education before enrolling in primary education. The principle in the 1981 Basic Education Act guiding this decision is that “education is designed to enable students to engage in an uninterrupted developmental process. It is tailored to the progress of the pupils’ development.” The institution in charge of ensuring this principle, the Inspectorate of Education in the Ministry of Education, Culture and Science, interprets it as the obligation to “make reasoned considerations about progression to the next *groep* based on the child’s development.” The Inspectorate “advocates that progression to *groep 3* should be based solely on developmental data” and explicitly discourages progression from kindergarten to primary education based solely on children’s birthdays (Huizenga and Damstra, 2016a).⁶ Elementary schools are allowed to establish their own protocol for the transition from kindergarten education (*groep 1* and *groep 2*) to primary education (*groep 3*). These so-called transition protocols generally include criteria based on social-emotional skills, the development of motor skills, task orientation, concentration span, and reading and math skills (Mulder et al., 2016).

2.2.2 Age at the Start of Primary Education

In Figure 1 we take a first look at the data. The data refers to all children born in the Netherlands between 2002 and December 2005 for whom we have the data required for our empirical analysis. See Appendix 2.A for details on

⁶Nevertheless, elementary schools are known to still put some weight on children’s birthdates when enrolling children in primary education (Huizenga and Damstra, 2016b).

the data and descriptive statistics. In Figure 2.1A we look at whether children leave kindergarten to start their primary education in the calendar year they turn 5, 6, or 7 years old.⁷ The share of children starting their primary education at different ages is plotted against children's birthdays.⁸ Over 90 percent of children born in January start their primary education in the calendar year they turn 6; the remaining children start in the calendar year they turn 5. As we consider children born later in the year, there is a gradual increase in the share of children starting in the calendar year they turn 7. On October 1—which used to be the cutoff date for enrollment in primary education before the 1981 Basic Education Act—there is a discontinuity in the relationship between children's birthdays and when they start their primary education. The share admitted in the calendar year they turn 7 jumps up by around 20 percentage points, with a corresponding drop in the share of children admitted age 6. This is an indication that the primary education enrollment policy of elementary schools still puts some weight on children's birthdates. After October 1, the share of children admitted age 7 increases more steeply than before October 1 as we consider children born later in the year. At the end of the calendar year, there is a second discontinuity in the relationship between children's birthdays and when they start their primary education. On December 31, 80 percent of children start their primary education in the calendar year they turn 7. If date of birth played no role for when children start their primary education, this should imply that around 80 percent of children born one day later—the first day of the next calendar year—should start their primary education in the calendar year they turn 6. However, the share of children born on January 1 starting their primary education in the calendar year they turn 6 is around 90 percent. The January 1 discontinuity is also visible in the share of children born on January 1 starting their primary education in the calendar year they turn 5. This share is around 10 percent. If date of birth played no role for when children start their primary education, around 10 percent of children born on December 31 should start their primary education in the calendar year they turn 6. However, this share is around 20 percent.

Figure 2.1B illustrates the role of family background for whether children start their primary education in the calendar year they turn 5, 6, or 7. In particular, we split children into 2 groups, those with parental income above the median and those with parental income strictly below the median. Starting from birthdays around April, the share of children enrolling in primary education at age 6 is always lower for children with parental income below the median. This means that children

⁷We have dropped the 0.14 percent of children who started their primary education age 4 or age 8.

⁸Children born on the 29th of February 2004 are reassigned to 28th of February 2004.

with parental income below the median stay in kindergarten education somewhat longer and start primary education somewhat older than children with parental income above the median. Results are very similar when we split the sample by maternal education instead of family income.⁹

Figure 2.1C plots the birthdays of children against average age on the day they start their primary education (we assume that the school year starts on the 1st of September). Children born on January 1 are on average 6.6 years old when they start their primary education. As we consider children born later in the year, the average age at the start of primary education decreases almost linearly to 6.1 years. There is a discontinuity on October 1, where average age at the start of primary education jumps from 6.1 to above 6.2 years. After October 1, average age at the start of primary education increases from around 6.2 years to around 6.5 years. This is around 0.1 years less than the average age at the start of primary education of children born on January 1. The October 1 and January 1 discontinuities in the relationship between birthdays and average age at the start of primary education are therefore of similar size. The maximum difference across birthdays in age at the start of primary education, as shown in Figure 2.1C, is around 0.5 years. For comparison, enrollment policies solely based on children's date of birth imply a maximum difference across birthdays in age at the start of primary education of one year.

2.2.3 Primary and Secondary Education Outcomes

In Figure 2 we examine the relationship between children's outcomes in primary and secondary education and their birthdays. The three education outcomes available are grade retention in primary education, test scores at the end of primary education, and the secondary school education track children end up attending. The test scores are from national standardized tests. We transform children's raw test scores into z-scores based on the mean and the standard deviation of all students taking the test in the same year.

In Figure 2.2A we look at the share of children not retained in any grade during their primary education. The V-shape we obtain as a result mirrors the V-shape in average age at the start of primary education in Figure 2.1C. Hence, across birthdays, children who start their primary education older are more likely to never be retained during their primary education.¹⁰ In Figure 2.2B we consider the average z-score of the standardized test at the end of primary education.

⁹See Appendix 2.B, Figure 1B.

¹⁰The two-way relationship between age at the start of primary education and the share of children never retained across birthdays is close to linear as we show in Appendix 2.B, Figure 2A.

We again obtain a V-shape that mirrors the V-shape in average age at the start of primary education in Figure 2.1C. Therefore, across birthdays, children who start their primary education older do better in the standardized test at the end of primary education.¹¹ Finally, in Figure 2.2C we look at the share of children going on to the most academic secondary school track (the so-called pre-university track). This yields the same V-shape we have found for never retained in primary school and test scores at the end of primary school.¹² Hence, all three primary and secondary education outcomes we observe point in the direction of a positive relationship between age at the start of primary education across birthdays and education outcomes.

2.3 The Model

We start with a model of elementary schools in the Netherlands where all children receiving a kindergarten education (in *groep 1* and *groep 2* of elementary school) are admitted to start their primary education (in *groep 3* of elementary school) in the calendar year they turn 6 or 7 years old. We then extend the model to account for the around 1 percent of children in our data who start their primary education in the calendar year they turn 5.

Elementary schools Children start elementary school when they turn 4 years old. There, they first receive a kindergarten education and then a primary education. All children in kindergarten education are admitted to start their primary education in the calendar year they turn 6 or 7 years old. Whether children start age 6 or age 7—and hence how much time they spend in kindergarten education—depends on their birthdays and skills as specified by the primary education enrollment policy of elementary schools. Primary education is organized in six grades. After elementary school, children enroll in different secondary school tracks partly depending on their performance in a standardized test.

Skill accumulation in kindergarten education Children start elementary school the day they turn 4 years old. We denote their skill endowments at this point by $\alpha + e$ where α denotes the average skill endowment and e the deviation from the average skill endowment. In kindergarten education, children accumulate skills linearly in time at rate $\beta + \lambda e$. As $e = 0$ for the average child, the average

¹¹The two-way relationship between age at the start of primary education and the standardized test score is also close to linear as we show in Appendix 2.B, Figure 2B.

¹²The two-way relationship between age at the start of primary education and the share of children attending the most academic track is again close to linear, as we show in Appendix 2.B, Figure 2C.

child accumulates skills at rate β . When $\lambda \neq 0$, the rate of skill accumulation in kindergarten education depends on the child's skill endowment. If child i has received a kindergarten education for l^K years, her skills are

$$SkillEoK(l^K, e) = \alpha + e + (\beta + \lambda e)l^K. \quad (2.1)$$

Enrollment in primary education On the first day of each school year, all children in kindergarten education turning 7 in the calendar year are admitted to primary education. Whether kindergarten children turning 6 in the calendar year are admitted to primary education depends on their birthdays and skills as specified by the primary education enrollment policy. Children's birthdays may matter for two different reasons. First, because the enrollment policy directly links enrollment to children's birthdays. Second, because children's birthdays are related to their age, and hence their skills, at the point the enrollment decision is made.

We denote the birthday of children by x and the first day of the school year by f , both measured as a fraction of a calendar year. On the first day of the school year, kindergarten children turning 6 in the calendar year are $6 + f - x$ years old. As they started elementary school at age 4, these children have spent $l^K = 2 + f - x$ years in kindergarten education and, using (2.1), have skills

$$skill_{\text{day-}f}^{\text{year-turn-6}}[x, e] = \alpha + e + (\beta + \lambda e)(2 + f - x). \quad (2.2)$$

The primary education enrollment policy of elementary schools $\mathcal{E}$ takes the form of birthday-specific minimum-skill threshold. Elementary schools enroll children in primary education in the calendar year they turn 6 if their skills are above a birthday-specific minimum-skill threshold for enrollment $minskill^6[x]$

$$skill_{\text{day-}f}^{\text{year-turn-6}}[x, e] > minskill^6[x]. \quad (2.3)$$

Children's birthdays enter both the left-hand and right-hand side of this inequality. On the left-hand side because of the link between birthdays and skills at the point in time when the enrollment decision is made (first day of the school year); On the right-hand side because the enrollment policy may specify annual cutoff dates. For example, children born after day q may never be admitted in the calendar year they turn 6 (in this case we have $minskill[x] = \infty$ if $x > q$). We assume throughout that $1 + \lambda(2 + f - x) > 0$, which implies that children are admitted to primary education age 6 if and only if their skill endowments exceed a threshold that depends on their birthday.

Skills at the end of kindergarten education We denote the time a child born at x with skill endowment e spends in kindergarten education given the primary-school enrollment policy $\mathcal{E}$ by $l^K[x, e; \mathcal{E}]$. Using (2.1), children's skills at the end of kindergarten education can then be written as a function of their birthdays, skill endowments, and the primary education enrollment policy of elementary schools

$$SkillEoK[x, e; \mathcal{E}] = \alpha + e + (\beta + \lambda e)l^K[x, e; \mathcal{E}]. \quad (2.4)$$

As children are admitted in the calendar year they turn 6 or 7, time in kindergarten can be written as

$$l^K[x, e; \mathcal{E}] = 2 + f - x_i + D^7[x, e; \mathcal{E}] \quad (2.5)$$

where $D^7[x, e; \mathcal{E}]$ is an indicator variable that takes the value 1 if and only if the child enrolls in primary education in the calendar year s/he turned 7. As this is the case if and only if (2.2) holds, we obtain

$$D^7[x, e; \mathcal{E}] = I\left[e < \frac{minskill^6[x] - \alpha - \beta(2 + f - x)}{1 + \lambda(2 + f - x)}\right]. \quad (2.6)$$

where $I[\cdot]$ is an indicator function and the term in square brackets is the condition for enrollment in the calendar year children turn 7.

Outcomes in primary and secondary education Our goal is to examine whether the skills that children have at the end of their kindergarten education affect outcomes in primary education. We will therefore examine the effect of $SkillEoK[x, e; \mathcal{E}]$ in (2.4) on retention in primary education, test scores at the end of primary education, and the secondary school track that children end up attending.

Adding primary education enrollment age 5 So far we assumed that children enroll in primary education in the calendar year they turn 6 or 7. We now extend the model to account for the (small) share of children in the Netherlands starting their primary education in the calendar year they turn 5. As a result, instead of (2.5), we have that the time a child with birthday x and skill endowment e will spend in kindergarten is

$$l^K[x, e; \mathcal{E}] = 2 + f - x_i + D^7[x, e; \mathcal{E}] - D^5[x, e; \mathcal{E}]. \quad (2.7)$$

D^7 continues to capture if a child enrolls in primary education in the calendar year s/he turns 7 ($D^7 = 1$) or not ($D^7 = 0$). D^5 is new and captures if a child enrolls

in primary education in the calendar year s/he turns 5 ($D^5 = 1$) or not ($D^5 = 0$). The skills of children at the end of kindergarten education can then be obtained by substituting (2.7) in (2.4).

To determine $D^5[x, e; \mathcal{E}]$, we need to detail how elementary schools decide whether to enroll children in primary education in the calendar year they turn 5. We assume that elementary schools make this decision based on the skills of these children on the first day of the school year. These skills can be obtained analogously to (2.2) as

$$skill_{\text{day-}f}^{\text{year-turn-5}}[x, e] = \alpha + (\beta + \lambda e)(1 + f - x) + e \quad (2.8)$$

where we assume that $1 + \lambda(1 + f - x) > 0$ and f continues to denote the day the school year starts. If the skills in (2.8) are above the age-5 admission threshold $minskill^5[x]$ specified by the primary education enrollment policy $\mathcal{E}$ of elementary schools, children will be admitted to primary education in the calendar year they turn 5. Hence

$$D^5[x, e; \mathcal{E}] = I\left[e \geq \frac{minskill^5[x] - \alpha - \beta(1 + f - x)}{1 + \lambda(1 + f - x)}\right] = I[e \geq \pi^5[x; \mathcal{E}]] \quad (2.9)$$

where $I[\cdot]$ continues to denote the indicator function taking the value 1 if and only if the condition in square brackets is satisfied and $\pi^5[x; \mathcal{E}]$ is defined implicitly by the second equality.

Children not admitted in the calendar year they turn 5 will be considered for admission to primary education on day f of the calendar year in which they turn 6. On that day, children with skills above the age-6-admission threshold that elementary schools apply, $minskill^6[x]$, will be admitted to primary education. All other children will be admitted to primary education in the calendar year they turn 7. Hence

$$D^7[x, e; \mathcal{E}] = I\left[e < \frac{minskill^6[x] - \alpha - \beta(2 + f - x)}{1 + \lambda(2 + f - x)}\right] = I[e \leq \pi^6[x; \mathcal{E}]] \quad (2.10)$$

where $\pi^6[x; \mathcal{E}]$ is defined implicitly by the second equality.

We assume that $\pi^6[x; \mathcal{E}] < \pi^5[x; \mathcal{E}]$. Hence, children with birthday x enroll in primary education in the calendar year they turn 5 if their skill endowments are above $\pi^5[x; \mathcal{E}]$; in the calendar year they turn 6 if their skill endowments are below $\pi^5[x; \mathcal{E}]$ but above $\pi^6[x; \mathcal{E}]$; and in the calendar year they turn 7 if their skill endowments are below $\pi^6[x; \mathcal{E}]$.

2.4 Age at Enrollment in Primary Education

We start by estimating the model for age at enrollment in primary education. Specifically, we want to understand the determinants of children enrolling in primary education in the calendar year they turn 5, 6, or 7 years old. To do so, we need to specify the policy that elementary schools use for enrollment in primary education and how children's skill endowments depend on their observable characteristics.

2.4.1 Skill Endowments, Skill Shocks, and Enrollment Policies

Children's skill endowments when they start their kindergarten education at age 4 are $\alpha + e_i$ where α is the average skill endowment and e_i the deviation of the skill endowment of child i from the average. Skill endowments reflect the effect of family background characteristics W_i and of skill endowment shocks v_i

$$e_i = \rho W_i + \sigma v_i \quad (2.11)$$

where ρW_i captures the effect of family background on skill endowments and v_i is an independent skill endowment shock that follows a standard logistic distribution;¹³ σ scales the effect of the standardized skill shock. Family background characteristics will be measured as deviations from the average across all children.

The primary education enrollment policy of elementary schools $\mathcal{E}$ consists of the birthday-specific minimum-skill thresholds $minskill^j[x]$ for $j = 5, 6$ in (2.9) and (2.10). We use two alternative functional forms for $minskill^j[x]$. The first functional form assumes that elementary schools may use different minimum-skill thresholds for age-5 and age-6 enrollment in primary education depending on whether children were born before or after October 1. However, children born before October 1 are subject to the same thresholds for age-5 admission and age-6 admission (independently of their birthdays) and the same holds for children born after October 1. We chose October 1 as the date where minimum-skill thresholds may change as this allows us to capture the upward jump in age-7 admission in the data (see Figures 2.1A and 2.1C). Hence, our first specification for the primary education enrollment policy is that for $j = 5, 6$

$$minskill^j[x] = minskill^j + \theta I[x \geq q] \quad (2.12)$$

¹³This will result in an ordered logit model for whether children enroll in primary education in the calendar year they turn 5, 6, or 7. We also considered a standard normal distribution for v , which results in an ordered probit model for whether children enroll in primary education in the calendar year they turn 5, 6, or 7. Overall, we found the ordered logit model to fit the data somewhat better than the ordered probit model.

where $minskill^j$ are constants, q is set to October 1, and θ is the difference in the minimum-skill thresholds used for children born after and before October 1. The policy in (2.12) is therefore defined by 3 parameters: $minskill^5$, $minskill^6$, and θ .¹⁴

We also consider a second specification for the primary education enrollment policy of elementary schools. This specification differs from our first specification in that the minimum-skill threshold for primary education enrollment may depend on children's birthdays for children born after October 1. As we will show below, the data seems to indicate that this is the case. The specification we use is that for $j = 5, 6$

$$minskill^j[x] = minskill^j + \theta I[x \geq q](1 + \kappa(x - q)). \quad (2.13)$$

For $\kappa = 0$, we are back to the primary education enrollment policy in (2.12). When $\kappa \neq 0$, the minimum-skill threshold for enrollment in primary education depends on children's birthdays for children born after October 1. A strictly positive κ , for example, implies that elementary schools apply a higher minimum-skill threshold for children born later in the calendar year. The primary education admission policy in (2.13) is defined by four parameters: $minskill^5$, $minskill^6$, θ , and κ .¹⁵ The implications of this primary education enrollment policy for children's age at enrollment in primary education as a function of their birthdays and skill endowments are illustrated in Appendix 2.B, Figure 3.

2.4.2 Ordered Logit Model for Age at Enrollment in Primary Education

We can now obtain the probabilities that children enroll in primary education in the calendar year they turn 5, 6, or 7. Using (2.9) we get the probability of admission age 5 as a function of children's birthdays x , family background W , and

¹⁴This specification also implies a change in the minimum-skill thresholds on January 1 (necessary to capture the jumps on January 1 in Figures 2.1A and 2.1C). The size of the change is a function of the policy parameters in (2.12). For example, consider children born on the 31st of December 2002. These children will be enrolled in primary education in the school year starting in 2008 only if their skills as defined in (2.8) are above $minskill^6 + \theta$. For children born one day later, on the 1st of January 2003, to be enrolled in primary education in the school year starting in 2008, their skills have to be above $minskill^5$. Hence, the minimum-skill requirement for enrollment in primary education in 2008 increases on January 1 if $minskill^5 > minskill^6 + \theta$.

¹⁵Just as in the case of the first specification for the primary education enrollment policy, the change in the minimum-skill thresholds on January 1 is a function of the policy parameters in (2.13).

the primary education enrollment policy $\mathcal{E}$

$$Prob^5(x, W; \mathcal{E}) = \mathbb{E}(D^5[x_i, e_i; \mathcal{E}] | x, W; \mathcal{E}) = 1 - F[\omega^5[x, W; \mathcal{E}]] \quad (2.14)$$

where $F[\cdot]$ is the standard logistic cumulative distribution function. The last equality uses $e = \rho W + \sigma v$, that v follows a standard logistic distribution, and the definition

$$\omega^5[x, W; \mathcal{E}] = \pi^5[x; \mathcal{E}]/\sigma - W\rho/\sigma \quad (2.15)$$

where $\pi^5[x; \mathcal{E}]$ is implicitly defined in (2.9). Similarly, the probability of admission age 7 is

$$Prob^7(x, W; \mathcal{E}) = \mathbb{E}(D^7[x_i, e_i; \mathcal{E}] | x, W; \mathcal{E}) = F[\omega^6[x, W; \mathcal{E}]] \quad (2.16)$$

where

$$\omega^6[x, W; \mathcal{E}] = \pi^6[x; \mathcal{E}]/\sigma - W\rho/\sigma \quad (2.17)$$

with $\pi^6[x; \mathcal{E}]$ is implicitly defined in (2.10). The probability of admission age 6 follows from the fact that all children are admitted age 5, 6, or 7. Our assumption that $\pi^6[x; \mathcal{E}] < \pi^5[x; \mathcal{E}]$ implies that (2.14)-(2.17) constitute a standard ordered logit regression model (Amemiya, 1985).

2.4.3 Evaluation of the Primary Education Enrollment Model

Figure 3 evaluates the ordered logit model for age at enrollment in primary education for all possible birthdays. We compare model prediction and data for the share of children starting primary education at different ages. Age 5, 6, 7 denotes children enrolling in primary education in the calendar year they turn 5, 6, or 7 years old respectively.

Figure 2.3A considers the baseline logit model. Elementary schools use the primary education enrollment policy in (2.13) with skill accumulation in kindergarten education depending on children's skill endowments ($\lambda \neq 0$). The model has a total of eight parameters.¹⁶ The model fits the data well.¹⁷ The only discernible

¹⁶There are four parameters defining the primary education enrollment policy and five parameters for children's skills in kindergarten education. The ordered logit model has eight parameters only as the average skill endowment of children is not identified separately from the minimum-skill parameters in (2.12). For the parameter estimates see the first column of Table 2.1.

¹⁷We obtain the model predictions by drawing a skill endowment shock v from a standard logistic distribution for every child in our sample, predicting the child's age at enrollment in primary education using the estimated ordered logit model, and calculating the share of children

difference between model prediction and data across the 365 possible birthdays is that the share of children admitted in the calendar year they turn 6 (turn 5) is overestimated (underestimated) by around 1 percentage point for children born in January.

Figures 2.3B and 2.3C evaluates the model when we consider a less general specification. In Figure 2.3B we assume that the rate of skill accumulation in kindergarten education does not depend on children's skill endowments ($\lambda = 0$). This does not affect the fit of the model. Figure 2.3C also assumes that the primary education enrollment policy applies the same minimum-skill threshold for all children born after October 1 ($\kappa = 0$). Now the model fits worse for children born in fall.

Figure 4 examines how well our baseline logit model explains the data for exact age at the start of primary education (assuming that the school year starts on September 1). Figure 2.4A compares predicted average age at the start of primary education with the data for all possible birthdays.¹⁸ Figure 2.4B relates the predicted variance of age at the start of primary education among children born on the same day for every birthday with the variance in the data.¹⁹ Figure 2.4C compares model prediction and data for average age at the start of primary education separately for children with parental income above and below median income.²⁰ The model fits the data well in all three dimensions.

Figure 2.5A evaluates our baseline ordered logit model using the individual-level data for all children (326,416 children). We do so using a binscatter plot with the actual age at the start of primary education of each child on the vertical axis. On the horizontal axis we have the model predicted age for each child given the child's family background and birthday.²¹ The binscatter dots are close to the (red) 45 degree line, which indicates that the model does well in capturing the variation in age at the start of primary education across children. Figure 2.5B evaluates our baseline logit model using the individual-level data for siblings (47,993 pairs). The binscatter plot has the actual between-sibling difference in the age at the start of primary education on the vertical axis.²² On the horizontal

enrolling in the calendar year they turn 5, 6, and 7 years old. We repeat this 50 times and obtain the model predictions in Figure 3 as the average share across the 50 draws for the skill endowment shocks.

¹⁸Age at the start of primary education is obtained as age on the 1st of September of the year children enroll in primary education. We obtain the model predictions using the approach described in footnote 17.

¹⁹We obtain the model predictions using the approach described in footnote 17. The model predicted variance for a given birthday is the average of the variances across the 50 draws for the skill endowment shocks.

²⁰We obtain very similar results when we split the sample by maternal schooling.

²¹We obtain the model predictions using the approach described in footnote 17. The model predicted outcome for each child is the average across the 50 draws for skill endowment shocks.

²²The difference is between the younger and the older sibling. We measure the difference

axis we have the model predicted between-sibling difference.²³ As siblings have the same family background, the model predictions for siblings differ only because of differences in their birthdays. The binscatter dots lie again close to the (red) 45 degree line. Hence, the model also does well in capturing the between-sibling variation in age at the start of primary education.

In sum, our baseline ordered logit model for children's age when they enroll in primary education fits the data well. We find this to be true for the share of children starting primary education at different ages across all possible birthdays; for the average and the variance of age at the start of primary school across birthdays; for the difference in the average age at the start of primary school across birthdays between families with above and below median income; for the variation in age at the start of primary school across all children; and also for the between-sibling variation in age at the start of primary school across all pairs of siblings.

2.5 Primary and Secondary Education Outcomes

We now turn to outcomes in primary and secondary education. The three outcomes we observe are: test scores of national standardized tests administered at the end of primary education; grade retention during primary education; and whether children end up attending the most academic secondary school track. Our models for these outcomes build on children's skills at the end of their kindergarten education as defined in (2.4).

We start with the model for test scores at the end of primary education *TestEoPE*. We assume that these depend on children's skills at the end of kindergarten education *SkillEoK*, family background W , and unobserved shocks η in primary education

$$TestEoPE[x, W, v, \eta; \mathcal{E}] = A_s + B_s SkillEoK[x, W, v; \mathcal{E}] + C_s W + \eta \quad (2.18)$$

where we assume that the primary education shock η is independent of x , W , v and has mean zero. This model for test scores involves four parameters, A_s , σB_s , C_s^S , C_s^Y , as we observe two family background variables (maternal education and family income). The model for retention in primary education builds on the same basic determinants. Children are assumed to be retained at least once in primary

relative to the sample average. We drop the 1547 families with 3 or more children in our data set.

²³We obtain the model predictions using the approach described in footnote 17. The model predicted outcome for each child is the average across the 50 draws for skill endowment shocks.

education if and only if

$$A_r + B_r SkillEoK[x, W, v; \mathcal{E}] + C_r W + u_r \leq RetentionThreshold \quad (2.19)$$

where u_r is a primary education shock that follows a standard logistic distribution. This model also involves four parameters: $A_r - RetentionThreshold$, B_r , C_r^S , C_r^Y (note that A_r cannot be identified separately from the *RetentionThreshold*). Similarly, we assume that children end up attending the most academic secondary school track if and only if

$$A_a + B_a SkillEoK[x, W, v; \mathcal{E}] + C_a W + u_a \geq AcademicTrackThreshold. \quad (2.20)$$

where u_a is a primary education shock that follows a standard logistic distribution. This model again involves four parameters: $A_a - AcademicTrackThreshold$, B_a , C_a^S , C_a^Y .

2.5.1 Calibrating the Models for Education Outcomes

The models for the outcomes in primary and secondary education in (2.18)-(2.20) involve four parameters each. We obtain these parameters using a calibration method. A key feature of our calibration is that we solely target average education outcomes related to family income and maternal education. Specifically, we target the following four moments: the average education outcome for children with family income above/below the median and the average education outcome for children with maternal education above/below the median. We do not target any moments related to children's birthdays.

The calibration of the models for the education outcomes in (2.18)-(2.20) builds on the ordered logit model for age at enrollment in primary education we estimated in Section 2.4. Our estimates of the parameters of the ordered logit model allow us to predict end-of-kindergarten skills (up to a scale factor) given skill endowment shocks v . This allows us to simulate average education outcomes of children with family income above/below the median and children with maternal education above/below the median using (2.18)-(2.20) and to calibrate the models for the education outcomes to match the targeted moments.

The starting point of our calibration is the prediction of end-of-kindergarten skills for given skill shocks v . Substituting children's skill endowments in (2.11) into the equation for their end-of-kindergarten skills in (2.4) and dividing by the scale factor σ yields our estimate of children's scaled skills at the end of kindergarten education as a function of x , W , v , and the primary education

enrollment policy $\mathcal{E}$

$$\frac{\widehat{SkillEoK}}{\sigma}[x, W, v; \mathcal{E}] = \widehat{\alpha/\sigma} + \widehat{\rho/\sigma}W + v + \left(\widehat{\beta/\sigma} + \widehat{\lambda}(\widehat{\rho/\sigma}W + v) \right) \widehat{l^K}[x, W, v; \mathcal{E}] \quad (2.21)$$

where all right-hand-side parameters with hats are estimates obtained using the ordered logit model for primary education enrollment in Section 2.4. The time in kindergarten $\widehat{l^K}$ as a function of x , W , and v can be obtained based on (2.7) and our estimates of the ordered logit model

$$\widehat{l^K}[x, W, v; \mathcal{E}] = 2 + f - x + \widehat{D^7}[x, W, v; \mathcal{E}] - \widehat{D^5}[x, W, v; \mathcal{E}] \quad (2.22)$$

where $\widehat{D^5}$ and $\widehat{D^7}$ are the ordered logit estimates of D^5 and D^7 in (2.9) and (2.10).

Equations (2.21) and (2.22) allow us to simulate the distribution of scaled end-of-kindergarten skills for all children in our sample by combining their characteristics x and W with draws v from a standard logistic distribution. We now discuss how the simulated scaled end-of-kindergarten skills are used to calibrate the models for primary and secondary education outcomes in (2.18)-(2.20).

Test Scores at the End of Primary Education

End-of-primary-education test scores can be expressed as a function of simulated scaled end-of-kindergarten skills by substituting (2.21) into (2.18)

$$TestEoPE[x, W, v, \eta; \mathcal{E}] = A_s + \sigma B_s \frac{\widehat{SkillEoK}}{\sigma}[x, W, v; \mathcal{E}] + C_s W + \eta. \quad (2.23)$$

As we assumed that η has mean zero and is independent of x , W , v , the model for expected end-of-primary-education test scores given $SkillEoK$ and W involves four parameters: A_s , σB_s , C_s^S , C_s^Y (σ and B_s cannot be identified separately).²⁴ We calibrate these four parameters using the average test scores of 4 different groups of children as targets: (i) children with parental income strictly below the median; (ii) children with parental income above the median; (iii) children with maternal education strictly below the median; (iv) children with maternal education above the median. Hence, the moments we target are based on children's family backgrounds. We do not target any moments related to the test results of children with different birthdays.²⁵

²⁴Hence, the scale parameter of the distribution of skill shocks σ cannot be identified separately from B_s .

²⁵The calibrated parameters are reported in Appendix 2.B, Table 1.

Retention in Primary Education

To calibrate the model for retention in primary education we first substitute (2.21) in (2.19). This yields a logit model where the probability of retention in primary education is a function of scaled end-of-kindergarten skills and family background characteristics. The model involves four parameters: σB_r , C_r^S , C_r^Y , and $A_r - \textit{RetentionThreshold}$.²⁶ We calibrate these parameters using as targets the share of retained children in the same 4 groups of children we employed to calibrate the model for test scores.

Secondary School Track

The calibration of the model for the secondary school track that children attend is analogous to the calibration of the retention model. We first substitute (2.21) in (2.20). This yields a logit model where the probability of attending the most academic secondary school track is a function of scaled end-of-kindergarten skills and family background characteristics. The model involves four parameters that we need to calibrate: σB_a , C_a^S , C_a^Y , and $A - \textit{AcademicTrackThreshold}$.²⁷ We calibrate these parameters using as targets the share of children attending the most academic secondary school track in the same 4 groups of children we employed to calibrate the model for test scores.

2.5.2 Evaluating the Model for Education Outcomes

We now evaluate our models for outcomes in primary and secondary education by comparing model predictions with the data.

Education Outcomes across Birthdays

Figure 2.6A examines how well our model for end-of-primary-education test scores captures the data across birthdays.²⁸ It can be seen that the model fits the test data well. It seems worthwhile noting that we do not use any moments involving test scores across birthdays for the model calibration. The four moments we use

²⁶The scale parameter of the distribution of skill shocks σ cannot be identified separately from B_r . Also, A_r cannot be identified separately from the *RetentionThreshold*.

²⁷The scale parameter of the distribution of skill shocks σ cannot be identified separately from B_a . Also, A_a cannot be identified separately from the *AcademicTrackThreshold*.

²⁸We obtain the model predictions by first drawing a skill endowment shock v from a standard logistic distribution for every child in our sample and then predicting the child's end-of-kindergarten skills using (2.21) and our estimates of the ordered logit model for age at enrollment in primary education. Then we substitute the child's predicted end-of-kindergarten skills into the calibrated test score model in (2.23) to obtain predicted end-of-primary-education test scores. We repeat this 50 times and obtain the model predictions in Figure 2.6A as the average across the 50 draws for skill endowment shocks. We implicitly assume that the average primary education shock η experienced by each child is zero (the expected value of η).

are average test scores for groups of children with different family background characteristics (above/below median parental income and maternal education). The variation across birthdays in Figure 2.6A comes therefore entirely from the variation in end-of-kindergarten skills implied by our ordered logit model for age at enrollment in primary education in Section 2.4.²⁹

Figure 2.6B examines how well our model captures the share of children never retained in primary education across birthdays. This model also fits well. Figure 2.6C shows that our model also captures the share of children attending the most academic secondary school track well. The variation across birthdays in both figures is again entirely driven by the variation in end-of-kindergarten skills implied by our ordered logit model for age at enrollment in primary education.

Variance of Education Outcomes among Children with the same Birthday

Figure 7 evaluates our models by examining the variance of outcomes in primary and secondary education among children born on the same day. We obtain the model predicted variance by simulating the education outcome for all children born on a given day and calculating the variance across these children. We then compare the variance in our model and in the data using two-way scatter plots at the birthday level. The vertical axis measures the variance in the data and the horizontal axis the variance in our model. Variances are measured relative to the average across birthdays.

Figure 2.7A compares the variance of end-of-primary-education test scores among children with the same birthday in our model and in the data.³⁰ For the variance of test scores in the data to be well captured by our model, the scatter plot should lie around the (red) 45 degree line. This turns out to be the case. Figure 2.7B shows the same comparison for whether children have or have not been retained in primary education. This outcome only takes two values for each child, 1 if the child is never retained in primary education and 0 if the child is retained once or more. The model variance among children with the same birthday simulates the binary outcome for children born on that day and then calculates

²⁹We obtain the model prediction for each child using an approach that is analogous to the one used for test scores in footnote 27. The main difference is that we also need to draw the primary education shocks u_r and u_a in (2.19) and (2.20) from a standard logistic distribution to predict whether the child is subject to grade retention or not and whether the child ends up attending the most academic secondary school track or not.

³⁰Variances are measured relative to the average across birthdays. The model predicted variance is calculated based on the model predicted end-of-primary-education test scores for each child obtained as explained in footnote 27 assuming $\eta = 0$. As we take η to be independent of $SkillEoK$ and W , this implies that we understate the true model predicted variance for every birthday by the (unknown) constant $Var(\eta)$. However, this constant is eliminated by measuring variances relative to the average across birthdays.

the variance of the simulated outcome across these children.³¹ The data variance is the variance of the actual binary outcome across the same children. The scatter plot lies around the (red) 45 degree line, which implies that our model variances capture the variances in the data well. Figure 2.7C compares the model variance for whether children with the same birthday attend the most academic secondary school track with the variance in the data. This outcome again takes two values only, 1 if the child attended the academic track and 0 if not. Our model does well in this dimension also.

In sum, the variance of outcomes in primary and secondary education among children with the same birthday in the data is captured well by our model.

Family Income and Education Outcomes across Birthdays

Figure 8 evaluates our model separately for the group of children with parental income above the median and the group of children with parental income strictly below the median. We examine the average of the education outcomes in these two groups for all possible birthdays. It is worthwhile recalling that our model calibration actually targets the average education outcomes in these groups. However, we do not use any moments involving the variation in outcomes across birthdays for either group. The variation across birthdays in each group is entirely driven by the variation in end-of-kindergarten skills implied by our ordered logit model for age at enrollment in primary education. It can be seen in Figure 8 that our model captures the variation in all three education outcomes well in both groups.

Between-Sibling Differences in Education Outcomes

Figure 9 evaluates how well our model does in explaining differences in education outcomes within families. The data for siblings are the same as in Figure 5 (50,908 pairs of siblings). The three binscatter plots have the actual between-sibling difference in education outcomes on the vertical axis.³² On the horizontal axis we have the between-sibling difference of the primary education outcome predicted by the model. Because siblings have the same family endowment, the model prediction for siblings only differs because of differences in their birthdays. Figure 2.9A contains the binscatter for the difference between the end-of-primary-education test scores of siblings, Figure 2.9B for the between-sibling difference in whether they were ever retained, and Figure 2.9C for the between-sibling difference in whether they end up attending the most academic track of secondary school.

³¹The model predicted variances in Figure 2.7B and 2.7C are calculated based on the model predicted outcomes for all children obtained as explained in footnote 28.

³²The difference is again between the younger and the older sibling. We measure the difference relative to the sample average.

The binscatter dots lie close to the (red) 45 degree line for all three education outcomes, which indicates that the models do well in explaining the between-sibling variation in education outcomes in the data.

Education Outcomes and Discontinuous Admission to Primary Education

We also evaluate our model for outcomes in primary and secondary education using a regression discontinuity design (RDD). We implement the RDD for the two birthdates where we found discontinuous changes in the minimum-skill thresholds entering the primary education enrollment policy, October 1 and January 1. For both birthdates, we first use a RDD design to estimate the change in end-of-primary-education test scores, the probability of grade retention during primary education, and the probability that children end up attending the most academic track of secondary education.³³ We then compare these estimates with RDD estimates using simulated data for the three education outcomes based our models in Sections 2.4 and 2.5.³⁴ It is worthwhile noting that our model implies that the RDD effect is heterogeneous at the individual level. Children with high and low skill endowments born on October 1 and January 1 are completely unaffected by the discontinuous changes in the minimum-skill thresholds. Moreover, the effect on children with intermediate skill endowments is heterogeneous because of $\lambda \neq 0$. As a result, the RDD estimates do not have a structural interpretation. The comparison of the RDD estimates in the actual data and the simulated data is nevertheless interesting as, if our model for education outcomes captures the data, this heterogeneity should be the same in the data and the model.

Figure 10 compares the RDD point estimates we get using the actual and the simulated data for the three primary and secondary education outcomes. For the RDD estimates using the data, which are on the vertical axis, we plot point estimates and 90 percent confidence bands. While the RDD estimates using the actual data are quite noisy, point estimates are close to the RDD estimates using the simulated data. It seems also noteworthy that the relative size of the RDD point estimates for grade retention and attendance of the most academic track differ for the two discontinuities, both using the actual and the simulated data.

³³We estimate separate RDD effects for October 1 and January 1 and pool the data across birth years, see Appendix 2.B, Table 2 for more information.

³⁴The RDD estimates using the simulated data on education outcomes are obtained as follows. We first predict the education outcome for all children using the approach in footnote 27 and 28 and then use the simulated data to obtain the RDD point estimate using the same approach as for the actual data. We repeat this for 50 draws from the shocks and calculate the average RDD point estimate across the 50 draws.

Instrumental Variables Estimation of Test Skill Model

Finally, we can also evaluate our model for end-of-primary-education test scores in Sections 2.4 and 2.5 against the causal effect of time in kindergarten education on test scores obtained using an instrumental-variables approach. Substituting (2.21) into (2.23) yields that child i 's end-of-primary-education test scores can be written as

$$TestEoPE_i = a + \beta B_s l_i^K + (\rho B_s + C_s)W_i + \rho \lambda B_s W_i l_i^K + \lambda \sigma B_s v_i l_i^K + \sigma B_s v_i + \eta_i \quad (2.24)$$

l_i^K is the time in kindergarten education.

The effect of particular interest in (2.24) is βB_s , the slope coefficient on l_i^K . It captures the effect of more time in kindergarten education on the test scores of children with average family background ($W = 0$) and skill endowment shock ($v = 0$). This effect is the product of two parameters entering the model used in previous sections to predict children's end-of-primary-education test scores. First, the scaled effect of more time in kindergarten on end-of-kindergarten skills (β/σ) we estimated using the ordered logit model for age at enrollment in primary education in Section 2.4. Second, the effect of an increase in scaled end-of-kindergarten skills on end-of-primary-education test scores (σB_s) we obtained by calibrating our model for test scores in Section 2.5. Hence, the product $(\beta/\sigma)(\sigma B_s)$ of these two parameters could be checked against causal estimates of βB based on (2.24).

Causal estimation of βB_s based on (2.24) requires addressing two (standard) issues. First, the endogeneity of l_i^K , as time in kindergarten education depends on the child's (unobserved) skill endowment shock v_i . Second, unobserved heterogeneity. If $\lambda \neq 0$, end-of-kindergarten skills—and hence end-of-primary-education test scores—depend on an interaction between time in kindergarten and the skill endowment shock ($v_i l_i^K$).

Both issues can be addressed by building on our ordered logit model for age at enrollment in primary education in Section 2.4. To see how, we define G_i as the expected value of $v_i l_i^K$ given the birthday x_i and family background W_i of child i , $G_i = \mathbb{E}(v_i l_i^K | x_i, W_i)$. Using this definition allows us to rewrite (2.24) as

$$TestEoPE_i = a + \beta B_s l_i^K + (\rho B_s + C_s)W_i + \rho \lambda B_s W_i l_i^K + \lambda \sigma B_s G_i + \epsilon_i \quad (2.25)$$

where $\epsilon_i = \lambda \sigma B_s (v_i l_i^K - \mathbb{E}(v_i l_i^K | x_i, W_i)) + \sigma B_s v_i + \eta_i$. Adam's law for conditional expectations implies that $\mathbb{E}(v_i l_i^K - \mathbb{E}(l_i^K v_i | x_i, W_i) | x_i) = 0$ which combined with

the properties of the shocks v_i and η_i yields $\mathbb{E}(\epsilon_i|x_i) = 0$ and hence

$$\mathbb{E}(\epsilon_i|\mathbb{E}(l^K|x_i)) = 0 \quad (2.26)$$

where $\mathbb{E}(l^K|x_i)$ is the expected time in kindergarten education given birthday x_i .

We can now obtain a causal estimate of βB_s using a 2SLS approach that combines (2.25) and (2.26) with our ordered logit model in Section 2.4. The first step is to replace the value of G_i for each child on the right-hand side of (2.25) by its estimate using the ordered logit model. Moreover, we also use the ordered logit model to estimate the expected time in kindergarten for every birthday x and then employ these estimates as an instrument for the actual time children with birthday x spent in kindergarten. This results in a 2SLS model with a generated regressor and instrument, which yields consistent 2SLS estimates as long as the generated regressor and instrument are estimated consistently (Wooldridge, 2002).³⁵

The 2SLS estimate of βB_s using (2.25) is around 0.27 with a standard error of 0.012.³⁶ The logit model yielded that β/σ , the scaled effect of time in kindergarten education on end-of-kindergarten skills, is 4.02. The calibration of the model for end-of-primary-education test scores yielded that σB_s , the effect of scaled end-of-kindergarten skills on end-of-primary-education test scores, is 0.074. Hence, $(\beta/\sigma)(\sigma B_s)$ is 0.29, close to our 2SLS estimate of βB_s .

2.6 Counterfactual Analysis of Enrollment Policies

We now analyze two counterfactual primary education enrollment policies and their effects on outcomes in primary and secondary education. We compare these policies with each other and with the primary education enrollment policy employed in the Netherlands.

The two counterfactual primary education enrollment policies are polar opposites in the weight put on children's birthdates versus their skills. The first policy enrolls children in primary education solely based on their birthdates. Conditional on children's birthdates, their skills are irrelevant for when they start their primary education. The second policy enrolls children in primary education solely based on their skills at the start of the school year. Conditional on children's skills, their birthdates are irrelevant.

We now describe the two counterfactual enrollment policies in more detail. Then we compare the primary and secondary education outcomes of the different enrollment policies.

³⁵Consistency of the standard errors requires additional assumptions (Wooldridge, 2002).

³⁶For the full 2SLS results, see Appendix 2.B, Table 3.

2.6.1 Two Counterfactual Primary Education Enrollment Policies

The first counterfactual policy we examine enrolls children in primary education based solely on their birthdates. The policy sets an annual cutoff date; children born before the cutoff date of calendar year T start their primary education in calendar year $T + 6$, while children born after the cutoff date start their primary education in calendar year $T + 7$. In our counterfactual analysis, we choose the annual cutoff date so that the average age of children at the start of primary education is the same as in the data. Put differently, the cutoff date is chosen such that the average age of children at the start of primary education under the counterfactual enrollment policy is the same as under the policy in the Netherlands. What will change compared to the policy in the Netherlands is that the distribution of age at the start of primary education across child birthdates more strongly and their skill endowments less strongly.

The second counterfactual policy enrolls children in primary education based solely on their skills at the start of the school year. Formally, children's age at primary education enrollment continues to be determined by (2.9) and (2.10). The key change compared to the policy in the Netherlands is that the counterfactual policy is based on minimum-skill threshold functions $minskill^j[x]$ for $j = 5, 6$ that are independent of children's birthday x . That is, $minskill^j[x] = minskill^j$ for $j = 5, 6$. In our counterfactual analysis, we choose these thresholds such that average age at the start of primary education is the same as under the policy in the Netherlands. What will change compared to the policy in the Netherlands is that the distribution of age at the start of primary education across children will reflect children's skill endowments more strongly and their birthdates less strongly.

2.6.2 Comparing Outcomes in Primary and Secondary Education

We start by comparing age and skills at the start of primary education under the different primary education enrollment policies. Then we turn to the primary and secondary education outcomes: retention in primary education, test scores at the end of primary education, and the secondary school track that children end up attending.

Age and Skills at the Start of Primary Education

Figure 2.11A summarizes our results for average age at the start of primary education across birthdays for the three different enrollment policies. The lines in this figure, and all subsequent figures, are unweighted lowess curves fit through

the simulated data points for 365 birthdays.³⁷ The annual cutoff date of the counterfactual policy where primary education enrollment is based solely on children's birthdates is October 28. This cutoff date implies that children's average age at the start of primary education is the same as under the policy in the Netherlands.

Figure 2.11A shows that the enrollment policy in the Netherlands implies a maximum difference across birthdays in average age at the start of primary education of 0.5 years. The counterfactual policy where enrollment is based solely on children's birthdates leads to differences of up to one year. The policy in the Netherlands leads to substantially smaller differences across birthdays mainly because children with high skill endowments born after October 28—who would have been the oldest in their class at the start of primary education under the counterfactual policy—start their primary education earlier than under the counterfactual policy based solely on birthdates. This frees up resources in kindergarten education, which are then mostly used for children born around the middle of the year who have low skill endowments.

The reallocation of time in kindergarten from high-skill children born after October 28 to low-skill children born earlier is strongest under the counterfactual policy where children are enrolled based solely on their skills at the start of the school year. It is therefore expected that, as shown in Figure 2.11A, differences in average age at the start of primary education across birthdays are smallest under this enrollment policy. As a matter of fact, the policy nearly eliminates any differences in average age at the start of primary education across birthdays.

Figure 2.11B shows the implications of the different enrollment policies for the standard deviation of age at the start of primary education among children born on the same day. The standard deviation is zero for all birthdays under the policy where primary education enrollment is based solely on birthdates. This is an immediate implication of this policy enrolling children born on the same day in the same school year. The standard deviation is largest, for every birthday, under the policy where enrollment in primary education is solely based on skills. The policy in the Netherlands implies a standard deviation between the two counterfactual enrollment policies.

Comparing Figure 2.11B with Figure 2.11A suggests a basic policy trade-off. While the policy enrolling children based solely on their skills leads to the smallest differences in average age at the start of primary education across birthdays, it

³⁷We use the default bandwidth. The underlying data consists of the model predictions for every birthday. The result for the enrollment policy used in the Netherlands is obtained by fitting separate lines before and after October 1. The line for the enrollment policy based solely on children's birthdates is obtained by fitting separate lines before and after October 28 (October 28 is the counterfactual annual cutoff date implying the same average age at the start of primary education as the enrollment policy used in the Netherlands).

leads to the largest differences in age at the start of primary education among children with the same birthday.

However, the comparison between the different enrollment policies will look very different when we examine children's skills, not their age, at the start of primary education. This is because of how the policy in the Netherlands, and the counterfactual policy enrolling children based solely on their skills, distribute time in kindergarten education across children with different skill endowments. Figure 2.11C summarizes this aspect of the different enrollment policies using the simulated covariance between children's implied ages at the start of primary education and their scaled skill endowments, which—using (2.11)—are obtained as $(\rho/\sigma)W + v$ with the estimate of ρ/σ taken from our ordered logit model in Section 2.4. Under the policy where enrollment is solely based on birthdates, this covariance is zero for all birthdays, as children's skills are irrelevant for when they start their primary education. Under the policy where enrollment is solely based on skills, the covariance is negative and varies little across birthdays. The primary education enrollment policy in the Netherlands implies a covariance between the two counterfactual policies. This is because, compared to the policy where enrollment is solely based on birthdates, the policy in the Netherlands reallocates time in kindergarten from children with high skill endowments who would have been the oldest in their class at the start of primary education to children with low skill endowments born earlier. This reallocation of time in kindergarten across children with different skill endowments is even stronger under the counterfactual policy where enrollment is solely based on skills at the start of the next school year.

Figure 12 shows the implications of the different primary education enrollment policies for scaled skills at the start of primary education or, equivalently, scaled skills at the end of kindergarten education. Figure 2.12A shows average skills at the start of primary education across birthdays. The main finding is very similar to Figure 2.11A. Differences are largest for the policy where enrollment is solely based on birthdates, followed by the policy in the Netherlands. Under the policy where enrollment is solely based on skills, there are basically no differences in average skills at the start of primary education across birthdays. Figure 2.12B shows the standard deviation of skills at the start of primary education among children born on the same day. Here the main finding is the opposite of what we obtained in Figure 2.11A. Now it is the policy where enrollment is solely based on skills that implies the *smallest* standard deviation in skills at the start of primary education among children born on the same day. This finding is driven by the way the enrollment policy based solely on skills allocates time in kindergarten education across children with different skill endowments, which we illustrated in

Figure 2.11C.

Primary and Secondary Education Outcomes Across Birthdays

Figure 13 shows average outcomes in primary education and in secondary education across birthdays for the three different primary education enrollment policies. We start with the comparison between the primary education enrollment policy in the Netherlands and the counterfactual policy based solely on birthdates. For children born in the first half of the year, the two policies lead to similar outcomes in primary and secondary education. This is true for average z-score of the standardized tests at the end of primary education (Figure 2.13A); for the share of children never retained in primary education (Figure 2.13B); and for the share of children attending the most academic secondary school track (Figure 2.13C).

For children born in the second part of the year, the enrollment policy in the Netherlands leads to substantially smaller differences in primary and secondary education outcomes across birthdays than the counterfactual policy based solely on birthdates. The policy in the Netherlands implies a maximum difference for the three education outcomes that is around half the difference under the counterfactual policy. For test scores at the end of primary education, the policy in the Netherlands leads to a maximum difference in average z-scores of the standardized test at the end of primary education of 0.125, compared to 0.25 under the policy based solely on birthdates (Figure 2.13A). For the share of children never retained in primary education, the policy in the Netherlands leads to a maximum difference of 6 percentage points across birthdays, compared to 15 percentage points under the counterfactual policy (Figure 2.13B). Finally, for the share of children attending the most academic secondary school, the policy in the Netherlands leads to a maximum difference of 6 percentage points across birthdays, compared to a difference 12 percentage points under the policy based solely on birthdates (Figure 2.13C).

Figure 13 also shows that differences in primary and secondary education outcomes across birthdays in the Netherlands could be reduced further by a primary education enrollment policy that puts more weight on children's skills. As a matter of fact, the counterfactual policy where enrollment in primary education is solely based on children's skills basically eliminates *all* differences across birthdays in standardized test scores at the end of primary education (Figure 2.13A); in the share of children never retained in primary education (Figure 2.13B); and in the share of children attending the most academic secondary school (Figure 2.13C).

Education Outcomes of Children with Low Skill Endowments

We now turn to outcomes in primary and secondary education of children with low skill endowments because of their family background. Figure 14 examines average education outcomes of children with skill endowments below the 10th percentile and the average skill endowment shock (the effect of skill endowment shocks will be examined in the next section).³⁸ We start with the comparison between the primary education enrollment policy in the Netherlands and the counterfactual policy based solely on birthdates. The differences for all three education outcomes are limited to the 4-week window between October 1 and 28. For all other birthdays, the two enrollment policies yield the same average outcomes.

The reason is straightforward. When compared to the counterfactual policy where enrollment is based solely on birthdates, the policy in the Netherlands reallocates time in kindergarten to children with low skill endowments. However, this reallocation effect is generally limited to children with very low skill endowments because of the joint effect of family background *and* adverse skill endowment shock. The exceptions are children born during the 4-week period from October 1 to October 28. Under the counterfactual policy, all these children start their primary education in the calendar year they turn 6 years old. By contrast, the policy in the Netherlands has around 40 percent of the children born during this 4-week period start their primary education in the calendar year they turn 7. As these children are those with the lowest skills, the enrollment policy in the Netherlands raises the average educational outcomes of children with low skill endowments compared to the policy based solely on birthdates. The policy in the Netherlands can assign more time in kindergarten education to children born in this 4-week window between October 1 and October 28, although time in kindergarten per child is the same as under the counterfactual policy, mainly because it assigns less time in kindergarten to children born after October 28 with high skill endowments.

While the difference between the enrollment policy in the Netherlands and the counterfactual policy based solely on birthdates is limited to birthdays between October 1 and October 28, the effect on children with skill endowments below the 10th percentile born during this 4-week period is substantial. Under the policy in the Netherlands, their average z-score in the end-of-primary-education standardized test improves by around 0.25; they are around 25 percentage points more likely to get through primary education without being retained; and they are around 4 percentage points more likely to end up attending the most academic secondary school track.

³⁸Children's skill endowments associated with family background are estimated as $(\rho/\sigma)W$ using the estimate of ρ/σ in our ordered logit model in Section 2.4.

Figure 14 also shows that the outcomes in primary and secondary education of children with low skill endowments in the Netherlands could be improved further. The counterfactual policy where enrollment in primary education is based solely on children's skills would substantially raise education outcomes of children with low skill endowments born in July, August, and September. Their average z-score in the end-of-primary-education test could be increased by up to almost 0.3 (Figure 2.14A); the share of these children never retained in primary education could be raised by up to 28 percentage points (Figure 2.14B); and the share of these children attending the most academic secondary school could be increased by up to more than 4 percentage points (Figure 2.14C).

Education Outcomes of Children with Low Skill Endowment Shocks

Figure 15 examines average outcomes in primary and secondary education of children with the average family background and skill endowment shocks below the 25th percentile. It can be seen that compared to the counterfactual policy based solely on birthdates, the policy in the Netherlands raises education outcomes substantially for children born between October 1 and 28. Their average z-score in the end-of-primary-education test is increased by up to almost 0.3 (Figure 2.15A); the share of children never retained is raised by up to 27 percentage points (Figure 2.15B); and the share of children attending the most academic secondary school is increased by up to almost 11 percentage points (Figure 2.15C). The reason is that the policy in the Netherlands enrolls children with low skills born during this 4-week period in primary education in the calendar year they turn 7, while the counterfactual policy enrolls these children in the calendar year they turn 6. Moreover, the policy in the Netherlands also leads to substantial improvements in education outcomes for children born before October 1. The reason is that children with skill endowment shocks below the 25th percentile tend to be a large share of the children with low skills at the time the enrollment decision is made. As a result, these children are a large fraction of the children the policy in the Netherlands enrolls in primary education in the calendar year they turn 7 years old.

Figure 15 also shows that there is potential for further improvement in the education outcomes of children with skill endowment shocks below the 25th percentile. For children born before October 1, the primary education enrollment policy in the Netherlands yields substantially worse education outcomes than the counterfactual policy where enrollment is solely based on skills. Children in the Netherlands with endowment shocks below the 25th percentile achieve an average z-score in the end-of-primary-education test that is up to around 0.2 worse than if enrollment were

solely based on skills (Figure 2.15A); the share of these children never retained is more than 10 percentage points higher than if enrollment were solely based on skills (Figure 2.15B); and the share of children attending the most academic secondary school is up to 7 percentage points lower than if enrollment were solely based on skills (Figure 2.15C).

2.7 Summary

In most countries, the default policy is to enroll children in primary education by date of birth. However, in the Netherlands, schools are asked by their supervisory institution to enroll children solely based on developmental skills. We analyze if and how enrolling children in primary education based on their skills affects equality of opportunity in education. Compared to enrollment by date of birth, enrollment policies based on children's skills have the potential to improve equality of opportunity along two interlinked dimensions. Enrolling children solely by date of birth implies that their skills at the start of primary education end up reflecting randomness in date of birth. Moreover, enrolling children by date of birth also preserves all differences in skill endowments at school entry. An enrollment policy based on children's skills may reduce the effect of randomness in children's dates of birth on education outcomes and raise education outcomes of children with skill endowments at the bottom of the distribution.


Our analysis yields two main results. First, compared to enrollment by date of birth, the policy in the Netherlands improves equality of opportunity in education by reducing differences in primary and secondary education outcomes across birthdays and by improving education outcomes of children with skill endowments at the bottom of the distribution. Second, despite the recommendation of their supervisory institution, schools in the Netherlands do not enroll children in primary education solely based on their developmental skills. We find that doing so would nearly eliminate differences in primary and secondary education outcomes across birthdays and further raise outcomes of children with low skill endowments.

2.8 Tables and Figures


Table 2.1: Ordered Logit Model for Age at Enrollment in Primary Education

Minimum-skill threshold	Step-function with slope		Step-function
Heterogenous learning?	Yes ($\lambda \neq 0$)	No ($\lambda = 0$)	No ($\lambda = 0$)
β/σ	4.064*** (0.079)	4.474*** (0.045)	5.162*** (0.043)
θ/σ	0.828*** (0.023)	0.918*** (0.0189)	1.163*** (0.0174)
κ/σ	3.726*** (0.109)	3.933*** (0.112)	
λ	-0.042*** (0.007)		
ρ_S/σ	0.081*** (0.002)	0.081*** (0.002)	0.081*** (0.002)
ρ_Y/σ	0.338*** (0.009)	0.337*** (0.01)	0.330*** (0.009)
cutoffA	2.649*** (0.027)	2.774*** (0.020)	2.620*** (0.184)
cutoffB	-4.498*** (0.078)	-4.947*** (0.026)	-5.326*** (0.026)


Notes: This table reports estimates of the ordered logit model for age at enrollment in primary education in 2.4. The sample includes all children born in the Netherlands between 2002 and 2005 for whom we have the relevant data. See Appendix 2.A for details on the data and summary statistics.

FIGURE 2.1A: Age at Enrollment in Primary Education


Notes: Share of children enrolling in primary education in the calendar year they turn 5, 6, or 7 years old by children's birthday. Enrollment shares sum to one. See Appendix 2.A for information on the data and summary statistics.

**FIGURE 2.1B: Parental Income and Primary Education Enrollment
Age 7**


Notes: Share of children enrolling in the calendar year they turn 7 years old by children's birthday, separately for children with parental income above and below the median. See Appendix 2.A for information on the data and summary statistics.

FIGURE 2.1C: Age at the Start of Primary Education


Notes: Age at the start of primary education measured as children's ages on September 1 of the year they enroll in primary education. See Appendix 2.A for information on the data and summary statistics.

FIGURE 2.2A: Non-Retention in Primary Education


Notes: Share of children not retained in any grade in primary education by children's birthday. Primary education consists of six grades. See Appendix 2.A for information on the data and summary statistics.

FIGURE 2.2B: End-of-Primary-Education Test Score


Notes: Z-score of the standardized test score at the end of primary education by children's birthday. See Appendix 2.A for information on the data and summary statistics.

FIGURE 2.2C: Share Attending Academic Secondary School


Notes: Share of children who end up attending the most academic secondary school track by children's birthday. We consider children to have enrolled in the academic secondary school track if they are enrolled in an academic secondary school in the second grade of secondary school.

FIGURE 2.3A: Predicted and Actual Age at Enrollment in Primary Education


Notes: Actual and predicted primary education enrollment shares at ages 5 and 6. The enrollment share age 7 is one minus the enrollment shares at ages 5 and 6. Model predicted shares are obtained using the ordered logit model in Section 2.4. We first simulate age at enrollment in primary education for all children given a simulated skill shock for each child, then obtain the enrollment shares at ages 5 and 6, and finally average enrollment shares across 50 skill shocks. See Section 4.3 for further information.

FIGURE 2.3B: Age at Primary Education Enrollment with Homogeneous Learning in Kindergarten


Notes: Actual and predicted primary education enrollment shares at ages 5 and 6. Model predicted shares are obtained using the ordered logit model in Section 2.4 assuming that learning in kindergarten does not depend on skill endowments ($\lambda = 0$). See the notes of Figure 2.3A for an explanation of how model predicted shares are obtained.

FIGURE 2.3C: Primary Education Enrollment with Identical Skill Threshold for October-December Birthdays


Notes: Actual and predicted primary education enrollment shares at ages 5 and 6. Model predicted shares are obtained using the ordered logit model in Section 2.4 assuming that the same minimum-skill threshold is applied to children born October-December ($\kappa = 0$). See the notes Figure 2.3A for an explanation of how model predicted shares are obtained.

FIGURE 2.4A: Predicted and Actual Age at the Start of Primary Education

Notes: Comparison of actual age and predicted age at the start of primary education by children's birthdays. Age at the start of primary education is measured as children's age on September 1 of the year they enroll in primary education. Model predicted age is obtained by using the ordered logit model in Section 2.4 to first obtain the age at the start of primary education of all children given a simulated skill shock for each child and then average across 50 skill shocks. See Section 4.3 for further information.


FIGURE 2.4B: Predicted and Actual Variance of Age at Start of Primary Education Among Children Born on the Same Day

Notes: Actual and predicted variance of age at the start of primary education among children with the same birthday. The model predicted variance is obtained by using the ordered logit model in Section 2.4. We first obtain age at the start of primary education for all children given a simulated skill shock for each child, then obtain the variance among children with the same birthday, and finally average the variance across 50 skill shocks. See Section 4.3 for further information.

FIGURE 2.4C: Predicted and Actual Age at Start of Primary Education and Income


Notes: Actual and predicted age at the start of primary education, separately for children with parental income above and below the median.

FIGURE 2.5A: Children's Predicted and Actual Age at the Start of Primary Education




Notes: Binscatter plot of the actual age of each child at the start of primary education against the expected age of the child predicted by the ordered logit model in Section 2.4. Model predicted expected age is obtained by calculating the predicted age at the start of primary education of each child for 50 simulated skill shocks and averaging across skill shocks. See Section 4.3 for further information. There are 333,465 children in the sample. Number of bins chosen optimally.


FIGURE 2.5B: Actual and Predicted Between-Sibling Difference in Age at the Start of Primary Education




Notes: Binscatter plot of between-sibling difference in actual age at the start of primary education against the expected between-sibling age difference predicted by the ordered logit model in Section 2.4. See the notes to Figure 2.5A for an explanation of how model predicted expected age is obtained. The sample consists of 50,908 families with two children. The between-sibling age difference is measured as the difference between the older and the younger sibling relative to the average difference across all families. Number of bins chosen optimally.

FIGURE 2.6A: Predicted and Actual End-of-Primary-Education Test Score

Notes: Predicted test scores are obtained by combining the models in Sections 2.4 and 2.5 to obtain the predicted test scores for all children given a simulated skill shock for each child and then average across 50 skill shocks. See Section 5.2.1 for further information.


FIGURE 2.6B: Predicted and Actual Non-Retention in Primary Education

Notes: Predicted shares are obtained by combining the models in Sections 2.4 and 2.5. We first obtain the predicted non-retention indicator variable for all children given a simulated skill endowment shock and primary education shock for each child and then average across 50 skill shocks. See Section 2.5 for further information on the primary education shock.

FIGURE 2.6C: Predicted and Actual Attendance of Academic Secondary School


Notes: Predicted shares are obtained by combining the models in Sections 2.4 and 2.5. We first obtain the predicted most-academic-secondary-school-track indicator variable for all children given a simulated skill endowment shock and primary education shock for each child and then average across 50 skill shocks. See Section 2.5 for further information on the primary education shock.

FIGURE 2.7A: Predicted and Actual Variance in End-of-Primary-Education Test Scores Among Children Born on the Same Day




Notes: Scatter plot at the birthday level. The model predicted variance is obtained combining the models in Sections 2.4 and 2.5. We first obtain the predicted test scores for all children given a simulated skill shock for each child, then calculate the variance among children with the same birthday, and finally average across 50 skill shocks. Variances are measured relative to the average across all birthdays.

FIGURE 2.7B: Predicted and Actual Variance in Non-Retention Among Children Born on the Same Day




Notes: Scatter plot at the birthday level. *Non-retention* is an indicator variable taking the value of 1 if and only if the child has not been retained in primary education. The predicted variance of the non-retention indicator is obtained combining the models in Sections 2.4 and 2.5. We first obtain the predicted non-retention indicator variable for all children given a simulated skill endowment shock and primary education shock, then calculate the variance among children with the same birthday, and finally average across 50 skill shocks. See Section 2.5 for further information on the primary education shock. Variances are measured relative to the average across all birthdays.


FIGURE 2.7C: Predicted and Actual Variance in Academic-Secondary-School Attendance Among Children Born on the Same Day




Notes: Scatter plot at the birthday level. *Academic-secondary-school attendance* is an indicator variable taking the value of 1 if and only if the child ended up attending the most academic secondary school. The predicted variance is obtained analogously to the variance of the non-retention indicator variable. See the notes to Figure 2.7B.

FIGURE 2.8A: Predicted and Actual Test Scores and Income

Notes: Predicted and actual end-of-primary-education test scores by birthday, separately for families with income above and below the median. Predictions are obtained combining the models in Sections 2.4 and 2.5. See the notes to Figure 2.6A for more information.


FIGURE 2.8B: Predicted and Actual Non-Retention in Primary Education and Income

Notes: Predicted and actual shares of children never retained in primary education by birthday, separately for families with income above and below the median. Predictions are obtained combining the models in Sections 2.4 and 2.5. See the notes to Figure 2.6B for more information.

FIGURE 2.8C: Predicted and Actual Enrollment in Academic Secondary School and Income


Notes: Predicted and actual shares of children enrolled in the most academic secondary school track by birthday, separately for families with income above and below the median. Predictions are obtained combining the models in Sections 2.4 and 2.5. See the notes to Figure 2.6C for more information.

FIGURE 2.9A: Actual and Predicted Between-Sibling Difference in Test Scores




Notes: Binscatter plot of between-sibling difference in the actual end-of-primary-education test scores against the expected between-sibling difference as predicted by the models in Sections 2.4 and 2.5. See the notes to Figure 2.7A for an explanation of how we obtain the model predicted expected outcome for each child and the notes to Figure 2.5B for details on the sample. Number of bins chosen optimally.

FIGURE 2.9B: Actual and Predicted Between-Sibling Difference in Non-Retention




Notes: Binscatter plot of between-sibling difference in the actual non-retention indicator variable against the expected between-sibling difference as predicted by the models in Sections 2.4 and 2.5. See the notes to Figure 2.7B for the definition of the non-retention indicator variable and an explanation of how we obtain the model predicted expected outcome for each child. Details on the sample are in the notes to Figure 2.5B. Number of bins chosen optimally.

FIGURE 2.9C: Actual and Predicted Between-Sibling Difference in Enrollment in Academic Secondary School




Notes: Binscatter plot of the between-sibling difference in the actual academic-secondary-school attendance indicator against the expected between-sibling difference predicted by the models in Sections 2.4 and 2.5. See the notes to Figure 2.7C for the definition of the academic-secondary-school attendance indicator variable, the notes to Figure 2.7B for an explanation of how we obtain the model predicted expected outcome for each child, and the notes to Figure 2.5B for details on the sample. Number of bins chosen optimally.

FIGURE 2.10A: Regression Discontinuities on October 1 in Data Versus Model




Notes: The estimates are obtained using a regression discontinuity design (RDD). The RDD estimates capture the effect of the discontinuities in the primary school enrollment policy on October 1 and January 1 on the three education outcomes. The outcomes are end-of-primary-education test scores; non-retention in primary education; and attendance of the most academic secondary school track. Estimates on the vertical axis are based on the actual data. Estimates on the horizontal axis are based on simulated data using the models in Sections 2.4 and 2.5. See Section 5.2.5 for more information.

FIGURE 2.10B: Regression Discontinuities on January 1 in Data Versus Model




Notes: The RDD estimates capture the effect on outcomes of being born on January 1 compared to being born on December 31 of the previous year. The outcomes are end-of-primary-education test scores; non-retention in primary education; and enrollment in the most academic secondary school track. Estimates on the horizontal axis are based on simulated data using the model in Sections 2.4 and 2.5. Estimates on the vertical axis are based on the actual data.

FIGURE 2.11A: Primary Education Enrollment Policies and Age at the Start of Primary Education




Notes: Average age at the start of primary education by birthday. Age at the start of primary education is measured as children's ages on September 1 of the year they start in primary education. The simulations for different primary education enrollment policies are based on the model in Sections 2.4 and 2.5. The lines are unweighted lowess curves fit through simulated data for 365 birthdays, see footnote 36 for more information.

FIGURE 2.11B: Standard Deviation in Age at the Start of Primary Education Among Children Born on the Same Day




Notes: Standard deviation of the age at the start of primary education among children born on the same day. The simulations for different primary education enrollment policies are based on the model in Sections 2.4 and 2.5. The lines are unweighted lowess curves fit through simulated data for 365 birthdays, see footnote 36 for more information.

FIGURE 2.11C: Correlation Between Age at the Start of Primary Education and Skill Endowments




Notes: Correlation between age at the start of primary education and skill endowments among children born on the same day. The simulations for different primary education enrollment policies are based on the model in Sections 2.4 and 2.5. The lines are unweighted lowess curves fit through simulated data for 365 birthdays, see footnote 36 for more information.

FIGURE 2.12A: Primary Education Enrollment Policies and Skills at the Start of Primary Education




Notes: Average skills at the start of primary education—or, equivalently, at the end of kindergarten education—by birthday. The simulations for different primary education enrollment policies are based on the model in Sections 2.4 and 2.5. The lines are unweighted lowess curves fit through simulated data for 365 birthdays, see footnote 36 for more information.

FIGURE 2.12B: Standard Deviation in Skills at the Start of Primary Education Among Children Born on the Same Day




Notes: Standard deviation of skills at the start of of primary education—or, equivalently, at the end of kindergarten education—among children born on the same day. The simulations for different primary education enrollment policies are based on the model in Sections 2.4 and 2.5. The lines are unweighted lowess curves fit through simulated data for 365 birthdays, see footnote 36 for more information.

FIGURE 2.13A: Primary Education Enrollment Policies and End-of-Primary-Education Test Scores




Notes: Average standardized test z-scores at the end of primary education by birthday. The simulations for different primary education enrollment policies are based on the model in Sections 2.4 and 2.5. The lines are unweighted lowess curves fit through simulated data for 365 birthdays, see footnote 36 for more information.

FIGURE 2.13B: Primary Education Enrollment Policies and Non-Retention in Primary Education




Notes: Share of children never retained in primary education by birthday under different enrollment policies. The simulations are based on the model in Sections 2.4 and 2.5.

FIGURE 2.13C: Primary Education Enrollment Policies and Secondary School




Notes: Share of children going to the most academic secondary school track by birthday under different primary education enrollment policies. The simulations are based on the model in Sections 2.4 and 2.5. The lines are unweighted lowess curves fit through simulated data for 365 birthdays, see footnote 36 for more information.

FIGURE 2.14A: End-of-Primary-Education Test Scores Among Children with Low Skill Endowments




Notes: Comparison of the primary education enrollment policy in the Netherlands with the policy based solely on birthdates and the policy based solely on skills. The outcome is the average z-score of the standardized test at the end of primary education of children with skill endowments below the 10th percentile and the average skill shock. The simulations are based on the model in Sections 2.4 and 2.5. The lines are unweighted lowest curves fit through simulated data for 365 birthdays, see footnote 36 for more information.

FIGURE 2.14B: Non-Retention in Primary Education Among Children with Low Skill Endowments




Notes: Comparison of the primary education enrollment policy in the Netherlands with the policy based solely on birthdates and the policy based solely on skills. The outcome is the share of children never retained in primary education among children with skill endowments below the 10th percentile and the average skill shock. The simulations are based on the model in Sections 2.4 and 2.5. The lines are unweighted lowest curves fit through simulated data for 365 birthdays, see footnote 36 for more information.


FIGURE 2.14C: Academic Secondary School Among Children with Low Skill Endowments




Notes: Comparison of the primary education enrollment policy in the Netherlands with the policy based solely on birthdates and the policy based solely on skills. The outcome is the share of children going to the most academic secondary school track among children with skill endowments below the 10th percentile and the average skill shock. The simulations are based on the model in Sections 2.4 and 2.5. The lines are unweighted lowest curves fit through simulated data for 365 birthdays, see footnote 36 for more information.

FIGURE 2.15A: End-of-Primary-Education Test Scores and Adverse Skill Shocks

Notes: Comparison of the primary education enrollment policy in the Netherlands with the policy based solely on birthdates and the policy based solely on skills. The outcome is the average z-score of the standardized test at the end of primary education of children with the average family background but skill endowment shocks below the 25th percentile. The simulations are based on the model in Sections 2.4 and 2.5. The lines are unweighted lowess curves fit through simulated data for 365 birthdays, see footnote 36 for more information.

FIGURE 2.15B: Non-Retention in Primary Education and Adverse Skill Shocks

Notes: Comparison of the primary education enrollment policy in the Netherlands with the policy based solely on birthdates and the policy based solely on skills. The outcome is the share of children never retained in primary education among children with the average family background and skill endowment shocks below the 25th percentile. The simulations are based on the model in Sections 2.4 and 2.5. The lines are unweighted lowess curves fit through simulated data for 365 birthdays, see footnote 36 for more information.

FIGURE 2.15C: Academic Secondary School and Adverse Skill Shocks

Notes: Comparison of the enrollment policy in the Netherlands with the policy based solely on birthdates and the policy based solely on skills. The outcome is the share of children going to the most academic secondary school track among children with the average family background and skill endowment shocks below the 25th percentile. The simulations are based on the model in Sections 2.4 and 2.5. The lines are unweighted lowess curves fit through simulated data for 365 birthdays, see footnote 36 for more information.

Appendices to Chapter 2

2.A Data Sources and Summary Statistics

The data is provided by the *Centraal Bureau voor de Statistiek* (CBS) of the Netherlands and consists of the administrative records from various branches of the administration. Records from different branches are merged using the CBS random personal identifier.

Our data on primary school enrollment covers students enrolled in elementary schools since the school year 2008-2009. For each school year, we observe the grade a child attends, which enables us to identify the calendar year children enroll in grade 3 (the beginning of what we refer to as primary education) and a complete history of progression through elementary school grades. We combine this with administrative records from municipal registries to obtain the gender, country of birth, and day of birth for each child³⁹.

Outcome measures: We assess children’s education outcomes using grade retention in primary education; standardized test scores at the end of primary education; and the secondary school education track children end up attending. In grade 8, at the end of elementary school, most elementary schools in the Netherlands administer a standardized test. The decision whether to test children is made by elementary schools. Standardized tests were first introduced in 1970 and are taken into consideration when elementary schools make individual recommendations regarding the secondary school track. The test is administered over a 3-day period and centrally scored. Since the school year 2014/15, elementary schools can choose between three tests. The last school year we use is 2018-19, as there was no testing in 2019-20 because of Covid-19. We transform raw test scores into z-scores within each test type and year. The data we use to determine which secondary school track children end up attending after completing elementary school comes from the secondary education registers. There are three tracks: VMBO (pre-vocational secondary education) is a four-year program that leads to vocational training; HAVO (senior general secondary education) is a five-year track that prepares students for higher professional education at universities of applied sciences; VWO (pre-university education), the most academic track, is a six-year program preparing students for university. In the first year of secondary school (grade 9), students may attend so-called bridge classes, which are not assigned to any of the three secondary school tracks. We therefore use the track students attend in the second year of secondary school (grade 10) as the secondary school track students end up attending.

Family Background: The administrative records from municipal registries link children to their legal parents. This allows us to obtain information about children’s

³⁹Day of birth requires a separate confidentiality statement.

family backgrounds, such as parental income and maternal schooling. We restrict the sample to children born in the Netherlands. Parental income comes from records that the tax authority provides to the CBS. We construct family income by linking mothers to records for gross household income in 2011.⁴⁰ Gross household income is defined as total income from all sources. It includes wage income, interest income, profits; income insurance benefits because of unemployment or disability; public and private pension income, including survivors' pensions; social security benefits including tied transfers for housing and study; and income transfers received from ex-spouses. The data on maternal education comes from the education degree register of the Ministry of Education. This register has data on the highest education degree attained by 2011 using the International Standard Classification of Education (ISCED) classification. We convert the data into years of schooling using the typical time it takes for different education degrees.⁴¹ We use maternal education only, as there are more missing links between children and their fathers than their mothers and there are more missing values for paternal education than maternal education.

Sample selection: Our main sample is chosen to ensure that we have data on the school year when the children enroll in grade 3 (the beginning of primary education); their test scores at the end of elementary school; parental income; and maternal schooling. This implies that we are limited to four birth years, 2002, 2003, 2004, and 2005. We drop children born in the first 4 months of 2002 as some enrolled in primary education in the school years 2007-2008 and we only have enrollment data starting in school year 2008-2009. We only consider children born in the Netherlands. We exclude the very small share of children who first enrolled in primary education in the calendar year they turned 4 or 8. We also exclude children enrolled in special education schools.

⁴⁰This is the year closest to when the children in our data started elementary school.

⁴¹See Luijkx and de Heus (2008).

Table 2.A.1: Data Sources

	CBS Dataset	Content
Municipality Register	Gbapersoontab (2019)	Personal ID, country of birth, and exact birthdates
Primary Education Register	Inschrwpotab (2008-2019)	Students enrolled in elementary schools, grade retention history, and test at the end of elementary school
Education Register	Onderwijsinschrtab (2008-2022)	Students enrolled in secondary education
Link Child-Parent	Kindoudertab (2011)	Link from children to legal parents
Education Level File	Hoogsteopltab (2011)	Highest education degree attained
Household Income	Inhatab (2011)	Gross Household Income

Table 2.A.2: Summary Statistics

	Mean	S.D.	Observations
A. Main Sample (With Maternal Education and Parental Income)			
Maternal Schooling in Years (2011)	13.551	3.616	326,416
Log Gross Household Income (2011)	11.213	0.661	326,416
Enrolled in Primary Education in Calendar Year Turn 5	0.012	0.111	326,416
Enrolled in Primary Education in Calendar Year Turn 6	0.787	0.410	326,416
Enrolled in Primary Education in Calendar Year Turn 7	0.201	0.401	326,416
End-of-Primary Education Test Score (z-Score)	0.156	0.953	301,144
Attendance Academic Secondary School	0.335	0.472	304,573
Non-Retention in Primary Education	0.900	0.300	300,117
B. Sibling Sample			
Maternal Schooling in Years (2011)	14.073	3.385	95,986
Log Gross Household Income (2011)	11.323	0.642	95,986
Enrolled in Primary Education in Calendar Year Turn 5	0.012	0.109	95,986
Enrolled in Primary Education in Calendar Year Turn 6	0.799	0.401	95,986
Enrolled in Primary Education in Calendar Year Turn 7	0.189	0.391	95,986
End-of-Primary Education Test Score (z-Score)	0.249	0.918	89,137
Attendance Academic Secondary School	0.376	0.484	90,123
Non-Retention in Primary Education	0.910	0.287	88,897
C. Largest Sample (With Parental Income Only)			
Log Gross Household Income (2011)	11.142	0.631	649,777
Enrolled in Primary Education in Calendar Year Turn 5	0.010	0.101	649,777
Enrolled in Primary Education in Calendar Year Turn 6	0.770	0.421	649,777
Enrolled in Primary Education in Calendar Year Turn 7	0.220	0.414	649,777
End-of-Primary Education Test Score (z-Score)	0.008	0.984	595,342
Attendance Academic Secondary School	0.260	0.439	604,512
Non-Retention in Primary Education	0.889	0.315	598,174

Notes: The main sample in Panel A refers to children (i) born in the Netherlands in 2002-2005; (ii) who enrolled in grade 3 (the beginning of primary education) in a Dutch elementary school in or after the school year 2008-2009; and (iii) with data on maternal education and parental income. We exclude children born in the first 4 months of 2002 as some enrolled in grade 3 before the school year 2008-2009. We also exclude children enrolled in special education schools. Panel B presents information about the sibling sample. Panel C drops the maternal education requirement of the main sample.

2.B Additional Tables and Figures

Table 2.B.1: Calibration Results for Education Outcome Models

	Test Scores	Never Retained	Academic Track
SkillEoK	0.067	0.438	0.170
S	0.056	0.037	0.137
Y	0.322	0.348	0.848
Constant	-0.489	-1.677	-2.527
Observations	305,793	306,586	302,456

Notes: Calibrated parameters of the models for primary and secondary education outcomes in Section 5 in the main paper.

Table 2.B.2: RDD Estimates for Primary and Secondary Education Outcomes

	October 1 Discontinuity			January 1 Discontinuity		
	Test Score	Never Retained	Academic Track	Test Score	Never Retained	Academic Track
Data	0.0383*** (0.0149)	0.0254*** (0.00724)	0.0171*** (0.00591)	0.0316* (0.0178)	0.0132*** (0.00404)	0.0187** (0.00832)
Simulated Data	0.0377*** (0.00144)	0.0286*** (0.00105)	0.0177*** (0.00077)	0.0347*** (0.0025)	0.0124*** (0.0012)	0.0197*** (0.001)
Bandwidth (in days)	38	29	34	28	26	41


Notes: Estimates using the regression discontinuity design (RDD) are presented in Section 5.2.5 in the main paper. Standard errors are in parentheses. The simulated data is obtained using the model in Sections 4 and 5 in the main paper. Bandwidths are selected optimally.

Table 2.B.3: Instrumental-Variables Estimation

<i>Dependent variable: Test Score at the End of Primary Education</i>			
	(1)	(2)	(3)
lk	0.264*** (0.0121)	0.269*** (0.0120)	0.268*** (0.0257)
lk×S	-0.00574 (0.00445)	-0.00620 (0.00442)	-0.00954 (0.00894)
lk×Y	-0.0187 (0.0223)	-0.0224 (0.0218)	-0.0197 (0.0426)
S	0.0754*** (0.0106)	0.0702*** (0.0105)	
Y	0.313*** (0.0520)	0.292*** (0.0509)	
G	0.000478 (0.00259)	-6.26e-05 (0.00258)	-0.00214 (0.00525)
Constant	-0.468*** (0.0283)		
Observations	301,144	301,125	88,419
Number of Schools		6,866	
School FE		yes	
Family FE			yes


Notes: This table reports instrumental-variable estimates of the equation for end-of-primary-education test scores in equation (25) in the main paper. See Section 5.2.6 for more information on the estimation method and Appendix A for more information on the data.

Figure 2.B.1A: PIRLS




Notes: Data from the 2011 and 2016 waves of the Progress in International Reading Literacy Study (PIRLS), see Foy (2013, 2017) for more information. The regression model we use includes fixed effects for months of birth, the PIRLS waves, and schools. The PIRLS study tested children aged 9 to 10 years. The variable for the age at the start of primary education is derived from the home questionnaire question: "How old was your child when he/she began primary/elementary school?" The possible answers were 5 years old or younger; 6 years old; 7 years old; and 8 years old or older. We code "5 years old or younger" as 5 years old and "8 years old or older" as 8 years old.

Figure 2.B.1B: Maternal Schooling and Primary Education Enrollment Age 7




Notes: Share of children enrolling in the calendar year they turn age 7 by children's birthday, separately for children with maternal years of education above the median and strictly below the median.

Figure 2.B.2A: End-of-Primary-Education Test Scores and Age at the Start of Primary Education




Notes: Scatter plot of the data at the birthday level. See Section 2.3 in the main paper for more information.

Figure 2.B.2B: Non-Retention in Primary Education and Age at the Start of Primary Education




Notes: Scatter plot of the data at the birthday level. See Section 2.3 in the main paper for more information.

Figure 2.B.2C: Attendance of Academic Secondary School and Age at the Start of Primary Education




Notes: Scatter plot of the data at the birthday level. See Section 2.3 in the main paper for more information.

Figure 2.B.3: Calendar Year of Enrollment in Primary Education of Children by Birthdates and Skill Endowments




Notes: The figure illustrates the implications of the primary education enrollment policy in equation (13) in the main paper using the parameter estimates in Section 4. Instead of age at enrollment in primary education, the figure contains information on the calendar year of enrollment in primary education (different calendar years are indicated by different colors), as a function of children's birthdates and their skill endowments. Age at enrollment in primary education is the difference between the calendar year of enrollment and calendar year of birth on the horizontal axis.

Figure 2.B.4A: End-of-Primary-Education Test Scores and Age at the Start of Primary Education




Notes: Scatter plot of the data and the model predictions at the birthday level. The model predictions are obtained using the model in Sections 4 and 5 in the main paper. They refer to the expected education outcomes and ages at the start of primary education given children's family backgrounds and birthdays.

Figure 2.B.4B: Non-Retention in Primary Education and Age at the Start of Primary Education




Notes: Scatter plot of the data and the model predictions at the birthday level. The model predictions are obtained using the model in Sections 4 and 5 in the main paper. They refer to the expected education outcomes and ages at the start of primary education given children's family backgrounds and birthdays.

Figure 2.B.4C: Attendance of Academic Secondary School and Age at the Start of Primary Education




Notes: Scatter plot of the data and the model predictions at the birthday level. The model predictions are obtained using the model in Sections 4 and 5 in the main paper. They refer to the expected education outcomes and ages at the start of primary education given children's family backgrounds and birthdays.

Figure 2.B.5A: End-of-Primary-Education Test Scores and Age at the Start of Primary Education




Notes: Scatter plot of the data and the model predictions at the birthday level, separately for families with income above and strictly below the median. See notes to Figure 4A for more information.

Figure 2.B.5B: Non-Retention in Primary Education and Age at the Start of Primary Education




Notes: Scatter plot of the data and the model predictions at the birthday level, separately for families with income above and strictly below the median. See notes to Figure 4B for more information.

Figure 2.B.5C: Attendance of Academic Secondary School and Age at the Start of Primary Education




Notes: Scatter plot of the data and the model predictions at the birthday level, separately for families with income above and strictly below the median. See notes to Figure 4C for more information.

Figure 2.B.6A: Actual and Predicted End-of-Primary-Education Test Scores at the Individual Level




Notes: Binscatter plot of the actual end-of-primary-education test score of each child against the expected test score as predicted by the model in Sections 4 and 5 of the main text given children's birthdays and family backgrounds. Number of bins chosen optimally.

Figure 2.B.6B: Actual and Predicted Non-Retention at the Individual Level




Notes: Binscatter plot of the actual non-retention indicator for each child against the expected non-retention indicator as predicted by the model in Sections 4 and 5 of the main text given children's birthdays and family backgrounds. The non-retention indicator is a variable taking the value of 1 if and only if the child has never been retained in primary education. Number of bins chosen optimally.

Figure 2.B.6C: Actual and Predicted Enrollment in Academic Secondary School at the Individual Level




Notes: Binscatter plot of the actual academic-secondary-school enrollment indicator for each child against the expected academic-secondary-school enrollment indicator predicted by the model in Sections 4 and 5 of the main text given children's birthdays and family backgrounds. The academic-secondary-school enrollment indicator is a variable taking the value of 1 if and only if the child ended up attending the most academic secondary school. Number of bins chosen optimally.

Figure 2.B.7A: Actual and Predicted End-of-Primary-Education Test Scores at the Individual Level—Excluding Implausible Values for Income and Education




Notes: Figure 6A excluding children from households with gross income in the bottom 2 percent of the distribution (below €17,985 per year) and whose mothers have not completed primary school. See Appendix A for more information on gross income and maternal education.

Figure 2.B.7B: Actual and Predicted Non-Retention at the Individual Level—Excluding Implausible Values for Income and Education




Notes: Figure 6B excluding children from households with gross income in the bottom 2 percent of the distribution (below €17,985 per year) and whose mothers have not completed primary school. See Appendix A for more information on gross income and maternal education.


Figure 2.B.7C: Actual and Predicted Attendance of Academic Secondary School at the Individual Level—Excluding Implausible Values for Income and Education




Notes: Figure 6C excluding children from households with gross income in the bottom 2 percent of the distribution (below €17,985 per year) and whose mothers have not completed primary school. See Appendix A for more information on gross income and maternal education.

Figure 2.B.8A: Age at the Start of Primary Education—Girls and Boys


Notes: Comparison of actual age and predicted (expected) age at the start of primary education by children's birthday, separately for boys and girls. Predicted age is obtained using the model in Section 4 of the main text but allowing children's skill endowments to also depend on their gender. Specifically, we re-estimate the model in Section 4 after changing the specification for skill endowment in equation (11) to $e_i = \psi F_i + \rho W_i + \sigma v_i$ where F_i is equal to 1 for girls and equal to 0 for boys.

Figure 2.B.8B: Test Scores at the End of Primary Education—Girls and Boys

Notes: Predicted test scores by birthday are obtained by augmenting the model in Sections 4 and 5 for gender, predicting children's expected test score given family backgrounds, birthdays, and gender, and average at the birthday level. For more on how the model in Section 4 is augmented see the notes to Figure 8A. The model for end-of-primary-education test scores in Section 5 is augmented by adding the term $\psi_s F_i$, where F_i is equal to 1 for girls and equal to 0 for boys, on the right-hand-side of equation (18). The additional parameter is calibrated by adding the average difference between the test scores of boys and girls as a target.


Figure 2.B.8C: Non-Retention Primary Education—Girls and Boys

Notes: Predicted shares are obtained by augmenting the model in Sections 4 and 5 for gender, predicting children's probability of non-retention in primary education given family backgrounds, birthdays, and gender, and average at the birthday level. For more on how the model in Section 4 is augmented see the notes to Figure 8A. The non-retention model in Section 5 is augmented by adding the term $\psi_r F_i$, where F_i is equal to 1 for girls and equal to 0 for boys, on the left-hand-side of equation (19). The additional parameter is calibrated by adding the average difference between the share of non-retained boys and non-retained girls as a target.

Figure 2.B.8D: Academic Secondary School—Girls and Boys


Notes: Predicted shares are obtained by augmenting the model in Sections 4 and 5 for gender, predicting children's probability of attending the most academic secondary school track given family backgrounds, birthdays, and gender, and average at the birthday level. For more on how the model in Section 4 is augmented see the notes to Figure 8A. The academic-track model in Section 5 is augmented by adding the term $\psi_a F_i$, where F_i is equal to 1 for girls and equal to 0 for boys, on the left-hand-side of equation (20). The additional parameter is calibrated by adding the average difference between the share of boys and girls attending academic secondary school as a target.

Figure 2.B.9A: End-of-Primary Education Test Scores and Age at the Start of Primary Education—Instrumental-Variables Results and Data




Notes: The predictions using the instrumental-variables results are based on the equation for end-of-primary-education test scores in equation (25) in the main paper. See Section 5.2.6 in the main paper for more information on the estimation method and Table 3 for the instrumental-variables estimates.

Figure 2.B.9B: End-of-Primary Education Test Scores and Age at the Start of Primary Education—Instrumental-Variables and Calibration Results




Notes: The predictions using the instrumental-variables results are based on the equation for end-of-primary-education test scores in equation (25) in the main paper. See Section 5.2.6 in the main paper for more information on the estimation method and Table 3 for the instrumental-variables estimates.

Figure 2.B.10A: Primary Education Enrollment Policies and Skills at the Start of Primary Education—Income Only Sample




Notes: See Figure 12A in the main text for more information. The difference with Figure 12A is that the model in Section 4 is implemented using parental income but not maternal education as this results in a substantially larger sample, see Appendix A for more information.

Figure 2.B.10B: Standard Deviation in Skills at the Start of Primary Education Among Children Born on the Same Day—Income Only Sample




Notes: See Figure 12B in the main text for more information. The difference from Figure 12B is that the model in Section 4 is implemented using parental income but not maternal education as this results in a substantially larger sample, see Appendix A for more information.

Figure 2.B.11A: Primary Education Enrollment Policies and End-of-Primary-Education Test Scores—Income Only Sample




Notes: See Figure 13A in the main text for more information. The difference from Figure 12A is that the model in Sections 4 and 5 is implemented using parental income but not maternal education as this results in a substantially larger sample. In addition to average outcomes above and below median parental income, the calibration in Section 5 uses the average outcomes for children born October-March and May-September as targets.

Figure 2.B.11B: Primary Education Enrollment Policies and Non-Retention in Primary Education—Income Only Sample




Notes: See Figure 13B in the main text for more information. The difference from Figure 13B is that the model in Sections 4 and 5 is implemented using parental income but not maternal education as this results in a substantially larger sample. In addition to average outcomes above and below median parental income, the calibration in Section 5 uses the average outcomes for children born October-March and May-September as targets.

Figure 2.B.11C: Primary Education Enrollment Policies and Academic Secondary School—Income Only Sample




Notes: See Figure 13C in the main text for more information. The difference from Figure 13C is that the model in Sections 4 and 5 is implemented using parental income but not maternal education as this results in a substantially larger sample. In addition to average outcomes above and below median parental income, the calibration in Section 5 uses the average outcomes for children born October-March and May-September as targets.

Figure 2.B.12A: End-of-Primary-Education Test Scores Among Children with Low Skill Endowments—Income Only Sample




Notes: See Figure 14A in the main text for more information. The difference from Figure 14A is that the model in Sections 4 and 5 is implemented using parental income but not maternal education as this results in a substantially larger sample. In addition to average outcomes above and below median parental income, the calibration in Section 5 uses the average outcomes for children born October-March and May-September as targets.

Figure 2.B.12B: Non-Retention in Primary Education Among Children with Low Skill Endowments—Income Only Sample




Notes: See Figure 14B in the main text for more information. The difference from Figure 14B is that the model in Sections 4 and 5 is implemented using parental income but not maternal education as this results in a substantially larger sample. In addition to average outcomes above and below median parental income, the calibration in Section 5 uses the average outcomes for children born October-March and May-September as targets.

Figure 2.B.12C: Academic Secondary School Among Children with Low Skill Endowments—Income Only Sample




Notes: See Figure 14C in the main text for more information. The difference from Figure 14C is that the model in Sections 4 and 5 is implemented using parental income but not maternal education as this results in a substantially larger sample. In addition to average outcomes above and below median parental income, the calibration in Section 5 uses the average outcomes for children born October-March and May-September as targets.

Figure 2.B.13A: End-of-Primary-Education Test Scores and Adverse Skill Shocks—Income Only Sample




Notes: See Figure 15A in the main text for more information. The difference from Figure 15A is that the model in Sections 4 and 5 is implemented using parental income but not maternal education as this results in a substantially larger sample. In addition to average outcomes above and below median parental income, the calibration in Section 5 uses the average outcomes for children born October-March and May-September as targets.

Figure 2.B.13B: Non-Retention in Primary Education and Adverse Skill Shocks—Income Only Sample



Notes: See Figure 15B in the main text for more information. The difference from Figure 15B is that the model in Sections 4 and 5 is implemented using parental income but not maternal education as this results in a substantially larger sample. In addition to average outcomes above and below median parental income, the calibration in Section 5 uses the average outcomes for children born October-March and May-September as targets.

Figure 2.B.13C: Academic Secondary School and Adverse Skill Shocks—Income Only Sample



Notes: See Figure 15C in the main text for more information. The difference from Figure 15C is that the model in Sections 4 and 5 is implemented using parental income but not maternal education as this results in a substantially larger sample. In addition to average outcomes above and below median parental income, the calibration in Section 5 uses the average outcomes for children born October-March and May-September as targets.

Chapter 3

Intra-Household Allocation of Parental Investment: Parents' Equity-Efficiency Preferences and Beliefs

with Katja Kaufmann and Pia Pinger

3.1 Introduction

It has been well documented that the resources parents allocate toward their children vary considerably across families, and that differences in parental investments are highly predictive of important life outcomes, such as educational attainment, earnings, health and family outcomes (see, e.g., Cunha and Heckman, 2007a). Also, higher socioeconomic status parents invest more into their children, both in terms of financial resources and time, thereby exacerbating preexisting inequalities in endowment. However, much less is known about how parents allocate their investments towards different children in the same household, how intra-household investment decisions vary across socio-economic groups, and about the underlying mechanisms.

Can parents' equity-efficiency preferences and their beliefs about the productivity of their investments help in explaining the allocation of time investments within and across families? This study investigates parents' intra-household investment decisions. It shows how parents' beliefs about the productivity of their time investment in a child of given ability as well as their equity-efficiency preferences influence their investment into the human capital of their children.

To assess the role of parental productivity beliefs and equity-efficiency preferences for human capital investments, we first introduce a unified theoretical framework that demonstrates how the interplay of equity-efficiency concerns and lower returns to parental involvement for more able children can moderate disparities in intra-family investments. We then provide empirical evidence based on novel individual-level household data containing innovative and theory-driven measures of parents' equity-efficiency preferences, their beliefs about the human capital production function (i.e. the perceived productivity of parental investment in terms of children's human capital), as well as the perceived probability that their child will enroll in or complete university education. Identification of the role of parental productivity beliefs and equity-efficiency preferences for parents' investment decisions is achieved by using direct measures of individual beliefs based on hypothetical investment scenarios that vary one input at a time to assess the perceived returns to human capital investments (as pioneered by Dominitz and Manski, 1996).¹ To elicit parents' equity-efficiency preferences, we provide parents with scenarios about a hypothetical family, where parents make different intra-household investment decisions between two children and how the resulting child outcomes look like, and ask them to choose between these scenarios. We allow for equality-in-investment preferences, equality-in-outcome preferences and efficiency preferences. Last, we

¹High quality belief data is important in this context, as observed choices (under uncertainty) can be consistent with different combinations of preferences and beliefs (Manski, 2004).

link our data to administrative records to assess, whether eliciting parents' beliefs and preferences helps to explain actual parental investments as well as actual child outcomes.

A better understanding of how parental beliefs and preferences shape within household investments can help explain important empirical patterns. Figure 3.1, for example, unveils a *negative* relationship between a child's academic potential and the parents' investment into this child relative to its siblings. *Higher* ability is associated with *less* investment. At first glance, this finding may seem counter-intuitive given the well-documented importance of parental investments for child development (e.g. Cunha and Heckman, 2007a; Attanasio, Meghir, and Nix, 2020a). Yet, our findings can reconcile this pattern with the existing evidence. First, we show that, on average, Dutch parents perceive higher marginal returns from investing in less academically able children when it comes to learning-related activities. Moreover, on average, parents exhibit equality-focused preferences in the treatment of siblings, i.e. they prefer to invest more into the less academically able child. By linking our survey data to administrative data from the Dutch Statistics Bureau (CBS), we show that parents' equality-focused preferences lead them to invest in a way that reduces the gap in academic outcomes (performance on a high-stakes standardized test) among their offspring. Data on parents' beliefs about the productivity of their time investment as well as their equity-efficiency preferences not only help to explain the magnitude of siblings' *outcomes* differences across families. They also predict differential *investments* between siblings across families. In particular, actual investment differences are smallest for parents with equality-in-investment preferences. Parents with equality-in-outcome preferences invest more in the less academically able child, while parents with efficiency preferences and the perception that investments into the more academically able child are more productive indeed invest more into this child.

Which parents' have equality-focused preferences? Around 50 percent of parents prefer investing equally into their children, while slightly more than 30 percent favor investing in a way that equalizes outcomes. The smallest group is the group with efficiency preferences. Interestingly, efficiency preferences are more relevant among non-native parents without college education and with more children. In combination with perceptions that investments are more productive for more able children, this will lead parents to invest in a way that exacerbates existing endowment differences, in particular among less privileged families. This is a finding that is not only directly policy-relevant, but also consistent with findings in the literature that in particular in developing countries parents invest in a way that reinforces pre-existing endowment differences.

Our study not only sheds light on the nuanced decision-making processes

within households, but also offers a new perspective on equality of opportunity in the context of educational attainment. As an example, our results imply that two children of similar ability will receive different amounts of parental investment depending on whether they happen to be born into families with more or with less academically able siblings compared to themselves. Similarly, high ability children from disadvantaged families might do worse than high ability children from advantaged families, because they are more likely to have a lower performing sibling. This is the case, if their parents have equality-focused preferences. If, on the other hand, they have efficiency preferences, which are more prevalent among less privileged families, then the lower performing siblings will suffer more and would experience a further disadvantage compared to a more privileged family.

This paper contributes to two key strands of the existing literature: the literature making use of subjective beliefs and preferences to shed light on how parents invest into their children's human capital, and a literature on intrahousehold resource allocation in response to differences in child endowments.

Investigating how parents invest time and resources into their children's human capital based on data on parents' subjective beliefs, our paper is closely related to Cunha, Elo, and Culhane (2013), Attanasio, Cunha, and Jervis (2015), Boneva and Rauh (2018) and Attanasio, Cunha, and Jervis (2019).²

Cunha, Elo, and Culhane (2013, 2022) develop a method that relies on the use of hypothetical investment scenarios to elicit parental beliefs about the returns to parental investments. In a sample of parents with low socioeconomic status, they document beliefs about the returns to parental investments when children are aged 0 to 2. Attanasio, Cunha, and Jervis (2015) and Attanasio, Cunha, and Jervis (2019) explore the importance of parental expectations regarding the returns to early childhood investments. They link subjective beliefs directly to observed outcomes, enriching our understanding of how parental expectations shape resource allocation decisions in the early stages of child development. Boneva and Rauh (2018) show that parents perceive the returns to investments in late childhood to be higher than the returns to early investments, and that investments in different time periods are perceived as substitutes. The authors show that parental beliefs are predictive of actual investment decisions and document that parental beliefs about the productivity of investments are higher for more privileged families. This paper applies similar belief-elicitation techniques to analyze not

²This literature is linked to and builds on the seminal paper by Manski (2004) and papers investigating how educational decisions are linked to perceived returns to education, such as Jensen (2010); Attanasio and Kaufmann (2014); Kaufmann (2014) who investigate how decisions about years of schooling or educational degrees are linked to parents' perceptions about the returns to schooling, and Wiswall and Zafar (2015) who analyze how students' college major choice is linked to their perceptions about returns to college majors.

the variation in educational investment decisions across households, but to shed light on how parents make intrahousehold decisions and invest into their different children weighing equality versus efficiency concerns.

The second related strand of the literature is on intrahousehold allocation decisions, which focused on how families distribute resources among children in response to differences in endowments, such as health or cognitive ability, and how these allocations affect long-term outcomes (see seminal papers by Behrman, Pollak, and Taubman (1982, 1986); Rosenzweig and Wolpin (1988)).³ In terms of recent papers, Yi et al. (2015) find that parents increase health investments in children who experience early-life health shocks, but reduce their educational investments suggesting that parents make trade-offs between different types of investments. Bharadwaj, Eberhard, and Neilson (2018) show that lower birth weight children receive greater compensatory investments from their parents, particularly in education-related activities. Falch, Fisher, and Nyhus (2021) use Norwegian data and find that while parents often increase health-related investments for children with early disadvantages, these compensatory efforts do not always fully close gaps in educational outcomes later in life. Kirchberger (2020) examine how parents allocate time and financial resources among siblings in developing countries such as India, Ethiopia, Peru, and Vietnam. The study finds evidence of gender-based differences in resource allocation, highlighting how cultural norms and gender preferences can influence intrahousehold resource allocation, leading to inequality in investments and outcomes.

While many of these papers focus on health endowments, which are of critical importance in particular in developing countries, we focus on intrahousehold investment decisions based on differences in children's *academic ability* in a developed country context. Also the main focus of these papers is generally on parental investments, which reflect the *interaction* of preferences and (generally unobserved) perceptions about the production function, and not to identify parents' preferences themselves. The goal of our paper is to complement this literature by aiming to disentangle the separate roles of preferences and beliefs, by eliciting them in hypothetical survey experiments.

Two recent papers aim to identify parents' preferences in terms of investing into their different children via eliciting beliefs and preferences or via conducting experiments and making indirect inferences about preferences (Giannola (2024) and Berry, Dizon-Ross, and Jagnani (2024), respectively). Giannola (2024) provides evidence on parental investment decisions within the household and makes use of

³Chiappori and Meghir (2015) develop a theoretical model, which suggests that parents may allocate more resources to children with higher expected returns on human capital investments, thereby reinforcing existing ability differences. However, the model also allows for the possibility of compensatory behavior when parents are motivated by equity concerns.

hypothetical survey methods of eliciting subjective beliefs and preferences (similarly to this paper) to shed light on intrahousehold investment decisions based on differences in children’s academic ability. The study finds that parents in India view investment and ability as complementary factors, leading them to allocate more resources to children who are perceived to have higher ability. They thus prioritize maximizing returns on investments over equity concerns.

Compared to Giannola (2024), we elicit a theory-driven survey measure on parents’ equity-efficiency preferences as well as parents’ perceptions about the human capital production function and link these measures to actual investments and actual child outcomes.⁴ Also we allow for parental preferences for equal *investments*. Complementing his approach by focusing on a developed country context, we find that in the Netherlands, the majority of parents prefer equal investments, then equal outcomes and the smallest group prefers efficiency. The preference towards the latter is however stronger for less-educated families with migration background and more children, consistent with the findings of Giannola (2024).

Another directly related recent paper on parental investment decisions within the household is Berry, Dizon-Ross, and Jagnani (2024), who explore how parents’ preferences for equality shape their decision-making in the context of educational investments. They find that, in both high- and low-income settings, parents are often willing to sacrifice household earnings to equalize educational opportunities between siblings, even when this comes at the cost of overall household returns. While Berry, Dizon-Ross, and Jagnani (2024) focus on short-run returns and vary those returns exogenously (and thus do not use parents’ actual beliefs about their own investments), we elicit parental preferences as well as their beliefs about the human capital production function and link these measures not only to parents actual investments, but also to children’s actual outcomes.⁵ Our findings are consistent with their results, in that also in our context, the majority of parents has preferences for equal investments, then for equal outcomes and then for efficiency, while the latter group becomes more important for families from lower socio-economic status.

Our study sheds light on the nuanced decision-making processes within households and offers a new perspective on equality of opportunity in the context of educational attainment. Since parental investments in childhood are critical for the development

⁴The literature eliciting both beliefs and preferences at the same time is still small, albeit growing (see, e.g. Delavande and Zafar (2019) or Azmat and Kaufmann (2024) on the decision about enrollment into different types of universities or enrolling into different educational degrees, respectively).

⁵One advantage of the approach by Berry, Dizon-Ross, and Jagnani (2024) is that the incentivize parental choices using monetary stakes.

of children's cognitive and noncognitive skills, intra-household allocations can have profound consequences for children's long-term outcomes in the labor market, marriage market, in terms of health and so forth. A better understanding of intra-household allocation decisions is therefore critical for being able to address the source of inequality (in outcomes and/or opportunity) and for the effective targeting of social programs.

The structure of this paper is as follows: In Section 2 we develop a theoretical framework for analyzing the role of parents' preferences and beliefs in investment decisions and derive predictions in terms of parents actual investment choices and children's actual outcomes. In Section 3 we discuss the survey design with respect to the elicitation of beliefs and preferences. In Section 4 we present the data and descriptive statistics, while Section 5 discussed our empirical strategy. Section 6 presents our results before we conclude in Section 7.

3.2 Theoretical Framework

In this section, we develop a theoretical framework of intra-household allocation decisions, where parents derive utility from the educational outcomes of their children and make decisions on time investment. We then study the interplay between investment decisions, beliefs about returns, and preferences for equality in equilibrium, focusing on three archetypal preference types: efficiency-oriented, equality-in-outcome-oriented, and equality-in-investment-oriented.

3.2.1 Model Set-Up

Consider a family i with two children of differing perceived ability levels s : high (h) and low (l). Let Y_i^s denote the perceived educational outcome for the child with ability level s , and x_i^s represent the resources allocated by the parents to this child.

The (perceived) education production function for the lower-ability child (l) is given by:

$$Y_i^l = a_i + b_i x_i^l \quad (3.1)$$

where a_i is the baseline outcome for the low-ability child without any parental investment, and b_i represents the perceived marginal return on investment for this child.

For the higher-ability child (h), the production function is:

$$Y_i^h = (a_i + \alpha_i) + (b_i + \beta_i)x_i^h \quad (3.2)$$

where α_i captures the ability premium in outcomes, and β_i represents the additional marginal return to investment for the high-ability child.

Parents are subject to the following constraints:

$$x_i^h + x_i^l = L_i \quad (3.3)$$

$$x_i^s \geq 0 \quad (3.4)$$

The former equation ensures that the total resources invested in both children do not exceed the available resources L_i , while the latter ensures non-negative investments.

The utility function for parents is given by:

$$U_i = \pi_1 \ln(Y_i^h + Y_i^l) - \pi_2 \ln(|Y_i^h - Y_i^l|) - \pi_3 \ln(|x_i^h - x_i^l|) \quad (3.5)$$

Here, π_1 , π_2 , and π_3 are parameters representing the weights parents place on efficiency, equality in outcomes, and equality in investment, respectively.

3.2.2 Interplay of Perceived Returns and Preferences

To illustrate how different parental preferences influence resource allocation, we consider three archetypal parent types, each prioritizing one component of the utility function.

Efficiency-Oriented Parents

We consider parents prioritizing the overall educational success of their children as efficiency-oriented type. They derive utility from:

$$U = \ln(Y^h + Y^l) \quad (3.6)$$

Substituting the production functions, we obtain:

$$U = \ln(2a_i + \alpha_i + b_i L_i + \beta_i x_i^h)$$

In equilibrium, parents maximize their utility and the equilibrium resource allocation is:

$$x_i^h = \begin{cases} L_i & \text{if } \beta_i > 0 \\ 0 & \text{if } \beta_i < 0 \end{cases} \quad (3.7)$$

These parents aim to maximize the total educational output in equilibrium. If the marginal return to investment is increasing with ability level, namely (β_i) is positive, they allocate all resources to this child. Conversely, if β_i is negative,

indicating a higher return for the low-ability child, all resources will go to that child.

- (i) For efficiency-oriented parents, the **difference in time investment** between children is $x_i^h - x_i^l = L_i$ when the marginal return is increasing in ability level ($\beta_i > 0$) and $x_i^h - x_i^l = -L_i$ when the marginal return is decreasing in ability level ($\beta_i < 0$).
- (ii) The **difference in child outcomes** is $Y_i^h - Y_i^l = \alpha_i + (b_i + \beta_i)L_i$ for increasing marginal returns and $Y_i^h - Y_i^l = \alpha_i - b_iL_i$ for decreasing marginal returns.

Equality-in-Outcome-Oriented Parents

Parents who prioritize equality in outcomes aim to minimize the differences in their children's educational achievements:

$$U = \ln(|Y_i^h - Y_i^l|) \quad (3.8)$$

In equilibrium, they will allocate resources such that:

$$a_i + b_i x_i^l = (a_i + \alpha_i) + (b_i + \beta_i)x_i^h$$

Solving for x_i^h , we get:

$$x_i^h = \frac{b_i L_i - \alpha_i}{2b_i + \beta_i} \quad (3.9)$$

These parents strive to ensure equal educational outcomes for both children. They adjust their investments to offset the differences in abilities, leading to an equilibrium where perceived outcomes are equal.

- (i) For equality-in-outcome-oriented parents, the **difference in time investment** is $x_i^h - x_i^l = -\frac{2\alpha_i + \beta_i L_i}{2b_i + \beta_i}$, which is negative for increasing marginal returns but can be positive if the marginal returns strongly decrease ($\beta_i < -\frac{2\alpha_i}{L_i}$).
- (ii) For these parents, the **difference in child outcomes** is zero ($Y_i^h - Y_i^l = 0$).

Equality-in-Investment-Oriented Parents

Parents who prioritize equality in investments aim to equalize the resources allocated to each child:

$$U = \ln(|x_i^h - x_i^l|) \quad (3.10)$$

In equilibrium, they will choose:

$$x_i^h = x_i^l \quad (3.11)$$

These parents value equity in resource allocation. They distribute resources equally, leading to outcome differences driven by the inherent ability premium (α_i) and additional marginal return (β_i).

For equality-in-investment-oriented parents, there is no **difference in investment** between the children ($x_i^h - x_i^l = 0$). The **outcome gap** ($Y_i^h - Y_i^l = \alpha_i + \beta_i \frac{L_i}{2}$) decreases as α_i and β_i decrease.

Building on the discussion above, we can now classify parents into distinct types based on two key dimensions: their preferences and their beliefs. The first dimension reflects parents' equality preferences, which may prioritize equality in outcomes, equality in investments, or efficiency. The second dimension relates to their beliefs about the marginal returns to investment—whether they believe these returns increase or decrease with the child's ability. By combining these two dimensions, we identify six distinct parent types. Table 3.1 summarizes the model's predictions for each type, outlining how different preferences and beliefs shape investment behaviors and the resulting outcomes.

3.3 Survey Design

To shed light on how parents decide to allocate their resources across their different children, we designed a survey using innovative data elicitation methods (such as the use of vignettes/ hypothetical scenarios) to elicit information about subjective measures, such as economic preferences and beliefs about the returns to parental investment (by child ability), as well as about parental investments and time use.⁶

The survey is divided into different survey blocks, described in detail below in the order in which they were presented to respondents. Appendix 3.A presents the exact wording of the survey questions. We collect detailed information on parental beliefs about the human capital production function, beliefs about educational outcomes, parents' equity-efficiency preferences, parental investment activities, and parent and child characteristics.

⁶The survey questions were part of a larger survey designed and implemented in the context of the CRC TR224 with the help of people responsible for running the Dutch household survey, LISS, the sample of which is a superset of our sample of households (for a short description of the dataset, see (<https://www.crctr224.de/research/data#data5>)).

3.3.1 Beliefs about the Human Capital Production Function

To elicit beliefs about the productivity of parental investments, we build on and extend the approach developed in Cunha, Elo, and Culhane (2013) and Boneva and Rauh (2018). In particular, we present parents with hypothetical investment scenarios that vary along two dimensions: (i) the level of parental investments and (ii) the initial human capital level of the child. For each scenario, parents are asked to state what the *likelihood of getting a university degree* will be (in percent) and what the *future earnings* of the child will be *at age 30*.

We focus on a particular type of parental investment that is relevant to all school-aged children: the number of hours parents spend every week helping their child with school work. We give examples and specify that helping the child with school work can mean "monitoring that the homework is done, going over the homework together, being available for questions, help/support in studying for school such as vocabulary, dictation or for tests".

We ask parents to imagine two different families, the De Jongs, and the Jansens, who make decisions about how much to help their child. Mr and Mrs De Jong have one child, Jan. Jan is 8 years old, and he is more academically able than the average child of his age group. In the following school years, Mr and Mrs De Jong can decide how much to help Jan with his school work. We ask parents to assume that apart from helping with school work, Jan's parents' behavior is exactly the same in the following two scenarios. We then ask what they expect Jan's likelihood of getting a university degree will be (from 0 to 100%) in each of the following scenarios:

- (A) if they help Jan 0 hours every week between age 8 and 9
- (B) if they help Jan 1 1/2 hours every week between age 8 and 9
- (C) if they help Jan 3 hours every week between age 8 and 9

Analogously, we ask parents to assume there is no inflation and that Jan will be working full-time, and elicit what they expect Jan's gross monthly earnings to be at age 30 in each of three scenarios.

We then ask parents to imagine a different family, the Jansens, who are very similar to the De Jongs in many respects (they also have a 8-year-old son, live in the same neighbourhood and they have similar levels of income and education). We tell them that there is one difference and that is that, unlike Jan, David is less academically able than the average child of his age group. Again Mr and Mrs Jansen can decide how much time to invest to help David with his school work. We ask respondents to indicate what they expect David's likelihood of getting a

university degree will be in each of the three scenarios (A), (B) and (C) and what David's future earnings will be at age 30.⁷

3.3.2 Equity-Efficiency Preferences

To elicit parents' equity-efficiency preferences with respect to parental investments into different children, we develop an innovative survey measure based on hypothetical scenarios where we vary how much parents invest into their two (hypothetical) children who are differently academically able, which affects the children's likelihood to graduate from university, i.e. the level and the difference between the two children. The three scenarios are as follows (where the average probability of graduating of the two children is always 45%):

- (A) Investing the same into both children, we specify that the more academically able child has a probability of 60% of graduating from university as opposed to 30% for the less academically able child.
- (B) Investing more into the less able child and thereby compensating the initial endowment differences, the likelihood for the more (less) able is now 50% (40%).
- (C) Investing more into the more able child instead, thereby reinforcing initial differences, the likelihoods are 70% and 20%.

In the hypothetical scenarios, we indicate the names of the children and randomize whether respondents receive a scenarios with two female, two male names or with both genders and which of latter is the more academically able. Moreover, we also randomize in which order the three possible answers are presented to avoid potential biases.

In a second step, we increase the cost for the option to which the respondent assigned the highest probability, by decreasing the average likelihood of graduating from university for both children for that particular scenario, while in the other two scenarios, the average probability of graduating for the two children remains at 45%. We then elicit again how the respondent would assign the probabilities across the three options given the new outcomes.

⁷The survey questions we elicit (in particular with respect to the production function and the equity-efficiency trade-off) are complex questions and respondents may be hesitant to answer for fear of answering wrongly. We therefore state before eliciting the questions that "[w]e know that the following questions are difficult hypothetical questions and that many aspects are relevant. It is normal that you are not certain about your answer, but for us it is very valuable if you take your best guess, i.e. let us know what you think is the likely outcome."

3.3.3 Parental Time Investment

Both in terms of actual investments and in terms of the hypothetical scenarios eliciting the perceived production function and equity-efficiency preferences, we focus on parental time investment. Focusing on time investments reduces the complexity of survey questions, while it has been shown that parental time investments are particularly important for child development.⁸

In terms of specific survey questions, we first elicit how much time parents spend on helping each of their children with school work. More specifically, we ask parents how many hours per week they spend on helping a particular child with her/his homework and specify that this can mean monitoring that the homework is done, going over the work, being available for questions or helping with studying (practicing vocabulary, practicing dictation or for tests) (see the Appendix ?? for the exact wording). We ask this separately for the years that the child is in elementary school and for when the child is in secondary school (if applicable based on the age of the child).

In addition to our main question/outcome on parental time investment above, as a supplementary question we also ask for how much time parents spend quality time with children other than helping with home work. In particular, we ask parents how many hours per week they spent on direct interactions other than helping with homework with each of their children. We specify that possibly examples are joint meals, joint activities (such as playing games/going on an excursion/to the zoo, museum, concert etc), reading to/with the child, talking about personal matters (see the Appendix for the exact wording). Again we ask this separately for the time that the child was in elementary school and for when the child was in secondary school (if applicable).

3.4 Data and Descriptive Statistics

The survey described in the previous section was conducted in the Netherlands in March 2019, targeting parents with at least one child aged 4 to 18. The respondents completed the survey online as part of the LISS Panel, a representative sample of Dutch households that participate in monthly internet surveys administered by CentERdata (Tilburg University, The Netherlands).⁹ The data collected includes

⁸See Cunha and Heckman (2007b), Guryan, Hurst, and Kearney (2008), Hsin and Felfe (2014), Del Bono et al. (2016), Baker and Milligan (2016), Kalil and Ryan (2017), Cobb-Clark, Salamanca, and Zhu (2019), and Attanasio, Meghir, and Nix (2020b).

⁹The survey questions were part of a larger survey designed and implemented in the context of the CRC TR224 with the help of people responsible for running the Dutch household survey, LISS (for a short description of the dataset, see <https://www.crcctr224.de/research/data#data5>). The LISS (Longitudinal Internet Studies for the Social Sciences) panel is a representative sample

information on parents' demographic background, household income, parental education, their perception regarding their children's likelihood of university enrollment and completion, and the answers to our designed hypothetical questions regarding equity-efficiency preferences and beliefs about the productivity of parental time investment.

3.4.1 Sample Selection

We begin with a "Full Sample" consisting of 985 parent respondents who participated in our survey section. These parents were included based on their complete responses to questions related to educational investment and equity-efficiency preferences. The dataset includes demographic information on the family, the parents and the children (including family structure and educational background).

To study the effects of beliefs and preferences on actual parental investment behavior, we linked this dataset with child-specific information on parental investments during elementary school. This merge process resulted in 2,154 matched parent-child links. Following the merge, the sample selection process was carefully implemented in several stages. First, 682 observations were excluded where the child had either not yet started primary education or had already progressed beyond secondary education. This step was crucial to ensure that parental investments during elementary school were assessed during the relevant educational period. Next, we removed 292 observations where the perceived probability of a child graduating from university (used as a proxy for parents' perceptions about children's relative ability ranking) was missing. This ensured that the analysis only included cases where parents' perceptions of their children's (relative) academic potential were available.

Following this, we excluded 72 observations due to missing data on parental involvement during elementary school, ensuring the completeness of the data used in subsequent analysis. Finally, 313 observations were dropped because they involved families, where only one child was observed, precluding the possibility of analyzing intra-household allocation of resources. The final sample, referred to as the "Sibling Sample," consisted of 344 families with at least two children, allowing us to examine how parents allocate educational resources among their children, with a particular focus on the role of parental beliefs and preferences.

Additionally, we link our survey data with Dutch administrative data maintained by Statistics Netherlands (Centraal Bureau voor de Statistiek, CBS), which covers the entire Dutch population.¹⁰ The linkage process involves collaboration with

of Dutch individuals who participate in monthly internet surveys which are administered by CentERdata (Tilburg University, The Netherlands). For more details see section 3.3.3.

¹⁰These microdata are available for statistical and scientific research under specific conditions.

the CBS team, using parents' names and birth dates to match LISS panel survey respondents to the administrative records. Through this process, we successfully matched 793 out of 985 parents in our full survey sample to the CBS administrative data. For our intra-household sample, We impose further sample restrictions to include only parents with more than one child in the administrative records and for whom CITO test scores -our primary measure of educational outcomes- are available for at least two children.¹¹ This ensures our analysis focuses on families where intra-household educational inequality can be assessed directly through multiple siblings' test scores. As a result, our final sample includes 236 families with CITO scores for at least two children.

3.4.2 Variable Description

In this subsection, we discuss the variables used in our primary analysis. We start with variables related to parental characteristics and family background.

Parental Education. Parental education is measured using the variable *highest level of education with diploma* from the general LISS panel. We categorize parents as college-educated if they have obtained a bachelor's degree or higher from either an applied science university (HBO) or a research university (WO).

Immigration Background. We classify parents' immigration backgrounds using a five-category variable in the LISS survey. The categories include native, first-generation western origin, second-generation western origin, first-generation non-western origin, and second-generation non-western origin¹². For our analysis, we group first- and second-generation individuals of western origin as "western immigrants" and those of non-western origin as "non-western immigrants."

Number of Children. The number of children in a household is determined based on the information provided by the parents. In the survey, parents are asked to list anonymously the age, gender, and phase of education of their children who are still in education, with a maximum number of ten children.

In the following, we discuss the variables related to child-specific beliefs, investments, and outcomes.

CITO Test Scores. We use the overall CITO standardized scores as a measure

Researchers can link the LISS panel to the administrative data using unique random IDs, with support from CBS. For more details, contact: microdata@cbs.nl.

¹¹The CITO test scores cover the period from 2006 to 2019.

¹²Statistics Netherlands (CBS) defines immigrants as individuals with at least one parent born abroad. A distinction is made between those born abroad (first-generation) and those born in the Netherlands (second-generation). Additionally, a differentiation exists between individuals with a western and a non-western migration background. Non-western backgrounds include origins from Africa, South America, Asia (excluding Indonesia and Japan), or Turkey. Western backgrounds include origins from Europe (excluding Turkey), North America, Oceania, Indonesia, or Japan. For more details, see the CBS website.

of children's educational outcomes at the end of primary school. In the Netherlands, primary education consists of six grade levels and children typically complete their primary education at age 12. Schools can choose the provider of the end-of-primary education test they administer. Approximately 85% of schools opt for the CITO test, renowned for its comprehensive assessment of key subjects such as mathematics, Dutch language skills, and study skills. Once a school decides to administer the CITO test, participation is mandatory for all enrolled students. End-of-primary education test scores are high-stakes as they are taken into consideration when teachers make recommendations for the secondary-school track. The standardized CITO scores, ranging from 501 to 550, are adjusted annually to maintain comparability across cohorts.

Hours Helping with Homework. Parental investment is measured by the reported average number of hours per week parents assist with homework or study-related tasks during primary school, including activities like checking homework, answering questions, and test preparation. This is reported by the parents for each child, including those who have already graduated from primary school.

3.4.3 Descriptive Statistics

Table 3.2 presents descriptive statistics on the two samples we use in our analysis, that is the Full Sample and the Sibling Sample. The Full Sample consists of 985 parent respondents who participated in the survey. 44.5% of them are male, and 43.0% have obtained a college degree. The average age of the parents in this sample is 48.37 years. The ethnic composition is predominantly Dutch, with 81.7% of the parents identifying as such, while 7.5% are from Western backgrounds and 10.3% from non-Western backgrounds. On average, families in this sample have 2.26 children. Parents reported spending an average of 2.006 hours per week helping their children with homework during elementary school.

In the Sibling Sample, which includes 344 parents, 44.2% are male, and 46.5% have a college degree. The average age of parents in this sample is 44.87 years. The ethnic composition is similar to the Full Sample, with 79.1% identifying as Dutch, 8.7% from Western backgrounds, and 11.6% from non-Western backgrounds. The average number of children in these households is 2.66. On average, parents in this group spend 1.696 hours per week helping their children with homework during elementary school.

T-tests comparing the Full Sample and the Sibling Sample indicate that, apart from age and the number of children, there are no statistically significant differences between the two groups in terms of gender distribution, college attainment, ethnic composition, or time spent helping with homework. This suggests that the

Sibling Sample is broadly representative of the Full Sample in most demographic characteristics, supporting the robustness of our analysis on how parents' beliefs and preferences influence the allocation of educational resources among their children.

3.5 Empirical Strategy

In this section, we outline the empirical strategy used to investigate how parental beliefs about the returns to time investments and preferences for equity and efficiency shape their decisions regarding educational investments in their children.

3.5.1 Parental Beliefs in Returns to Time Investment

We begin by examining parental perceptions of the returns to time investments in their children's education, focusing on variation based on children's academic ability. We employ a simple OLS regression to quantify these perceptions:

$$\begin{aligned} Outcome_{ij} = & \alpha + \beta_1 Investment_{ij} + \beta_2 MoreAble_{ij} \\ & + \beta_3 (Investment_{ij} \times MoreAble_{ij}) + \beta_4 X + \epsilon_{ij} \end{aligned} \quad (3.12)$$

Where $Outcome_{ij}$ represents the perceived likelihood of a *hypothetical* child obtaining a university degree or expected earnings at age 30, depending on the level of parental investment in scenario j by parent i . $Investment_{ij}$ refers to weekly hours of parental time dedicated to helping with schoolwork in scenario j . $MoreAble_{ij}$ is a dummy variable indicating whether the child in scenario j is perceived to be more academically able than average. X represents additional control variables of the parent.-

In this framework, β_1 captures the average perceived return to investment for less academically able children, reflecting how parents believe that additional time investment improves outcomes for these children. β_2 represents the perceived baseline advantage of more academically able children when no parental investment is made, indicating the benefit of higher endowment. The interaction term β_3 is of particular interest, as it measures the differential in perceived returns between less and more academically able children. A negative β_3 suggests that parents perceive diminishing returns for investing in more able children, implying that they believe time investment is more effective for children who are less academically advantaged.

We further explore the role of parental education by interacting a college-

educated parent dummy with the investment variable and the interaction term ($Investment_{ij} \times MoreAble_{ij}$). This allows us to assess whether college-educated parents perceive different returns based on their child's academic ability.

3.5.2 Analyzing Parental Equity-Efficiency Preferences

In this section, we examine parental preferences in balancing equity and efficiency when making educational investments in their children, building on the survey design detailed in Section 3.3. Parents are presented with a choice between three investment strategies, each offering different probabilities of educational outcomes for their children.

To analyze how parents' choices relate to the implied inequality of each investment option, we first calculate the inequality in educational outcomes for each choice. Let $Ineq_j$ represent the inequality in educational outcomes for scenario j , measured as the difference in the probabilities of university graduation between the academically more able and less able child:

$$Ineq_j = Q_j(Grad_{more\ able}) - Q_j(Grad_{less\ able}) \quad (3.13)$$

The log-odds of choosing each scenario j for respondent I is then modeled as a function of the inequality measure $Ineq_j$ and other relevant control variables X_I (if applicable):

$$\log\left(\frac{P_{ij}}{1 - P_{ij}}\right) = \alpha + \beta Ineq_j + \gamma X_i + \epsilon_{ij} \quad (3.14)$$

Here, β captures the effect of inequality on the likelihood of parent i choosing scenario j , while γ represents the coefficients for any control variables X_i , and ϵ_{ij} is the error term. This approach allows us to assess how outcome disparities influence parents' decisions, providing a direct measure of their preference for equality compared to their preference for maximizing returns.

Additionally, based on the probabilities that parents assign to each investment option, we classify parents into three distinct preference types—"equality in outcome," "equality in investment," and "efficiency"—as described in Section 3.2. A parent is classified into a preference type based on the highest probability assigned to a specific option. To ensure clarity in categorization, we retain only those observations where a single preference type dominates, excluding cases where ties occur among the probabilities.¹³ We then conduct a multinomial logit analysis on these types,

¹³This also includes cases where the parent inputs probabilities like 33% and 33%, and the system mechanically adjusts one to 34%; we eliminate these cases as they do not reflect a strict preference.

using parental characteristics such as education level and other demographic variables as predictors, with the "efficiency" option (maximizing return) as the reference category. This analysis allows us to understand how different parental attributes correlate with preferences for equality, providing insights into the factors that influence the likelihood of choosing equality-oriented strategies over an efficiency-driven approach.

At last, we study the trade-off between efficiency concerns and equality preferences. We give parents an efficiency cost by reducing total outcomes for the option to which they initially assigned the highest probability in the baseline question. This option, referred to as the previously most favorable option, allows us to observe how parents adjust their choice probabilities when faced with a reduction in efficiency. Specifically, we regress the percentage change in the choice probability for this previously most favorable option, measured as the difference between the new probability after the efficiency cost and the initial probability, relative to the initial probability, on the parents' preference type.

The regression equation is expressed as:

$$\frac{P_{\text{new},if} - P_{\text{initial},if}}{P_{\text{initial},if}} = \alpha + \beta \text{Pref}_i + \gamma X_i + \epsilon_{if} \quad (3.15)$$

Here, $\frac{P_{\text{new},if} - P_{\text{initial},if}}{P_{\text{initial},if}}$ represents the percentage change in the choice probability for respondent i for the previously most favorable option f after the introduction of the efficiency cost. Pref_i captures the parent's preference types or weights. X_i includes other control variables related to the respondent, and ϵ_{if} is the error term.

3.5.3 Equity-Efficiency Preferences and Child Outcome Gap

Establishing a relationship between parents' equity-efficiency preferences and within-family inequalities in children's outcomes is not straightforward, largely due to challenges in obtaining reliable outcome measures. First, standardized, high-stakes outcome measures that accurately reflect children's development are rare. Second, comparing the performance of children within the same family can be difficult because they often take standardized tests at different ages.

To address these issues, we link our household survey data to administrative records from the CBS, which include CITO test scores. The CITO test, a high-stakes school exit exam in the Netherlands, is standardized across pupils and years, enabling cross-pupil and cross-year comparisons. Since most families enroll their children in the same primary school, we can observe consistent learning outcomes within families. Moreover, children typically take the CITO test around age 12, which allows for precise comparisons within families.

We use the mean log deviation (MLD) measure to quantify within-household inequality in CITO scores:

$$\text{CITOMLD}_{ij} = \ln(\text{CITO}_{ij}) - \ln(\overline{\text{CITO}}_i), \quad (3.16)$$

where CITOMLD_{ij} represents the log deviation in CITO test scores for child j in family i , and $\overline{\text{CITO}}_i$ is the family average CITO score.

We then estimate the following regression model:

$$\text{CITOMLD}_{ij} = \alpha + \beta_1 \text{Equal-Outcome-Type}_i + \mathbf{X}'_i \gamma + \epsilon_{ij}, \quad (3.17)$$

where $\text{Equal-Outcome-Type}_i$ is a dummy variable indicating whether parents have a strong preference for equal outcomes among their children. The control vector $\mathbf{X}_i$ includes variables such as parental education, family size, and family origin. This regression tests whether parents who prefer equal outcomes tend to have children with more comparable academic achievements. To explore potential heterogeneity in these effects, we conduct separate regressions for different socioeconomic groups, particularly focusing on variations by parental education levels.

3.5.4 Equity-Efficiency Preferences and Parental Investment Gap

Next, we explore whether parents' equity-efficiency preferences are reflected in their actual investment behaviors. We assess, whether a preference for equal investment among children predicts differences in the time parents invest in each child, irrespective of the children's ability ranks.

The key variable of interest is the time parents spend helping their children with homework during primary school, as discussed in Section 3.3.3. To eliminate family-specific factors, we focus on within-family differences in investment. We estimate the following model:

$$\Delta \text{Investment}_{ij} = \alpha + \beta_2 \text{Equal-Investment-Type}_i + \mathbf{X}'_i \gamma + \epsilon_{ij}, \quad (3.18)$$

where $\Delta \text{Investment}_{ij}$ represents the difference in hours spent on homework help between siblings in family i . The variable $\text{Equal-Investment-Type}_i$ is a dummy that equals 1 if the parent expresses a strong preference for equal investment.

We also extend the analysis by replacing the $\text{Equal-Investment-Type}$ dummy with a continuous score ranging from 0 to 1, representing the strength of the parents' preference for equal investment. This approach allows us to explore, whether the investment gap decreases as the preference for equality in investments

strengthens.

3.5.5 Equity-Efficiency Preferences, Perceived Returns, and Parental Investments

Our final analysis aims at exploring how parents' equity-efficiency preferences and their beliefs about the returns to investment interact to shape the allocation of parental resources among siblings. Specifically, we focus on the within-family differences in time investment between children perceived as more versus less academically able.

To conduct this analysis, we first rank children within families based on their perceived academic ability. This step is crucial because our model relies on the concept of an ability differential within families, but we cannot directly observe actual or perceived ability. Therefore, we use parents' estimates of each child's probability of graduating from university, conditional on hypothetical enrollment, as a proxy for academic ability. This ranking is derived from parents' responses to a survey question framed as a hypothetical scenario: "If your child were able to enroll, what is the probability that they would graduate?" This framing is designed to mitigate the issue of children already being on different educational tracks, which might be influenced by prior parental investments. By focusing on this hypothetical scenario, we aim to capture parents' underlying beliefs about each child's potential, independent of their current educational pathway. Our assumption is that this proxy accurately reflects the true *ranking* of a parent's perceived academic ability across all of their children.¹⁴

We retain only those observations where there is a clear distinction in parents' perceptions of which child has a higher chance of graduating. This results in the exclusion of about one-third of sibling pairs where such a distinction could not be made. After establishing this ranking, we classify parents according to their expressed preferences for equality in outcomes and their beliefs about whether the returns to investment are higher for children with greater or lesser ability.

We then estimate the following regression model to examine the interaction between these preferences and perceived returns:

$$\begin{aligned} \Delta \text{Investment}_{ij} = & \alpha + \beta_1 \text{Equal-Outcome-Type}_i + \beta_2 \text{Decreasing-Return}_i \\ & + \beta_3 (\text{Equal-Outcome-Type}_i \times \text{Decreasing-Return}_i) + \mathbf{X}_i' \gamma + \epsilon_{ij}, \end{aligned} \quad (3.19)$$

where $\Delta \text{Investment}_{ij}$ represents the difference in hours spent helping with

¹⁴The assumption implies that parents did not invest in a way that changed the ranking of the children based on their endowment.

schoolwork during elementary school between the more and less academically able child in family i . The variable $\text{Equal-Outcome-Type}_i$ is a dummy that equals 1 if the parent has a strict preference for equal outcomes, and 0 if the parent prioritizes efficiency. The variable $\text{Decreasing-Return}_i$ indicates whether the parent believes marginal returns to investment are lower for the more academically able child than for the less academically able child. The interaction term $\text{Equal-Outcome-Type}_i \times \text{Decreasing-Return}_i$ captures the combined effect of these preferences on investment allocation.

We first test whether - as predicted by our theoretical framework- the equal-investment-type leads to a smaller investment gap than the two other types (which should be independent of perceiving increasing or decreasing returns to investment in ability). In a second step, we drop the equal-investment-type to analyze in which direction parents invest, which according to our theoretical framework depends on the interaction between the perceived production function (returns to investment decreasing or increasing in ability) and the preference type (i.e. equal-outcome type or efficiency type)

We further extend our analysis by examining how these effects vary across families with different numbers of children and by including additional controls for parental education and other socioeconomic factors. The robustness is assessed by including additional specifications that account for the presence of more than two children in the family and by incorporating controls for the "equality-in-investment" preference type.

3.6 Results

In this section, we aim to link the different parental types we identified based on their elicited preferences and beliefs to outcomes of and investment into the different children in the family. We first present descriptive evidence and link different parent types (based on their beliefs about the human capital production function and on their equity-efficiency preferences) to demographic characteristics (such as parents' education, migration background and number of children) (see Subsections 3.6.1 and 3.6.2).

In a second step, the goal is to link the preference-belief types to children's actual outcomes and to parental investments into their different children. We thereby test the predictions of our theoretical framework, as derived in Section 3.2, and analyze if it can predict real-world outcomes. First, we assess whether parents who favor equal outcomes for their children actually witness more comparable academic achievements among their offspring at the end of primary school than parents of different preference types (such as equal-investment type or efficiency

type). Second, we analyze how parents' preferences correlate with their actual time spent investing in their children.

3.6.1 Descriptive Evidence on the Perceived Production Function

We start by classifying parents into different types based on their beliefs about the human capital production function. In particular, we analyze, whether parents believe that children's educational achievement increases with increasing parental time investment and how (marginal returns) and whether marginal returns to time investment are higher or lower for more/less academically able children. As discussed in Sections 3.3.3 and 3.4.2, parental time investment is conceptualized with the time parents help children with their schoolwork during elementary school. In terms of outcomes, we elicited the perceived effect both on the child's *probability of obtaining a university degree* and on a child's *income at the age of 30*.

Figures 3.2 and 3.3 show parents' perceptions about the *probability of obtaining a university degree* and about the *expected earnings at age 30*, respectively, of a hypothetical child as a function of *parental time investment* in helping with schoolwork during elementary school (ranging from 0 over 1,5 to 3 hours). We display these perceptions separately for the hypothetical child with higher and lower than average academic ability (in light and dark grey, respectively). Moreover, we show the patterns for college-educated parents (right panel) and those without college education (left panel). For further details and variable definitions, see Section 3.3 and 3.4.2.

Figure 3.2 shows that parents with and without college education perceive positive returns to *parental time investment* both for the less and the more academically able child, i.e. parents believe that the probability of the child obtaining a university degree is higher the more time parents spend supporting children with their schoolwork. Also both college-educated and not college-educated parents believe that the likelihood of obtaining a university degree is substantially higher for the more academically able child, but the difference is larger for the college-educated parents (about 25 percentage points versus less than 15 percentage points). Interestingly, only college-educated parents believe that marginal returns are (strongly) decreasing for the hypothetical child that is more academically able, i.e. the additional return from 1.5 hours to 3 hours is small and substantially smaller than from 0 to 1.5 hours. Instead college-educated parents believe that marginal returns for the less academically able child are nearly linear (or only slightly decreasing). Thus, on average, returns to investment tend to be smaller (in

particular in relative terms) for the more academically able child. For not college-educated parents, marginal are only slightly or not decreasing in time spend, but also in their case parents perceive returns to be smaller for the more academically able child.

Figure 3.3 displays extremely similar patterns with respect to the outcome *expected income at age 30*. On average, parents perceive returns to time investment to be positive with respect to the child's future income. Also, they expect higher earnings for the more academically able child (see light grey bars). Importantly, as in the case of the *probability of obtaining a college degree*, parents believe –on average– that the returns to investment are smaller (at least in relative terms) for the more academically able child. However, these average beliefs are hiding an important amount of heterogeneity. Our theoretical framework in Section 3.2 predicts important differences in investment patterns based on perceiving higher/lower returns for the more or less academically able child, which is what we will investigate further in Section 3.6.4 below.

Table 3.3 presents regression estimates of the perceived effect of *parental time investment* on parents' perceptions about a child's *probability of obtaining a university degree* (Columns (1) and (2)) and about a child's log income at age 30 (Columns 3 and 4) for hypothetical children under different hypothetical scenarios. The variable *Investment* takes three values: 0, 1.5 and 3 hours, and is treated as continuous. The dummy variable *More Able Child* indicates whether the hypothetical child in the scenario is more academically able than the average child in their age group. In Columns (2) and (4), we interact *Investment* and *More Able Child* with a dummy variable indicating whether parents are college-educated.

As was shown graphically above, parental time investment is positively related to parents' perceptions about a child's *probability to obtain a university degree* (Columns (1) and (2)) and to their *expectations about the child's log income at age 30* (Columns (3) and (4)). More specifically, given the linear specification, one additional hour of parental investment time per week during elementary school increases the probability of a university degree by 6 percentage points and increases expected income at age 30 by 29 percent (significant at the 1-percent level). Also, on average, parents expect the academically more able child to have a substantially higher probability of obtaining a university degree (by 13 percentage points, see Column (1)) and they expect higher income (by 85 percent, see Column (3)). Interestingly, however, the coefficient on the interaction between investment and more academically able child is negative, i.e. on average, parents perceive a lower return to investment for more academically able children (by 1 percentage point per hour invested in terms of probability of a university degree and by 6 percent in terms of expected earnings, both significant at the 1-percent level).

Columns (2) and (4) add a further interaction with a dummy for whether parents are college educated. College-educated parents perceive a higher return to ability per se, i.e. college educated parents perceive an 19 percentage points higher likelihood for the more able child to obtain a university degree (compared to a 8 percentage points for not college-educated parents). Also they perceive a higher return to ability in terms of log income. They expect the more academically able child to have earnings that are more than twice as large (110 percent more) as the less able child compared to 64 percent for the not college-educated parents. Interestingly, college-educated parents perceive the return to investment for the more able child to be lower (1.4 p.p. for likelihood of university degree and 5 percent for expected income, significant at 5 and 1 percent, respectively).

As we discussed in Sections 3.2 and 3.5, whether parents perceive returns to investment to be higher or lower for the more academically able child (in interaction with parents' equity-efficiency preferences) is critical for predictions about how parents invest into their different children. This will be investigated further below.

3.6.2 Descriptive Results on Equity-Efficiency Preferences

In this section, we classify parents into different types based on equity-efficiency preferences. In particular, we split them into three categories, as discussed above, parents with preferences for equal investment (i.e. investing into their children equally, independent of their endowment), parents with preferences for equal outcomes and parent with efficiency preferences.

Figure 3.4 shows the weight parents place on choosing different scenarios as a function of investment and outcomes (in)equality between children of different academic ability in the same family. These questions are based on the following hypothetical scenarios. In Scenario 1, parents invest equally in both children. In Scenario 2, parents invest more in the less academically able child. In Scenario 3, parents invest more in the more academically able child. We display results separately for college educated parents (right panel) and those who have no college education (left panel). For further details and variable definitions, see Section 3.3 and 3.4.2.

According to Figure 3.4, the dark grey bars indicate that, on average, Dutch parents are most likely to choose equality-in-investment scenarios. The average probability of selecting this scenario is around 50 percent for college-educated parents and approximately 46 percent for parents without a college education. The second most likely scenario, represented by the light grey bars, is equality-in-outcome, with an average probability of around 33 percent for both education

groups. The least likely scenario, shown by the grey bars, is efficiency-based, with an average probability of about 15 percent for college-educated parents and close to 20 percent for parents without a college education.

We then categorize the parents into preference types based on their strict preferences. The procedure is defined in Section 3.5. Appendix Figures 3.B.1, 3.B.2 and 3.B.3 display the distribution of the three preference types for different subgroups, based on parents' education, number of children and migrant status, respectively. Appendix Figure 3.B.1 shows the preference type *Equality-in-Investment* is the most common type for all education groups. This type is however most prevalent among the university-educated (with 60 percent) compared to the other three education groups (with around 50 percent). The least frequent category is the *Efficiency* preference type, which is particularly rare among the university-educated (with less than 10 percent), while it is most frequent among those parents with a vocational degree (about 16 percent).

According to Appendix Figure 3.B.2 again the Equality-in-Investment Type is more prevalent with around 50 percent, while the prevalence of the other two types is related to the number of children a family has. In particular, the larger the number of children, the more frequent becomes the *Efficiency* Type, ranging from less than 10 percent for families with one child to around 23 percent for families with four or more children.

Lastly, Appendix Figure 3.B.3 shows that families with migrant background are less likely to have *Equality-in-Investment* preferences *Equality-in-Outcome* preferences, but substantially more likely to have *Efficiency* preferences (25 percent as opposed to 10 percent for natives).

Table 3.4 summarizes these descriptive results. In particular, we investigate how the likelihood to choose a particular scenario is linked in a multivariate framework to parents' background characteristics. We presents regression estimates from a multinomial choice model, where the outcome variables consist of three categories based on strict preference types: the *Equality-in-Outcome* Type denotes parents who assign the strongest preference (highest weight) to the scenario where the difference in the probability of getting into university between the more and the less academically able child is the smallest; the *Equality-in-Investment* Type represents parents who give the highest weight to the scenario where the investment is equal for the more and the less academically able child; and the *Efficiency* Type denotes parents who assign the highest weight to the hypothetical scenario where the total probability of enrolling in university for both children combined is the highest.¹⁵

¹⁵The number of observations (parents) used in this regression are fewer than 985, because we drop those respondents who give equal weight to all three scenarios (33-34%).

Consistent with the figures described above, Table 3.4 shows that college-educated parents are more likely to be classified as *Equality-in-Investment* and *Equality-in-Outcomes* types (as opposed to the reference group, which is the *Efficiency* type), while parents with migrant background or a larger number of children are less likely classified as those two types.

Table 3.5 presents regression estimates of the effect of *Scenario-Implied Outcome Inequality* on the *log odds ratio of choosing that scenario*. More specifically, in the hypothetical scenarios, the *Scenario-Implied Outcome Inequality* is the difference in the probability of enrolling in a university between the more and the less academically able child, which takes three possible values 10%, 30% and 50%, and is treated as a continuous variable. Column (1) of Table 3.5 shows that on average the log-odds ratio of choosing a particular scenario decreases with the level of outcome inequality between the hypothetical children that is implied by that scenario. Converting the log-odds to probabilities, a 10 percentage points increase in inequality reduces the probability of choosing the scenario by approximately 3.35 percentage points. This suggests that parents prefer scenarios that lead to more equal outcomes between their children.

In Column (2), we add a dummy variable *Equal Investment Scenario*, which indicates whether the scenario implies investing into the two children equally (independently of their academic ability). As can be seen, being a scenario where investment into the two hypothetical children is equal increases the likelihood of choosing that scenario by 56 percentage points.

In Columns (3)-(5), we interact the *Scenario-Implied Outcome Inequality* with dummy variables indicating whether parents are college-educated, parents are migrants, and the total number of children in the family. Consistent with the results above, we find that scenario-implied inequality decreases the log-odds ratio of choosing that scenario even more for college-educated parents, but less for parents with migrant background or with a larger number of children.

Figure 3.5 shows the results of a follow-up survey question, in which we increase the costs for the most preferred option, to see whether/how parents change their most preferred scenario. More specifically, based on each parent's strictly preferred scenario, we decreased the probability of graduating for both the more and less able children by 5 percentage points. In the other two scenarios, the average graduation probability for the two children remains constant at 45%. The figure plots the average change in choice probability by preference type, measured in percent relative to the initial choice probability before introducing the additional cost. Results are displayed separately for parents with a college education (right panel) and those without a college education (left panel). All three preference types decrease the choice probability for their previously most preferred but now more

costly option. However, the *Equality-in-Investment* and the *Equality-in-Outcome* types (dark and light grey bar, respectively) respond the least, in particular among college-educated parents (decrease in choice probability of between 14 to 18 percent). The *Efficiency* type responds more strongly, with decrease in choice probability of 25 and close to 35 percent for not college-educated and college-educated parents, respectively. This result is consistent with this preference type caring the most about efficiency costs.

Table 3.6 investigates the effect of additional efficiency costs for the different preference types in a regression framework. Column (1) of Table 3.6 shows that –relative to the *Equal-Investment* Type (excluded category)– the *Efficiency* type responds significantly more strongly (by 11 percentage points, at the 1-percent level) in terms of reducing the probability for the previously preferred scenario, while there is no difference between the *Equal-Investment* and *Equal-Outcome* Type. Table 3.6, Column (2), shows that results are robust to including controls for parents being college-educated, number of children and migrant background.

The last two columns of Table 3.6 display the same relationship in a continuous way, i.e. regressing the relative change in choice probabilities on the weight given to/*probability of choosing the equal outcome scenario* or the *efficiency scenario* (instead of on dummies for the preference types). Columns (3) and (4) show that the higher the weight given to the efficiency scenario, the larger the decrease in the probability of choosing the previously preferred scenario after the additional efficiency cost (without and with additional controls for parental background, respectively).

To summarize, *Equality-in-Investment* and *Equality-in-Outcomes* Types are the most prevalent types (in descending order) and are more prevalent among highly-educated families with few children and without migration background. The *Efficiency* Type is more prevalent among less-educated families with more children and migrant background. Even for this group it is, however, the least likely type. Moreover, consistent with the notion of efficiency, it is the Efficiency Type who responds most strongly to increased costs.

3.6.3 Equity-Efficiency Preferences and Outcome Gap

In this section, we investigate whether parents' stated equity-efficiency preferences are linked to the actual realized outcome inequality between the different children in the family. We therefore focus from now on on parents with at least two children, leaving us with 344 parents (see Table 3.2 for descriptive statistics, discussed in Section 3.4). Child outcomes are measured using test scores from the standardized CITO test taking place at the end of primary school, which is a high-stakes test

used to allocate students to different tracks in high school (see Section 3.4 for more details and definition of variables).¹⁶

Table 3.7 presents regression estimates of the effects of having a preference for *Equality in Outcomes* on actual within-family outcome inequality. Inequality is quantified using *within-family log deviations in the CITO test*, as described in Section 3.5. The variable *Equality in Outcome* is a dummy variable that takes the value 1, if parents assign the highest probability to the scenario where the difference in the probability of getting into university between the more and the less academically able child is the smallest, and 0 otherwise. Column (1) of Table 3.7 displays the pooled results, while Columns (2) to (4) report results for parents with a secondary schooling degree or less, a vocational degree or at least a college degree, respectively. According to Column (1), there is a clear link between parents' preference for equal outcomes to the actual inequality in child outcomes (test scores) within the family, in that having *Equality in Outcome* preferences decreases within-family deviations in the CITO test by 10 percent. This relationship is mostly driven by the intermediate-education group of parents, the ones with vocational degree, but the negative sign and similar coefficient is also there for parents with college degree or more. More specifically, among parents with vocational degree, having *Equality in Outcome* preferences decreases the within-family deviations in CITO test by 13 percent (significant at the 5-percent level).

Thus, we have shown that the *Equality in Outcome* type is linked to less inequality in outcomes between their children, as predicted by our model (see Section 3.2).

3.6.4 Equity-Efficiency Preferences, Perceived Production Function and Investment Gap

In this section, we test the other predictions of our theoretical framework (see Section 3.2), which is linking parents' preferences and beliefs to the level of investment into the different children within the family. We start by testing the first clear prediction of the model, which is that the *Equal Investment Type* is the one investing most equally into the different children. For the two other preference types, the direction of investment (investing more into the more or into the less academically able child) depends not only on the preference type, but also the perception of whether marginal returns to investment are higher or lower for

¹⁶While basically all schools administer a mandatory standardized test at the end of primary school, most but not all schools choose the CITO test. Therefore our number of observations (number of families/parents) is reduced to 236.

the more able child. For the *Equal Investment Type* instead, the model always predicts investing into different children more equally than the other two types, independently of the endowment of the children.

In Table 3.8, we therefore regress actual within-family investment inequality on having a preference for equality in investment. In Columns (1) and (2) (without and with controls, respectively), the outcome variable is a dummy variable, which takes the value 1 if parents invest equally into their children, i.e. the actual hours helping with school work during elementary school is the same for the two children of a parent, and 0 otherwise. The variable *Equality Invest Type* is a dummy variable that takes the value 1 if a parent assigns the highest weight to the scenario where the investment into the more and less academically able child is the same and 0 otherwise. In Columns (3) and (4) (without and with controls, respectively), the outcome variable is the *absolute difference in investment*, i.e. the absolute difference between two children in the family in terms of hours parents help each of them with school work during elementary school. The explanatory variable *Equality Invest Weight* is the probability parents assign to the scenario where the investment into the more and less able child is the same.

Columns (1) and (2) of Table 3.8 show that parents being the *Equal Invest Type* according to stated preferences based on hypothetical scenarios is linked to a significantly higher probability of investing equally into the different children. More specifically, it increases the likelihood by 12 percentage points (significant at the 1-percent level). Results remain basically unchanged after including control variables for parental background. According to Columns (3) and (4) of Table 3.8, the same is true when looking at the relationship between the weight parents assign to the hypothetical scenario in which both children are invested into equally, i.e. the *Equal Invest Weight*, and the actual (absolute) difference in parental investments between the different children in a family. A ten percentage points increase in the weight given to the equal-invest scenario decreases the absolute difference in investment between the two children by nearly one hour (0.7 hours after including controls), significant at the 5-percent level.

Lastly, we test the more complex prediction of the model in terms of the actual direction of investment (i.e., investing more into the more or the less able child), which depends on the interaction between parents' preference type and their beliefs about the marginal returns to investment, i.e. whether they are higher or lower for the academically more able child (see Table 3.1 for a summary of the predictions of the model with respect to these interactions).

Table 3.9 presents regression estimates of how parents' preference type and perceived returns to investment and the interaction thereof are linked to the difference in investment between the more and less academically able child within

the family. The analysis focuses on families with at least two children, ranking children within each family pairwise based on parents' perceptions about children's academic ability. We exclude child pairs that are indistinguishable based on parental perceived ability and limit the sample to families with no more than three children to ensure accurate parental rankings. The outcome variable is the *difference in hours of parental investment between the more and less academically able child*. The regression model includes a dummy variable *Equal Outcome Type*, which is equal to 1, if the parent has a strict preference for equality in outcomes and 0 if the parent has a strict preference for efficiency. Also, we include a dummy variable *Decreasing Return Type*, which takes the value 1 if the parent perceives decreasing returns to investment with respect to ability and 0 otherwise. Lastly, we include an interaction term of the two dummies. Columns (1)-(2) report results for families with exactly two children, Columns (3)-(4) include families with three children, and Columns (5)-(6) incorporate the Equal Invest Type, who prefers to invest equally into the children, so that in this case *Equal Outcome Type* takes the value zero if the parent has a strict preference for efficiency or for equality in investment. Columns (2), (4), and (6) include controls as specified in Section 3.5.

According to Table 3.9 being an *Equal Outcome Type* is linked to a smaller (and possibly negative) difference in investment between the more and less academically able child, as predicted by the model (see Section 3.2 and Table 3.1 for a summary of the model predictions). The coefficient is similar in magnitude (-2.1 to -2.4) and always significant at least at the 5-percent level. Perceiving lower marginal returns for the more academically able child (*Decreasing Return Type*) is also linked to a less positive/more negative difference in investment between the more and less able child. The coefficient is significantly negative (after adding parental controls), but only when the excluded category only contains the Efficiency Type, as in Columns (1) to (4), as predicted by the model. After including also the *Equal Invest Type* into the excluded category (as in Columns (5) and (6), the coefficient is smaller and not significant, consistent with perceived marginal returns not playing a role for this type. For the *Equal Outcome Type* instead, perceiving decreasing returns goes in the opposite direction (or even reverses) the main effect, again as predicted by the model if returns are perceived as being clearly lower for the more able type. The intuition is that the *Equal Outcome Type* wants their children to have similar outcomes and they therefore tend to aim to counteract endowment differences (by actually investing more into the less able child). However, since this preference type does not aim to reverse the original ability ranking of the two children, perceiving strongly decreasing returns implies that they would not want to invest (much) more into the less able child.

3.7 Conclusion


In this paper, we investigated parents' intra-household investment decisions, focusing on how parents' beliefs about the productivity of their time investment (conditional on child ability) and their equity-efficiency preferences influence their investment in the human capital of their children. Based on a unified framework that demonstrates how the interplay of equity-efficiency concerns and diminishing returns to parental involvement can moderate disparities in intra-family investments, we designed and implemented a survey to elicit parents' beliefs and preferences using innovative survey instruments. Analyzing responses from parents in the Netherlands, we uncovered a negative correlation between a child's academic potential and the parents' investment in this child relative to its siblings. To explain this finding, we showed that—on average—Dutch parents perceive higher marginal returns from investing in less academically able children in learning-related activities. Moreover, on average, parents exhibit equality-focused preferences in the treatment of siblings, i.e., they prefer to invest more in the less academically able child.

By linking our survey data to administrative data from the Dutch Statistics Bureau (CBS), we showed that parents' equality-focused preferences lead them to invest in a way that reduces the gap in academic outcomes (performance on a high-stakes standardized test) among their offspring. Data on parents' beliefs about the productivity of their time investment, as well as their equity-efficiency preferences, not only help to explain the magnitude of siblings' outcome differences across families, but also predict differential investments between siblings across families. In particular, actual investment differences are smallest for parents with equality-in-investment preferences. Parents with equality-in-outcome preferences invest more in the less academically able child, while parents with efficiency preferences and higher perceived productivity of investments in the more academically able child invest more in this child.

Our study not only sheds light on the nuanced decision-making processes within households but also offers a new perspective on equality of opportunity in the context of educational attainment. Since parental investments in childhood are critical for the development of children's cognitive and non-cognitive skills, intra-household allocations can have profound consequences for their long-term outcomes in the labor market, marriage market, health, and so forth. A better understanding of intra-household allocation decisions is, therefore, critical for addressing the sources of inequality (in outcomes and/or opportunity) and for the effective targeting of social programs.


3.8 Figures

Figure 3.1: Within-Family Correlation between Parental Investment and Child Ability




Notes: This figure shows a binscatter plot of deviations in parental time investment from the family average against deviations in parental perceived child ability from the family average. Parental time investment is measured by hours spent helping with schoolwork during elementary school, while perceived child ability is measured by the reported probability that the child would graduate from university, conditional on hypothetical enrollment. The sample is restricted to families with at least two children for whom we have data on parental investment and where the perceived ability is distinguishable (see Section 3.4 for more details and variable definitions).

Figure 3.2: Production Function: Perceived Return to Parental Investment (Probability University Degree)



Notes: This figure shows parents' perceptions about the probability of obtaining a university degree of a hypothetical child as a function of parental time investment in helping with schoolwork during elementary school (ranging from 0 over 1,5 to 3 hours). These perceptions were elicited for hypothetical child with higher and lower than average academic ability. Also, we display results separately by whether parents are college educated (right panel) or not (left panel). For further details and variable definitions, see Section 3.4.

Figure 3.3: Production Function: Perceived Return to Parental Investment (Expected Income)



Notes: This figure shows parents' expectations about the income (at age 30) of a hypothetical child as a function of parental time investment in helping with schoolwork during elementary school (ranging from 0 over 1,5 to 3 hours). These expectations were elicited for hypothetical child with higher and lower than average academic ability. Also, we display results separately by whether parents are college educated (right panel) or not (left panel). For further details and variable definitions, see Section 3.4.

Figure 3.4: Equity-Efficiency Preferences in Choice Probabilities




Figure Notes: This figure shows the weight parents place on choosing different scenarios as a function of investment and outcomes (in)equality between children of different academic ability in the same family based on a hypothetical survey question. In Scenario 1, parents invest equally in both children (*Equality in Investment*). In Scenario 2, parents invest more in the less academically able child (*Equality in Outcome*). In Scenario 3, parents invest more in the more academically able child (*Efficiency*). We display results separately for college educated parents (right panel) and those who have no college education (left panel). For further details and variable definitions, see Section 3.4.

Figure 3.5: Change in Likelihood of Response




Figure Notes: This figure illustrates the parent response to a reduction in total outcomes for the option that parents assigned the highest probability in the baseline question. Specifically, based on each parent's strictly preferred scenario, we decreased the probability of graduating for both more and less able children by 5 percentage points. In the other two scenarios, the average graduation probability for the two children remains constant at 45%. The figure then plots the average change in choice probability by preference type, measured in percentage (between 0 and 1) relative to the initial choice probability before introducing the additional cost. Results are displayed separately for parents with a college education (right panel) and those without a college education (left panel). For further details and variable definitions, see Section 3.4.

3.9 Tables

Table 3.1: Model Predictions for Different Parent Types: Investment and Outcome Gaps

	Decreasing Returns (DR)	Increasing Returns (IR)
Equality-in-Investment	No investment gap Outcome gap remains	No investment gap Outcome gap remains
Equality-in-Outcome	Invest more in less able (unless strongly DR) Reduce outcome gap	Invest more in less able Reduce outcome gap
Efficiency	Invest more in less able Reduce outcome gap	Invest more in more able Increase outcome gap

Notes: This table presents model predictions of different parents types based on two key dimensions: their preferences and their beliefs. The first dimension reflects parents' equality preferences, which may prioritize equality in outcomes, equality in investments, or efficiency. The second dimension relates to their beliefs about the marginal returns to investment—whether they believe these returns increase or decrease with the child's ability. For each type we summarize model predictions for within-family investment pattern and the outcome pattern of their children.

Table 3.2: Summary Statistics of Parents

Parent Respondent	Full Sample		Sibling Sample	
	Mean	Std. Dev.	Mean	Std. Dev.
Male	0.445	0.497	0.442	0.497
Age	48.372	8.060	44.869	5.541
College Educated	0.430	0.495	0.465	0.500
Dutch Background	0.817	0.387	0.791	0.407
Western Background	0.075	0.264	0.087	0.283
Non-Western Background	0.103	0.304	0.116	0.321
Number of Children	2.260	0.972	2.657	0.950
Hours Helping Homework Per Child	2.006	4.965	1.696	2.609
Number of Parents	985		344	

Notes: This table presents summary statistics for the full sample of parents who responded to the survey questions on the perceived production function and on equality-efficiency preferences, as well as a subsample of parents who have at least two children for whom we have data on parental investment (helping with homework in elementary school).

Table 3.3: Perceived Return to Parental Investment

	Probability University Degree	Expected Log Income		
	(1)	(2)	(3)	(4)
Investment (Hours)	0.0624*** (0.00386)	0.0685*** (0.00543)	0.287*** (0.0150)	0.308*** (0.0193)
More Able Child	0.126*** (0.0239)	0.0801** (0.0350)	0.854*** (0.0551)	0.638*** (0.0620)
Investment * More Able	-0.0108*** (0.00363)	-0.0108*** (0.00363)	-0.0633*** (0.0147)	-0.0633*** (0.0147)
More Able * Parent College		0.107** (0.0423)		0.471*** (0.100)
Investment * Parent College		-0.0141** (0.00700)		-0.0454* (0.0262)
Constant	7.319*** (0.0129)	7.319*** (0.0129)	-1.191*** (0.0334)	-1.191*** (0.0331)
Number of Parents	985	985	960	960
Number of Parent-Scenario Pairs	5,910	5,910	5,760	5,760
Parent Fixed-Effects	yes	yes	yes	yes
R-squared	0.247	0.268	0.050	0.055

Notes: This table presents regression estimates of the perceived effect of parental investment (hours spent helping with homework during elementary school) on parents' perceptions about a child's probability of obtaining a university degree (Columns (1) and (2)) and about a child's log income at age 30 (Columns 3 and 4) for hypothetical children under different hypothetical scenarios. Robust standard errors are reported in parentheses. The variable *Investment* takes three values: 0 hours, 1.5 hours, and 3 hours, and is treated as continuous. The dummy variable *More Able Child* indicates whether the hypothetical child in the scenario is more academically able than the average child in their age group. In Columns 2 and 4, we interact *Investment* and *More Able Child* with a dummy variable indicating whether parents are college-educated.

Table 3.4: Preference Types and Parental Characteristics

Preference Types	Coefficient
Equality-in-Outcome	
Parent College Educated	0.5347*** (0.2365)
Parent Migrant	-1.0564*** (0.2592)
Number of Children	-0.3309*** (0.1095)
Equality-in-Investment	
Parent College Educated	0.4856** (0.2272)
Parent Migrant	-1.1543*** (0.2452)
Number of Children	-0.2892*** (0.1020)
Efficiency Type	(Reference Group)
Number of Parents:	846

Notes: This table presents regression estimates from a multinomial choice model. The outcome variables consist of three categories based on strict preference types: the *Equality-in-Outcome Type* denotes parents who assign the strongest preference (highest weight) to the scenario where the difference in the probability of getting into university between the more and the less academically able child is the smallest; the *Equality-in-Investment Type* represents parents who give the highest weight to the scenario where the investment is equal for the more and the less academically able child; and the *Efficiency* denotes parents who assign the highest weight to the hypothetical scenario where the total probability of enrolling in university for both children combined is the highest. The number of observations (parents) used in this regression is fewer than 985 because we drop those respondents who give equal weight to all three scenarios (33-34%). Robust standard errors are reported in parentheses.

Table 3.5: Preferences for Outcome Inequality

	Log-odds Ratio Choosing the Scenerio				
	(1)	(2)	(3)	(4)	(5)
Scenario-Implied Outcome Inequality	-1.344*** (0.170)	-1.344*** (0.170)	-1.344*** (0.169)	-1.344*** (0.169)	-1.344*** (0.168)
Equality-in-Investment Type		0.558*** (0.0693)	0.558*** (0.0693)	0.558*** (0.0693)	0.558*** (0.0693)
Scenario Inequality * Parent College			-0.667* (0.345)		
Scenario Inequality * Parent Migrant				1.071** (0.449)	
Scenario Inequality * Number Children					0.533*** (0.172)
Constant	-0.430*** (0.0509)	-0.616*** (0.0578)	-0.616*** (0.0578)	-0.616*** (0.0577)	-0.616*** (0.0574)
Number of Parents	521	521	521	521	521
Number of Parent-Scenario Pairs	1,563	1,563	1,563	1,563	1,563
Parent Fixed-Effects	yes	yes	yes	yes	yes
R-squared	0.045	0.110	0.112	0.114	0.117

Notes: This table presents regression estimates of the effects of scenario-implied outcome inequality on the log odds ratio of choosing that scenario. In the hypothetical scenarios, the difference in the probability of enrolling in a university between the more and the less academically able child takes three possible values 0.1, 0.3 and 0.5, and is treated as a continuous variable. The dummy variable *Equality-in-Investment type* indicates whether the scenario implies equal investment. In Columns (3)-(5), we interact the scenario-implied inequality with dummy variables indicating whether parents are college-educated, parents are migrants, and the total number of children in the family. Robust standard errors are reported in parentheses.

Table 3.6: Response to Additional Efficiency Cost

	Relative Change in Probability Choosing Scenario			
	(1)	(2)	(3)	(4)
Equality-in-Outcome Type	0.00966 (0.0240)	0.00990 (0.0240)		
Efficiency Type	-0.107*** (0.0383)	-0.106*** (0.0393)		
Equality-in-Outcome Weight			-0.0148 (0.0420)	-0.0147 (0.0421)
Efficiency Weight			-0.191*** (0.0653)	-0.190*** (0.0673)
Parent College Educated		0.00378 (0.0225)		0.00261 (0.0226)
Parent Migrant		-0.00891 (0.0294)		-0.00786 (0.0294)
Number of Children		0.00481 (0.0106)		0.00402 (0.0105)
Constant	-0.161*** (0.0153)	-0.172*** (0.0298)	-0.134*** (0.0204)	-0.143*** (0.0323)
Number of Parents	845	845	845	845
R-squared	0.013	0.013	0.013	0.013

Notes: This table presents regression estimates of the effects of adding an additional efficiency cost to the hypothetical scenarios on the relative change in the probability of parents choosing those scenarios. Specifically, if a scenario was previously considered the most favorable, the probability of graduating from university for both the more and the less academically able child is reduced by 5%. The outcome variable is the ratio of the change in the probability of selecting the scenario relative to the initial probability. Robust standard errors are reported in parentheses.

Table 3.7: Preferences and Actual Outcome Inequality

	Within-Family Log Deviations in CITO Test By Parents' Education			
	Pooled (1)	Secondary or Less (2)	Vocational (3)	College and More (4)
Equality-in-Outcome type	-0.0971* (0.0533)	0.0474 (0.203)	-0.130** (0.0575)	-0.0982 (0.0906)
Constant	0.180* (0.0710)	0.0493 (0.0743)	0.243 (0.119)	0.790** (0.0390)
Observations	236	51	80	105
Controls	yes	yes	yes	yes
R-squared	0.022	0.005	0.051	0.016

Notes: This table presents regression estimates of the effects of having a preference for equality in outcomes on actual within-family outcome inequality. The child outcome is measured using CITO test scores at the end of primary school (see Data Section for more details and definition of variables), and inequality is quantified using log deviations, as described in Section 3.5. The variable *Equality-in-Outcome* is a dummy variable that takes the value 1 if a parent assigns the highest probability to the scenario where the difference in the probability of getting into university between the more and the less academically able child is the smallest, and 0 otherwise. Robust standard errors are reported in parentheses.

Table 3.8: Preference Types and Actual Investment Gap

	Probability Invest Equally		Absolute Difference in Investment	
	(1)	(2)	(3)	(4)
Equality-in-Investment Type	0.129*** (0.0413)	0.115*** (0.0417)		
Equality-in-Investment Weight			-0.900** (0.355)	-0.699** (0.326)
Constant	0.483*** (0.0280)	-0.421 (0.738)	1.778*** (0.261)	0.996 (5.024)
Number of Parents	344	344	344	344
Number of Child Pairs	579	579	579	579
Controls		yes		yes
R-squared	0.017	0.043	0.007	0.055

Notes: This table presents regression estimates of the effects of having a preference for equality in investment on actual within-family investment inequality. In Columns (1) and (2), the outcome variable is a dummy variable, which takes the value 1 if parents invest equally into their children, i.e. the actual hours helping with school work during elementary school is the same for the two children of a parent, and 0 otherwise. The variable *Equality-in-Investment Type* is a dummy variable that takes the value 1 if a parent assigns the highest weight to the scenario where the investment into the more and less academically able child is the same, and 0 otherwise. In Columns (3) and (4), the outcome variable is the *absolute difference in investment*, i.e. the absolute difference between two children in the family in terms of hours parents help each of them with school work during elementary school. The variable *Equality-Investment-weight* is the probability parent assigns to the scenario where the investment into the more and less able child is the same. Robust standard errors are reported in parentheses.

Table 3.9: The Interaction of Preference and Beliefs on Actual Investment

	Difference in Investment Between High and Low Ability Child					
	(1)	(2)	(3)	(4)	(5)	(6)
Equality-in-Outcome Type	-2.100** (0.932)	-2.196** (0.942)	-2.286*** (0.788)	-2.361*** (0.717)	-2.429** (1.069)	-2.157** (1.067)
Decreasing-Return Type	-1.624 (1.018)	-1.686** (0.825)	-0.821 (0.737)	-1.258* (0.713)	-0.721 (0.969)	-0.624 (1.002)
Equality-in-Outcome type * Decreasing-Return-type	2.757** (1.214)	2.787** (1.119)	2.396** (0.938)	2.476*** (0.860)	2.296** (1.126)	2.001* (1.136)
Constant	0.800 (0.740)	1.379 (0.972)	0.429 (0.579)	1.464* (0.741)	0.571 (0.929)	0.538 (0.782)
Observations	62	62	95	95	208	208
Sample: Include 3 Children			yes	yes	yes	yes
Sample: Include Equal Invest Type					yes	yes
Controls		yes		yes		yes
R-squared	0.068	0.138	0.095	0.176	0.029	0.060

Notes: This table presents regression estimates of the effects of equality-efficiency preferences and perceived returns to investment on within-family investment heterogeneity concerning ability differences. The analysis focuses on families with more than one child, ranking children within each family pairwise based on parents' perceptions about children's academic ability. We exclude child pairs that are indistinguishable based on parental perceived ability and limit the sample to families with no more than three children to ensure accurate parental rankings. The outcome variable is the difference in hours of parental investment between the more and less academically able child. The regression model includes a dummy variable *Equal Outcome Type*, which is equal to 1, if the parent has a strict preference for equality in outcomes and 0 if the parent has a strict preference for efficiency. Also, we include a dummy variable *Decreasing Return Type*, which takes the value 1 if the parent perceives decreasing returns to investment and 0 otherwise. Lastly, we include an interaction term of the two dummies. Columns (1)-(2) report results for families with exactly two children, Columns (3)-(4) include families with three children, and Columns (5)-(6) incorporate the *Equal Invest Type*, who prefers to invest equally into the children, so that in this case *Equal Outcome Type* takes the value zero if the parent has a strict preference for efficiency or for equality in investment. Columns (2), (4), and (6) include controls as specified in Section 3.5. Robust standard errors are reported in parentheses.

Appendices to Chapter 3

3.A Survey Questions: Exact Wording

3.A.1 Equity-Efficiency Preferences

Imagine you have two children, [Name 1 and Name 2] [randomize: Daan and Thijs OR Anne and Vera OR Daan and Vera OR Vera and Daan], in elementary school. [child] is academically more able. Imagine that you as a parent could influence the long-run prospects of the child by your “investment” choices (for example, by helping your child with the school work (such as monitoring whether homework is done, being available for questions, helping to study for tests), by reading with the child, by spending money on tutoring etc). Which of the following options would you be more likely to choose? Indicate the probability with which you would choose each of the options indicated below.

Choice scenario 1: [randomize order in which options A), B), C) are presented]

A) By investing the same into both children, the more academically able child has a probability of 60% graduating from university and the less academically able child of 30%.

B) By investing more into the academically less able child, the more academically able child has a probability of 50% graduating from university and the less academically able child of 40%.

C) By investing more into the more academically able child, the more academically able child has a probability of 70% graduating from university and the less academically able child of 20%.


3.A.2 Parental Time Investment

Main question: ”How many hours per week do/did you spend on helping [child] with homework (such as monitoring that homework is done, going over the work, being available for questions etc) or helping with studying (practicing vocabulary, practicing dictation or for tests) during elementary/secondary school?”

Additional question: ”How many hours per week did you spend on direct interaction with [child] other than helping with homework? Examples are joint meals, joint activities (such as playing games/going on an excursion/to the zoo, museum, concert etc), reading to/with the child, talking about personal matters (also when other family members were present) when [child] was in elementary/secondary school?”


3.B Equality Preferences Types

Figure 3.B.1: Equality Preferences Type: by Parental Education




In this figure parents are classified into three preference types based on their choices in hypothetical scenarios that trade off between equality and efficiency: Equality-in-Investment, where parents select the scenario in which parents invest equally in both children; Equality-in-Outcome, where parents select the scenario that minimizes the difference in outcomes between the more and the less academically able child; and Efficiency, where parents select the scenario that gives the highest overall probability of successful outcomes by potentially allocating more resources to the more academically able child. The figure shows the distribution of these preference types across different levels of parental education, including university, college, vocational, and secondary education or below. For further details and variable definitions, see Section 3.4.

Figure 3.B.2: Equality Preferences Type: by Number of Children



In this figure parents are classified into three preference types based on their choices in hypothetical scenarios that trade off between equality and efficiency: Equality-in-Investment, where parents select the scenario in which parents invest equally in both children; Equality-in-Outcome, where parents select the scenario that minimizes the difference in outcomes between the more and the less academically able child; and Efficiency, where parents select the scenario that gives the highest overall probability of successful outcomes by potentially allocating more resources to the more academically able child. The figure shows the distribution of these preference types across families with different numbers of children: one, two, three, and four. For further details and variable definitions, see Section 3.4.

Figure 3.B.3: Equality Preferences Type: by Family Origin

In this figure parents are classified into three preference types based on their choices in hypothetical scenarios that trade off between equality and efficiency: Equality-in-Investment, where parents select the scenario in which parents invest equally in both children; Equality-in-Outcome, where parents select the scenario that minimizes the difference in outcomes between the more and the less academically able child; and Efficiency, where parents select the scenario that gives the highest overall probability of successful outcomes by potentially allocating more resources to the more academically able child. The figure shows the distribution of these preference types across families of different origins, specifically distinguishing between immigrant and native families. For further details and variable definitions, see Section 3.4.

Bibliography

- Aizer, Anna and Janet Currie. 2019. “Environmental Factors in Disadvantaged Neighborhoods and Behavioral and Cognitive Outcomes.” *American Economic Review* .
- Amemiya, Takeshi. 1985. *Advanced Econometrics*. Harvard University Press.
- Attanasio, Orazio, Flavio Cunha, and Paulina Jervis. 2015. “Eliciting beliefs, choice probabilities, and preferences: Experimental evidence from education.” *Journal of Human Capital* 9 (3):259–306.
- . 2019. “Early investments and child development.” *American Economic Review* 109 (3):255–285.
- Attanasio, Orazio and Katja Kaufmann. 2014. “Education choices and returns to schooling: Mothers’ and youths’ subjective expectations and their role by gender.” *Journal of Development Economics* 109 (C):203–216.
- Attanasio, Orazio, Costas Meghir, and Emily Nix. 2020a. “Human capital development and parental investment in India.” *The review of economic studies* 87 (6):2511–2541.
- . 2020b. “Parental time investments and child development: The role of human capital.” *Review of Economic Studies* 87 (4):2180–2204.
- Azmat, Ghazala and Katja Maria Kaufmann. 2024. “Formation of College Plans: Expected Returns, Preferences, and Adjustment Process.” *Journal of the European Economic Association* 22 (2):669–711.
- Baker, Michael and Kevin Milligan. 2016. “Parental time and child development: Evidence from the Canadian time-use survey.” *Journal of Labor Economics* 34 (1):B24–B67.
- Bassok, Daphna and Sean F. Reardon. 2013. “Academic Redshirting in Kindergarten: Prevalence, Patterns, and Implications.” *Educational Evaluation and Policy Analysis* 35 (3):283–297.

BIBLIOGRAPHY

- Bedard, Kelly and Elizabeth Dhuey. 2006. "The Persistence of Early Childhood Maturity: International Evidence of Long-Run Age Effects." *Quarterly Journal of Economics* 121 (4):1437–1472.
- Behrman, Jere R, Robert A Pollak, and Paul Taubman. 1982. "Parental preferences and provision for progeny." *Journal of political economy* 90 (1):52–73.
- . 1986. "Do parents favor boys?" *International Economic Review* :33–54.
- Berry, James, Rebecca Dizon-Ross, and Maulik Jagnani. 2024. "Parental preferences for equality and efficiency in educational investments." *Journal of Development Economics* Forthcoming.
- Bharadwaj, Prashant, Juan Pedro Eberhard, and Christopher A. Neilson. 2018. "Health at Birth, Parental Investments, and Academic Outcomes." *Journal of Labor Economics* 36 (2):349–394.
- Black, Sandra E., Paul J. Devereux, and Kjell G Salvanes. 2011. "Too Young to Leave the Nest? The Effects of School Starting Age." *The Review of Economics and Statistics* 93 (2):455–467.
- Boneva, Teodora and Christopher Rauh. 2018. "Parental Beliefs about Returns to Educational Investments—The Later the Better?" *Journal of the European Economic Association* 16 (6):1669–1711.
- Card, David, Ciprian Domnisoru, and Lowell Taylor. 2018. "The Role of School Quality in Regional Differences in Upward Mobility." *Journal of Economic Perspectives* 32 (3):167–194.
- Cascio, Elizabeth U. and Diane W. Schanzenbach. 2016. "First in the Class? Age and the Education Production Function." *Education Finance and Policy* 11:225–250.
- Chen, Kelly, Nicole Fortin, and Shelley Phipps. 2015. "Young in Class: Implications for Inattentive/Hyperactive Behaviour of Canadian Boys and Girls." *Canadian Journal of Economics* 48:1601–1634.
- Chetty, Raj, John N. Friedman, Nathaniel Hendren, Maggie R. Jones, and Sonya R. Porter. 2020. "The Opportunity Atlas: Mapping the Childhood Roots of Social Mobility." *Nature* 577:462–468.
- Chetty, Raj and Nathaniel Hendren. 2018a. "The Impacts of Neighborhoods on Intergenerational Mobility I: Childhood Exposure Effects." *The Quarterly Journal of Economics* 133 (3):1107–1162.

- . 2018b. “The Impacts of Neighborhoods on Intergenerational Mobility II: County-Level Estimates.” *The Quarterly Journal of Economics* 133 (3):1163–1228.
- Chiappori, Pierre-André and Costas Meghir. 2015. “The family: A theoretical framework and evidence from micro data.” *Annual Review of Economics* 7:553–583.
- Chyn, Eric and Lawrence F. Katz. 2021. “A Summary of Recent Literature on Neighborhood Effects.” *Journal of Economic Literature* .
- Cobb-Clark, Deborah A., Nicolas Salamanca, and Anna Zhu. 2019. “Time use and child outcomes: Results from the Longitudinal Study of Australian Children.” *Economic Journal* 129 (619):F332–F358.
- Crawford, Claire, Lorraine Dearden, and Ellen Greaves. 2013. *When You Are Born Matters: Evidence for England*. R80. IFS Report.
- Cunha, Flavio, Irma Elo, and Jennifer Culhane. 2013. “Eliciting Maternal Expectations about the Technology of Cognitive Skill Formation.” NBER Working Paper 19144, National Bureau of Economic Research.
- Cunha, Flávio, Irma Elo, and Jennifer Culhane. 2022. “Maternal subjective expectations about the technology of skill formation predict investments in children one year later.” *Journal of Econometrics* 231 (1):3–32.
- Cunha, Flavio and James Heckman. 2007a. “The technology of skill formation.” *American economic review* 97 (2):31–47.
- Cunha, Flavio and James J Heckman. 2007b. “Formulating, identifying and estimating the technology of cognitive and noncognitive skill formation.” *Journal of Human Resources* 43 (4):738–782.
- Del Bono, Emilia, Marco Francesconi, Yvonne Kelly, and Amanda Sacker. 2016. “The economics of parental time investments in children.” *Economic Journal* 126 (596):F96–F123.
- Delavande, Adeline and Basit Zafar. 2019. “University choice: The role of expected earnings, nonpecuniary outcomes, and financial constraints.” *Journal of Political Economy* 127 (5):2343–2393.
- Deutscher, Nathan. 2020. “Place, Peers, and the Teenage Years: Long-Run Neighborhood Effects in Australia.” *American Economic Journal: Applied Economics* 12 (2):220–49.

BIBLIOGRAPHY

- Dhuey, Elizabeth, David Figlio, Krzysztof Karbownik, and Jeffrey Roth. 2019. "School Starting Age and Cognitive Development." *Journal of Policy Analysis and Management* 38 (3):538–578.
- Dhuey, Elizabeth and Stephen Lipscomb. 2008. "What Makes a Leader? Relative Age and High School Leadership." *Economics of Education Review* 27:173–183.
- Dominitz, J and CF Manski. 1996. "Eliciting student expectations of the returns to schooling." *Journal of Human Resources* 31 (1):1–26.
- Dustmann, Puhani Patrick A., Christian and Uta Schönberg. 2016. "The Long-Term Effects of Early Track Choice." *The Economic Journal* 127:1348–1380.
- Elder, Todd E. and Darren H. Lubotsky. 2009. "Kindergarten Entrance Age and Children's Achievement." *Journal of Human Resources* 44:641–683.
- Epple, Dennis and Richard Romano. 2003. "Neighborhood Schools, Choice, and the Distribution of Educational Benefits." In *The Economics of School Choice*, edited by Caroline Hoxby. University of Chicago Press, 227–286. Accessed: 2024-10-05.
- European-Commission. 2019. "Education and Training Monitor 2019 - Netherlands." Accessed: 2024-10-05.
- Falch, Rune, Mats Fisher, and Ole H. Nyhus. 2021. "Compensatory investments in early childhood development: Evidence from Norway." *Economics of Education Review* 84:102–124.
- Foy, Pierre. 2013. "TIMSS and PIRLS 2011 User Guide for the Fourth Grade Combined International Database." *TIMSS & PIRLS International Study Center, Lynch School of Education, Boston College and International Association for the Evaluation of Educational Achievement (IEA)* .
- . 2017. "PIRLS 2016 User Guide for the International Database." *TIMSS & PIRLS International Study Center, Lynch School of Education, Boston College and International Association for the Evaluation of Educational Achievement (IEA)* .
- Fredriksson, Peter and Björn Öckert. 2014. "Life-Cycle Effects of Age at School Start." *The Economic Journal* 124:977–1004.
- Giannola, Michele. 2024. "Parental investments and intra-household inequality in child human capital: evidence from a survey experiment." *The Economic Journal* 134 (658):671–727.

- Gibbons, Stephen and Olmo Silva. 2018. "Urban Density and Pupil Attainment." *Journal of Urban Economics* 103:96–115.
- Guryan, Jonathan, Sara B. Heller, Jens Ludwig, and Harold A. Pollack. 2021. "The Effect of Parental Resources and Neighborhood Violence on Cognitive Skills." *American Economic Review* .
- Guryan, Jonathan, Erik Hurst, and Melissa Kearney. 2008. "The quantity and quality of parental time investments: Evidence from the American Time Use Survey." *Journal of Labor Economics* 26 (1):93–124.
- Hsin, Amy and Christina Felfe. 2014. "Parental education and parental time with children." *Journal of Economic Perspectives* 28 (3):109–132.
- Huizenga, Janneke and Geertje Damstra. 2016a. "Doorstroom van Kleuters – Is het Kind Klaar voor Groep 3 of Is Groep 3 Klaar voor het Kind?" OC&W/PO-raad Den Haag.
- . 2016b. "Flow of Preschool Children: Is the Child Ready for Group 3, or Is Group 3 Ready for the Child?" OC&W/PO-raad Den Haag.
- Jackson, C. Kirabo. 2020. "Public School Funding and Student Outcomes." *Brookings Papers on Economic Activity* 2020 (1):65–138.
- Jensen, Robert. 2010. "The (Perceived) Returns to Education and the Demand for Schooling*." *The Quarterly Journal of Economics* 125 (2):515–548.
- Kalil, Ariel and Rebecca Ryan. 2017. "Family income, parenting practices, and child development." *Future of Children* 27 (2):163–184.
- Kaufmann, Katja Maria. 2014. "Understanding the income gradient in college attendance in Mexico: The role of heterogeneity in expected returns." *Quantitative Economics* 5 (3):583–630.
- Kawaguchi, Daiji. 2011. "Actual Age at School Entry, Educational Outcomes, and Earnings." *Journal of The Japanese and International Economies* 25:64–80.
- Kirchberger, Martin. 2020. "Parental time allocation and gender inequalities in education." *Journal of Development Studies* 56:963–978.
- Laliberté, Jean-William. 2021. "Long-term Contextual Effects in Education: Schools and Neighborhoods." *American Economic Journal: Economic Policy* 13 (2):336–377.

BIBLIOGRAPHY

- Luijkx, Ruud and Manon de Heus. 2008. “The Educational System of the Netherlands.” In *The International Standard Classification of Education*. MZES, Mannheim, 47–75.
- Manski, Charles F. 2004. “Measuring expectations.” *Econometrica* 72 (5):1329–1376.
- Matsubayashi, Tetsuya and Michiko Ueda. 2015. “Relative Age in School and Suicide among Young Individuals in Japan: A Regression Discontinuity Approach.” *PloS one* 10 (8).
- McEwan, Patrick J. and Joseph S. Shapiro. 2008. “The Benefits of Delayed Primary School Enrollment.” *Journal of Human Resources* 93 (2):1–29.
- Mulder, Bianca Leest, Lia, Annemiek Veen, Iris Bollen, Janneke Huizinga, and Geertje Damstra. 2016. “Doorstroom van Kleuters: Achtergrondrapportage bij de Brochure voor Scholen en Besturen.”
- Mühlenweg, Andrea M. 2010. “Young and Innocent: International Evidence on Age Effects within Grades on Victimization in Elementary School.” *Economic Letters* 109:157–160.
- Mühlenweg, Andrea M. and Patrick A. Puhani. 2010. “The Evolution of the School-Entry Age Effect in a School Tracking System.” *Journal of Human Resources* 45:407–438.
- OECD. 2023. *Education at a Glance 2023: OECD Indicators*. OECD Publishing. Accessed: 2024-10-05.
- Oosterbeek, Hessel, Simon ter Meulen, and Bas van der Klaauw. 2021. “Long-Term Effects of School-Starting-Age Rules.” *Economics of Education Review* 84.
- Page, Lionel, Dipanwita Sarkar Sarkar, and Juliana Silva-Goncalves. 2019. “Long-Lasting Effects of Relative Age at School.” *Journal of Economic Behavior and Organization* 168:166–195.
- Patrinos, Harry Anthony. 2011. “School choice in the Netherlands.” *CESifo DICE Report* 9 (2):55–59.
- Rosenzweig, Mark R and Kenneth I Wolpin. 1988. “Migration selectivity and the effects of public programs.” *Journal of Public Economics* 37 (3):265–289.
- Rossin-Slater, Maya. 2018. “The Importance of Early Childhood Education.” *Annual Review of Economics* 10:269–295.

BIBLIOGRAPHY

- Rothstein, Jesse. 2019. “Inequality of Educational Opportunity? Schools as Mediators of the Intergenerational Transmission of Income.” *Journal of Labor Economics* 37 (S1):85–123.
- Schneeweis, Nicole and Martina Zweimüller. 2014. “Early Tracking and the Misfortune of Being Young.” *Scand. Journal of Economics* 116:394–428.
- Sykes, Brooke and Sako Musterd. 2011. “Examining Neighbourhood and School Effects Simultaneously: What Does the Dutch Evidence Show?” *Journal of Urban Studies* 48 (4):1–18.
- Valdés, Manuel T. and Miguel Requena. 2024. “The Effect of the Age at School Entry on Educational Attainment and Field of Study: An Analysis Using the Spanish Census.” *Higher Education* 87 (4):1061–1083.
- Webbink, Dinand, Bas ter Weel, and Céline Odding. 2023. “On the Estimation of Neighbourhood Exposure Effects.” Unpublished manuscript.
- Wiswall, Matthew and Basit Zafar. 2015. “Determinants of college major choice: Identification using an information experiment.” *Review of Economic Studies* 82 (2):791–824.
- Wodtke, Geoffrey T. 2018. “Parental Resources and Children’s Cognitive Development.” *Sociology of Education* 91 (1):32–53.
- Wooldridge, Jeffrey M. 2002. *Econometric Analysis of Cross-Section and Panel Data*. MIT Press.
- Yamaguchi, Shintaro, Hirotake Ito, and Makiko Nakamuro. 2023. “Month-of-Birth Effects on Skills and Skill Formation.” *Labour Economics* 84.
- Yi, Junjian, James J. Heckman, Shuaizhang Zhang, and Gabriella Conti. 2015. “The dynamics of parental investments in response to health shocks in China.” *Journal of Labor Economics* 33 (2):323–363.

This dissertation is the result of my own work, and no other sources or means, except the ones listed, have been employed.

19.10.2024, Mannheim

Xi Lin

Curriculum Vitae

University of Mannheim

08.2018 – 11.2024

Ph.D. Economics

University of Mannheim

08.2016 – 08.2018

M.Sc. Economics

University of Mannheim

08.2013 – 08.2016

B.Sc. Economics

University of International Business and Economics (Beijing)

08.2009 – 08.2013

B.A. Economics