

Schüle, Paul Valentin

Research Report

Essays in public economics and on equality of opportunity

ifo Beiträge zur Wirtschaftsforschung, No. 108

Provided in Cooperation with:

Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

Suggested Citation: Schüle, Paul Valentin (2024) : Essays in public economics and on equality of opportunity, ifo Beiträge zur Wirtschaftsforschung, No. 108, ISBN 978-3-95942-138-6, ifo Institut - Leibniz-Institut für Wirtschaftsforschung an der Universität München, München

This Version is available at:

<https://hdl.handle.net/10419/300803>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Essays in Public Economics and on Equality of Opportunity

Paul Valentin Schüle



ifo
BEITRÄGE
zur Wirtschaftsforschung

108
2024

**Essays in Public Economics and
on Equality of Opportunity**

Paul Valentin Schüle

Herausgeber der Reihe: Clemens Fuest

Schriftleitung: Chang Woon Nam

ifo INSTITUTE

Leibniz Institute for Economic Research
at the University of Munich

Bibliografische Information der Deutschen Nationalbibliothek

Die Deutsche Nationalbibliothek verzeichnet diese Publikation in der Deutschen Nationalbibliografie; detaillierte bibliografische Daten sind im Internet über <https://dnb.d-nb.de> abrufbar.

ISBN Nr. 978-3-95942-138-6

Alle Rechte, insbesondere das der Übersetzung in fremde Sprachen, vorbehalten. Ohne ausdrückliche Genehmigung des Verlags ist es auch nicht gestattet, dieses Buch oder Teile daraus auf photomechanischem Wege (Photokopie, Mikrokopie) oder auf andere Art zu vervielfältigen.
© ifo Institut, München 2024

Druck: Pinsker Druck und Medien GmbH, Mainburg

ifo Institut im Internet:
<https://www.ifo.de>

Essays in Public Economics and on Equality of Opportunity

Inaugural-Dissertation
zur Erlangung des Grades Doctor oeconomiae publicae (Dr. oec. publ.)
an der Ludwig-Maximilians-Universität München

2024

vorgelegt von
Paul Valentin Schüle

Referent:	Prof. Dr. Andreas Peichl
Koreferent:	Prof. Dr. Dominik Sachs
Promotionsabschlussberatung:	10. Juli 2024

Acknowledgements

I am indebted to many wonderful people without whose support this thesis would not have been possible. First and foremost, I want to express my deepest gratitude to my supervisor Andreas Peichl for his trust and support during the last years. Andreas has been an excellent academic mentor who provided me with the freedom to independently pursue my research interests and the necessary opportunities to make use of this freedom.

Furthermore, I want to thank my second supervisor, Dominik Sachs, for his continuous support, for the successful collaboration, and for many insights in the field of public economics. I thank Clemens Fuest for serving as my third supervisor, for helpful comments, and for his efforts to make the ifo Institute a great place for pursuing a PhD.

I learned a lot from my co-authors Majed Dodin, Sebastian Findeisen, Lukas Henkel, Paul Hufe, Martyna Kobus, Sebastian Link, and Manuel Menkhoff, without whom the first three chapters of the thesis would not exist. It has been a pleasure to work with you. I am also grateful to Florian Englmaier and Yves Le Yaouanq: without your encouragement, I would not have started the PhD in the first place. I am indebted to Danny Yagan for hosting me at UC Berkeley, and I have benefited from financial support granted by the Deutsche Forschungsgemeinschaft, Fondazione Compagnia di San Paolo and the Volkswagen Foundation, CESifo, and the Joachim Herz Foundation.

I would like to thank all my colleagues at CES and ifo for creating a friendly and motivating workplace environment, and for the time spent together. I will miss both the coffee breaks at CES and the after-work beers at ifo. I am especially grateful to Mehmet Ayaz, Julia Baarck, Maximilian Blömer, Mathias Dolls, Oana Garbasevschi, Elena Herold, Leonie Koch, Carla Krolage, Max Lay, Jonas Löbbing, Jakob Miethe, Sascha Möhrle, Ann-Christin Rathje, Marcel Schlepper, Thomas Schwab, Anne Steuernagel, Marc Stöckli, Daniel Weishaar, and Lisa Windsteiger for joint collaboration or advice. Moreover, I thank Philipp Ludwig for many conversations around all aspects of the PhD, and Peter Naschold for encouraging my interest in economics early on.

This thesis is dedicated to my parents and to Marisa, from whom I received love and unconditional support at all times.

Contents

List of Figures	VII
List of Tables	XI
Preface	XIII
1 Downward Revision of Investment Decisions after Tax Hikes	1
1.1 Introduction	1
1.2 Institutional Background and Data	5
1.2.1 The German Local Business Tax	5
1.2.2 Firm-level Data on Revisions of Investment Plans	7
1.3 Empirical Strategy	11
1.4 Results	14
1.4.1 Revision of Investment Plans after Tax Hikes: Main Results	14
1.4.2 Magnitude of Effect Size in Comparison to the Literature	20
1.4.3 State Dependence and Heterogeneity	27
1.5 Conclusion	32
2 Social Mobility in Germany	35
2.1 Introduction	35
2.2 Related Literature and Institutional Background	37
2.2.1 Related Literature	37
2.2.2 Institutional Background	39
2.3 Data and Measurement Strategy	40
2.3.1 Variable Definition	41
2.3.2 Mobility Statistics	43
2.3.3 Sample Definition and Limitations	45
2.4 National Estimates	47
2.4.1 Subgroup Estimates	51
2.4.2 Time Trends	54
2.5 Regional Estimates	60
2.5.1 States	61

2.5.2	Cities	63
2.5.3	Local Labor Markets	65
2.6	Conclusion	71
3	Multidimensional Equality of Opportunity in the United States	73
3.1	Introduction	73
3.2	Measurement of Multidimensional Equality of Opportunity	76
3.3	Data	79
3.4	Results	81
3.5	Sensitivity Analysis	87
3.6	Conclusion	90
4	Career Preferences and Socio-Economic Background	91
4.1	Introduction	91
4.2	Conceptual Framework	95
4.3	Data	97
4.3.1	The German Socio-Economic Panel (SOEP)	97
4.3.2	The British Cohort Study	100
4.4	Empirical Evidence	104
4.4.1	Career Preferences Vary with Parental Background	104
4.4.2	Preferences Are Intergenerationally Persistent	105
4.4.3	Career Preferences Predict Future Earnings	108
4.5	Career Preferences and Income Mobility	114
4.6	Conclusion	118
A	Appendix to Chapter 1	121
A.1	Business Taxation in Germany	121
A.2	Data Appendix	124
A.2.1	Administrative Data at the Municipality Level	124
A.2.2	The ifo Investment Survey	129
A.2.3	Construction and Descriptive Statistics of the Merged Dataset	133
A.3	Supplementary Figures and Tables	140
A.4	Back-of-the-Envelope Calculation	154
A.5	Calculation of Effective Tax Rates	156
B	Appendix to Chapter 2	163
B.1	Additional Information on the Mikrozensus	163

B.2	Additional Figures and Tables	165
B.3	Data Appendix SOEP and PISA	178
B.3.1	Data Appendix SOEP	178
B.3.2	Data Appendix PISA	180
B.3.3	Adjustment of Ability Trends	181
B.4	Regional Predictors of Mobility	185
B.5	Robustness of Regional Estimates	197
C	Appendix to Chapter 3	201
C.1	Measurement—Some Simple Examples	201
C.2	Attribute Decomposition	203
C.3	Additional Figures and Tables	206
D	Appendix to Chapter 4	215
D.1	Data Appendix	215
D.1.1	Additional Data Description SOEP	215
D.1.2	Additional Data Description British Cohort Study	220
D.2	Additional Figures and Tables	223
	Bibliography	235

List of Figures

Figure 1.1:	Variation in Local Business Tax Rates (1980-2018)	6
Figure 1.2:	Relationship between Planned and Realized Investment	10
Figure 1.3:	Event Study: Investment Revision Effect after a Tax Hike	18
Figure 1.4:	Effect of Tax Hike on Investment Plans, Realizations, and Revisions	21
Figure 1.5:	Investment Revisions after a Tax Hike: State Dependence	28
Figure 1.6:	Testing for Further Heterogeneity	31
Figure 2.1:	Move-out Frequency by Parental Income Rank	46
Figure 2.2:	Social Mobility at the National Level	48
Figure 2.3:	Differences by Parental Education	53
Figure 2.4:	A-Level Share by Cohort	55
Figure 2.5:	Parental Income Gradient by Cohort	56
Figure 2.6:	Quintile Measures by Cohort	56
Figure 2.7:	A-Level Share by Cohort Quintile	57
Figure 2.8:	Time Trend Math Grades and Test Scores	59
Figure 2.9:	Social Mobility in Hamburg and Leipzig	65
Figure 2.10:	A-Level Share and Q1 Measure by Local Labor Market	66
Figure 2.11:	Parental Income Gradient by Local Labor Market	67
Figure 2.12:	Sorting: Conditional and Unconditional Rank Gradients	70
Figure 3.1:	Intergenerational Income Mobility and the Distribution of Monetary Resources in the United States	74
Figure 3.2:	Inequality of Opportunity in the US - Intergenerational Sample	82
Figure 3.3:	Inequality of Opportunity in the US, 1983-2016 Baseline Estimates	84
Figure 3.4:	Inequality of Opportunity in the US, 1983-2016 Decomposition by Background Characteristic and Outcome Dimension	86
Figure 3.5:	Inequality of Opportunity in the US, 1983-2016 Sensitivity to Parameter Choices	88
Figure 4.1:	Causal Model	95
Figure 4.2:	Correlation of Career Preferences	98
Figure 4.3:	Principal Component Analysis Career Preferences	99

List of Figures

Figure 4.4:	Correlation of Career Preferences in the BCS	102
Figure 4.5:	Variation in Career Preferences by Parental Income Rank	105
Figure 4.6:	Variation in Career Preferences by Parental Income Rank in the British Cohort Study	106
Figure 4.7:	Intergenerational Persistence in Preferences and Values	107
Figure A.1:	Timing of Tax Hike News	122
Figure A.2:	Time Series of Local Scaling Factors by Municipality	126
Figure A.3:	Share of Municipalities Increasing the LBT over Time	127
Figure A.4:	Number of Tax Hikes and Distribution of Tax Changes	127
Figure A.5:	Predictability of Tax Hikes as a Function of Past Tax Hikes in the Same Municipality	128
Figure A.6:	Distribution of Firms by Number of Employees	136
Figure A.7:	Time Series of Investment Plans and Realizations	136
Figure A.8:	Time-Series of Investment Revisions	137
Figure A.9:	Time-Series of Share of Large Revenue Drops	137
Figure A.10:	Distribution of the Log Revision Ratio	139
Figure A.11:	Long Event Study: Effect of Tax Hike on Investment Plans, Realizations, and Revisions	140
Figure A.12:	Investment Revision Effect after a Tax Hike: Alternative Estimators	141
Figure A.13:	Event Study: Expenditures and Revenues of Municipalities	142
Figure A.14:	Investment Revisions after a Tax Hike: Permutation Test	143
Figure A.15:	Collectively Bargained Wage Growth in Manufacturing	144
Figure A.16:	Obstacles to Investment by Firm Size	144
Figure A.17:	Time Series of Average Interest Rate on Loans for Firms	159
Figure A.18:	Present Discounted Value of Depreciation: 7% vs Time-Varying Interest Rate	159
Figure A.19:	Relation of Changes in LBT Rate and Changes in Effective Tax Rates	160
Figure A.20:	Relation of Changes in Effective Tax Rates and Changes in User Cost of Capital	161
Figure B.1:	Illustration of the Microcensus Survey Design	163
Figure B.2:	Household Income by Percentile Rank	164
Figure B.3:	A-Level Wage Premium, Years 1997-2016	166
Figure B.4:	National Estimates under Different Equalization Schemes	167
Figure B.5:	Social Mobility for Subgroups	168
Figure B.6:	Time Trend A-Level Share for Subgroups	169
Figure B.7:	Time Trend Parental Income Gradient for Subgroups	170

Figure B.8: Social Mobility at the State Level	171
Figure B.9: Social Mobility for Cities	172
Figure B.10: Q5/Q1 Ratio by Local Labor Market	173
Figure B.11: Mobility Estimates by Spatial Planning Region	175
Figure B.12: Mean Parental Income Rank by Local Labor Market	176
Figure B.13: Robustness to State and Region Specific Parental Income Ranks	177
Figure B.14: Social Mobility in the SOEP	179
Figure B.15: Predicting the A-Level Share by Parental Income Quintile	193
Figure B.16: Shrinkage Estimates and Parametric EB Confidence Intervals	199
Figure B.17: Shrinkage Estimates and Robust EB Confidence Intervals	200
Figure C.1: Wealth in PSID and SCF, 1983-2016	208
Figure C.2: Parental Income by Family Background Type	209
Figure C.3: Inequality of Opportunity in the US Re-weighted Sample	210
Figure C.4: Equality of Opportunity in the US Varying Age Restrictions	211
Figure C.5: Equality of Opportunity in the US Extended Family Background Character- istics	212
Figure C.6: Inequality of Opportunity in the US, 1983-2016 Sensitivity to Data Choices	213
Figure D.1: Ordered Career Preferences	215
Figure D.2: Career Question in the SOEP Questionnaire	217
Figure D.3: Ordered Career Preferences in the BCS	221
Figure D.4: Principal Component Analysis Career Preferences in the BCS	221
Figure D.5: Career Question in the BCS Questionnaire	222
Figure D.6: Variation in Career Preferences by Parental Years of Education	223
Figure D.7: Variation in Career Preferences by Parental Occupation (Standard Interna- tional Occupational Prestige Scale)	224
Figure D.8: Intergenerational Persistence in Preferences and Values – By Quintile of Parental Income	225
Figure D.9: Career Preferences Predict Future Earnings	226

List of Tables

Table 1.1:	Tax Hikes across Municipalities and Firms: Summary Statistics	8
Table 1.2:	Difference-in-Differences: Investment Revisions after a Tax Hike	16
Table 1.3:	Investment Revisions after a Tax Hike: Effective Tax Rates	24
Table 1.4:	Summary of Estimated (Semi-)Elasticities	26
Table 2.1:	Co-Residence Rate by Child Age	42
Table 2.2:	Monthly Child-related Expenditures of Single Child Households	43
Table 2.3:	Average Characteristics of Children by Age of Observation	46
Table 2.4:	Intergenerational Income Mobility in the US, Denmark and Germany . . .	50
Table 2.5:	Mobility Statistics for Subgroups	52
Table 2.6:	Math Grades and Test Scores of Marginal Children	60
Table 2.7:	Social Mobility at the State Level	62
Table 2.8:	Social Mobility in the 15 Largest Urban Labor Markets	64
Table 4.1:	Labor Earnings and Career Preferences in the SOEP	109
Table 4.2:	Controlling for Mediating Factors	112
Table 4.3:	Labor Earnings and Career Preferences in the BCS	113
Table 4.4:	Career Preferences and Intergenerational Income Mobility	116
Table 4.5:	Career Preferences and Intergenerational Income Mobility in the British Cohort Study	117
Table A.1:	Distribution of Firms in the IVS by Industry and Size	130
Table A.2:	Summary Statistics of Firms in the Sample	135
Table A.3:	Balance Statistics of Firms in the Treatment and Control Group	135
Table A.4:	Information Content of Investment Plans for Realized Investment	139
Table A.5:	Robustness: Baseline Estimates Excl. Reunification Period	145
Table A.6:	Treatment Effect Heterogeneity: State Dependence	146
Table A.7:	Difference-in-Differences: Investment Revisions after a Tax Hike	147
Table A.8:	Treatment Effect Heterogeneity: Volatility of Revenue Growth	148
Table A.9:	Treatment Effect Heterogeneity: Current Revenue Growth I	149
Table A.10:	Treatment Effect Heterogeneity: Current Revenue Growth II	150
Table A.11:	Treatment Effect Heterogeneity: Financial Constraints	151

List of Tables

Table A.12: Treatment Effect Heterogeneity: Firm Size and Settlement Structure . . .	152
Table A.13: Treatment Effect Heterogeneity: Tax Hike Dynamics	153
Table B.1: National Estimates for Different Age-Restrictions	165
Table B.2: Correlation between Mobility Measures	174
Table B.3: Math Grades and Test Scores of Inframarginal Children	183
Table B.4: Math Grades and Test Scores of Marginal Children - Adjusted for Changes among Inframarginal Children	184
Table B.5: List of Regional Indicators	188
Table B.6: The 15 Most Informative Predictors of Relative Mobility	190
Table B.7: Social Mobility and Regional Characteristics	191
Table B.8: The 15 Most Informative Predictors by LLM Size	196
Table B.9: Reliability Ratios of Mobility Measures	198
Table C.1: Descriptive Statistics	206
Table C.2: Observation Loss due to Sample Restrictions	207
Table C.3: Attribute Decomposition Alternative Parameterization	214
Table D.1: Correlation Table of Career Preferences	216
Table D.2: Big Five Scales in the SOEP	219
Table D.3: Annual Work Hours and Past Career Preferences	227
Table D.4: Years of Education and Past Career Preferences	228
Table D.5: Career Preferences Mediating Intergenerational Income Mobility	229
Table D.6: Career Preferences and Intergenerational Income Mobility – Robustness to Different Weighting Schemes	230
Table D.7: Career Preferences and Intergenerational Income Mobility – Controlling for Childhood Characteristics	231
Table D.8: Career Preferences Mediating Intergenerational Income Mobility in the British Cohort Study	232
Table D.9: Career Preferences and Intergenerational Income Mobility in the British Cohort Study Controlling for Childhood Characteristics	233

Preface

This thesis consists of four self-contained chapters that provide empirical evidence on firm responses to business taxation (Chapter 1) and the measurement of equality of opportunity (Chapters 2-4), covering the two main areas of research that I have been working on for the last four years.

In the first area of research, I use quasi-experimental evidence from Germany to measure the impact of public policies on the behavior of firms. Firms' behavioral responses to changes in taxes or subsidies can be quite complex. For example, their location decisions and the allocation of factor inputs across existing plants tend to be more sensitive to taxes and subsidies than those of individuals. Firms furthermore face many margins of adjustment. In response to a tax increase, firms may theoretically change prices, wages, production, or investment. In Chapter 1 of this thesis, co-authored with Sebastian Link, Manuel Menkhoff, and Andreas Peichl (Link et al., forthcoming), I evaluate the impact of increases in business taxes on firms' investment responses. Credible empirical evidence on this question is still scarce, as quasi-experimental variation in business tax rates is difficult to find. To make progress in this direction, we combine the specific system of business taxation in Germany with unique data on firm level investment plans and their realizations to study how an increase in business taxes affects investment.

Our identification strategy builds on two pillars. First, we take advantage of the decentralized design of the German local business tax: While tax base and liability criteria are set by the federal government, municipalities each year autonomously decide on the statutory tax rates. We can therefore distinguish the variation in tax rates from potential changes in the tax base. Second, we estimate the investment response of corporate firms to these tax changes by using panel data on both planned and realized investment volumes. The unique feature of our data is that each fall, firms report the planned volume of investment for the subsequent year. Municipalities announce tax changes for the subsequent year typically in December, that is after firms have reported their investment plans. Consequently, firms are surprised by the tax changes and have not yet included this information in their investment plans. At the same time, investment plans arguably incorporate all other pieces of firms' (partially unobserved) private and public information that determine investment in the subsequent

year. Focusing on the revision of investment plans, i.e. the difference between the investment volume planned prior to the tax change and the investment volume ultimately realized, thus eliminates many potential confounding factors of firms' investment choices.

Our results show economically large and statistically significant investment responses for firms experiencing a tax increase. On average, the share of firms that invest less than previously planned increases by 3 percentage points after a tax hike. The magnitude of the investment response varies substantially over the business cycle. Compared to our baseline estimates, the share of firms that invest less than previously planned in response to a tax hike is twice as large if taxes are increased during a recession. Potential explanations for this state dependence of tax shocks relate to uncertainty about expected returns to investments, cashflow sensitivity, and tax incidence.

In ongoing work not included into the dissertation, I build on the results of this chapter by asking how local shocks such as business tax hikes affect the allocation of resources within firms with multiple establishments, and study the implications for the spatial allocation of economic activity in Germany.

The second area of my research concerns the measurement of equality of opportunity. The fairness principle of equality of opportunity states that individuals should only be held responsible for factors within their control. It implies that inequality resulting from circumstances beyond personal control, such as parental income, is considered unfair and ought to be compensated. To assess the extent to which this principle is realized in practice, it is necessary to map the distribution of opportunities against the distribution of circumstances.

Most discussions in the literature have so far centered around the correct measurement of circumstances. For example, while the literature in economics, reviewed in Black and Devereux (2011), has traditionally measured equality of opportunity by the intergenerational correlation of income, another strand of the literature reviewed in Roemer and Trannoy (2016) argues that next to parental income, other circumstances such as race, gender, or parental education, should be taken into account as well. The measurement of opportunities, on the other hand, has received less attention. In fact, many applied papers in the literature directly equate opportunities with realized income without further discussion. Yet, while I agree that income constitutes the natural first outcome to look at, there are several reasons why it is valuable to also consider other measures of opportunity. Chapters 2, 3, and 4 of this thesis are dedicated to this task.

In Chapter 2, co-authored with Majed Dodin, Sebastian Findeisen, Lukas Henkel, and Dominik Sachs (Dodin et al., 2024), I characterize intergenerational social mobility in Germany. Social mobility is an important indicator for both equality of opportunity and economic efficiency in a society. Despite its importance, reliable mobility statistics are not available for many countries, as measuring social mobility requires data that allow to link parental outcomes to a measure of opportunities for children. Household panel studies may contain this information but are typically too small to deliver sufficiently precise estimates for regional comparisons or the analysis of time trends. An attractive alternative are administrative data sources, such as linked tax records. As in many other countries, however, such data is not available in Germany, where to date no large-scale empirical study of social mobility across time and space exists.

In order to fill this gap, we implement a new measurement strategy for social mobility in Germany. Motivated by Germany's early tracking system in secondary education, our mobility statistics measure the association between parental income and the educational opportunities of children. Our measure of opportunities captures whether a child will obtain the A-Level (Abitur), the highest secondary schooling degree in Germany. Using census data, we are able to link 526,000 children to their parents.

We show that on average, a 10 percentile increase in parental income rank is associated with a 5.2 percentage point increase in the probability of obtaining an A-Level degree. For the birth cohorts 1980-1996, this parental income gradient has not changed despite a large-scale expansion of upper secondary education in Germany. We furthermore are the first to document geographic variation in social mobility across German states, cities, and local labor markets. For example, the top-bottom gap in the probability of obtaining an A-Level degree is 20 percentage points larger in Bremen than in Hamburg, two city states approximately 100 kilometers apart. Finally, we can show that household characteristics can explain only a small fraction of the variation in mobility measures across local labor markets, leading us to reject the hypothesis that sorting is the major driver behind the regional variation in mobility.

In Chapter 3, co-authored with Paul Hufe, Martyna Kobus, and Andreas Peichl (Hufe et al., 2022a), I analyze the association between family background characteristics and the joint distribution of income and wealth. Based on a novel multidimensional measure (Kobus et al., 2020), we characterize the development of equality of opportunity in the US for the time period 1983-2016. By focussing exclusively on either income or wealth, many prior studies neglect important drivers of individual consumption possibilities, which arguably are the relevant metric to assess the financial well-being of individuals. For example, unidimensional

analyses will misrepresent both the financial well-being of income-poor heirs who supports their lifestyle by selling assets and asset-poor persons with high and stable income flows. Therefore, if society cares for the economic well-being of individuals more broadly, it is useful to move from unidimensional analyses of monetary resources to analyses that target the joint distribution of income and wealth.

We find that multidimensional inequality of opportunity is consistently and significantly higher than inequality of opportunity in incomes. This finding entails that unidimensional analyses that focus on income only underestimate the extent to which monetary resources are associated with family background characteristics. Furthermore, inequality of opportunity in 2016 is 77% higher than in 1983; hence, the playing field in the US has become more tilted in recent decades. Time trends are different when accounting for the multidimensionality of monetary resources. For example, an exclusive focus on income suggests only moderate increases in unequal opportunities after the year 2000. This relative stability, however, is accompanied by strong increases in the wealth dimension. As a consequence, when accounting for the multidimensionality of monetary resources, it is much harder to reject the hypothesis that social mobility in the US has declined in recent years.

In the single-authored Chapter 4, I take yet another look at the intergenerational persistence of socio-economic status. To explain why incomes and other markers of socio-economic status are correlated between children and parents, researchers in economics have typically pointed to differential circumstances for children from advantaged and disadvantaged households, such as credit constraints, school and peer quality, or parental time investments. However, even conditional on these factors, economic outcomes between children from advantaged and disadvantaged households may differ if these children hold systematically different views about what is valuable in a career.

In the chapter, I show for the UK and Germany that career preferences indeed differ significantly by socio-economic status. For example, children from advantaged backgrounds report that they place less value on income and job security, but more on having an interesting job than children from disadvantaged backgrounds. The household panel data that are used in the analysis further enable me to link career preferences to actual labor market outcomes. I can therefore provide rich evidence on the interaction between own income, parental income, and career preferences. I find that career preferences elicited at age 16 or 17 are highly

predictive of the labor market outcomes of these children when they are 28 to 46 years old. Overall, my results show for the first time that career preferences are a likely mediator of the intergenerational persistence of socio-economic status.

1 Downward Revision of Investment Decisions after Tax Hikes

This chapter is based on co-authored work with Sebastian Link, Manuel Menkhoff, and Andreas Peichl, and is forthcoming in the American Economic Journal: Policy. See Link et al. (forthcoming) for the full reference.

1.1 Introduction

The effect of corporate taxes on firm investment is a central question in macroeconomics and public finance. Corporate tax reforms like the US Tax Cuts and Jobs Act (TCJA) are often motivated by the argument that high corporate tax rates inhibit firm investment and growth (CEA, 2017). Standard theories of corporate taxation indeed predict that firms cut on investment projects if their after-tax net present value is reduced by tax increases (Hall and Jorgenson, 1967). To what degree corporate taxation affects investment, however, is ultimately an empirical question. Credible evidence on it is still scarce, as estimating the causal effect of corporate taxes on investment is challenging.

On the one hand, attributing cross-country discrepancies in investment behavior to differences in corporate tax rates is difficult to justify, as the timing of tax reforms often correlates with other macroeconomic determinants of firm investment. On the other hand, studies exploiting within-country variation need a valid control group and face the problem that many national-level tax reforms such as the TCJA change several parameters of the tax system simultaneously. For these reasons, quasi-experimental evidence on the response of investment to changes in the corporate tax burden originates predominantly from targeted tax deductions, which provide exogenous variation in exposure to tax decreases across firms of different size or in different industries (e.g., Zwick and Mahon, 2017; Garrett et al., 2020; Ohn, 2018). However, to what extent the effects of such specific policies generalize to changes in the corporate tax rate remains unclear.

This paper addresses this gap by combining the specific system of business taxation in Germany with unique data on firm-level investment plans and their realizations. Our identification strategy builds on two pillars. First, we exploit the decentralized design of the German local

1 Downward Revision of Investment Decisions after Tax Hikes

business tax (LBT): While tax base and liability criteria are set by the federal government, municipalities each year autonomously decide on the statutory tax rates.¹ We can therefore distinguish tax rate variation from potential changes in the tax base. Furthermore, municipalities adjust their taxes frequently. Restricting the analysis to tax increases, which are much more common than tax cuts, our identifying variation consists of 1,443 tax hikes between 1980 and 2018. The large number of tax hikes allows us to control for potentially heterogeneous time trends across regions or industries.

Second, we estimate the investment response of firms to these tax changes by leveraging panel data on both planned and realized investment volumes among a large, representative survey of on average 1,500 German manufacturing firms. The unique feature of our data is that each fall, firms report the planned volume of investment for the subsequent year. Municipalities announce tax changes for the subsequent year typically in December, i.e., after firms have reported their investment plans. In consequence, firms are surprised by the tax changes and have not included this information in their investment plans. At the same time, investment plans arguably incorporate all other (partially unobserved) private and public information of the firms that determine investment in the subsequent year.

Focusing on the revision of investment plans, i.e., the difference between the investment volume planned prior to the tax change and the investment volume ultimately realized, is advantageous from several perspectives.² Most importantly, investment revisions allow us to estimate the effect of corporate taxes on firm investment under weaker assumptions than usually possible. Because investment plans incorporate all relevant firm-level information, our results would still be unbiased if, for example, the occurrence of tax hikes were endogenous to local economic conditions. Moreover, considering revisions avoids problems with sensitivity in estimates due to the lumpy nature of investment, and hedges against potential bias in two-way fixed effects models (de Chaisemartin and D'Haultfœuille, 2022).

Our results show economically large and statistically significant investment responses for firms experiencing a tax increase. On average, the share of firms that invest less than previously planned increases by approximately 3 percentage points after a tax hike. In terms of magnitudes, a 1 percentage point increase in the LBT rate is associated with a decrease in the ratio of realized over planned investment by 2.3-3.8 percent, depending on the empirical

¹ This variation has been used by Fuest et al. (2018) to study the wage incidence of corporate taxation and Isphording et al. (2021) to assess the effects on R&D spending.

² Comparing planned to realized quantities connects to the macro literature exploiting deviations from forecasts for identification (e.g., Romer and Romer, 2004).

specification. As firms on average invest approximately as much as previously planned, this maps into a semi-elasticity of investment with respect to the LBT rate of around 3. The corresponding elasticity of investment with respect to the net-of-tax rate is of similar magnitude. We verify our identification approach with an event study design, demonstrating that firms only deviate from the baseline probability for revising an investment decision in the year of the tax hike. While our baseline specification exploits variation in statutory tax rates (as previous literature on LBT in Germany did, see, e.g., Fuest et al., 2018; Isphording et al., 2021), we find similar effects when relying on effective tax rates that are more common in studies for other countries and settings.

The magnitude of the investment response varies substantially over the business cycle. Compared to our baseline estimates, the share of firms that invest less than previously planned in response to a tax hike is twice as large if taxes are increased during a recession. We discuss three potential explanations for this state dependence of tax shocks, relating to uncertainty about expected returns to investments, cashflow sensitivity, and tax incidence.

Our main contribution is to investigate the impact of hikes in the corporate tax rate on firm investment.³ While we are not the first to study this important question, we add to the literature along two dimensions. First and foremost, by using a novel identification strategy based on revisions of investment plans to investigate firms' investment response, we can eliminate concerns about omitted variable bias that have not been fully resolved in most of the previous literature. When using realized investments as outcome variable instead, results could be biased if tax policy responds to economic conditions. For example, Giroud and Rauh (2019) and Ivanov et al. (2022) investigate the effects of changes in US state level taxes on firm level outcomes.⁴ To be precise, the former paper studies the effects of tax changes on the reallocation of labor and capital across states while the latter looks at corporate leverage as the main outcome. In additional analyses, both papers also investigate—among others—effects on firms' capital stock. For identification, both papers rely on parallel trends between US states.⁵ While in our context the variation is on a more local level (municipalities within

³ Other studies investigate firm-level responses to the corporate income tax along other margins (e.g., Auerbach, 2006; Suárez Serrato and Zidar, 2016; Fuest et al., 2018; Garrett et al., 2020; Ljungqvist and Smolyansky, 2018). We add to this literature by providing new evidence on the investment response.

⁴ Mertens and Ravn (2013) use aggregate data and combine a narrative approach with a structural VAR model to exploit changes in US federal corporate taxes.

⁵ Both studies provide an extension using a narrative approach in the spirit of Romer and Romer (2010) that classifies arguably exogenous tax changes. Yet, these approaches only exclude a small number of potentially endogenous tax changes from the analysis and hence some concerns remain. Furthermore, the sample in Giroud and Rauh (2019) is restricted to large multi-state firms, and Ivanov et al. (2022) study tax decreases, whereas we focus on tax hikes.

1 Downward Revision of Investment Decisions after Tax Hikes

states), the key advantage is that assuming parallel trends is well justified by the fact that firms' ex ante planned volume of investment—i.e., the counterfactual level of investment in absence of a tax hike—should incorporate all firm-level information besides the tax shock that is relevant for investment in the subsequent year. That is, we do not require flat pre-trends in realized investment levels but only in terms of revisions of investment plans, which is much less demanding.

An alternative way to overcome endogeneity concerns is to focus on targeted tax deductions or accelerated depreciation allowances, giving rise to arguably exogenous variation in exposure to tax decreases for firms in different sectors and industries (Zwick and Mahon, 2017; Xu and Zwick, 2022; Garrett et al., 2020; Ohn, 2018; House and Shapiro, 2008; Curtis et al., 2021; Maffini et al., 2019; Guceri and Albinowski, 2021). However, the extent to which the effects of such specific policies generalize to changes of the tax rate remains unclear. Studying these targeted policies can therefore not substitute for a direct evaluation of the investment effects of changing the corporate tax rate, which affects all corporate firms at the same time and independently of their investment behavior. To the best of our knowledge, the only other paper using firm-level data and quasi-experimental variation to study the investment responses to a change in the universal corporate income tax rate is Harju et al. (2022).⁶ However, as the Finnish corporate tax cut also entailed an increase in dividend taxation, they cannot consistently disentangle the effects of both channels. Moreover, the German setting has the advantage to offer substantially larger variation, especially in terms of the number of tax rate changes.

In addition, our findings of higher investment responses during recessions relate to an ongoing debate about the state dependence of fiscal multipliers (Auerbach and Gorodnichenko, 2013; Ghassibe and Zanetti, 2022; Ramey and Zubairy, 2018) and the state dependence of investment effects in response to tax changes more specifically (Demirel, 2021; Hayo and Mierzwa, 2021; Jones et al., 2015; Ljungqvist and Smolyansky, 2018; Winberry, 2021). We complement this macroeconomic evidence by means of firm-level microdata and a distinct research design, showing that investment reacts much stronger to tax increases during recessions.

The remainder of the paper is structured as follows. Section 1.2 describes the municipality-level data on local business tax rates and the survey data on firm-level investment plans and their realizations. Section 1.3 presents our empirical strategy, while Section 1.4 documents the results. Section 1.5 concludes.

⁶ In the German context, Dobbins and Jacob (2016) compare the differential investment responses of domestically and foreign-owned firms after a cut in the federal corporate tax rate in 2008. Lerche (2022) estimates the effects of an investment tax credit in East Germany on firms' production behavior.

1.2 Institutional Background and Data

To investigate the effects of corporate tax rate changes on firm investment, we merge municipality-level data on local business tax rates with unique data on firm-level investment plans and their realizations.

1.2.1 The German Local Business Tax

Institutional Background. The local business tax (LBT) is one of three types of taxes on business income in Germany. It is applied to the operating profits of both corporate and non-corporate firms. While tax base and liability criteria of the LBT are set at the federal level, municipalities decide autonomously on the tax rate. The tax rate consists of two components: a basic rate, which is determined by the federal government, and a local scaling factor, which is set at the municipal level. Each year, the municipal council has to vote on next year's scaling factor, even if it remains unchanged. As it is common practice to decide on next year's local scaling factor jointly with the adoption of the budget in the year's last meeting of the municipal council, tax changes are typically announced in December.⁷ Municipalities in our sample are approximately ten times more likely to increase rather than decrease their local scaling factor. In consequence, the identifying variation in our setting is too weak to consistently estimate the effect of tax decreases on investment.⁸ We thus restrict the analysis to tax changes induced by municipalities increasing their local scaling factors, henceforth referred to as a tax hike. This implies that the tax reforms exploited in this paper affect investment exclusively via increases in the tax rate, and not via changes in the tax base. Taxable profits of firms with establishments in more than one municipality are divided between municipalities according to formula apportionment based on the payroll share. Appendix A.1 provides additional details on the institutional setting.

Variation in Business Tax Rates. We use information on municipal tax scaling factors from the Statistical Offices of the German Federal States for the years 1980 to 2018. We enrich these data with information on municipality budgets and local economic conditions from several administrative data sources, leaving us with a panel of all German municipalities with extensive

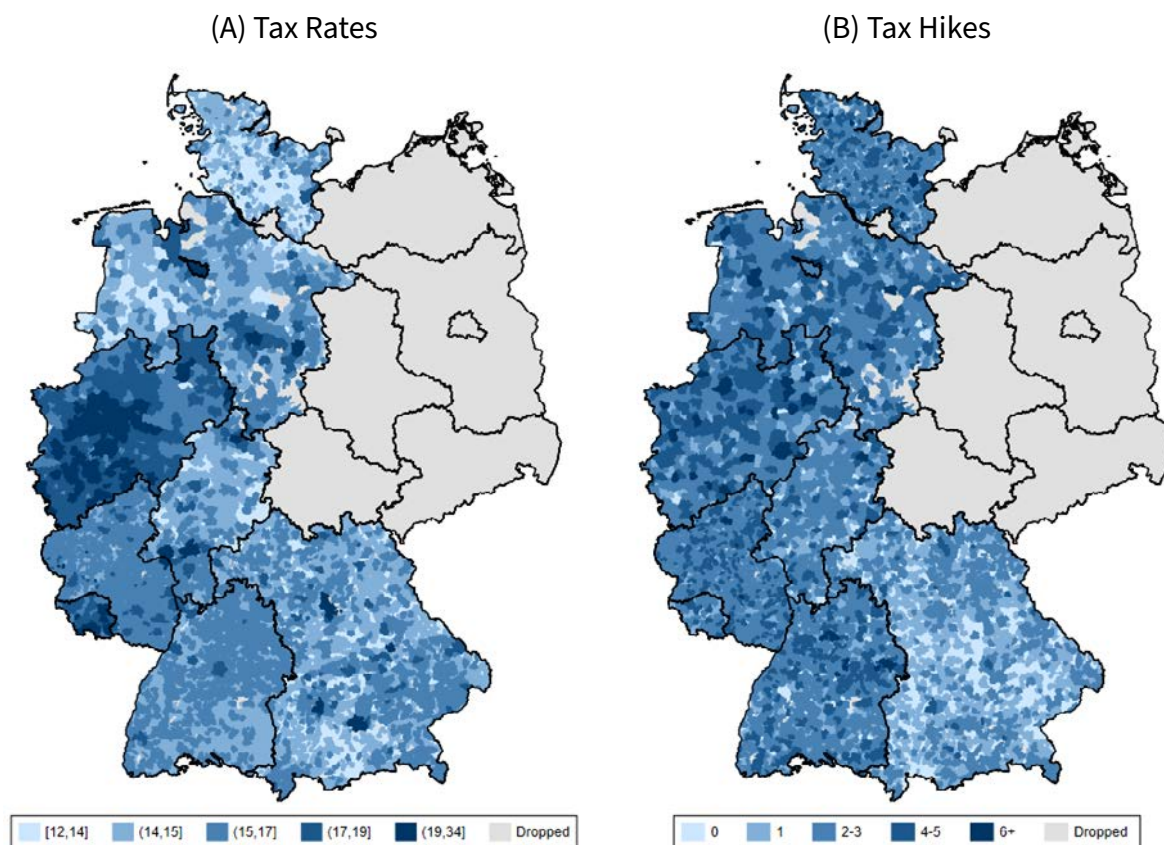
⁷ Appendix Figure A.1 substantiates this empirically, showing that newspaper coverage of LBT hikes indeed peaks each year in December.

⁸ The number of tax decreases that could in principle be used in the analysis is very low. Combining the municipality-level data on LBT rates and the firm-level data from the IVS, our analysis could only exploit 236 firm-year observations (0.7% of all observations) that face a tax drop in a given year despite spanning a time frame of almost four decades.

1 Downward Revision of Investment Decisions after Tax Hikes

information on taxes, revenues and expenditures. To avoid capturing structural changes of the German reunification, and as data for East Germany are only available since 1990, we restrict our sample to West German municipalities (excl. West-Berlin). We furthermore exclude the few municipalities that underwent a municipal merger during the period of consideration, as we cannot determine their exact tax rates.⁹

Figure 1.1: Variation in Local Business Tax Rates (1980-2018)



Notes: This figure shows the cross-sectional and time variation in municipal scaling factors of the German local business tax (LBT). Panel (A) plots the average LBT rate (in percent) induced by different scaling factors for the period 1980-2018. Panel (B) indicates the number of tax hikes, defined as an increase of the scaling factor. Municipalities in light grey areas are dropped from the sample as they are either located in East Germany or underwent a change of boundaries due to a merger. Moreover, we exclude observations where a tax hike was followed or preceded by another tax hike in the next or last two years.

There is substantial variation in LBT rates across municipalities and over time.¹⁰ As shown in Panel (A) of Figure 1.1, average tax rates differ strongly between municipalities, ranging from 12 to 34 percent. Panel (B) displays the identifying variation we rely on, i.e., the number of tax

⁹ Municipal mergers were very frequent in East Germany after 1990 and this rule would also lead to an exclusion of many municipalities in East Germany.

¹⁰ See Appendix A.2.1 for a more detailed description and investigation of the variation in LBT rates.

hikes between 1980 and 2018. Only few municipalities never increased the LBT in this period, while the median municipality increased the LBT rate three times and the median duration between two tax hikes in our sample is 13 years. The distribution of tax hikes is rather stable over time in terms of average size and dispersion (see Appendix Figure A.4). Importantly, past increases in the LBT contain very little predictive power for future tax hikes, as shown in Appendix Figure A.5.

After combining the municipality-level data on LBT rates with the firm-level investment data described in Section 1.2.2, we can exploit large parts of this variation in LBT rates. As summarized in the left panel of Table 1.1, our empirical strategy outlined in Section 1.3 relies on 1,443 tax hikes in 802 municipalities. The average tax hike amounts to 0.92 percentage points, corresponding to a 6 percent increase on average. The right panel summarizes the variation in tax hikes across firms. On average, approximately 7 percent of firms are exposed to a tax hike each year.

1.2.2 Firm-level Data on Revisions of Investment Plans

We use micro data on firms' investment behavior from the ifo Investment Survey (IVS, 2019). The IVS is conducted biannually (spring and fall) by the ifo Institute on behalf of the European Commission and covers a representative sample of incorporated firms in the German manufacturing sector.¹¹ The main purpose of the IVS is to obtain timely information on investment activity at disaggregated industry levels.¹² To achieve this goal, the IVS does not only elicit quantitative information on ex post realizations, but also on the planned volume of investment for the subsequent year. Thus, the panel structure of the IVS allows measuring how firms have revised their investment plans. In addition, survey participants provide quantitative information on revenues and the number of employees. The survey is usually completed by high-level management personnel at the firms' controlling departments (Sauer and Wohlrabe, 2020).¹³

¹¹ Appendix A.2.2 and Sauer and Wohlrabe (2020) provide additional information on the purpose and design of the survey, its representativeness, data access, and the wording of the survey questions used in the paper. The IVS micro data have been extensively used in recent research, e.g., Bachmann et al. (2017); Bachmann and Zorn (2020); Link et al. (2023).

¹² The German Federal Statistical Office releases information on realized investment at the levels of disaggregated industries only with a time lag of two years.

¹³ As noted on the cover letter of the survey, the ifo Institute guarantees compliance with strict data security criteria, that the data is evaluated in anonymous form, and that the survey results are only made available at the aggregate industry level. The participants therefore know that the firm-specific information reported to the


1 Downward Revision of Investment Decisions after Tax Hikes

Table 1.1: Tax Hikes across Municipalities and Firms: Summary Statistics

	Municipalities			Firm Observations			
	with Tax Hikes			with Tax Hikes		without Tax Hikes	
	N	Mean	SD	N	Share of Downward Revisions	N	Share of Downward Revisions
1980-1984	119	1.05	0.83	265	0.50	2655	0.45
1985-1989	131	0.92	0.53	340	0.49	4940	0.45
1990-1994	266	1.09	0.54	546	0.58	4831	0.54
1995-1999	228	0.94	0.47	385	0.51	4560	0.51
2000-2004	178	1.06	0.55	269	0.60	4711	0.58
2005-2009	106	0.78	0.51	161	0.66	4446	0.59
2010-2014	263	0.74	0.42	413	0.58	4118	0.58
2015-2018	152	0.69	0.38	248	0.63	2422	0.60
Full Sample	1443	0.92	0.54	2627	0.56	32683	0.54

Notes: This table reports summary statistics of the final sample used in the main analysis, i.e., after combining the municipality-level data on LBT rates and the firm-level data from the IVS. The left panel depicts the number of tax hikes at the municipality level that can be exploited in the empirical analysis along with the average size and standard deviation of these hikes. The right panel summarizes the number of firm observations that face a tax hike in a given year or not, as well as the average share of downward revisions of investment plans ($(I_{i,t})/(E_{i,t-1}[I_{i,t}]) < 1$) for each of these groups.

The timing of the survey is as follows:



In the fall of year t_{-1} , firms report how much they plan to invest in equipment and buildings (in Euro) in the subsequent year t_0 , denoted $E_{t-1}[I_{t0}]$. The realized investment volume of year

IVS neither is available to their stakeholders, nor can be related to the municipality they are located in—and where the LBT is set. Therefore, incentives to strategically misreport the company's future investment activity are limited and unrelated to the variation in the LBT.

t_0 , I_{t_0} , is elicited in the spring and fall survey of year t_1 .¹⁴ By comparing planned investment $E_{t-1}[I_{t_0}]$ to realized investment I_{t_0} , we observe whether firms in year t_0 invested more, less, or the same amount as previously planned. As municipalities announce the LBT rate for year t_0 at the end of year t_{-1} , i.e., after the fall survey, firms' investment plans for year t_0 reported to the IVS do not include information about changes in the LBT.


The investment data of the IVS have been shown to be very accurate. For instance, Bachmann and Zorn (2020) show that aggregate investment growth calculated from the microdata of the IVS is highly correlated with manufacturing investment growth reported by the Federal Statistical Office, and Sauer and Wohlrabe (2020) report that the average absolute deviation of the former from the latter is less than two percentage points. Moreover, Bachmann et al. (2017) present a series of stylized facts on the cross-sectional and time-series properties of revisions of investment plans, i.e., the difference between ex ante planned and ex post realized investment volumes, showing that these deviations are meaningful along many dimensions. For example, they document that the overall distribution of revisions is not systematically skewed, while their cross-sectional average is procyclical.¹⁵ This indicates that participants provide accurate investment plans given their current level of knowledge at the time of the survey.

We restrict the sample to firms that report both planned and ex post realized volumes of investment, referring to all years t_0 for which we have information on LBT rates in the municipality of their location available. Following the protocol deposited in Appendix A.2.3, we keep those firms for which we can observe revisions in investment plans for at least five years and, following Fuest et al. (2018), drop firms with legal forms that are exempted from the LBT. Our final sample consists of 35,310 firm-year observations in years $t_0 \in \{1980, 2018\}$ that are spread across 1,192 municipalities in West Germany. According to the descriptive statistics presented in greater detail in Appendix A.2.3, the median firm in our sample is a typical representative of the “German Mittelstand” employing 264 workers, generating annual revenues of 45 million Euro (CPI inflation-adjusted and—if denominated in German marks—converted to 2015 Euros), and investing 1.4 million Euro each year. For each firm, we can rely on information on reported planned and realized investment volumes in, on average, 17 years. In the final sample, firms report zero investment in only 0.7% of all observations.

¹⁴ Following Bachmann et al. (2017), we take the average of I_{t_0} if firms report it in both waves of year t_1 and drop the observation if these reports deviate more than 20% from the mean (see Appendix A.2.3 for details). The results are similar once we restrict the analysis to I_{t_0} reported in the fall wave.

¹⁵ Relatedly, Appendix Figure A.8 shows that the investment plans are more frequently and more strongly revised downward during recessions.

Figure 1.2: Relationship between Planned and Realized Investment



Notes: This figure shows a binned scatter plot between ex ante planned and ex post realized levels of investment in year t_0 (both in logs) as reported by firms to the IVS in years t_{-1} and t_1 , respectively. The red line depicts a linear fit of the data. The sample is restricted to observations in years without tax changes.

Importantly, investment plans for the next year contain a large amount of information that is highly predictive for the level of investment that is subsequently realized and that is changing within firms from year to year. The binned scatter plot depicted in Figure 1.2 demonstrates that the relationship between ex ante planned and ex post realized volumes of investment, i.e., $E_{t-1}[I_{t0}]$ and I_{t0} (both in logs), is highly linear and virtually corresponding to the 45 degree line. As depicted in Appendix Table A.4, 84% of the unconditional variation in (log) realized investment is explained by the investment plans for the respective year. Appendix A.2.3 presents a more detailed investigation of this relationship that, *inter alia*, demonstrates that investment plans, $E_{t-1}[I_{t0}]$, are much more strongly correlated with ex post realized investment, I_{t0} , than the realized level of investment in the previous year, I_{t-1} , and that these patterns even hold when controlling for firm fixed effects. Taken together, investment plans contain accurate information on subsequent year's investment that goes beyond the extrapolation of the level of investment that was realized in the year these plans are reported to the IVS.

The raw data provide a first indication of the main result of the paper, i.e., that firms revise investment decisions downwards after tax hikes. For each five-year interval of the data, the

right panel of Table 1.1 depicts the average share of downward revisions of investment plans separately for firms in municipalities with and without tax hikes. The share of downward revisions is—at least weakly—larger among treated firms than untreated firms in each time interval.¹⁶ We investigate this effect more systematically in the remainder of the paper.

1.3 Empirical Strategy

Research Design. We seek to identify the average treatment effect of an increase in the statutory LBT rate on firm investment. We consider a firm as treated in year t_0 if residing in a municipality that increased its LBT scaling factor from year t_{-1} to t_0 . The hypothesis guiding our analysis is that firms surprised by the announcement of a tax hike in December of t_{-1} will on average invest less in year t_0 than previously planned. We therefore expect downward revisions of planned investment to be more frequent in municipalities that increased their local scaling factors. At the same time, firms' investment plans elicited in the fall should incorporate all other, potentially unobserved, information influencing investment in the subsequent year.

Our identification strategy thus eliminates concerns about omitted variable bias. When using realized investment as outcome variable, results could be biased if tax policy responded to economic conditions, even after controlling for unit and time fixed effects, violating the parallel trends assumption. In our context this is quite different, as we observe the ex ante planned volume of investment—i.e., the counterfactual level of investment in absence of a tax hike—in addition to the ex post realized level of investment directly in our data. Using investment revisions instead of realized investment, we have a strong theoretical argument why we can extrapolate a (flat) pre-trend into the post-treatment period.

Hence, compared to using realized investment as dependent variable, our analysis only requires the weaker assumption that there are no unobserved factors that are both (i) correlated with investment and local tax policy in year t_0 , and (ii) not in the information set of the firm when forming investment plans in the fall of year t_{-1} . The only scenario that could violate this assumption would be a *local* shock that hits after firms have reported their investment plans, and that induces municipalities to implement a tax hike *within a few weeks*. Given the municipal decision structures and the “speed” of German bureaucracy, however, such an

¹⁶ In the pre-treatment year t_{-1} , the averages of both main outcome variables (Log Revision Ratio and Downward Revision Indicator) are not statistically different for firms that eventually are affected by a tax hike in year y_0 and firms ending up in the control group, see Panel B of Appendix Table A.3.

1 Downward Revision of Investment Decisions after Tax Hikes

immediate response is highly unlikely. Relatedly, Fuest et al. (2018) show that changes in the LBT are typically not triggered by shocks to economic variables, and Blesse et al. (2019) demonstrate that tax setting of the municipalities substantially deviates from theoretically optimal behavior. As in the US (Robinson and Tazhitdinova, 2022), regional variation in corporate tax rates seems to be to a large extent idiosyncratic and not readily explained by standard theories of tax setting. Overall, we are therefore confident that omitted variables do not threaten identification in our setting.

Instead, a potential limitation of our identification strategy is that some firms may put a positive probability on the scenario that taxes will be increased in the subsequent year, whereas our analysis implicitly assumes that firms expect taxes to remain constant. To the extent that this was not true, there would exist two potential sources of bias pointing in opposite directions. A downward bias originating from the treatment group (where some firms revise investment less strongly), and an upward bias originating from the control group (where some firms upward-revise investment if taxes are *not* increased). As long as the expected probability of a tax hike does not differ between treatment and control group, both biases will cancel out on average. Furthermore, any systematic and time-constant differences in the expected probability to be treated between firms and municipalities are inconsequential once we include firm fixed effects. However, private information about the likelihood of tax hikes may lead to systematically different beliefs in t_{-1} , the year before the LBT is increased (Riedel and Simmler, 2021). In particular, some firms may receive signals on the likely occurrence of a tax hike in the subsequent year even before investment plans are reported in the fall. In this case, they will—at least partially—incorporate this information into their investment plans and hence revise their investment decisions less strongly on average thereafter. If the private information helped firms to better predict the occurrence or absence of tax hikes, we hence should, if anything, tend to underestimate the investment response to a tax hike.¹⁷ In our data, however, we do not find evidence for a downward bias: Private information about future tax hikes should be more prevalent in smaller municipalities, where social ties to the municipality council are more likely, but treatment effects are not significantly different between cities and rural municipalities (see Figure 1.6, Panel B).

¹⁷ Moreover, the investment response might be underestimated due to the fact that firms pay the LBT according to the payroll share attributable to each municipality. As firm investment reported in the IVS refers to all domestic plants, the tax hike variation is measured with error for firms with plants in multiple municipalities. Although the IVS data lacks information on the prevalence and the payroll share of multi-establishments, the resulting attenuation bias is arguably small as only 7% of firms in the manufacturing sector were operating in multiple municipalities in 2017 according to aggregated, administrative LBT data (*Gewerbesteuerstatistik*). Furthermore, Panel A of Figure 1.6 demonstrates that treatment effects do not differ by firm size, a proxy for being a multi-plant company.

Measurement and Estimation. We use two variables to measure investment revisions. The first is an indicator for revising investment decisions downwards, defined as:

$$\text{Downward Revision:} \quad \mathbb{1} \left(\frac{I_{i,t}}{E_{i,t-1}(I_{i,t})} < 1 \right)$$

The downward revision indicator is attractive due to its robustness against outliers and non-linear investment responses. The second variable is the log revision ratio and takes the magnitude of each revision into account. It is defined as the natural logarithm of the ratio between realized and planned investment volumes:

$$\text{Log Revision Ratio:} \quad \ln \left(\frac{I_{i,t}}{E_{i,t-1}(I_{i,t})} \right)$$

We choose the logarithmic form due to the lumpy nature of investment, which means that the distribution of investment volumes is skewed and the revision ratio can get very large for small denominators. Moreover, the resulting estimates directly translate into the semi-elasticity of investment with respect to the tax rate, the relevant quantity of interest that we can directly compare to other estimates in the literature. As firms invest approximately as much as previously planned, the ratio of realized over planned investment is equal to one on average. As furthermore the revision ratio and realized investment are measured in logarithmic form, a tax hike which decreases the log revision ratio by 0.01 implies that both the revision ratio and realized investment decrease by 1%.¹⁸

In our main analysis, we estimate the following linear model by OLS:

$$\text{InvestmentRevision}_{i,m,t} = \gamma \text{TaxHike}_{m,t} + \mu_i + \phi_{l,t} + \psi_{s,t} + \varepsilon_{i,t}, \quad (1.1)$$

which explains the investment revision of firm i in municipality m and year t by municipality level tax hikes $\text{TaxHike}_{m,t}$ that take one of the following two forms:

$$\text{Tax Hike Indicator:} \quad \mathbb{1} (\Delta \text{tax}_{m,t} > 0)$$

$$\text{Tax Hike in Percentage Points:} \quad \Delta \text{tax}_{m,t}$$

¹⁸ In more recent years, firms tend to invest on average slightly less than previously planned (compare Appendix Figure A.8).

For this reason, the constant for the log revision ratio in Table 1.2 is not exactly zero, and the constant for the share of downward revisions is slightly larger than 0.5. As such, the semi-elasticity of the revision ratio with respect to an LBT hike will slightly understate the semi-elasticity of investment with respect to an LBT hike.

1 Downward Revision of Investment Decisions after Tax Hikes

The tax hike indicator equals one if at time t municipality m increased the LBT. In addition, $\Delta tax_{m,t}$ denotes the tax change in percentage points. As discussed above, the focus on deviations of realized investment from the planned value should by itself rule out omitted variable bias.¹⁹ Still, some specifications additionally include firm fixed effects (μ_i) and year fixed effects at the level of industries ($\psi_{s,t}$) and federal states ($\phi_{l,t}$) to flexibly control for any time-invariant heterogeneity or systematic time trends in the probability of investment revisions and the frequency of tax hikes. In these specifications, we obtain a (generalized) Difference-in-Differences (DiD) estimate.²⁰ Standard errors are clustered at the municipality level.

1.4 Results

We present our results in three steps: first, we show our baseline results and their robustness along various dimensions. Second, we discuss how the magnitude of the effects relates to other estimates in the literature, before third, documenting effect heterogeneity over the business cycle.

1.4.1 Revision of Investment Plans after Tax Hikes: Main Results

The baseline results presented in Table 1.2 reveal that firms affected by a tax hike strongly downward revise their investment decisions in the year this change is enacted. Panel (A) displays the estimates for the downward revision indicator. In Column (1), we compare the share of firms investing less than previously planned between municipalities where a tax hike is enacted and municipalities where the LBT rate did not change, without including any controls. We find that the share of firms that revise their investment decisions downwards is 2.7 percentage points higher in affected municipalities (Panel A1). The estimates presented in the remaining columns demonstrate that the point estimates for the tax hike indicator are

¹⁹ While Section 1.2.2 demonstrates that the investment plans reported to the IVS contain valuable information that is highly predictive for ex post realized investment volumes, these variables might be elicited imprecisely. The resulting measurement error in the dependent variable should thus decrease the precision of our estimates without resulting in attenuation bias.

²⁰ Note that using investment revisions as the outcome of interest implies that treatment effects realize exclusively in the treatment period. Due to this lack of treatment effect dynamics, the recent concerns about bias in two-way fixed effects models (e.g., de Chaisemartin and D'Haultfœuille, 2022) do not apply in this setting, as discussed below in more detail. It is, however, relevant in a setting when using realized investment (instead of revisions) as the outcome variable in Section 1.4.2. Inspired by Dube et al. (2023), firms are then assigned to another firm identifier in the middle between two tax hikes in order to ensure that there is only one treatment for each unit and to allow for different long-run trends.

barely affected by sequentially adding fixed effects at various dimensions, indicating that firms' investment plans already largely absorb regional and industry-specific shocks. For the size of the tax change (Panel A2), the inclusion of fixed effects tends to slightly increase the estimated coefficients. In Column (5), where we impose the most restrictive set of fixed effects, the effects of the tax hike indicator and the percentage change in the LBT on the probability of downward revising investment decisions are estimated at 3.3 and 2.4 percentage points, respectively.

Panel (B) repeats the analysis using the log revision ratio as dependent variable. The estimated coefficients are negative in all specifications of Panel B1, indicating that firms invest less than previously planned in response to a tax hike. While the effects are estimated less precisely compared to Panel A1, the point estimates are largely unaffected by the choice of the control vector. Again focusing on the most restrictive specification in Column (5), we find that the ratio of realized over planned investment decreases by 3.6 percent in response to a tax hike. Taking the magnitude of tax changes into account in Panel B2, the estimate in Column (5) implies that a 1 percentage point increase in the LBT rate is associated with a decrease in the revision ratio by 3.8 percent. Since in the absence of a tax hike firms invest approximately as much as they have planned, the ratio of realized over planned investment is close to one (and the log of the ratio is close to zero, as visible from the constant). Hence, our estimates directly map into a semi-elasticity of investment with respect to the LBT of around 3.

Overall, we find a clear and statistically significant negative investment response of firms to increases in corporate tax rates in all estimated models.

Economic Size of the Investment Response. The estimated investment response is economically sizable. To illustrate this, we conduct a back-of-the-envelope calculation, described in detail in Appendix A.4. According to our estimated semi-elasticity of 3, each additional Euro of tax revenues raised comes with a loss in firm investment of 2.12 Euro in the first year after a tax hike. If we also consider that lower firm investment reduces tax revenues in the medium run due to lower profits, the approximated investment loss for each additional Euro of tax revenue increases to a range between 2.14 Euro and 2.28 Euro, depending on the assumed strength of the second-round effect. While these projections rely on a series of simplifying assumptions, they still illustrate that the foregone volume of investment is non-negligible.

1 Downward Revision of Investment Decisions after Tax Hikes

Table 1.2: Difference-in-Differences: Investment Revisions after a Tax Hike

	(1)	(2)	(3)	(4)	(5)
<i>Panel (A): Downward Revision</i>					
A1: Tax Hike Indicator: $\mathbb{1}(\Delta tax_{m,t} > 0)$					
	0.027	0.028	0.026	0.028	0.033
	(0.011)	(0.010)	(0.011)	(0.011)	(0.011)
Constant	0.536	0.536	0.536	0.536	0.535
	(0.005)	(0.005)	(0.001)	(0.001)	(0.001)
A2: Tax Hike in Percentage Points: $\Delta tax_{m,t}$					
	0.012	0.018	0.017	0.021	0.024
	(0.009)	(0.009)	(0.010)	(0.009)	(0.010)
Constant	0.537	0.536	0.536	0.536	0.536
	(0.005)	(0.005)	(0.001)	(0.001)	(0.001)
Observations	35310	35310	35310	35310	35310
<i>Panel (B): Log Revision Ratio</i>					
B1: Tax Hike Indicator: $\mathbb{1}(\Delta tax_{m,t} > 0)$					
	-0.031	-0.033	-0.025	-0.029	-0.036
	(0.016)	(0.015)	(0.017)	(0.016)	(0.017)
Constant	-0.033	-0.033	-0.033	-0.033	-0.032
	(0.006)	(0.006)	(0.001)	(0.001)	(0.001)
B2: Tax Hike in Percentage Points: $\Delta tax_{m,t}$					
	-0.023	-0.032	-0.028	-0.034	-0.038
	(0.014)	(0.013)	(0.014)	(0.014)	(0.016)
Constant	-0.033	-0.033	-0.033	-0.033	-0.032
	(0.006)	(0.006)	(0.001)	(0.001)	(0.001)
Observations	34421	34421	34421	34421	34421
Firm FE	-	-	✓	✓	✓
Year FE	-	✓	-	✓	-
Year × State FE	-	-	-	-	✓
Year × Industry FE	-	-	-	-	✓

Notes: This table reports estimates from linear regressions of Equation (1.1). “Downward Revision” is an indicator that is one if the fraction of realized investment over planned investment is below one. “Log Revision Ratio” is the natural logarithm of this ratio. “Tax Hike Indicator” is an indicator that is one if the local corporate tax rate is higher than in the year before. “Tax Hike” is the change in the local corporate tax rate in percentage points compared to the previous year. Industry fixed effects refer to the ifo industry classification, comparable to two-digit NACE industries. Standard errors in parentheses are clustered at the municipality level.

From this approximation of the behavioral response, we can also derive the marginal value of public funds (MVPF) in the spirit of Hendren and Sprung-Keyser (2020), given as:

$$MVPF = \frac{\text{Beneficiaries' Willingness to Pay}}{\text{Net Cost to Government}}$$

In our setting, firms are the beneficiaries and their willingness to pay is equal to the change of the tax burden. The net cost of the government equals the change of tax revenues plus the additional revenue changes via the behavioral response. According to this, our estimates point at a MVPF in the range between 1.01 and 1.08, i.e., slightly above one.²¹ However, given that investment is not the only margin of adjustment through which firms may react to increases in the LBT, the true MVPF is presumably larger. For example, Fuest et al. (2018) find that workers bear about half of the total burden of the LBT via lower wages, while employment is unaffected. The corresponding loss in payroll tax revenues would show up in the denominator of the MVPF formula and thus further raise the MVPF. Likewise, the cross-base elasticity with respect to the corporate income tax base is presumably positive, so that corporate income tax revenues accruing to the federal government will decline as well. The MVPF should be even higher once these fiscal externalities were taken into account.

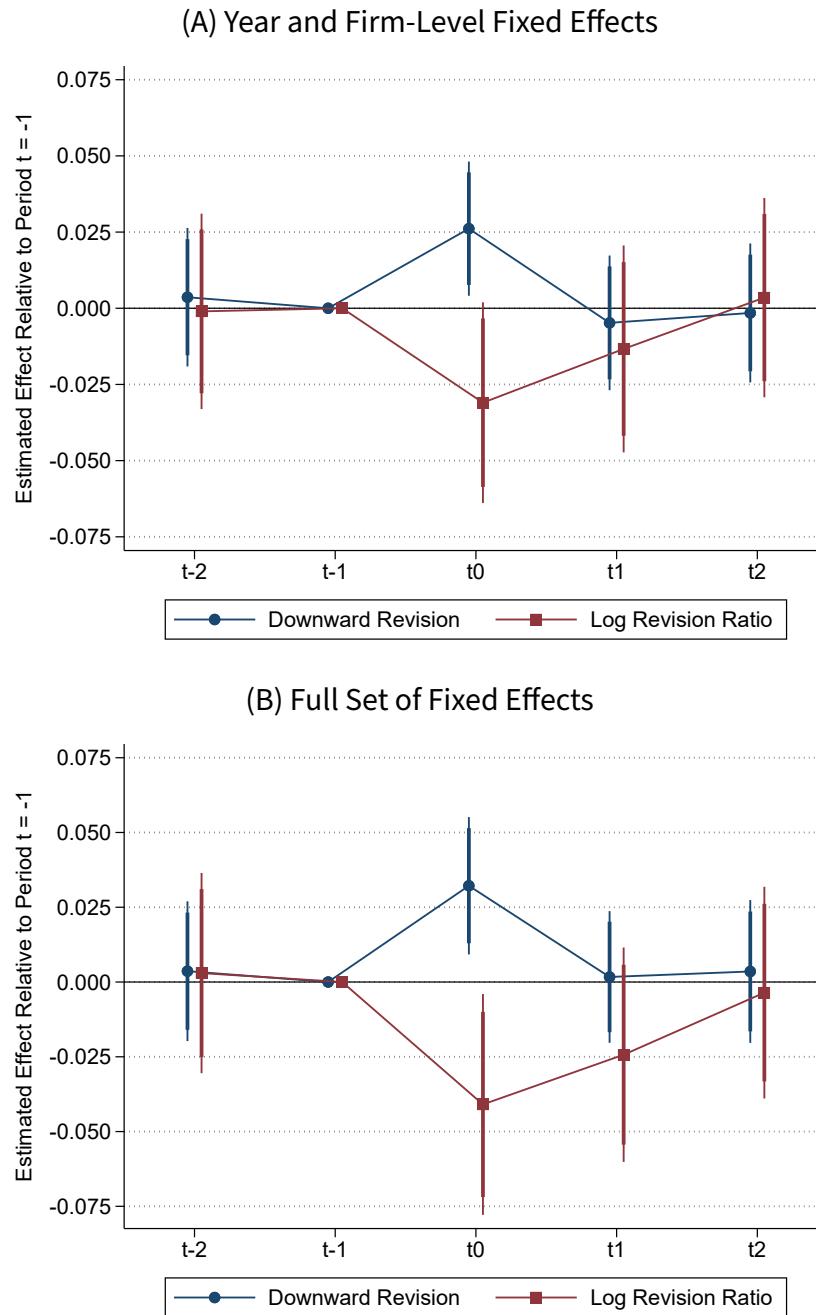
Validity of the Identifying Assumptions. Next, we estimate an event study to test a central implication of our identifying assumptions: an increase in investment revisions should only occur in the year of the tax hike (t_0), while no effect should be visible in the years before, when the tax hike could not have been anticipated, i.e., we should observe parallel trends before the tax change. Moreover, a tax hike implemented in January of year t_0 should be known to the firm (a) when reporting its investment plans for year t_1 to the IVS (in the fall wave of year t_0) and (b) when reporting the actual volume of investment for year t_1 (in the spring or fall wave of year t_2). Hence, investment revisions should also not be systematically higher in any period after year t_0 .

The results of the event study regression presented in Figure 1.3 confirm that investment revisions occur immediately in t_0 when the tax hike is enacted. In contrast, the point estimates

²¹ The calculation is: $(19,800)/(19,800 - 139) = 1.01 \geq MVPF \leq 1.08 = (19,800)/(19,800 - 1,386)$, where the value of 19,800 EUR refers to the increase in overall tax revenues paid by the median firm in response to a one percentage point increase in the LBT and the values of 139 EUR and 1,386 EUR denote the assumed upper and lower bound of the behavioral response, as described in Appendix A.4.

1 Downward Revision of Investment Decisions after Tax Hikes

Figure 1.3: Event Study: Investment Revision Effect after a Tax Hike



Notes: This figure shows the estimates of the following event-study regression: $\text{InvestmentRevision}_{i,t} = \sum_{j=-2}^2 \gamma_j \text{TaxHike}_{m,t}^j + \varepsilon_{i,t}$. In Panel (A), we additionally include year and firm fixed effects. In Panel (B), industry-year, state-year, and firm fixed effects are included. The reference period is t_{-1} . The dependent variable is based on the ratio of realized investments over planned investments (elicited in the fall of the previous year). “Downward Revision” is an indicator that is one if the ratio is below one. “Log Revision Ratio” is the natural logarithm of this ratio. Industry fixed effects are at the ifo industry classification level that is comparable to two-digit NACE industries. The confidence intervals refer to the levels of 90% (thick line) and 95% (thin line).

are (close to) zero in all other years, supporting the validity of our identifying assumptions. Panel A of Appendix Figure A.11 shows that these patterns also hold when extending the time window covered by the estimation to four years prior and post treatment.²²

Moreover, recent research in econometrics calls for caution when estimating two-way fixed effects models in generalized DiD settings with multiple treatment groups and periods (see, e.g., the survey by de Chaisemartin and D'Haultfœuille, 2022), as these estimators only provide an unbiased DiD estimate if the treatment effect is constant between groups and over time. This problem is less relevant in our setting given that the estimated treatment effects are not dynamic, i.e., do not evolve over time (as shown above). Still, to demonstrate that the recent critique does not apply here, we repeat the event study using the imputation estimator proposed by Borusyak et al. (forthcoming) as well as the interaction-weighted estimator by Sun and Abraham (2021). As shown in Appendix Figure A.12, results are very similar to Figure 1.3.

Robustness of Main Results. Next, we demonstrate that our main results are robust along various dimensions. We start by highlighting that the fact that the estimates summarized in Table 1.2 are barely affected by sequentially adding fixed effects at various dimensions provides a first indication for the robustness of our main results. If confounding local shocks were important, estimates should vary across these different specifications, which they do not.²³ This pattern suggests that firms' investment plans already incorporate shocks along various dimensions that might simultaneously affect firm investment and the municipalities' decisions to increase the LBT. Hence, focusing on deviations of realized investment from investment plans reported prior to the tax hike should by itself rule out many potential channels of omitted variable bias.

Nevertheless, attributing investment revisions to increases in the LBT could be problematic if tax hikes were accompanied by changes in municipality expenditures. If municipalities re-invested the additional tax revenue in local infrastructure, tax hikes would not only lead to higher tax payments on profits, but could also increase the value of local amenities for

²² The reason why we restrict ourselves to two pre- and post-event periods in Figure 1.3 is sample size. Because we always require that no other tax change happened in the pre- and post-event periods, extending the number of periods would shrink the size of the estimation sample considerably.

²³ Relatedly, Appendix Table A.9 shows that treatment effects are not heterogeneous for firms experiencing large revenue drops (compared to those firms who do not) as a proxy for a local (or even firm-specific) shock. This provides further empirical evidence that local shocks are not driving our results. This is not surprising, given that we analyze manufacturing firms whose products are tradable across municipalities and hence are less reliant on local markets.

1 Downward Revision of Investment Decisions after Tax Hikes

firms. If this created incentives for investment, this would counteract the direct effect of the tax hike and the true investment response would be underestimated. While this scenario is not implausible in general, we cannot detect concurrent expenditure shocks in our data. In line with evidence from Fuest et al. (2018) and Isphording et al. (2021), Appendix Figure A.13 shows that, on average, municipalities do not increase their expenditures jointly with the LBT.

Moreover, our results are robust to excluding the years after the German reunification from our sample. Although we only focus on firms located in West Germany, many of these firms were affected by this particularly turbulent economic time and their investment decisions were potentially affected by many investment subsidies that were introduced with the aim to foster investment in East Germany. Indeed, the estimated effect size is slightly, but not substantially, larger when excluding the time period after the German reunification (Appendix Table A.5).

As a final, more general robustness check, we conduct a permutation test by randomly assigning tax hikes to municipalities and, for each permutation, estimate Model (1) with both dependent variables, the downward revision indicator and the log revision ratio, along with the full set of fixed effects. Appendix Figure A.14 plots the cumulative distribution function of these placebo treatment coefficients. The non-parametric p-values obtained from this exercise are 0.0005 for the downward revision indicator and 0.0115 for the log revision ratio, and thus in the same order of magnitude as in our baseline regression.


1.4.2 Magnitude of Effect Size in Comparison to the Literature

While previous literature focuses on the effect of tax changes on the realized level of investment, a key novelty of our paper is studying the revision of investment plans. In order to facilitate the comparison of our result to previous findings, this section first demonstrates that our results regarding the downward revision of investment decisions can indeed be interpreted in terms of a reduction in realized investment of equal size. In a second step, we convert the identifying variation in the statutory LBT rate to changes in the net-of-tax rate, the effective tax rate, or the user cost of capital to show how our results compare to studies that rely on these frequently used specifications.

Effect of Tax Hikes on Realized Investment. As argued in Section 1.3, we can directly interpret a one percent decrease in the log revision ratio as a one percent decrease in the realized level of investment because firms on average invest as much as previously planned. In order to

demonstrate this empirically, Figure 1.4 plots the coefficients of three event study regressions using either the log revision ratio (red), the log realized investment volume (yellow), or the log planned investment volume (green) as dependent variable. As expected, the point estimates of both realized investment and the revision ratio are of comparable size in t_0 , i.e., the year of the tax hike, and indicate a drop by approximately 3%. Accordingly, investment plans fall to a persistently lower level only one period later in t_1 , when firms have incorporated the tax hike of year t_0 into their information set.

Figure 1.4: Effect of Tax Hike on Investment Plans, Realizations, and Revisions



Notes: This figure shows event-study estimates of log planned investment (green, short dashed lines), log realized investment (orange, long dashed lines), and the log revision ratio (red, solid lines) on the tax hike indicator and fixed effects at the levels of firm identifiers and years. The reference period is $t - 1$. In addition, end-periods $t - 3$ and $t + 3$ are binned and not shown. The sample is trimmed outside the event window. “Log Revision Ratio” is the natural logarithm of the ratio between the ex-post realized and ex-ante planned volume of investment. Inspired by Dube et al. (2023), when estimating the effects with respect to log planned and realized investment, firms are assigned to another firm identifier after the year that is in the middle between two tax hikes in order to ensure that there is only one treatment for each unit and to allow for different long-run trends. The confidence intervals refer to the significance levels of 90% (thick lines) and 95% (thin lines).

Note that Figure 1.4 is estimated on a different sample than our baseline estimates: Because treatment effects on investment levels are persistent and do not return to zero after one year, heterogeneous treatment effects may bias the estimates without further restrictions (e.g., de Chaisemartin and D’Haultfoeuille, 2022). Inspired by Dube et al. (2023), we therefore assign

1 Downward Revision of Investment Decisions after Tax Hikes

firms that experienced several tax hikes into distinct episodes to distinct firm identifiers, such that each spell contains a window around one tax hike only. For each of these spells, we assume that treatment effects have stabilized after three years and trim more distant periods. We further require that no other tax hikes took place during these periods, which—in total—reduces the sample size by 40%. In consequence, the coefficients are estimated less precisely than in our baseline specification. Note that despite these necessary sample restrictions, the coefficient for the log revision ratio in Figure 1.4 is still identified under much weaker assumptions than the coefficients for investment plans and realizations.

Overall, the results suggest that increases in the LBT have lasting effects on firms' investment decisions. While—given that the tax hike is also incorporated into firms' ex ante investment plans for years t_1 and thereafter—the coefficient on the log revision ratio returns to zero in the years following t_0 , the point estimates regarding the level of investment remain negative. Panel B of Appendix Figure A.11 shows that these patterns also hold when extending the time window covered by the estimation to four years before and after the treatment, which, by construction, relies on substantially fewer observations.

Effect Sizes Expressed in Terms of Net-of-Tax Rates. A common quantity of interest in the public finance literature is the elasticity of investment with respect to the net-of-tax rate. To interpret our finding through the lens of this literature, Appendix Table A.7 re-estimates our baseline results regarding the log revision ratio after expressing the variation in the LBT in terms of changes in the net-of-tax rate, defined as $\log(1 - \tau_t) - \log(1 - \tau_{t-1})$. The resulting elasticity of investment with respect to the net-of-tax rate ranges between 2 and 3, depending on the choice of the control vector. In the most restrictive specification, which includes the full set of fixed effects, a 1% increase in the net-of-tax rate increases the log revision ratio and thus investment by 3%.

Effect Sizes Expressed in Terms of Effective Tax Rates. Our main specification estimates the investment response to changes in statutory marginal tax rates, i.e., the LBT parameter that is directly set by municipal policymakers and that hence can be evaluated empirically without imposing further assumptions. However, large parts of the literature estimate treatment effects in relation to changes in effective marginal tax rates, i.e., also accounting for deductions – including depreciation rules – and other exemptions or tax credits (of which there are very few in the German context). To better compare our estimates with these studies, we thus also run such alternative specifications of our baseline estimation.

Our procedure to calculate effective marginal tax rates τ_{eff} , which is described in detail in Appendix A.5, follows the framework of Hall and Jorgenson (1967), as, e.g., recently applied by Furno (2022). Under this framework, the effective marginal tax rate is given by $\tau_{eff} = 1 - ((1 - \tau)/(1 - z * \tau))$, and depends only on the present discounted value (PDV) of the depreciation z and the statutory LBT rate τ .²⁴ To obtain z , we rely on information on depreciation schedules for machinery and buildings obtained from the Oxford Corporate Tax Database. Due to lack of the respective information, expressing changes in the LBT in terms of changes in effective marginal tax rates requires additional assumptions on, *inter alia*, (i) firms' discount rate and (ii) the distribution of the total volume of investment among categories subject to different depreciation schedules, i.e., investment in machinery or buildings. The choice of the adequate discount rate is not innocuous in our setting, given that our analysis covers a period of almost four decades during which interest rates have fluctuated strongly (see Appendix Figure A.17). Moreover, the composition across investment categories can only be roughly approximated by either relying on yearly aggregates of the entire manufacturing sector or using time-invariant shares of firm-level investment in machinery and buildings.

To account for this, we compute four different versions of τ_{eff} based on two sets of assumptions regarding the discount rate and the relative share of investment in machinery and buildings, each. In the first and second specification, we follow Zwick and Mahon (2017) in assuming a time-constant discount rate of 7% when calculating the present discounted value of depreciation, while the remaining specifications use time-varying interest rates on loans for discounting. Further, the first and third specification rely on information on the average share of investment in machinery and buildings obtained from aggregate data from the Federal Statistical Office of Germany, while the others use the firm-specific share of investment in machinery and buildings reported to the ifo Investment Survey whenever available.²⁵ Appendix Figure A.19 shows that across all specifications, the variation captured by changes in effective tax rates is strongly associated with the underlying changes in the LBT rate. This is not surprising, as all tax base rules of the LBT are set at the federal level (and potential changes of those over time are largely absorbed by year fixed effects) and no specific tax credits exist. Hence, apart from its scale, the identifying variation exploited in the empirical estimation does not differ strongly between the different approaches.

²⁴ This simplified version of the formula can be applied in our setting because there are no relevant tax credits in the German LBT that would complicate the calculation.

²⁵ We use the firm-specific mean across all years if firms reported machinery and building investments at least three times to the IVS and replace missing values by the aggregate data used in the first specification, see Appendix A.5.

1 Downward Revision of Investment Decisions after Tax Hikes

Table 1.3: Investment Revisions after a Tax Hike: Effective Tax Rates

	Log Revision Ratio			
	(1)	(2)	(3)	(4)
<i>Panel A: Variation in Effective Tax Rate</i>				
Effective Tax Hike	-0.132 (0.059)	-0.124 (0.057)	-0.136 (0.067)	-0.133 (0.066)
Constant	-0.033 (0.001)	-0.033 (0.001)	-0.033 (0.001)	-0.033 (0.001)
<i>Panel B: Variation in User Cost of Capital</i>				
User Cost Hike	-0.120 (0.054)	-0.107 (0.050)	-0.121 (0.061)	-0.111 (0.057)
Constant	-0.033 (0.001)	-0.033 (0.001)	-0.033 (0.001)	-0.033 (0.001)
Assumptions:				
Interest Rate	0.07	0.07	time-varying	time-varying
Spec. Investment Share	I	II	I	II
Observations	34421	34421	34421	34421
Firm FE	✓	✓	✓	✓
Year × State FE	✓	✓	✓	✓
Year × Industry FE	✓	✓	✓	✓

Notes: This table reports estimates from linear regressions of the log revision ratio on the size of the tax changes. In Panel A, the estimation is based on variation in effective tax rates (τ_{eff}) calculated as described in Appendix A.5. Panel B runs separate regressions exploiting changes in the user cost of capital (multiplied by 100). The respective specifications rely on different assumptions regarding the calculation of the present discounted value of depreciation, either assuming a time constant interest rate of 7% (as, e.g., in Zwick and Mahon, 2017), or based on time-varying interest rates on firm loans as depicted in Appendix Figure A.17. Further, τ_{eff} (and $UserCost$) is either calculated based on the average share of investment in machinery and buildings based on aggregate data from the Federal Statistical Office of Germany (Specification “I”) or on the firm-specific share of investment in machinery and buildings reported to the IVS whenever available (Specification “II”). All regressions apply firm fixed effects, as well as industry-by-year and state-by-year fixed effects. Standard errors in parentheses are clustered at the municipality level.

The results presented in Panel A of Table 1.3 show that the estimated effect sizes are largely comparable across the different specifications of τ_{eff} . For better comparison with our baseline results, the estimates have to be rescaled, as the effective tax rates are on average much smaller than the statutory ones. In the first specification, rescaling takes into account that the average statutory marginal tax rate of 16.79% is considerably larger than the average effective

marginal tax rate (assuming a discount rate of 7%). The size and precision of the estimated coefficient shown in Column (1) is remarkably close to the respective baseline specification using variation in the statutory LBT rate (compare $-0.132 * 3.82/16.79 = -0.030$ vs. -0.038). When relying on time-varying interest rates, e.g., in Column (3), the rescaled point estimate implies a slightly lower effect size, which is, however, still in the same order of magnitude ($-0.136 * 2.9/16.79 = -0.023$).

Panel B of Table 1.3 repeats this exercise using hikes in the tax term of the user cost of capital as explanatory variable, again aiming to produce estimates comparable to other parts of prior literature. Given the linear relationship between the user cost of capital and the effective tax rate depicted in Appendix Figure A.20 ($\tau_{eff} = 1 - UserCost^{-1}$), this approach does not impact the results apart from rescaling the coefficients. Across all specifications, the estimated effects of a hike in the user cost of capital on the log revision ratio are statistically significant and range between -0.107 and -0.121 .

Comparison to the Literature. How do the investment effects documented in our paper compare to findings in other studies? The earlier public finance literature (e.g., surveyed by Hassett and Hubbard, 2002) typically estimated the effect of changes in the tax term of the user cost of capital on investment, measured relative to the lagged capital stock (I/K). For the sake of comparability, the coefficient of -0.12 depicted in Panel (B) of Table 1.3, which can be interpreted as the effect on log investment, as demonstrated above, hence needs to be expressed in terms of I/K . As information on the capital stock is not available in our data, we use the information on I/K documented by Zwick and Mahon (2017) to rescale our estimate. Accordingly, a one unit change in the user cost of capital is associated with a decrease in the ratio of investment over the lagged capital stock by 1.2 percentage points in our setting.²⁶ Table 1.4 summarizes the main (semi)-elasticities reported in this section. The I/K transformation in the last column suggests that the investment response documented in our paper is slightly stronger compared to the estimates summarized in Hassett and Hubbard (2002) that range between -0.5 and -1 , but smaller than the comparatively large estimate of -1.6 found by Zwick and Mahon (2017).²⁷ Given that the investment stimulus studied in the latter paper was an explicitly recessionary policy, our evidence on state-dependence in the effect of tax hikes on firm investment delivers a potential explanation for the larger

²⁶ Mean I/K in Zwick and Mahon (2017) amounts to 0.1. Hence, our estimate can be expressed in terms of I/K as follows: $-0.12/0.1 = -1.2$.

²⁷ Ohn (2018, p. 296) also derives, after rescaling his estimates (Appendix J), an effect size for the same policy that is remarkably close to Zwick and Mahon (2017).

Table 1.4: Summary of Estimated (Semi-)Elasticities

	Log Revision Ratio	I/K
Change in Tax Rate	-0.038 (0.016)	-
Change in Net-of-Tax Rate (in %)	3.032 (1.312)	-
Change in Effective Tax Rate	-0.133 (0.066)	-
Change in User Cost of Capital	-0.111 (0.057)	-1.11 (0.57)

Notes: This table summarizes the main (semi)-elasticities reported in Section 1.4.2, relating the change in corporate tax rates to changes in the log revision ratio, and, as explained in the main text, to changes in realized investment. We always report the results of the most restrictive specification that controls for firm fixed effects as well as industry-year and state-year fixed effects. The I/K estimate is not directly estimated in our data, but obtained by rescaling the estimate for the log revision ratio with mean I/K of 0.1 reported by Zwick and Mahon (2017).

effect size found in their setting. Zwick and Mahon (2017) also find that loss-making firms are less responsive to tax shocks, in line with our evidence on smaller treatment effects among firms facing large revenue drops documented below (Appendix Table A.10). Furthermore, our semi-elasticity of 3 can be compared to Ohrn (2018) who reports a semi-elasticity of 4.7 for a tax decrease induced by a specific tax provision for the manufacturing sector. Besides differences in research design and institutional setting, these larger responses documented in more recent studies might be due to targeted policies being more effective at stimulating investment than statutory tax rate cuts as rationalized by Chen et al. (2022).

Two recent studies also estimate how firms respond to changes in universal corporate tax rates. Giroud and Rauh (2019) study how firm-level variables react to changes in US state level corporate taxes, including changes in the capital stock. They focus on a selected sample of large multi-state firms, for which they find a semi-elasticity of 0.24. As pointed out by the authors, this small elasticity might arise due to measurement error. Furthermore, the elasticity of capital can be temporarily lower than the elasticity of investment due to adjustment costs along the transition path in the firms' dynamic response. Mertens and Ravn (2013) use aggregate data and combine a narrative approach with a structural VAR model to exploit changes in US federal corporate taxes. They find semi-elasticities between 2.1 and 4, comparable to the responses documented in this paper.

In general, the elasticities reported in our paper are of course specific to the institutional setting of corporate taxation in Germany and the sample of firms affected. For instance, the elasticities we estimate for the manufacturing sector might differ from elasticities in other sectors, as manufacturing is more capital intensive than the overall economy. For the US, Cloyne et al. (2023) argues that tax cuts stimulate real investment activity mainly in the manufacturing sector, whereas firms in the services sector are more prone to adjust dividends, instead.

Finally, our results complement evidence from Isphording et al. (2021), showing that tax hikes in the German LBT reduce plant-level R&D spending by around 2 to 3 percent in the year of implementation. As R&D spending constitutes a (small) part of firm investment, we can directly compare the estimate to our semi-elasticity of investment of around 3. While both estimates suggest comparable effect sizes, our results are obtained for a different sample of firms and under less restrictive identifying assumptions.


1.4.3 State Dependence and Heterogeneity

State Dependence. Next, we exploit the long time dimension of our data to analyze potential heterogeneity in effect sizes over the business cycle. While a large literature in macroeconomics studies the state dependence of fiscal policy, there is not yet a consensus on whether effects of corporate tax changes are state dependent (Demirel, 2021; Hayo and Mierzwa, 2021; Jones et al., 2015; Ljungqvist and Smolyansky, 2018; Winberry, 2021). As most quasi-experimental evaluations of the effect of corporate taxes on investment behavior rely on few tax changes or just a single tax reform, the treatment variation is typically not large enough to distinguish effect size heterogeneity along the business cycle. In contrast, the long time dimension of our data in combination with the occurrence of multiple local tax changes in each given year allows us to evaluate whether the treatment effect is state dependent.²⁸

The effect of tax hikes on revisions of investment plans are substantially stronger during recessions compared to normal times. Figure 1.5 presents the estimation results of interacting the tax hike treatment with indicators capturing periods of recession and normal times. To this end, Panel A classifies t_0 as a recession year if at least one quarter of that year is defined as a recession by the German Council of Economic Experts. The average effect that we estimated

²⁸ As shown in Appendix Figure A.3, municipalities are as likely to raise taxes in recessions as in normal times. The reasons why municipalities increase taxes (also in recession) are diverse, ranging from growing budget requirements to electoral cycles (Foremny and Riedel, 2014) and rent extraction (Langenmayr and Simmler, 2021); see also the discussion in Section 1.2.

Figure 1.5: Investment Revisions after a Tax Hike: State Dependence



Notes: This figure estimates how the probability of investing less than previously planned (left panel) or the log revision ratio (right panel) change in response to a tax hike separately for recession and non-recession years by including respective interaction terms in Equation (1). In Panel A, recession years are defined following the classification of the German Council of Economic Experts and refer to 1980-1982, 1992-1993, 2001-2003, and 2008-2009. Panel B classifies recessions as years with negative real GDP growth according to World Bank data (<https://data.worldbank.org/indicator/NY.GDP.MKTP.KD.ZG?locations=DE>), resulting in a smaller set of recession years (1982, 1993, 2002, 2003, and 2009). The estimation purges for firm fixed effects, as well as year fixed effects at the levels of federal states and industries. Standard errors are clustered at the municipality level. Confidence intervals refer to the 90% level. The full regression output is disclosed in the even columns of Appendix Table A.6.

in Table 1.2 masks substantial heterogeneity over the business cycle. For instance, while in normal times the share of firms that invest less than previously planned increases by 2 percentage points in years with a tax hike, this figure triples to 6 to 7 percentage points in recessions. The same pattern also holds for the remaining specifications and the results tend to become even stronger when using a narrower classification of recession periods, defined as years with negative real GDP growth in Panel B.²⁹

²⁹ Despite the fact that the effects during the recession period can only be estimated relatively imprecisely due to the small sample size, the estimated effects during recessions are statistically different from those during expansions in half of the specifications, while being close to approaching significance in the remaining specifications (see Appendix Table A.6).

Mechanisms of State Dependence. While our baseline estimates are in line with the predictions of standard theories of investment (Hall and Jorgenson, 1967), theory fails to explain why the effect of tax hikes should be state dependent. In the following, we discuss three channels which could explain the stronger effect during recessions.

The first channel relates to the fact that investment projects are risky. As investments are only partially deductible from the tax base, profits and losses are treated unequally by the tax authorities.³⁰ In expectation, tax hikes thus lead to stronger decreases in the net present value of those investment projects with a higher variance of expected returns as first formalized by Domar and Musgrave (1944).³¹ During recessions, the expected return to many investment projects becomes more uncertain, as it is unknown when the economy will recover again. Tax hikes should therefore lead to stronger behavioral responses in economic downturns when a higher share of planned investments is risky. While we cannot test this conclusively in our data, we can assess whether firms with more volatile revenue paths react stronger to tax hikes. For this purpose, we calculate the standard deviation of yearly revenue growth for each firm and construct an indicator for having above median volatility. Appendix Table A.8 shows the regression results when the tax hike effect is interacted with this volatility indicator. While the effects are estimated imprecisely and are sensitive to the specified model, they indeed show slightly larger responses of firms with more volatile revenue paths, suggesting that one reason for the state dependence of tax shocks may be the heightened uncertainty about returns to investment during recessions.

Second, firm investment is sensitive to cashflow (Almeida et al., 2004). Corporate taxes decrease the cashflow for profitable firms and therefore lower investment. At the same time, Almeida et al. (2004) show that cashflow sensitivity is higher in recessions. During recessions, firms expect a higher probability of being cash constrained in the future and therefore retain more earnings for profitable investment opportunities. Taken together, this could give rise to an interaction effect, which reduces investment disproportionately if taxes are increased during recessions. Two regularities in our data support such a mechanism. First, we find that profitable firms react stronger to tax hikes during recessions. We use an indicator for a revenue drop by more than 10 percent compared to the previous year as a proxy for no

³⁰ As discussed by Fuest et al. (2018), costs of debt financing are usually fully deductible from the LBT, while costs of equity financing are not and loss offset is restricted. Moreover, due to depreciation rules, investment costs are split over several years while the revenues are fully taxed in each year.

³¹ While Domar and Musgrave (1944) refer to the personal income tax, the same logic applies to the corporate tax and has been tested in the data. For state-level corporate tax rates in the US, Ljungqvist et al. (2017) show that in response to a tax increase the average firm reduces risk as measured by their earnings volatility. Langenmayr and Lester (2017) find similar results in a cross-country panel and among small Spanish firms.

1 Downward Revision of Investment Decisions after Tax Hikes


longer being profitable. While firms that experience a large revenue drop in general revise investments downwards, the revision effect after a tax hike is smaller compared to firms without a large revenue drop during a recession (Appendix Table A.9). Firms with a large decline in revenues might still be profitable if they reduce their labor costs significantly. Appendix Table A.10 shows that the results hold when we exclude firms with a reduction in the number of employees by more than 5 percent as a robustness check. Second, if an adverse financing situation is reported to be a factor for a strong slowdown in investment volumes, the revision effect tends again to be larger in recessions (Appendix Table A.11). Both findings provide suggestive evidence that the stronger investment response in recessions may relate to cashflow sensitivity.

Finally, the stronger investment response to tax hikes in recessions could result from diminished possibilities to shift the tax burden to third parties. Fuest et al. (2018) show that workers bear approximately half of the incidence of the LBT in Germany. However, as wages are nominal downward rigid (e.g., Barattieri et al., 2014), firms often cannot decrease wages in response to adverse economic conditions. This lower bound bites predominantly in recessions, especially given that collective bargaining agreements are still the norm in the German manufacturing sector and bargained wages usually slow down only with a considerable time lag as depicted in Appendix Figure A.15. This could suggest that during recessions, (cashflow sensitive) firms reduce their investment disproportionately as downward rigid wages do not allow shifting the tax burden on workers. Consistent with such a channel, Fuest et al. (2018) report a lower wage incidence for less profitable firms. In our data, we do not observe wages, preventing us from investigating this issue further.

Testing for Further Heterogeneity. While the sample size of our data does not permit a comprehensive heterogeneity analysis, we perform the main estimation for a number of additional sample splits emphasized in the literature. For example, the investment effects of accelerated depreciation allowances in the corporate tax code have often been found to be much stronger among small (liquidity-constrained) firms (e.g., Zwick and Mahon, 2017). Figure 1.6 summarizes the results, showing that treatment effects do not differ by firm size when splitting our sample into firms with more or less than 250 employees (Panel A).³² We also fail

³² Given that firm size is often been used as a measure of cashflow sensitivity, this finding might seem inconsistent with the aforementioned suggestive evidence that the stronger investment response in recessions may relate to cashflow sensitivity. There are at least two possible explanations for this discrepancy. First, the share of very small manufacturing firms, which arguably suffer most from adverse financing conditions, in our sample is small (Appendix Table A.2). Second, Appendix Figure A.16 demonstrates that small firms are not only more likely to be

Figure 1.6: Testing for Further Heterogeneity



Notes: This figure estimates how the probability of investing less than previously planned (left panel) or the log revision ratio (right panel) change in response to a tax hike separately for different groups of firms by including respective interaction terms in Equation (1). Panel A provides separate estimates for small and large firms (split at the threshold of 250 employees). Panel B sorts firms according to their location using the definition of urban and rural areas of the Federal Institute for Research on Building, Urban Affairs, and Spatial Development (BBSR) that is mainly based on population density. In Panel C, we split the sample into municipalities with few (≤ 3) and many (> 3) tax hikes over the entire sample period. In Panel D, the tax hike treatment is split into cases where at least one tax hike has already occurred in the previous five years and where no tax hike occurred in the previous five years in the respective municipality. The estimation purges for firm fixed effects, as well as year fixed effects at the levels of federal states and industries. Standard errors are clustered at the municipality level. Confidence intervals refer to the 90% level. The full regression output is disclosed in the even columns of Appendix Tables A.12 and A.13.

to detect significant differences between rural and urban municipalities (Panel B), suggesting that a downward bias in our estimates due to private information of the firms—which should be more relevant in rural municipalities—is not a major concern.

cashflow sensitive, but also more likely to underinvest due to a weak earnings situation. As the LBT taxes profits and less profitable firms should thus react less sensitively to tax hikes, the latter mechanism could counteract the former.

1 Downward Revision of Investment Decisions after Tax Hikes

In Panel C, we split the sample into municipalities with few (≤ 3) and many (> 3) tax hikes over the entire sample period. If firms in municipalities, which often increase the LBT, expected a tax hike with a higher probability, then downward revisions of investment should be less likely among these firms. However, effect sizes are again very similar for both groups. An alternative way to investigate whether tax setting dynamics at the municipality level correlate with the effect sizes is to split the sample by the occurrence of a tax hike in the last years. The results depicted in Panel D suggest that having experienced a tax hike in the last five years is plausibly associated with a larger investment response, although the estimates are not statistically different from each other in any specification (Appendix Table A.13). This result could be consistent with the notion that higher policy uncertainty triggers a stronger response after tax hikes.

Overall, Figure 1.6 demonstrates that other than the strong effect heterogeneity with respect to the business cycle, the effect of tax hikes on the investment behavior of firms is rather homogeneous across other important partitions of our data.

1.5 Conclusion

This paper provides novel empirical evidence on the effect of corporate taxation on firm investment. Our research design allows us to address several concerns that often complicate identification of an investment response. By considering 1,443 tax changes of the German local business tax between 1980 and 2018, we draw on extensive treatment variation and average out idiosyncratic characteristics of single tax reforms. By observing both planned and realized investment volumes, we can control for ex ante investment plans when estimating the effect of tax hikes on firm investment, eliminating a wide set of further potentially confounding factors.

We find significant and economically large investment responses for firms experiencing a tax shock. The share of firms that invest less than previously planned increases by 3 percentage points after a tax hike, with strong heterogeneity along the business cycle. While in normal times the share of firms that revise their investment decisions downwards increases by 2 percentage points in response to a tax hike, this figure triples to over 6 percentage points if taxes are increased during a recession. These findings have direct policy implications that support the countercyclical Keynesian notion of “do not increase taxes during recessions”. While we

find suggestive evidence that the state dependence of tax shocks could plausibly be related to uncertainty about expected returns to investments, cashflow sensitivity, and tax incidence, more research is needed to disentangle the channels behind this finding.

Overall, our results confirm the view that investment decreases substantially in the corporate tax burden. While our estimates were obtained for increases in the statutory corporate tax rate, prior studies have often evaluated targeted tax policies which were deliberately designed to stimulate investment. We look forward to future research comparing the effects of both types of policies within a unified framework.

2 Social Mobility in Germany

This chapter is based on co-authored work with Majed Dodin, Sebastian Findeisen, Lukas Henkel, and Dominik Sachs, and has been published in the Journal of Public Economics. See Dodin et al. (2024) for the full reference.

2.1 Introduction

Intergenerational social mobility is an important indicator for both fairness and economic efficiency in a society. Next to violating widely held fairness ideals, low mobility can lead to the misallocation of resources, as talented children from disadvantaged backgrounds are impeded from realizing their potential. Despite its importance, reliable mobility statistics are not available for many countries. Measuring social mobility across generations is challenging, as it requires data that links parental outcomes to a measure of opportunities for children.¹ Household panel studies may contain this information but are typically too small to deliver sufficiently precise estimates for regional comparisons or the analysis of time trends (Lee and Solon, 2009; Mazumder, 2018). An attractive alternative are administrative data sources, such as linked tax records (e.g. Chetty et al., 2014a). As in many other countries, however, such data is not available in Germany, where to date no large-scale empirical study of social mobility across time and space exists.

In order to fill this gap, this paper implements a new measurement strategy for social mobility in Germany and provides estimates across time and regions. Motivated by Germany's early tracking system in secondary education, our mobility statistics measure the association between parental income and the educational opportunities of children. Our measure of opportunities captures whether a child will obtain the A-Level (Abitur), the highest secondary schooling degree in Germany. We are able to link 526,000 children to their parents, using census data spanning the years from 1997 to 2018.

Our first finding is that relative mobility, defined as the percentage point difference in the probability to obtain an A-level degree between children with different parental income

¹ In the literature, the expression social mobility refers to inter-generational social mobility and in other cases also to intra-generational mobility (i.e. social mobility between different periods of a lifetime). In this paper, we focus on the relationship across generations.

2 Social Mobility in Germany

ranks, has remained constant for recent birth cohorts. On average, a 10 percentile increase in parental income rank was associated with a 5.2 percentage point increase in the probability of obtaining an A-Level degree. For the birth cohorts 1980-1996, this parental income gradient has not changed despite a large-scale expansion of upper secondary education in Germany, the *Bildungsexpansion*. This long-term expansion was in parts a policy response to a public debate on social mobility (Dahrendorf, 1966; Hadjar and Becker, 2006) and increased the A-level share from 39% for children born in 1980 to 53% for the 1996 birth cohort. We document that the *Bildungsexpansion* took place uniformly across the income distribution, with almost identical increases in the share of A-Level educated children in all quintiles of the parental income distribution. This enhanced the odds ratio for disadvantaged children, but left the parental income gradient unaffected. The same pattern emerges when estimating mobility trends for population subgroups typically emphasized in social mobility policies, such as children in single parent households or children of parents with low levels of formal education. Complementing our main analysis with data on test scores and grades, we find no evidence that measured ability was better for marginal students from disadvantaged backgrounds than for marginal students from affluent households. We therefore cannot draw a positive conclusion about the *Bildungsexpansion* in the sense that it revealed more hidden talent among children at the bottom of the income distribution than among those at the top.

We also document geographic variation in social mobility across German states, cities, and local labor markets. For example, the top-bottom gap in the probability of obtaining an A-Level degree between children at the top and the bottom of the income distribution is 20 percentage points larger in Bremen than in Hamburg, two city states approximately 100 kilometers apart. We also find significant and meaningful differences within states. For example, the top-bottom gap is 8 percentage points larger in Cologne than in Düsseldorf, two large cities in North Rhine-Westphalia located approximately 40 kilometers apart. Overall, the within-state component of the variance in the parental income gradient across local labor markets or cities is around six times higher than the between-component. This is remarkable, as education policies, which prior literature has suspected to be a key determinant of mobility, vary mainly at the state level in Germany.²

We show that household characteristics can explain only a small fraction of the variation in mobility measures across local labor markets. Differences in mobility estimates can arise either due to structural differences between places or due to systematic sorting of households

² Helbig and Nikolai (2015) provide a comprehensive account of state level school reforms in Germany since 1949. Studies trying to evaluate their effects on social mobility include Betthäuser (2017), Büchler (2016), and Jähnen and Helbig (2015).

into different local labor markets (Chetty et al., 2014a). The census data employed in this paper contains rich information on the structure and characteristics of households, allowing us to directly test the importance of sorting by conditioning on an extensive set of household characteristics. We find that the mobility ranking between local labor markets is largely unchanged when conditioning on household characteristics, leading us to reject the hypothesis that sorting is the major driver of the regional variation in mobility.

Our paper is the first to provide a comprehensive account of social mobility in Germany, characterizing its evolution over time, heterogeneity across regions, estimates for many subgroups, and disentangling sorting versus place effects. Due to its early-age tracking system, Germany is particularly suited for studying social mobility through the lens of educational opportunities. Only completion of the highest track grants the A-Level degree and thus direct access to the tuition-free national university system, opening up the full range of career prospects. As a result, the A-Level wage premium amounts to more than 40%. Besides the economic benefits, having obtained an A-Level is also an important sign of social distinction in the German society. More broadly, a large literature shows that educational attainment has intrinsic value and predicts a wide range of non-pecuniary outcomes (Oreopoulos and Salvanes, 2011; Lochner, 2011). Educational attainment as a measure of opportunity is thus a strong and comprehensive indicator for the opportunities of an individual in the German context. Beyond Germany, this approach to measure mobility may also prove useful in other countries where the highest secondary school degree plays a similarly important role in shaping future career options.

The remainder of this paper is organized as follows. Section 2.2 discusses the related literature and relevant aspects of the German institutional framework. In Section 2.3, we describe data and measurement strategy. Section 2.4 reports our results at the national level. Regional estimates, including the analysis of local labor markets, are presented in Section 2.5. Section 2.6 concludes.

2.2 Related Literature and Institutional Background

2.2.1 Related Literature

The study of intergenerational social mobility has a long tradition in economics, sociology and educational research. While early sociological studies focused on measuring occupational transitions between generations, educational research studied intergenerational correlations

2 Social Mobility in Germany

in educational attainment. The literature in economics has traditionally measured social mobility by the intergenerational elasticity of (lifetime) earnings, or, more recently, by rank-rank correlations in lifetime income, making use of linked administrative tax data (e.g. Chetty et al., 2014a; Acciari et al., 2022; Corak, 2020).

In Germany, it is not possible to link individual tax returns. For that reason, most empirical evidence on income mobility is based on the German Socio-Economic Panel (SOEP), a household survey with limited sample size. Time trends or more fine-grained geographic variation in social mobility hence cannot be documented in the SOEP with a sufficient degree of statistical confidence. Schnitzlein (2016) shows that estimates of the national IGE based on the SOEP are sensitive to small variations in sampling criteria, resulting in a wide range of plausible estimates. It is therefore not surprising that the empirical evidence regarding the level of social mobility in Germany is not coherent. Studies that investigate intergenerational income mobility in the SOEP include Eisenhauer and Pfeiffer (2008), Riphahn and Heineck (2009), Eberhardter (2013) and Bratberg et al. (2017). These studies typically find higher levels of income mobility in Germany than in the US, and lower levels of mobility in East than in West Germany, albeit with high statistical uncertainty. On the other hand, sibling correlations (Schnitzlein, 2014) or measures of educational mobility have placed Germany closer to the immobile end of the scale in an international comparison.

Our measurement approach focuses on children's educational opportunities, while retaining the interpretability advantages of income based measures of parental socioeconomic status. This allows us to draw on the German census data, providing us with the statistical power necessary to conduct a more comprehensive study of social mobility in Germany.³ At the same time, we can document social mobility for very recent cohorts, because – unlike lifetime income – the A-Level degree can be measured already relatively early in the lifecycle. An additional advantage of our measurement approach is that it works great even in the presence of non-labor force participation or zero earnings in the child generation. Therefore, while much of the intergenerational mobility literature focuses on men, our method is well suited at including women.

Hilger (2015) employed a comparable approach for the US, examining mobility statistics based on census data linking children's years of schooling to parental income. Unlike our study, their focus on later-life outcomes raises sample selection concerns, requiring an imputation proce-

³ A less comprehensive version of the German Census data has previously been used to document differences in the intergenerational correlation in educational attainment between East and West Germany (Riphahn and Trübswetter, 2013; Klein et al., 2019).

due to most children leaving the parental household. Emphasizing years of schooling is justified in the US, where almost all children attend academic high school programs. In contrast, the German system's academic and vocational tracks make it ideal for our outlined census-based social mobility analysis.

2.2.2 Institutional Background

The salient feature of Germany's system of secondary education is early age tracking, where only the successful completion of the highest track results in the award of an A-Level degree (Abitur) and grants direct access to the tuition-free national university system. After finishing the four-year⁴ elementary school around the age of 10, children are allocated into one of three tracks. While the highest track, the Gymnasium (grades 5-12/13), provides general academic education that aims to prepare children for college, the lower two tracks (grades 5-9/10) provide vocational training with a focus on preparing students for an apprenticeship.⁵ The specific design of the tracking system in secondary education can vary across the 16 federal states, which bear the main responsibility for the education system. However, there exist only minor differences in state-provided financing. In addition, the Standing Conference of State Education Secretaries has the stated goal to ensure a high degree of comparability of educational qualifications across German states and there are no legal differences between the A-Level degrees issued from different states.

Since the early educational careers of children have important consequences for the choices available to them at later stages, and early track choices are heavily influenced by parental characteristics (Dustmann, 2004), the German institutional framework is particularly suited for studying social mobility through the lens of educational opportunities. The importance of track choices for social mobility is reinforced by the fact that almost all primary and secondary schools as well as universities are state-funded, mostly based on student headcounts, resulting in a comparatively large equality in the endowments and quality between different schools and universities.

Consequently, the A-Level degree is by far the most important qualification in the German education system, and individuals who obtain it enjoy substantially above-average economic

⁴ In the states of Berlin and Brandenburg, elementary school lasts six years.

⁵ The rigor of the tracking system is mediated by the possibility of switching tracks. In particular, it is common that talented students from the medium track switch to the general high track or attend a specialized high track after they finish their vocational degree when they are around 16 years old. A more detailed overview of the tracking system and track switching in Germany is provided in Biewen and Tapalaga (2017) and Dustmann et al. (2017).

outcomes. Using data on full-time workers aged 30-45, we find an A-Level wage premium of 42% for monthly net income.⁶ This estimate mirrors Schmillen and Stüber (2014) who report a 44% A-Level wage premium for total gross lifetime earnings. An A-Level degree is also associated with a lower risk of being unemployed (Hausner et al., 2015) and a higher life expectancy (Gärtner, 2002). Furthermore, it constitutes a beneficial factor for obtaining vocational training in certain white-collar occupations (Klein et al., 2019) and marks an important sign of social distinction in the German society. Overall, this illustrates that, for children in Germany, the A-Level degree is a compelling measure of their social and economic opportunities.

2.3 Data and Measurement Strategy

Our analysis is based on data of the German Microcensus (Mikrozensus, hereafter MZ), a large-scale annual representative survey of the German population administered by the Federal Statistical Office of Germany (FDZ der Statistischen Ämter des Bundes und der Länder, 1997-2018). The survey was first administered in West Germany in 1957 and includes East Germany since 1991. The MZ has several features that make it particularly suited for our research question. First, it allows us to reliably match children to their parents as long as they are still registered at their parents' household. By law, it is compulsory for individuals living in Germany to register at their place of residence, and the sampled households are obliged to provide information on each person registered in their household. Second, it contains fine-grained geographic information and is sufficiently large to permit the estimation of mobility statistics for single cohorts and regions.

Each year, a randomly selected 1% sample of the population living in Germany is asked to participate in the survey. By law, participation is mandatory for members of the selected households, which remain in the survey for at most four subsequent years. The primary sampling units consist of clusters of neighboring buildings, and all households belonging to a sampled cluster are interviewed. The unit non-response rate is approximately 3%.⁷ Each year, one quarter of the initially sampled clusters are replaced by new clusters, resulting in partial overlap of sampling units. Appendix B.1 contains additional information on the survey and

⁶ We use the waves 1997-2018 of the German Microcensus (described below) and compute the A-Level wage premium by regressing the log of net monthly personal income of full-time workers aged 30-45 on an A-Level dummy, controlling for a full set of age and year fixed effects to implicitly account for job experience.

⁷ The non-response rate is driven by households that could not be reached and residents in shared accommodations (Statistisches Bundesamt, 2018), which we exclude from our sample. The item non-response rate in our sample for the survey questions that we utilize is typically below 1%.

sampling design of the MZ. The detailed nature of the questionnaire together with the low non-response rate and the large sample size allow us to mitigate measurement and sample selection concerns often brought forward in the context of survey data.

2.3.1 Variable Definition

Measuring Opportunities of Children. Motivated by the importance of the A-Level degree for children's future educational and labor market opportunities in the German institutional framework, we measure opportunities by a binary variable Y_i that is equal to one if a child has obtained, or is on track to obtain, a degree that is equivalent to an A-Level, and zero otherwise. Specifically, our outcome variable is equal to one if (i) a child has obtained a degree that qualifies for tertiary education⁸ or if (ii) a child is enrolled in the last 2-3 years of a track which leads to such a degree at the successful completion of school.⁹ In the following, we refer to this outcome as an A-Level degree and characterize intergenerational mobility in terms of the conditional probabilities of obtaining an A-Level degree for children of different parental backgrounds.

Our outcome definition takes into account three considerations. First, while the MZ survey is conducted on a rolling basis, A-Level degrees are typically awarded in the second quarter of the calendar year. Back of the envelope calculations suggest that, if we only count children who have already obtained an A-Level degree, we would miss-measure our outcome for around 40% of the graduating cohort in each survey year. Second, since the share of children failing the final examination in a given year is low¹⁰, including upper stage students allows us to capture children that can reasonably be expected to obtain an A-Level degree but rotate out of the survey before they do so. Finally, including younger children disproportionately increases sample size, as younger children are more likely to live with their parents. Table 2.1 displays the share of children living with at least one parent by age of the child, calculated from our data. Virtually all children younger than 15 still co-reside with at least one parent. However, the share of children living with their parents is decreasing with child age, especially after the

⁸ We classify educational qualifications as equivalent to an A-Level if they grant access to the tuition-free national university system. This includes *Allgemeine Hochschulreife (Abitur)*, *Fachgebundene Hochschulreife* and *Fachhochschulreife*.

⁹ The MZ data contains information on the type of school and grade level attended by all sampled children. Our definition subsumes all students on *Allgemeinbildende Schulen* enrolled in the *Gymnasiale Oberstufe* as well as students from specialized tracks like *Berufliches Gymnasium* or *Fachoberschule* which award an A-Level degree.

¹⁰ The national average failure rate is approximately 3 percent on average for the years 2010-2020. For an overview of the share of children failing the final examination see <https://www.kmk.org/dokumentation-statistik/statistik/schulstatistik/abiturnoten.html>

Table 2.1: Co-Residence Rate by Child Age

Child Age	15	16	17	18	19	20	21	22	23
Share Living with Parents	0.99	0.98	0.97	0.92	0.84	0.72	0.62	0.52	0.44

Notes: This table reports the fraction of individuals which live in the same household as at least one of their parents in the MZ waves 1997 to 2018 by age at observation.

legal age of 18. While 92% of the 18 year olds are living with at least one of their parents, this fraction drops to 44% for individuals at the age of 23. In Section 2.3.3, we discuss how the co-residency and move-out patterns observed in the MZ data affect the interpretation of our results.

Measuring Parental Background. We measure parental background by a household’s self-reported monthly net income, excluding the income of all dependent children. Our income measure covers all sources of income, including labor income, business profits and social security transfers. To account for differences in need and standard of living by household composition, we scale all household incomes by the modified OECD equivalence scale.¹¹ We then compute the households’ percentile ranks in the sample distribution of equivalized household income, and assign each child the rank of their respective household, which we refer to as the parental income rank R_i .¹²

Parental income ranks are a conceptually attractive measure of family circumstances, as the relevance of financial resources and costly enrichment activities for different aspects of child development is widely recognized and there exists empirical evidence of significant disparities in child-related expenditures across the income distribution in Germany. Table 2.2 reports estimates of monthly child-related expenditures in different categories based on data of the 2018 Income and Consumption Survey (EVS) for dual parent households with single children in the top and bottom decile of the national income distribution. The estimates reveal substantial gaps in monthly expenditures on child-enrichment activities in categories

¹¹ Figure B.4 demonstrates that the choice of the scaling factor is not influential for our results at the aggregate level. However, the empirical conditional expectation function of our A-Level indicator can be better approximated linearly when computing ranks based on equivalized incomes.

¹² In Appendix B.1 we provide information on the sample income distributions and details on the construction of the rank variable.

such as education, health as well as culture and leisure activities, suggesting that parental income ranks are a suitable measure of parental background for the construction of mobility statistics in Germany.

Table 2.2: Monthly Child-related Expenditures of Single Child Households

Category	Total	Education	Health	Food	Culture	Mobility	Other
Top Decile	1212	83	113	156	205	85	244
Bottom Decile	424	28	11	104	47	29	65
Ratio	2.85	2.96	10.27	1.5	4.36	2.93	3.75

Notes: This table reports estimates of the monthly child-related expenditures in Euro of dual parent, single child households in the top and bottom decile of the German national income distribution for different expenditure categories. The data is reported in the 2018 Income and Consumption Survey (EVS) of the Federal Statistical Agency (Statistisches Bundesamt, 2021).

The continuous measure of household income provided in the MZ data that we use to compute parental income ranks is not asked for directly in the survey but imputed by the Statistical Office. The survey respondents report their personal income in 24 predefined bins. The Statistical Office then transforms the personal binned income into a continuous variable, essentially randomizing individuals uniformly within each bin. In a second step, these values are summed up to a continuous measure of household income. We discuss potential implications of this procedure for the external validity of our mobility statistics in Section 2.3.3.

2.3.2 Mobility Statistics

The central building block of all mobility statistics reported in this paper are estimates of the probability of children attaining an A-Level degree conditional on parental income rank $E[Y_i|R_i]$. Following the recent literature, we define two sets of mobility statistics with the aim of distinguishing between two mobility concepts: absolute and relative mobility. While measures of absolute mobility are informative about the level of opportunities for disadvantaged children, relative mobility measures seek to capture differences in opportunities between children of disadvantaged backgrounds relative to those of more advantaged backgrounds.

Absolute Mobility. We measure absolute mobility by the probability of obtaining an A-Level degree for a child from the bottom quintile of the parental income distribution:

$$Q1 = E(Y_i | R_i \leq 20). \quad (2.1)$$

2 Social Mobility in Germany

We refer to this estimand as the Q1 measure. A high value of the Q1 measure implies high absolute mobility, as it indicates that a large share of disadvantaged children are eligible to enter the university system.

Relative Mobility. We define two measures of relative mobility, both concerned with the difference in opportunities between children from low and high-income families. A simple non-parametric measure of relative mobility is the Q5/Q1 ratio:

$$Q5/Q1 = \frac{E(Y_i | R_i > 80)}{E(Y_i | R_i \leq 20)}, \quad (2.2)$$

which captures the odds ratio of obtaining an A-Level degree for children from the top quintile relative to those in the bottom quintile of the parental income distribution. A high value of the Q5/Q1 ratio implies low relative mobility. For example, a ratio of $Q5/Q1 = 2$ means that children from the top quintile of the income distribution are twice as likely to obtain an A-Level degree as children from the bottom quintile of the income distribution.

Next to the Q5/Q1 ratio, we also estimate a parametric statistic of relative mobility. As demonstrated in the results section of this paper, the empirical conditional expectation function, $E[\widehat{Y_i} | R_i]$, of our outcome given the parental income rank is close to linear in various partitions of our data. As a consequence, we can approximate the respective conditional expectation function (CEF) by its best linear predictor, which is defined as

$$\theta_{LP} = \arg \min_{\theta} E[(Y_i - Z_i' \theta)^2],$$

with $Z_i = (1, R_i)'$ and $\theta = (\alpha, \beta)$. In practice, we estimate the model parameters by running an OLS regression of our outcome indicator on the parental income rank variable. The slope coefficient β measures the gap in the probability of obtaining an A-Level between children at the top and the bottom of the income distribution. We refer to the slope coefficient as the *parental income gradient* and report estimates of $\beta \times 100$, which captures the gap in percentage points, for improved readability. While the Q5/Q1 ratio measures the relative outcome difference between children at the top and the bottom of the income distribution, the parental income gradient characterizes the absolute outcome difference and is therefore not sensitive to the baseline probability of obtaining an A-Level in the underlying population of interest.

Note that both of our measures of relative mobility are relative only in the sense that parental income is measured in ranks, whereas opportunities of children are measured with the A-Level degree, which is an absolute, rather than relative outcome.

2.3.3 Sample Definition and Limitations

We use the MZ survey waves from 1997 to 2018, for which a consistent definition of all relevant variables is available. For our national and regional estimates, we restrict our sample to the survey waves 2011-2018 (231,000 children) to produce recent mobility statistics and avoid ambiguities caused by a series of administrative reforms that changed county boundaries. The mobility statistics by birth cohort reported in Section 2.4.2 are computed based on the 1980-1996 birth cohorts (526,000 children).

Our primary sample contains all children aged 17-21 which are observed in the same single-family household as at least one of their parents. The age range is chosen to balance the following trade-off: For older children, our outcome is measured more precisely, i.e. we do not need to rely on upper-stage enrollment but are more likely to observe the completed degree. At the same time, the fraction of children in our sample that has already moved out of the parental household, and thus can not be matched to their parents, increases with age, which guides our choice for the upper bound. The lower bound is chosen as children enrolled in the upper stage of an A-Level track are typically at least 17 years old. In the following, we discuss potential concerns regarding the external validity of our mobility estimates.

Sample Selection. An immediate concern caused by the observed move-out patterns in the MZ data relates to the representativeness of our sample. If the observed move-out decisions were systematically related to both parental income and the educational attainment of children, the external validity of our estimates would be undermined as our statistics would not measure social mobility in the population of interest. While we acknowledge that dependencies of this type are generally plausible, we do not find evidence of sample selection in our data. Table 2.3 documents how time-constant characteristics of the children in our sample change with the age at observation. If move-out were to occur randomly, we should not see systematic changes in these statistics for older children for which the co-residency rate is lower. While move-out varies with social characteristics like gender, the average parental income and the associated income rank of children in the age range 17-21 are essentially constant. In addition, we can exploit the partial panel dimension of the MZ to investigate selection


2 Social Mobility in Germany

Table 2.3: Average Characteristics of Children by Age of Observation

Child Age	Share Female	Mean Parental Inc. (Equiv.)	Parental Inc. Rank	Share Parents with A-Level
17	0.49	1367	50	0.33
18	0.48	1367	50	0.32
19	0.47	1367	50	0.32
20	0.44	1359	50	0.31
21	0.42	1360	50	0.31

Notes: This table reports average attributes of children in the MZ waves 1997 to 2018 that are observed in the same household as at least one of their parents by age of observation. The ranks are computed based upon the sample distribution of equivalized household income as described in Section 2.3.3.

Figure 2.1: Move-out Frequency by Parental Income Rank



Notes: This figure shows the relative frequency of move-outs of children aged 17-20 by parental income rank. It is computed based on a sample of 265,229 children in the years 2012-2018 where we observe the partial panel dimension of the MZ and can identify households surveyed for more than one wave. We define households with “lost children” as households which report a lower number of children aged 17-20 than in the previous year.

patterns more directly. Figure 2.1 displays the share of observed move-outs of children by parental income rank for the subsample of households in our data that is observed in the survey in multiple years. It shows that move-outs occur near uniformly across the income

distribution and are thus uncorrelated with parental income rank. Both exercises suggest that sample selection is not a major concern for our analysis. In addition, we demonstrate in the next section that choosing alternative age ranges barely affects our results.

Standard Errors. The standard errors reported alongside our estimates in the results section of this paper abstract from the fact that we estimate the cutoffs defining the percentile ranks. For the parental income gradient as well as the Q1 and Q5 measure, we cluster standard errors at the level of the sampling district, the primary sampling unit of the MZ. For the Q5/Q1 ratio, we report plug-in standard errors based on a delta-method argument.¹³

2.4 National Estimates


We begin our empirical analysis by characterizing social mobility at the national level. Figure 2.2 shows the share of children with an A-Level degree by parental income rank in our data, as well as the best linear approximation to the empirical CEF. As can be seen, a linear model provides a reasonable approximation to the CEF, a regularity that we observe in essentially all considered partitions of our data. In the national data, we estimate the parental income gradient at $\beta \times 100 = 0.52$, implying a gap of roughly 50% in the probability of obtaining an A-Level degree between children from the top and the bottom of the income distribution.¹⁴ Our measure of absolute mobility in the national data suggests that one third of children from the bottom quintile of the income distribution complete an A-Level degree, with Q1 estimated at 0.34. Both parametric and non-parametric mobility statistics imply that the odds ratio in the probability of obtaining an A-Level degree between children from the top quintile relative to the bottom quintile is greater than 2, with Q5/Q1 estimated at 2.25.¹⁵

¹³ The MZ data allows for consistent identification of primary sampling units across waves following the 2011 survey. For the estimates in Section 2.4.2, where we also use prior waves, we instead cluster standard errors at the household level. For the delta method, we linearize the ratio of averages which yields the following approximation for the variance of the sampling distribution of the $\overline{Q5}/\overline{Q1}$ sample ratio: $V(\overline{Q5}/\overline{Q1}) \approx \frac{1}{(\overline{Q1})^2} \left(V(\overline{Q5}) + \left[\frac{\overline{Q5}}{\overline{Q1}} \right]^2 V(\overline{Q1}) - 2 \frac{\overline{Q5}}{\overline{Q1}} Cov(\overline{Q5}, \overline{Q1}) \right)$.

¹⁴ For the national estimates, we pool our data over the period 2011-2018 to ensure consistency with the regional estimates in Section 2.5, for which obtaining results before 2011 is difficult due to frequent reforms of local administrative boundaries.

¹⁵ Appendix Table B.1 summarizes the estimates and shows that they are robust to variations in the age restriction defining our sample. Furthermore, results are unchanged when averaging parental income over several years before assigning the ranks, strongly suggesting that transitory income shocks in parental income do not bias our estimates.

Figure 2.2: Social Mobility at the National Level



Notes: This figure shows the fraction of children aged 17-21 that are either enrolled in the upper stage of an A-Level track or have already attained an A-Level degree by percentile rank of their parents in the national income distribution for the MZ waves 2011-2018. The income ranks are computed with respect to the national income distribution of households with children aged 17-21 in each survey year. The reported slope coefficient of 0.0052 (SE 0.004) is estimated by OLS using the underlying micro data.

Do these estimates depict Germany as a country of high or low relative mobility? While a cross-country comparison of our results is not straightforward, as the German system of upper secondary education and university funding is unusual, we are aware of two US studies which report comparable mobility statistics. Using data from the Census 2000, Hilger (2015) reports a parental income rank gradient of 3.6 percentage points in attending college for children aged 19-21. A higher point estimate is reported in Chetty et al. (2014a), who estimate the rank gradient in college enrollment at 6.7 percentage points for children aged 18-21 based on tax registry data. Under the assumption that college enrollment conditional on having obtained an A-Level degree is weakly increasing in parental income rank, our estimate of 5.2 percentage points implies a college enrollment gradient that falls into the range of point estimates reported for the US. Abstracting from differences in the distributions of college quality and the selection of students of different parental backgrounds into colleges

of different quality, our estimates suggest that educational mobility in Germany is similar to the US. We consider this finding noteworthy, as (after tax) income inequality is more pronounced in the US than in Germany, suggesting that one could expect steeper rank gradients in the US.¹⁶

The similar gradients between parental income and higher education in Germany and the US could imply two different things with respect to the transmission of income from parents to their children. On the one hand, intergenerational income mobility in Germany might be similarly low as in the US. On the other hand, the gradient between own and parental income in Germany could be less steep than the gradient between the A-level and parental income.¹⁷

To shed light on this question, we compute measures of intergenerational income mobility in the German Socio-economic Panel (SOEP), and compare them to the US and Denmark, two countries with recent available estimates and typically viewed at opposing ends of income mobility among high income countries. We use the studies by Chetty et al. (2014a) and Helsø (2021) as comparisons. Both focus on child incomes early in the lifecycle around age 30. To ensure comparability, we restrict the income observation window to ages 29-33.¹⁸ Parental income is measured as gross family income, child income either as individual labor earnings or as gross family income.¹⁹ Our analysis is then limited in sample size with around 800 to 1000 linked parent-child pairs. Somewhat reassuringly, as shown in Appendix Figure B.14, the gradient between obtaining an A-Level degree and parental income rank in the SOEP is estimated at 0.52, which is the same number we obtain in our main estimates based on the MZ.

Table 2.4 shows the results for estimated income mobility. We consider both, rank-rank coefficients and the IGE. The estimates based on the SOEP suggest that income persistence

¹⁶ Rauh (2017), for example, finds a negative correlation between inequality and public education expenditures across countries. If public education expenditures benefit lower-income children more, one expects a steeper rank gradient in the US. Our results do not support this conclusion.

¹⁷ Compare, for example, the insights from Landersø and Heckman (2017), who find that Denmark, a society that is characterized by high levels of income mobility, is similar to the US in terms of measures of educational social mobility.

¹⁸ We can only cover around half of the cohorts included in the main analysis, since for the younger ones we do not observe earnings at age 29-33 yet. More information on sample restrictions and some descriptive evidence is disclosed in Appendix B.3.1.

¹⁹ The reason why we use two different child income definitions is as follows. More than 20% of the children in our linked parent-child sample are still living in the parental household at the age of 29-33. We address this with two alternative approaches. First, we drop all cohabiting children from our sample. Second, we compute child family income as the sum of individual labor earnings of the child and its cohabiting partner, missing out on non-labor income since this is only measured at the household level. The first approach has the advantage to account for other sources of income than labor income, the second one the advantage to avoid sample selection.

Table 2.4: Intergenerational Income Mobility in the US, Denmark and Germany

Child Income	Parental Income	Estimand	US	DK	DE
Individual Labor Earnings (excl. 0)	Gross Family Income	IGE	-	-	0.276 (0.052)
					0.278 [†]
Gross Family Income	Gross Family Income	IGE	0.344 (0.000)	0.171 (0.004)	(0.057) 0.360 ^{††} (0.080)
Individual Labor Earnings (excl. 0)	Gross Family Income	Rank-rank	0.282 (0.000)	0.223 (0.003)	0.341 (0.037)
					0.320 [†]
Gross Family Income	Gross Family Income	Rank-rank	0.341 (0.000)	0.203 (0.003)	(0.039) 0.354 ^{††} (0.043)

Notes: This table shows estimates of intergenerational elasticities (IGE) and rank-rank slopes of intergenerational income mobility in the US and Germany (DE). The US estimates are taken from Table I in Chetty et al. (2014a), the estimates for Denmark from Table 1 in Helsø (2021). The German estimates are own calculations based on the SOEP. To ensure the highest possible degree of comparability between estimates, the German sample is restricted to children between 29-33 years old, and parental income is measured when children are 15-19 years olds. [†] indicates that *child* family income is measured as the sum of individual labor earnings of the child and its cohabiting partner (excluding zero incomes), whereas ^{††} indicates that child family income is measured as gross household income among all children are no longer cohabiting with their parents. The sample sizes underlying the German estimates range from 834 to 1041 individuals, depending on the specification. Robust standard errors in parentheses.

in Germany and the US is similar.²⁰ Remarkably, the estimates for Germany are outside the confidence bands of the reported numbers for Denmark. As such, we find no evidence that Germany should be considered as having high levels of income mobility, as observed in the Scandinavian countries. If anything, the estimates suggest similar magnitudes as for the US comparing similar age cohorts.²¹ In light of the strong sample size limitations encountered in

²⁰ The association between individual earnings rank and parental income is actually higher in Germany, while the comparison of the association between child and parent family income depends on the way we measure family income of children.

²¹ Our analysis updates previous work by Bratberg et al. (2017) with the SOEP, which focuses on cohorts around 20 years older (birth years 1956-1976). Their study finds income mobility in Germany to be more comparable to Scandinavia, suggesting a decline in income mobility compared to the birth cohorts preceding our sample. Closer investigation of these trends has to be left for future research, however, as it would require different data sources and much larger sample sizes than currently available.

the German Socio-Economic Panel, we now shift the focus back to the examination of social mobility patterns within the MZ data. This offers the most robust and reliable assessment of social mobility in Germany.

2.4.1 Subgroup Estimates

A natural question to ask is whether the national estimates mask meaningful differences in mobility measures across subpopulations. We focus on selected subgroups typically emphasized in the analysis of social mobility. Next to parental income, parental education is the second main measure of socio-economic background in the literature. We are therefore interested in the change of our mobility measures when conditioning on A-level degrees in the parental household. As intergenerational transmission mechanisms are further dependent on the family structure, we split by gender, parenting status (i.e. whether the child grew up with one or both parents in the household), the number of siblings, and the birth order. Our measurement approach is in particular suited to study how mobility varies between men and women, as our outcome measure is not affected by differential labor market participation, which complicates the analysis of gender differences in intergenerational income mobility. Specific to Germany, we want to additionally distinguish mobility between the eastern and western part of the country, which still differ widely in many socio-economic characteristics 30 years after the reunification. Finally, we focus on migration status, since we know that mobility patterns can differ substantially between migrants and natives (e.g. Abramitzky et al., 2021).


Table 2.5 reports mobility statistics separately for these groups. We document several interesting patterns. Most importantly, we find substantial differences by parental education. Figure 2.3 displays the A-Level share of children by parental income rank and the associated parental income gradient separately for children from households where no parent has an A-Level degree and for children from households where at least one parent has an A-Level degree. The A-Level share among children of parents without an A-Level degree at the top of the income distribution is comparable to the A-Level share among children with at least one A-Level educated parent at the bottom of the income distribution. Roughly speaking, the empirical distribution for children of A-Level educated parents is shifted upwards by approximately 30 percentage points, uniformly across ranks. The conditional rank gradients are attenuated due to the positive correlation between parental education and income ranks, with point estimates of approximately 0.3 in both groups. The intergenerational correlation in A-Level attainment

Table 2.5: Mobility Statistics for Subgroups

		Gradient	Q1	Q5	Q5/Q1	A-Level	N
Parental Education	No A-Level	0.33 (0.006)	0.28 (0.003)	0.55 (0.006)	1.94 (0.028)	0.39	145,892
	A-Level	0.29 (0.007)	0.61 (0.007)	0.84 (0.003)	1.36 (0.016)	0.75	85,080
Parenting Status	Single Parent	0.50 (0.010)	0.34 (0.004)	0.72 (0.009)	2.13 (0.037)	0.47	50,622
	Two Parents	0.54 (0.005)	0.34 (0.004)	0.76 (0.003)	2.26 (0.027)	0.54	179,715
Parents Married	Not Married	0.46 (0.010)	0.33 (0.004)	0.69 (0.008)	2.12 (0.037)	0.47	51,018
	Married	0.54 (0.005)	0.35 (0.004)	0.77 (0.003)	2.22 (0.025)	0.54	172,999
Gender	Male	0.53 (0.006)	0.29 (0.004)	0.72 (0.004)	2.49 (0.033)	0.47	123,649
	Female	0.50 (0.006)	0.40 (0.004)	0.81 (0.003)	2.02 (0.023)	0.58	107,323
Migration Status	Native	0.55 (0.005)	0.32 (0.004)	0.76 (0.003)	2.35 (0.028)	0.54	164,018
	Migrant	0.47 (0.009)	0.36 (0.004)	0.75 (0.007)	2.11 (0.032)	0.48	60,908
Region	West	0.50 (0.005)	0.34 (0.003)	0.76 (0.003)	2.19 (0.022)	0.52	201,684
	East	0.60 (0.011)	0.31 (0.007)	0.80 (0.007)	2.61 (0.062)	0.51	29,288
Siblings	Yes	0.55 (0.005)	0.35 (0.003)	0.79 (0.003)	2.29 (0.024)	0.52	156,960
	No	0.49 (0.007)	0.32 (0.005)	0.72 (0.004)	2.27 (0.039)	0.52	74,012
Birth Order	1st Child	0.51 (0.005)	0.34 (0.003)	0.76 (0.003)	2.22 (0.023)	0.53	165,336
	2nd Child	0.52 (0.008)	0.34 (0.005)	0.77 (0.005)	2.27 (0.036)	0.51	56,996
	Later Child	0.57 (0.021)	0.31 (0.009)	0.78 (0.017)	2.48 (0.092)	0.45	8,640

Notes: This table reports mobility statistics for selected groups of children observed in the MZ survey waves 2011-2018. The gradient measures the gap in the probability of obtaining an A-Level between children at the top and the bottom of the parental income distribution. Q1 and Q5 denote the share of children obtaining an A-Level in the first and fifth quintile of parental income; Q5/Q1 is the ratio between both measures. Migration background subsumes all individuals who immigrated to Germany after 1949, as well as all foreigners born in Germany and all individuals born in Germany with at least one parent who immigrated after 1949 or was born in Germany as a foreigner. The standard errors reported in parentheses below each point estimate are computed as described in Section 2.3.3.

Figure 2.3: Differences by Parental Education



Notes: This figure shows the fraction of children aged 17-21 observed in the MZ survey waves 2011-2018 that are either enrolled in the upper stage of an A-Level track or have already attained an A-Level degree by parental income rank, separately for children of parents who have not obtained an A-Level degree and children of parents where at least one of the parents has obtained an A-Level degree. The ranks are computed based upon the sample distribution of equivalized household income as described in Section 2.3.3. The reported estimates of the parental income gradient are based on the underlying micro data. Standard errors are reported in the first panel of Table 2.5.

in our data is 0.54. This finding highlights that the interpretability advantages of income-only based measures of parental background come at the cost of missing observable attributes of households that could be used to characterize social mobility more comprehensively.

The estimates reported in Table 2.5 reveal a few more interesting discrepancies. At the bottom of the income distribution, females and children with migration background are approximately 11 and 4 percentage points more likely to obtain an A-Level degree than their respective male and native counterparts. While the gender-gap is close to constant across the income distribution, the difference between migrant and native children vanishes in the top quintile. Moreover, we document larger income rank gradients for children of married and cohabiting couples, as well as for natives and children living in East Germany. The East-West gap in

parental income gradients is 0.1, implying a 10 percentage points larger top-bottom gap in the probability of attaining an A-Level degree in East Germany as compared to West Germany. We investigate such regional patterns in more detail in Section 2.5.

2.4.2 Time Trends

We next ask how social mobility has evolved over time. While our descriptive approach does not allow us to attribute changes in mobility measures to specific policies, our measurement strategy enables us to provide novel evidence on the evolution of social mobility in Germany for relatively recent birth cohorts. The period we study is particularly interesting, as it covers the second half of the arguably most significant educational reform in post-war Germany, the *Bildungsexpansion*, a large-scale policy of expanding upper secondary and higher education that, starting in the early 1970s, increased the A-Level share from around 20% to approximately 50% for the birth cohorts since the mid 1990s. This expansion was a policy response to a heated public debate on social mobility (Dahrendorf, 1966) and the increasing importance of education for economic growth at the time (Picht, 1964; Hadjar and Becker, 2006). We ask whether the large-scale expansion of upper-secondary education in Germany was accompanied by changes in social mobility as defined by our mobility measures.

To this end, we focus on a sample of 526,000 children born between 1980-1996.²² At the time of writing, the children of the respective birth cohorts are 25-40 years old and constitute a significant part of the German working population. Including relatively young cohorts in our analysis is feasible, as, in contrast to traditional measures that rely on the labor market incomes of children, our education-based measure of opportunities does not suffer from life-cycle biases. Figure 2.4 depicts the evolution of the A-Level share among 17-21 year old children in the MZ data for the birth cohorts under consideration. Our data covers roughly the second half of the expansion, with an observed increase in the A-Level share of 14 percentage points from 39% for the 1980 birth cohort to 53% for children born in 1996.²³ At the same time, income inequality increased only moderately,²⁴ and we do not find evidence that the


²² We restrict our attention to these cohorts to rule out that our estimates are affected by differences in the distribution of age at measurement. For the considered cohorts, the share of 17, 18-, 19-, 20- and 21-year-olds in our data is constant.

²³ The *Bildungsexpansion* featured a parallel increase of tertiary education and did not decrease the share of A-Level graduates taking up university studies. In the years 2002-2015, where most of our birth cohorts graduate, it fluctuated around 70% (<https://www.datenportalbm.bf.de/portal/de/Tabelle-2.5.74.html>).

²⁴ While wage inequality rose in the 1990s and early 2000s when most children in our sample grew up, Fuchs-Schündeln et al. (2010) document that inequality in consumption and disposable income, the income concept used in this paper, increased only moderately.

expansion was accompanied by a decline in annually measured A-Level wage premia, as documented in Appendix Figure B.3. However, as the children under consideration have only partially entered the labor market even today, we note that with the currently available data it is not possible to rule out that the A-Level premium may eventually differ for these cohorts. Furthermore, the counterfactual development of the A-Level wage premium in absence of the *Bildungsexpansion* is inherently unobserved.

Figure 2.4: A-Level Share by Cohort




Notes: This figure shows the fraction of children born between 1980 and 1996 and observed at ages 17-21 that are either enrolled in the upper stage of an A-Level track or attained an A-Level degree in the MZ data. The shaded area displays pointwise 95% confidence intervals based on standard errors as described in Section 2.3.3.

Figures 2.5 and 2.6 display estimates of our mobility measures for the same cohorts. While the odds ratio captured by the Q5/Q1 ratio decreased by approximately one third, from around 3 for the 1980 birth cohort to slightly above 2 for the 1996 cohort, the parental income gradient has remained constant at around 0.52, the point estimate that we report at the national level based on more recent data. At the same time, absolute mobility as measured by the Q1 measure increased substantially, from approximately 0.22 in 1980 to 0.35 in 1996. The same overall pattern emerges when estimating mobility trends for the subgroups studied in Section 2.4.1 as reported in Figures B.6 and B.7 in the Appendix.


2 Social Mobility in Germany

Figure 2.5: Parental Income Gradient by Cohort




Notes: This figure shows for children aged 17-21 the evolution of the parental income gradient by birth cohort. The shaded area displays pointwise 95% confidence intervals based on standard errors as described in Section 2.3.3.

Figure 2.6: Quintile Measures by Cohort



Notes: This figure shows for children aged 17-21 the evolution of the quintile based measures of social mobility by birth cohort. While the left axis corresponds to the Q5/Q1 ratio, the right axis corresponds to the Q1 measure. The shaded areas display pointwise 95% confidence intervals based on standard errors as described in Section 2.3.3.

Figure 2.7: A-Level Share by Cohort Quintile



Notes: This figure shows the share of children born between 1980 and 1996 who obtained an A-Level degree by birth cohort and quintile of the parental income distribution in the MZ data. The shaded area displays pointwise 95% confidence intervals based on standard errors as described in Section 2.3.3.

The connection between these findings is best summarized in Figure 2.7, which depicts the A-Level share by quintile across birth cohorts: The *Bildunsexpansion* took place uniformly across the income distribution, with increases of about 14 percentage points in the A-Level share in all parts of the distribution. Did the *Bildunsexpansion* achieve its goal of fostering social mobility in Germany? While the expansion unquestionably increased absolute mobility as we measure it, the time trend in relative mobility is less straightforward to interpret. On the one hand, the attenuation of the Q5/Q1 ratio caused by the uniform increases in A-Level shares could suggest an increase in relative mobility according to a proportional notion of the concept. On the other hand, a less optimistic angle to interpret the same development is to consider the inverse odds ratio, that is the ratio between the probability *not* to obtain an A-Level for children in both quintiles. In the birth cohort 1980, children in Q1 were 2.2 times more likely not to obtain an A-Level degree than children in Q5. For children born in 1996, this inverse odds ratio has increased to 2.8, meaning that the relative gap in not obtaining an A-Level has actually widened. In contrast, the unaltered top-bottom gap in the probability of attaining an A-Level captured by the parental income gradient emphasizes stagnation in absolute differences. As the parental income gradient is insensitive to the chosen reference

point, we tend to interpret the evidence primarily as a stagnation of relative mobility. However, as both absolute and relative disparities often form the normative basis for interventions, all readings can be justified.

Trends in Ability and Selection Patterns by Parental Income. An interesting question concerns the selection of students who were marginal with respect to the *Bildungsexpansion* policy – meaning they would not have entered the A-level track without this education expansion. If marginal students from low income families are more talented than marginal children from high income families, this could suggest that the school system itself discriminates against children from disadvantaged backgrounds at the costs of overall efficiency of the system and that the *Bildungsexpansion* was partially a remedy in that respect.

We turn to an additional data source to obtain measures of ability for the cohorts in question. The well-known Programme for International Student Assessment (PISA) administered by the OECD provides test scores at age 15. It is generally accepted as a measure which displays a high correlation with e.g. IQ tests and other skill assessments (e.g. Rindermann, 2007; Pokropek et al., 2022). It only covers the more recent birth cohorts 1990-1996 considered in our paper because parental income is only collected since the 2006 PISA wave. To complement this, we employ the German Socio-Economic Panel (SOEP), which has annually collected school grades at age 17 for the birth cohorts 1982-1996. While the SOEP does not offer test-score data, it contains information about grades. Following the literature in the economics of education (e.g. Hanushek et al., 2022; Gneezy et al., 2019; Jensen and Rasmussen, 2011; Brunello and Rocco, 2013), we use grades and test scores in math to obtain an ability proxy which can be compared consistently across social groups.²⁵

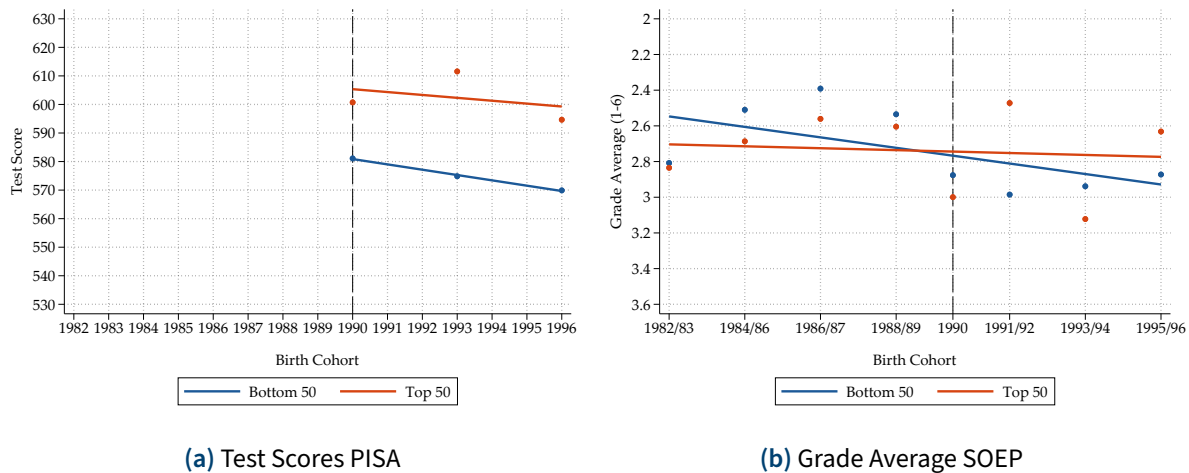
Figure 2.8 shows time trends of averaged grades (Panel B) and test scores (Panel A) for students attending the highest school track. The red line refers to above- and the blue line to below-median parental income.²⁶ Both measures suggest a slight deterioration of test scores and grades over time for both parental income groups.

The interesting question is about the differences in test scores and grades for marginal students from high versus low parental income. Marginal here refers to these students who only entered the highest track because of the educational expansion and the increase in the number of

²⁵ We obtain similar but slightly noisier results when averaging over all available grade and test score information. In the SOEP, this additionally includes grades in German and the first foreign language, in PISA test scores for German and “Science”.

²⁶ The sample sizes do not permit finer parental income splits, unfortunately.

Figure 2.8: Time Trend Math Grades and Test Scores



Notes: The figure shows averages math grades in the German Socio-Economic Panel (SOEP) in Panel (A), and average math test scores in the PISA-I data in Panel (B) by birth cohort. Math performance increases in PISA test scores, and decreases in SOEP grades, which range from 6 (worst) to 1 (best). The lines show the corresponding linear OLS fits. The PISA-I sample includes around 1,000 15-year old students on Gymnasium and the Gymnasialzug of Gesamthochschulen per cohort, the SOEP sample covers 1,061 children in total. Additional information on the underlying data is disclosed in Appendix B.3.

students in the highest track. Since “being marginal” is, naturally, an unobservable state, we present two different ways to make assumptions that enable us to learn about ability differences between marginal students of both parental income groups. First, we assume that test scores (and grades) have no trend for inframarginal students, such that the changes seen in Figure 2.8 can be attributed to entering marginal students. Appendix B.3.3 shows how, under this assumption, the ability for marginal students from both parental income groups can be inferred in a straightforward way by accounting for the increase of students in each group. In a second approach, we conduct a prediction exercise based on observables of children and parents to classify students as inframarginal. Then we consider how grades/test scores changed over time for students with those observables and impose these trends on inframarginal students. This procedure is described and results are shown in B.3.3. Although the assumptions behind the two approaches are rather different, they yield consistent results.

Table 2.6 shows the results of the first exercise. According to the PISA data, marginal children among birth cohorts 1990 to 1996 from the bottom half of the income distribution displayed lower test scores than marginal children from the top 50%. The difference of 21 test points corresponds to 31% of a standard deviation. For the same birth cohorts, the grade averages obtained in the SOEP data also suggest higher ability among marginal children from the upper

Table 2.6: Math Grades and Test Scores of Marginal Children

	SOEP			PISA		
	Bottom 50	Top 50	Δ	Bottom 50	Top 50	Δ
1982-1990	2.9	2.5	0.45 SD	-	-	-
1990-1996	3.5	2.6	0.84 SD	552	573	0.31 SD
1982-1996	3.1	2.6	0.54 SD	-	-	-

Notes: This table shows the average math grades in the German Socio-Economic Panel (SOEP) and the average PISA math test scores among “marginal” children in birth cohorts 1982, 1990 and 1996, separately for children below and above median parental income. The grades are computed using Equation B.2 in Appendix B.3.3, which also contains more details about the calculation. The third column expresses the differences between both groups in terms of the standard deviations, which is 1.06 for math grades in the SOEP, and 69 points for PISA test scores. Due to the small sample size of the SOEP, three year averages around the actual birth cohort are used to compute grade averages (1982: 1982-1984, 1990: 1989-1991, 1996: 1994-1996).

half of the income distribution. This pattern is also there for older birth cohorts. Results are similar for the second approach. In Appendix B.3.3, we report that for the early birth cohorts (82-90) there is a small advantage for lower income students. However, for later cohorts (90-96) this reverses and the evidence suggest more favorable test scores and grades for high income students.

Summing up, over the whole period considered (birth cohorts 1982 to 1996), there is no evidence that marginal students with below-median parental income perform better than marginal students with above-median parental income. There is some evidence, however, that among the more recent cohorts (1990-1996) test scores and grades for marginal students from higher parental income backgrounds are better compared to lower parental income backgrounds.

2.5 Regional Estimates

An interesting regularity documented in the recent empirical literature on social mobility is that there exists substantial geographic variation in social mobility measures within politically homogeneous entities, suggesting that regional comparisons can be used to gain a better understanding of the causes of social mobility (e.g. Chetty et al., 2014a; Acciari et al., 2022; Corak, 2020; Deutscher and Mazumder, 2020; Chuard and Grassi, 2020). This idea is appealing, as attributing cross-country discrepancies in social mobility to differences in single

characteristics or policies is difficult to justify. Complementary to well-designed evaluations of political reforms that rely on variation across time (e.g. Bertrand et al., 2021), within-country geographic variation can be helpful in understanding the causal mechanisms fostering or impeding social mobility by identifying exposure effects (Chetty and Hendren, 2018; Bütikofer and Peri, 2021). Moreover, pronounced regional differences can suggest mechanisms that warrant investigation.

The regional analysis conducted in this section is motivated by these considerations. In a first step, we present evidence of meaningful geographic variation in our mobility measures across regions in Germany. In a second step, we then ask what we can learn from the observed differences. We structure our regional analysis by disaggregating our data in a stepwise fashion, lending credence to our parametric mobility statistics while taking into account the political and economic landscape of Germany.

2.5.1 States

A natural starting point for our regional analysis are the 16 federal states of Germany. By constitutional law, the responsibility for the design and implementation of the education system falls under the jurisdiction of the German states and not under the jurisdiction of the federal government. As a consequence, the states have considerable discretion in the design of their education systems, leading to distinctions in the rigor of the tracking system, the capacities of each track, the types of schools and curricula and other important features of the education system. In particular, states differ with respect to the duration of primary school after which all children are allocated into the different tracks, the number of tracks (2 or 3) and the importance of teacher recommendations for admitted track choices. While in all states teachers recommend a track for each child at the end of primary school, track recommendations are binding only in some states. These parameters of the state education systems and their suspected consequences for social mobility are often at the center of the public debate on educational mobility in Germany.

Table 2.7 reports our mobility estimates for the 16 states, sorted by the point estimate of the parental income gradient in ascending order. We document significant and economically meaningful differences in both absolute and relative mobility measures between states. For example, the top-bottom gap in the probability of attaining an A-Level degree is approximately 20 percentage points larger in Bremen than in Hamburg, two city states in north-west Germany approximately 100 kilometers apart.

Table 2.7: Social Mobility at the State Level

State	Gradient	Q1	Q5	Q5/Q1	A-Level Share	Tracks	Binding Rec.
Hamburg	0.45 (0.033)	0.43 (0.023)	0.80 (0.017)	1.86 (0.109)	0.60	2	No
Rhineland-Palatinate	0.50 (0.019)	0.36 (0.013)	0.76 (0.011)	2.12 (0.086)	0.53	2	No
North Rhine-Westphalia	0.51 (0.009)	0.41 (0.006)	0.82 (0.005)	2.02 (0.032)	0.59	3	Ref
Hesse	0.52 (0.015)	0.39 (0.011)	0.81 (0.007)	2.07 (0.061)	0.59	3	Ref
Baden-Württemberg	0.52 (0.011)	0.34 (0.008)	0.76 (0.006)	2.24 (0.056)	0.53	3	Ref
Saarland	0.53 (0.040)	0.33 (0.024)	0.74 (0.025)	2.28 (0.186)	0.54	2	Ref
Schleswig-Holstein	0.53 (0.023)	0.32 (0.015)	0.76 (0.014)	2.34 (0.117)	0.52	2	No
Lower Saxony	0.54 (0.013)	0.29 (0.008)	0.73 (0.009)	2.52 (0.077)	0.48	3	Ref
Bavaria	0.54 (0.011)	0.24 (0.007)	0.67 (0.006)	2.75 (0.084)	0.42	3	Yes
Berlin	0.56 (0.021)	0.39 (0.013)	0.85 (0.011)	2.20 (0.082)	0.59	2	No
Brandenburg	0.57 (0.027)	0.35 (0.019)	0.84 (0.014)	2.37 (0.134)	0.60	2	Ref
Saxony-Anhalt	0.58 (0.034)	0.25 (0.017)	0.72 (0.026)	2.88 (0.227)	0.43	2	Ref
Saxony	0.61 (0.025)	0.28 (0.014)	0.78 (0.016)	2.83 (0.156)	0.48	2	Yes
Mecklenburg-Vorpommern	0.63 (0.041)	0.25 (0.020)	0.76 (0.028)	3.00 (0.256)	0.45	2	No
Bremen	0.64 (0.044)	0.32 (0.025)	0.86 (0.026)	2.65 (0.220)	0.55	2	No
Thuringia	0.65 (0.032)	0.25 (0.017)	0.76 (0.023)	3.07 (0.234)	0.46	2	Yes

Notes: This table reports mobility statistics for each federal state of Germany based on all children observed in the MZ waves 2011-2018. The gradient measures the gap in the probability of obtaining an A-Level between children at the top and the bottom of the parental income distribution. Q1 and Q5 denote the share of children obtaining an A-Level in the first and fifth quintile of parental income; Q5/Q1 is the ratio between both measures. The standard errors reported in parentheses below each point estimate are computed as described in Section 2.3.3. The classification of the state education systems is based on the description of educational reforms in Helbig and Nikolai (2015). In the last column, “Ref” indicates that teacher recommendations were reformed during the time period relevant for our analysis.

Similarly, the share of children obtaining an A-Level degree from the bottom quintile of the parental income distribution is 10 percentage points larger in Baden-Württemberg than in Bavaria, the two southernmost states of Germany. The estimated differences between states do not result from differences in the shape of the empirical CEFs, as we find that the linearity assumption underlying our parametric mobility estimates is supported by the data (compare Figure B.8). The table also reiterates the east-west gap documented in Section 2.4.1: except for Bremen, the least mobile states are all located in East Germany.

While we find that the differences in our measure of absolute mobility can be well explained by differences in the states' A-Level shares, that is the relative capacity of the highest track, there is no clear pattern in our estimates with respect to the aforementioned characteristics of the state education systems displayed in the last two columns of the table. Our findings suggest that, while certainly important, the design of the tracking system cannot readily explain the pronounced differences in our mobility measures between states.

2.5.2 Cities

A similar picture emerges when we restrict our analysis to urban regions of Germany. Table 2.8 reports our mobility estimates for the 15 largest labor markets of Germany, consisting of cities and their catchment areas as defined by commuting flows.

Compared to the national average, the largest urban regions of Germany show lower levels of relative, but higher levels of absolute social mobility. At the same time, the table shows that the regional differences observed at the state-level can also be found within states. For example, the top-bottom gap is approximately 8 percentage points larger in Cologne than in Düsseldorf, two large cities in North Rhine-Westphalia located approximately 40 kilometers apart. Similarly, our estimates of absolute mobility differ by 8 percentage points between Nuremberg and Munich, two large cities in Bavaria.


Table 2.8: Social Mobility in the 15 Largest Urban Labor Markets

City	State	Gradient	Q1	Q5	Q5/Q1	A-Level Share
Hamburg	HH/SH	0.47 (0.025)	0.41 (0.018)	0.79 (0.012)	1.94 (0.090)	0.58
Düsseldorf	NW	0.47 (0.029)	0.45 (0.023)	0.84 (0.014)	1.87 (0.100)	0.65
Münster	NW	0.47 (0.041)	0.47 (0.030)	0.84 (0.021)	1.78 (0.120)	0.62
Gelsenkirchen	NW	0.50 (0.035)	0.40 (0.018)	0.81 (0.029)	2.01 (0.116)	0.57
Stuttgart	BW	0.50 (0.024)	0.34 (0.017)	0.75 (0.012)	2.19 (0.114)	0.55
Bonn	NW	0.50 (0.039)	0.44 (0.030)	0.86 (0.016)	1.94 (0.135)	0.65
Duisburg	NW	0.51 (0.033)	0.42 (0.022)	0.84 (0.017)	2.02 (0.113)	0.58
Frankfurt	HE	0.52 (0.025)	0.42 (0.019)	0.83 (0.011)	1.97 (0.093)	0.62
Munich	BY	0.54 (0.025)	0.31 (0.021)	0.71 (0.011)	2.32 (0.162)	0.53
Dortmund	NW	0.55 (0.033)	0.40 (0.022)	0.86 (0.017)	2.16 (0.125)	0.59
Cologne	NW	0.55 (0.027)	0.38 (0.019)	0.85 (0.014)	2.25 (0.120)	0.60
Hanover	NI	0.56 (0.036)	0.30 (0.022)	0.76 (0.021)	2.51 (0.195)	0.53
Berlin	BE	0.56 (0.021)	0.39 (0.013)	0.85 (0.011)	2.20 (0.082)	0.59
Nuremberg	BY	0.60 (0.035)	0.23 (0.022)	0.70 (0.023)	3.01 (0.297)	0.43
Leipzig	SN	0.68 (0.044)	0.26 (0.026)	0.80 (0.028)	3.11 (0.335)	0.48

Notes: This table reports mobility statistics for the 15 largest urban local labor markets in Germany, as measured by their total population in 2017, based on the MZ waves 2011-2018. The gradient measures the gap in the probability of obtaining an A-Level between children at the top and the bottom of the parental income distribution. Q1 and Q5 denote the share of children obtaining an A-Level in the first and fifth quintile of parental income; Q5/Q1 is the ratio between both measures. The local labor markets are sorted, in ascending order, by the point estimate of the parental income gradient. Standard errors are computed as described in Section 2.3.3. The point estimates for the city-states can differ from those reported in Table 2.7, as the urban labor markets typically also include surrounding towns and villages.

The most striking discrepancy between cities in our data is observed for Hamburg and Leipzig, with a difference of approximately 20 percentage points in the estimated top-bottom gap, as well as 15 percentage points in our estimate of the Q1 measure. Figure 2.9 displays our raw data for the two cities. Similar to the previously considered partitions of our data, we show in Figure B.9 that the empirical CEFs are well approximated by a linear function. Overall, our city-level findings suggest that the relative opportunities of children can differ meaningfully across politically similar and geographically close regions of Germany.²⁷

Figure 2.9: Social Mobility in Hamburg and Leipzig




Notes: This figure shows the fraction of children aged 17-21 observed in the MZ survey waves 2011-2018 that are either enrolled in the upper stage of an A-Level track or have already attained an A-Level degree in Hamburg (Panel A) and Leipzig (Panel B). The reported slope coefficients are estimated by OLS using the underlying micro data. Standard errors are reported in Table 2.8.

2.5.3 Local Labor Markets

We finally disaggregate our data once more to the level of local labor markets (LLMs). The 258 LLMs in Germany represent aggregations of counties based on commuting flows, comparable to the commuting zones in the US. Except for five local labor markets (Bremen, Bremerhaven, Hamburg, Mannheim and Ulm), all counties aggregated into LLMs belong to a single state. The median number of children in our sample (observations) per LLM is 552 (mean: 895). The lowest number of observations across all LLMs is 100 (LLM Sonneberg) and the largest number of observations is 8159 (LLM Stuttgart).

²⁷ What cannot be inferred from Table 2.8 is the individual rank of each city. To obtain valid inference on rankings in terms of the parental income gradient or other mobility statistics, it is necessary to apply the methods developed in Mogstad et al. (2024).

Figure 2.10: A-Level Share and Q1 Measure by Local Labor Market



Notes: This figure presents heat maps of the A-Level share (Panel A) and the Q1 measure (Panel B) by LLM. Children are assigned to LLMs according to their current residence. The estimates are based on children aged 17-21 in the years 2011-2018 for which we have non-missing information on educational attainment and parental income. The A-Level share is defined as the fraction of children aged 17-21 that are either enrolled in the upper stage of an A-Level track or have already attained an A-Level degree. The Q1 measure reports this same share for children in the bottom 20% of the parental income distribution.

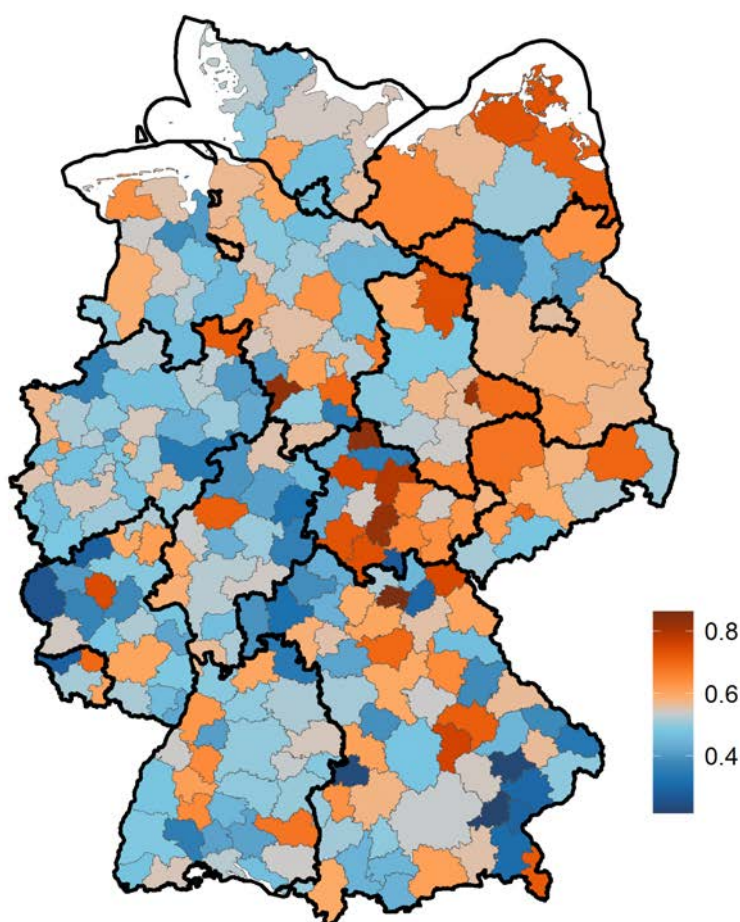
Regional Patterns in Absolute Mobility. We begin our local labor market-level analysis by studying regional variation in absolute mobility. Figure 2.10 shows the A-Level Share (Panel A) and our estimate of the Q1 measure (Panel B) in each of the 258 LLMs. Red areas correspond to regions with low, and blue areas to regions with high values of the respective statistic. For both statistics, state-level clusters are clearly visible. Panel (A) shows that the A-Level share is uniformly higher in the local labor markets of states with high average A-Level capacities, such as North Rhine-Westphalia or Hesse. Comparing the two panels demonstrates that, unsurprisingly, our measure of absolute mobility is closely linked to the local A-Level share ($\rho = 0.76$). Consequently, we observe lower levels of absolute mobility in regions with low A-Level shares, such as Bavaria.

Overall, there exists substantial variation in absolute mobility. In some regions, less than 15% of children from the bottom quintile of the national income distribution obtain an A-Level

degree, whereas in other regions this number exceeds 50%. We find that 44% of the variation in the Q1 measure and 57% of the variation in the A-Level share can be attributed to state level differences.

Regional Patterns in Relative Mobility. While the variation in absolute mobility can be well explained by state A-Level shares, regional patterns in relative mobility are less obvious. Figure 2.11 presents a heat map of our estimates of the parental income gradient.²⁸ Blue areas

Figure 2.11: Parental Income Gradient by Local Labor Market



Notes: This figure presents a heat map of the parental income gradient by LLM. Children are assigned to LLMs according to their current place of residence. The estimates are based on children aged 17-21 in the years 2011-2018 for which we have non-missing information on educational attainment and parental income. The parental income gradient is obtained as the slope coefficient of a regression of the A-Level dummy on a constant and the parental income rank, multiplied by 100.

²⁸ The corresponding heat map for the Q5/Q1 ratio is displayed in Figure B.10 in the Appendix. The correlations between our mobility measures are reported in Appendix Table B.2.

2 Social Mobility in Germany

represent regions of high mobility (low gradients), whereas red areas indicate low mobility. In some rural labor markets, the parental income gradient is estimated below 0.3, whereas in the least mobile areas the gradient exceeds 0.8. While LLMs in the East exhibit lower mobility on average, clusters of high and low mobility are spread out across all of Germany. In contrast to our estimates of absolute mobility, some of the observed clusters extend beyond state borders. The LLMs with the highest gradient (Lichtenfels) and the lowest gradient (Mühldorf) are both located in Bavaria. Indeed, we find that only 13% of the variation across LLMs can be explained by state level differences.

Robustness of Regional Estimates. While disaggregating our data to the LLM level allows us to ask several interesting questions, it makes it harder to distinguish meaningful variation from sampling error, as our mobility estimates are based on fewer observations. In Appendix B.5, we employ empirical Bayes methods to address this concern in a principled manner. Reassuringly, we find evidence of substantial overdispersion. Moreover, the main patterns described above also become evident when computing mobility statistics at the level of spatial planning regions, a higher-level aggregation of LLMs. The median number of observations per spatial planning region is 1741 (mean: 2406). Figure B.11 displays heat maps of our mobility statistics for all 96 spatial planning regions of Germany. By construction, dispersion in mobility estimates is more muted as we move to a higher level of aggregation. Yet, we still find substantial variation in mobility estimates and clusters of high and low relative mobility crossing state borders (Panel C). Moreover, it is again the case that state level differences explain more of the variation in absolute than relative mobility (72% vs 37%). Furthermore, while average parental income ranks naturally vary across Germany (Figure B.12), we show in Figure B.13 that mobility estimates for local labor markets remain virtually unchanged when computing parental income ranks not with respect to the national income distribution but with respect to the income distribution in the respective state or region type.

Sorting. What can we learn from the estimated regional differences across local labor markets? A first insight relates to the debate on the potential of place-based mobility policies. An active literature argues that places shape economic outcomes and that place-based policies can be an effective and cost-efficient tool to improve outcomes by amending local conditions (Kline and Moretti, 2014; Neumark and Simpson, 2015). In the context of educational policies and social mobility, it is often argued that the government should allocate additional resources to the local public school systems of socially immobile regions to enhance mobility. However,

such a policy is unlikely to achieve its objective if social mobility in the respective regions is low for reasons other than the quality of local schools. For example, if a region exhibits a high degree of inequality in parental educational attainment, the patterns we document in Section 2.4.1 would likely result in low levels of relative mobility as measured by the parental income gradient.

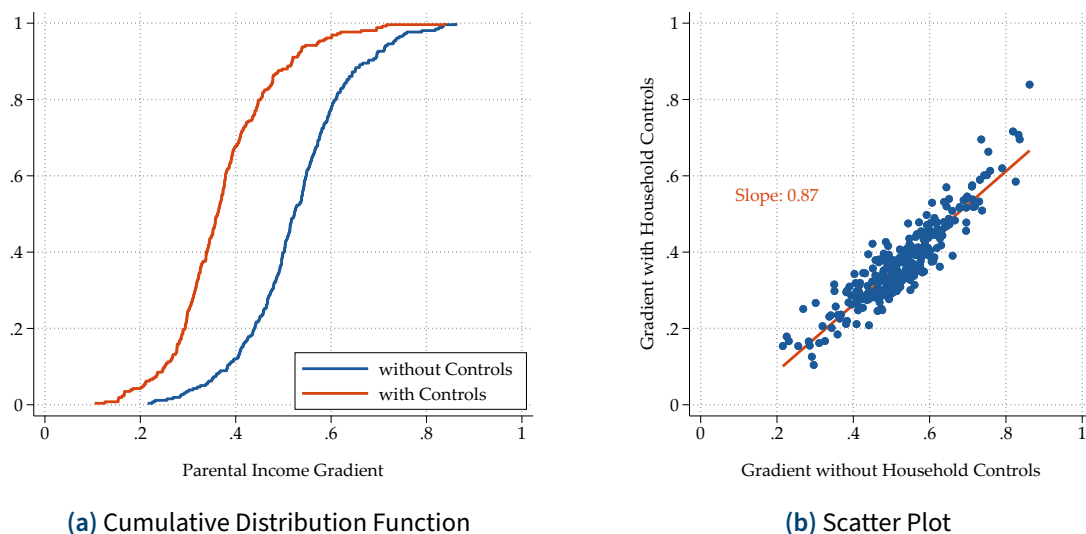
Such systematic sorting mechanisms are at the center of the academic debate regarding the interpretation of the regional differences in estimated mobility measures within countries.²⁹ The German census data allows us to directly test whether regional differences are muted once we account for household characteristics. We do so by computing conditional rank gradients, which we then compare to our parental income gradient. The set of conditioning variables we use for this exercise includes age and gender of the child, migration background, age and marital status of the parents, the number of siblings, a dummy for single parents and the highest parental education level in four categories. Figure 2.12, Panel (A) plots the marginal distributions of conditional and unconditional rank gradients. It shows that the CDF of the unconditional gradient first order stochastically dominates the CDF of the conditional gradient, which is expected given the patterns documented in Table 2.5. At the same time, the variance of the distribution of conditional rank gradients is approximately the same as the variance of the unconditional gradient. Moreover, as reported in Panel (B) we find that, despite the predictive power of the included household attributes, the relative ordering of gradients is largely unaffected by conditioning, which suggests that regional sorting of households cannot explain the regional variation in relative social mobility as we measure it. Conditional and unconditional gradients are strongly correlated, with a Pearson correlation of 0.91 and a Spearman rank correlation of 0.89. The same pattern emerges when repeating this analysis for higher levels of regional aggregation.^{30 31}

²⁹ For example, Rothbaum (2016) and Gallagher et al. (2018) suggest that in the US a substantial share of the geographic variation in the intergenerational mobility measures reported in Chetty et al. (2014a) can be explained by differences in household characteristics across commuting zones. In Chetty et al. (2014a), this was not tested directly, whereas in later work, Chetty and Hendren (2018) draw on a movers design to overcome this problem. By comparing outcomes of children who move across commuting zones, they can separate place effects from sorting patterns. Compared to our approach, the movers design utilizes only a subset of children, but has the advantage that it can control for a large share of potential sorting on (unobserved) household characteristics not captured by our set of variables.

³⁰ At the level of spatial planning regions, the Pearson correlation is 0.90 and the Spearman rank correlation 0.86. At the state level, the Pearson correlation amounts to 0.91 and the Spearman rank correlation to 0.84.

³¹ Note that, while this finding suggests that sorting does not play a major role, the same pattern would emerge if our regional estimates were dominated by sampling error, in the sense that the between local labor market variation in gradients was negligible relative to the estimation uncertainty. We address this concern in Appendix B.5.

Figure 2.12: Sorting: Conditional and Unconditional Rank Gradients



Notes: This figure compare unconditional and conditional estimates of the parental income gradient by local labor market. The conditioning variables include age and gender of the child, migration background, age and marital status of the parents, the number of siblings, a dummy for single parents and the highest parental education level in four categories. Panel (A) plots the Cumulative Distribution Function (CDF) of the conditional and unconditional parental income gradients, Panel (B) shows a scatter plot of the point estimates as well as their linear fit.

Predictors of Mobility. If sorting cannot account for most of the spatial variation in mobility, the question remains why some regions of Germany exhibit a higher degree of social mobility than others. Similar to previous studies which document geographic variation in intergenerational mobility, we are not able to test existing theories of intergenerational transmission processes which could explain these patterns conclusively in our data. To nevertheless learn from our estimates, we conduct a prediction exercise to characterize mobile regions in more detail. In Appendix B.4, we describe the methodology underlying the prediction exercise and present the results, with Appendix Table B.6 displaying the 15 most informative predictors of mobility differences between local labor markets. Overall, our selection procedure highlights social characteristics, the local organization of the education system and labor market conditions. These correlational findings are consistent with causal studies that emphasize the importance of local characteristics for child and adolescent outcomes (Chetty and Hendren, 2018; Damm and Dustmann, 2014).

2.6 Conclusion

This paper provides novel empirical evidence on the level, evolution and geography of social mobility in Germany. Our measurement strategy allows for the use of large-scale census data and characterizes mobility using robust statistical measures of the association between the educational attainment of a child and its parents' relative position in the national income distribution. We find that on average a 10 percentile increase in parental income rank is associated with a 5.2 percentage point increase in the probability to obtain an A-Level degree, implying a top-bottom gap of approximately 50 percentage points. This gap remained stable for the 1980-1996 birth cohorts, despite a concurrent massive roll-out of higher secondary education. An expansion in access to higher education alone may therefore not be sufficient to reduce the opportunity gap between children from high and low income households. At the same time, we find that absolute mobility increased substantially.

We further document variation in mobility measures across regions and show that household characteristics cannot account for these differences. As such, our findings are consistent with place-based rather than sorting-type explanations of geographic dispersion in mobility measures. Obtaining an optimal set of mobility predictors based on our disaggregated estimates, we find that social characteristics, the local organization of the education system and labor market conditions best predict mobility at the regional level. More research is needed to understand whether these correlations reflect structural relationships.

The measurement approach described in this paper provides a timely and feasible way to monitor the development of social mobility in Germany for recent cohorts. This framework may also prove useful in other countries where the highest secondary school degree is crucial for future career options. Education systems with secondary school degrees of comparable importance to the Abitur in Germany include Italy (Maturità), Austria (Matura) and the UK (A-Level).

3 Multidimensional Equality of Opportunity in the United States

This chapter is based on co-authored work with Paul Hufe, Martyna Kobus, and Andreas Peichl. A previous version has been published as a CESifo Working Paper. See Hufe et al. (2022a) for the full reference.

3.1 Introduction

In a fair economy, people act on a level playing field to acquire monetary resources. This idea—oftentimes labeled as *equality of opportunity*—is widely reflected in fairness conceptions of academic philosophers and the general public (Fong, 2001; Alesina et al., 2018; Rawls, 1971; Cohen, 1989; Roemer, 1998; Almås et al., 2020; Arneson, 2018; Cappelen et al., 2007). As a consequence, there is an active literature in economics that assesses the satisfaction of the opportunity-egalitarian ideal in different countries at different points in time. We contribute to this literature by providing the first analysis of the association between family background characteristics and the joint distribution of income and wealth in the US.

Existing studies on inequality of opportunity and intergenerational mobility focus on income—and to a lesser extent on wealth—to measure monetary resources.¹ By excluding either income or wealth from the analysis, these studies neglect important information on individual consumption possibilities which arguably are the relevant metric to assess the financial well-being of individuals. For example, unidimensional analyses will misrepresent the financial well-being of income-poor heirs who support their lifestyle by selling assets or of asset-poor persons with high incomes. Therefore, if society cares for the financial well-being of individuals more broadly, we should move from unidimensional analyses of monetary resources to analyses of the joint distribution of income and wealth.


The focus on unidimensional analyses would be innocuous if income and wealth were perfect substitutes as indicators for monetary resources. There are at least two reasons why this

¹ For the US, see Solon (1992); Davis and Mazumder (2017); Chetty et al. (2014a) for intergenerational income mobility; Pfeffer and Killewald (2018); Charles and Hurst (2003) for wealth mobility, and Pistoiesi (2009); Niehues and Peichl (2014); Hufe et al. (2022b) for inequality of opportunity in incomes.

3 Multidimensional Equality of Opportunity in the United States

is implausible. First, well-off parents transmit monetary resources to the next generation through bequests and inter vivo gifts (Elinder et al., 2018; Boserup et al., 2016; Wolff, 2002). In turn, expected wealth transfers distort the education and labor supply decisions of children (Kopczuk, 2013; Kindermann et al., 2020). Such behavioral responses create a wedge between the relative positions of individuals in income and wealth distributions: individuals that receive a lot of wealth from their parents are not necessarily those who earn high incomes. This observation is particularly relevant for the analysis of time trends as inheritances have grown in many Western societies in recent decades (Piketty and Zucman, 2015). Second, changes in wealth are a function of savings and asset price changes. While the savings channel depends on income, the price channel depends on portfolio compositions. Therefore, changes in asset prices are another force that drives a wedge between the relative positions of individuals in income and wealth distributions. Again, this observation is particularly relevant for the analysis of time trends as wealth-to-income ratios—and therefore the sensitivity of wealth to asset price fluctuations—has grown over time (Kuhn et al., 2020).

Figure 3.1: Intergenerational Income Mobility and the Distribution of Monetary Resources in the United States



Data: PSID.

Note: Panel (A) shows a binned scatter plot of average child income ranks in the period 2010-2016 by income rank of their parents in the period 1983-1988. All individuals are aged 25-60. Panel (B) shows a heatmap of year-specific income and wealth ranks for the pooled sample of individuals aged 25-60 in the period 1983-2016. Each data point shows the share of individuals in a fixed two-percentile income (wealth) bin that belong to a particular two-percentile wealth (income) bin. See Section 3.3 for detailed definitions of income and wealth.

In Figure 3.1, we use data from the Panel Study of Income Dynamics (PSID) to show that these concerns are relevant for the analysis of equal opportunities in the United States. In Panel (A), we replicate the well-known finding that child incomes increase with the income of their

parents during childhood: an increase of parental income by 10 percentile ranks is associated with an average increase of 3.7 percentile ranks in child income. This estimate is very similar to the slope estimate of 0.34 in Chetty et al. (2014a). In Panel (B), a heatmap of income and wealth ranks demonstrates that income and wealth are far from perfect correlates (Rank correlation $\rho = 0.55$).² Taken together, these patterns suggest that unidimensional analyses of equality of opportunity and intergenerational mobility provide a distorted image of the importance of family background for individual consumption possibilities and financial well-being.

In this paper, we address these shortcomings by analyzing the association between family background and the joint distribution of income and wealth. We use the PSID to implement and extend a novel measure of multidimensional equality of opportunity (Kobus et al., 2020). Our analysis proceeds in two steps. First, we construct an *intergenerational sample* in which we measure equality of opportunity in monetary resources by using parental income ranks as the only proxy for socioeconomic background. This practice is consistent with the literature on intergenerational mobility; however, the sparsity of data links across generations prevents meaningful analyses of time trends. Second, we construct an *individual sample* in which we substitute parental income ranks by a vector of alternative socioeconomic background characteristics. These data are available on an annual basis and allow us to assess trends over the period 1983-2016.

Our findings can be summarized as follows. First, multidimensional inequality of opportunity is consistently and substantially higher than inequality of opportunity in income. Hence, unidimensional analyses that focus on income only underestimate the extent to which monetary resources are associated with family background. Second, the playing field in the US has become more tilted in recent decades: inequality of opportunity in 2016 is 56% higher than in 1983. Furthermore, time trends are markedly different when accounting for the multidimensionality of monetary resources. For example, an exclusive focus on income suggests small increases in unequal opportunities after the year 2000. This relative stability, however, is accompanied by strong increases in the wealth dimension leading to an overall increase in unequal opportunities.

The contribution of this paper is threefold. First, we complement recent literature that characterizes the joint distribution of income and wealth in the US (Kuhn et al., 2020; Berman and Milanovic, 2023). This literature focuses on inequalities in outcomes but remains silent on opportunities and intergenerational transmission processes. Second, we provide novel insights

² The moderate rank correlation is not due to idiosyncratic fluctuations in income or wealth. Using 5-year moving averages for income and wealth yields $\rho = 0.59$.

regarding the development of equality of opportunity in the United States. While existing literature documented relative stability of equality of opportunity in terms of income after 2000 (Chetty et al., 2014b; Hartley et al., 2022), we show that decreases emerge once we account for the wealth dimension. Third, we provide a novel decomposition of the multidimensional measure into inequality of opportunity in income, inequality of opportunity in wealth, and the association of both outcomes across family background types. Association is a distinctive feature of joint distributions that cannot be captured by unidimensional analyses. It indicates whether individuals of a given family background are more likely to fare better or worse in both dimensions simultaneously. We use a multidimensional framework to combine these dimensions and to obtain an overall conclusion regarding the extent of unequal opportunities in the US.

3.2 Measurement of Multidimensional Equality of Opportunity

Consider a population $\mathcal{N} := \{1, \dots, N\}$ and a set of outcomes $\mathcal{K} := \{1, \dots, K\}$ that capture monetary resources. Individuals $i \in \mathcal{N}$ receive utility from $q \in \mathcal{K}$. We can summarize the distribution of monetary resources by outcome matrix X of dimension $N \times K$, where an element x_{iq} denotes i 's outcome in dimension q . Outcomes are determined by two sets of factors: a set Ω that captures family background characteristics and a set Θ that captures individual choices. We define $\omega_i \in \Omega$ as a comprehensive description of family background and $\theta_i \in \Theta$ as a comprehensive description of the choices made by $i \in \mathcal{N}$. For each q , there is an outcome-generating function defined as follows:

$$x_{iq} = f_q(\omega_i, \theta_i), \forall i \in \mathcal{N}. \quad (3.1)$$

In an equal-opportunity society, outcome differences are determined by individual choices θ_i but are invariant to family background ω_i (Roemer, 1998). There are different ways of translating this idea into measures. Most empirical literature relies on an *ex-ante* approach, which broadly consists of two steps. First, one partitions the population into types $T = \{t_1, \dots, t_M\}$. Individuals belong to a type if they share the same set of family background characteristics: $i, j \in t_m \Leftrightarrow \omega_i = \omega_j$. For example, in rank-rank measures of intergenerational mobility, types are defined by parental income ranks. Second, one assesses differences in average outcomes across types by regressing child outcomes on a measure of family background:

$$x_{iq} = \alpha_q + \beta_q \omega_i + \epsilon_{iq}. \quad (3.2)$$

There are two prominent ways of summarizing the resulting information in measures of inequality of opportunity: (i) β_q , which is the standard statistic in the literature on *intergenerational mobility* (Black and Devereux, 2011). (ii) $I(X) = I(\mathbb{E}[x_{iq}])$, where $I(\cdot)$ is any inequality index, and which is the standard statistic in the literature on *equality of opportunity* (Roe-mer and Trannoy, 2016). It defines inequality of opportunity as inequality between types: all within-type variation is removed and inequality reflects only inequality that arises due to family background. Clearly, both measures are isomorphic and capture the opportunity-egalitarian idea: the higher β_q , the more life outcomes x_{iq} are predicted by family background ω_i , and the higher the corresponding measure of inequality of opportunity.³

In this paper, we follow the tradition of the equality of opportunity literature and summarize outcome differences across types with an inequality index. In particular, we use the measure of Kobus et al. (2020), which allows us to account for the multidimensionality of monetary resources. For the sake of simplicity and in line with our empirical application, we focus on the case of two outcome dimensions and set $K = 2$. In this case, the index is given by

$$I(X) = 1 - \left(\sum_{t=1}^M \frac{N_t a_t}{\sum_{t=1}^M N_t a_t} \frac{(\mu_1^t)^{r_1} (\mu_2^t)^{r_2}}{(\mu_1)^{r_1} (\mu_2)^{r_2}} \right)^{\frac{1}{r_1 + r_2}} \quad \forall_t a_t < 0, r_1, r_2 < 0, \quad (3.3)$$

where N_t denotes the number of individuals in type t and μ_q^t (μ_q) the type (population) means in outcome q . In the following, we will describe the roles of r_q and a_t which are weights for outcome dimension q and types t , respectively. However, before doing so we note that if $r_q = 0$ for either outcome, $I(X)$ boils down to a unidimensional measure of inequality of opportunity which is the well-known Atkinson (1970) index applied to types t .⁴

Dimension weights r_q govern the sensitivity of the measure to between-type inequality in outcome q . The more negative r_q , the more convex the measure in q , and the higher its sensitivity to between-type inequality in this dimension. For example, if $r_1 < r_2$, $I(X)$ is more sensitive to inequality in the first than in the second outcome, i.e., the former is then relatively more important in the inequality assessment. r_q is also related to the degree of inequality aversion ϵ_q via $r_q = 1 - \epsilon_q$. As ϵ_q rises, the index becomes more sensitive to inequality at the bottom of the distribution than at the top. Note that ϵ_q is a parameter chosen by the researcher. In his seminal work, Atkinson (1970) arbitrarily set ϵ_q equal to 1, 1.5 and 2. Subsequently,

³ One advantage of the equality of opportunity approach is that it allows for a flexible accommodation of family background characteristics other than family income or wealth.

⁴ Generally, the index in Equation (3.3) is a multidimensional generalized entropy measure. Similarly, the empirical literature on equality of opportunity often uses unidimensional generalized entropy measures like the mean log deviation to summarize inequality between family background types.

3 Multidimensional Equality of Opportunity in the United States

empirical research has tried to infer plausible values of ϵ_q from economic policy design and tax schedules (Young, 1990; Gouveia and Strauss, 1994; Aristei and Perugini, 2016). These estimates range between 1 and 2 depending on the country and period of interest. In our baseline calculations, we choose $\epsilon_q = 1.2$ ($r_q = -0.2$) for both income and wealth. However, in section 3.5, we show that our conclusions on time trends do not change for a wide range of plausible choices for ϵ_q .

Type weights a_t determine how much the social planner values respective types. The more negative a_t , the higher the weight attached to type t . To ensure that $I(X)$ measures inequality, higher weight is assigned to types that have lower values of $(\mu_1^t)^{r_1}(\mu_2^t)^{r_2}$. Note that a_t is also a parameter chosen by the researcher. In our benchmark calculations, we choose type weights that decrease linearly with type ranks in the values of $(\mu_1^t)^{r_1}(\mu_2^t)^{r_2}$. However, in section 3.5, we show that our conclusions on time trends do not change when type weights are concave or convex in these ranks.

The functional form $(\mu_1^t)^{r_1}(\mu_2^t)^{r_2}$ is the same as the Cobb-Douglas utility function but the parameterization of weights differs. In particular, both a_t and r_q are negative ensuring that the index is convex and supermodular. Convexity ensures sensitivity to Pigou-Dalton transfers between types, i.e., that $I(X)$ increases after transfers that increase between-type inequality in dimension q . This is a fundamental property for ex-ante measures of inequality of opportunity. Supermodularity ensures sensitivity to correlation-increasing transfers, i.e., that $I(X)$ increases after transfers that increase the cross-type correlation of income and wealth. This is a fundamental property for multidimensional measures of inequality. It is important to note that the functional form is not adopted arbitrarily but that it is derived from first principles: $I(X)$ is the only index fulfilling the fundamental principles of ex-ante equality of opportunity in a multidimensional setting while satisfying standard properties of inequality measures such as *monotonicity*, *utilitarian aggregation*, and *ratio scale invariance*.⁵

Beyond its normative foundations, the index has several useful properties for empirical analyses. First, it can be decomposed to distinguish between the impact of inequality of opportunity in each outcome and the association of outcomes across types. We will use this property in our empirical analyses to understand the drivers of inequality of opportunity in the US. Second, the index is bounded in the interval $[0, 1]$. If $\mu_q^t = \mu_q$ for every type t and outcome q , the index will be zero. Finally, the index is a welfare-based inequality measure in line with the pioneering work by Atkinson (1970). For example, a value of 0.25 (0.5) means that existing

⁵ See also Appendix C.1 for an illustration of its core properties.

inequality of opportunity imposes a welfare cost of 25% (50%) of the population average of each outcome. In other words, if there was perfect equality of opportunity, society would achieve the same level of welfare using only 75% (50%) of the available monetary resources in income and wealth (Atkinson, 1970; Kolm, 1969; Sen, 1973).

3.3 Data

Data Source. We assess the evolution of equal opportunities in the US while accounting for the multidimensionality of monetary resources. Therefore, we require data with information on income, wealth, and family background for a long time period. In the US, the Panel Study of Income Dynamics (PSID) is the only publicly available data source that satisfies these criteria. For example, while the Survey of Consumer Finances (SCF) offers a long time series on household income and wealth, it contains limited information on the family background of its respondents. Since 1968 the PSID collects rich information on income and family background for a nationally representative sample of US households. Since 1984 it also collects data on wealth.⁶ Children who leave the parental household become independent sampling units in the PSID. Therefore, it is possible to link data across generations. Income information is collected for the year predating the survey year. Hence, we use information from the income reference (survey) period 1983-2016 (1984-2017). We now turn to a description of relevant variables.

Monetary Resources. We consider two dimensions of monetary resources: income and wealth. We measure income as annual disposable household income. It comprises total household income from labor, asset flows, windfall gains, private transfers, public transfers, private retirement income and social security pensions net of total household taxes. We measure wealth as household net worth. It comprises the sum of home equity, other real estate, private businesses, vehicles, transaction accounts, corporate equities, annuities/IRAs and other savings net of any debt.

We scale household incomes and wealth by the modified OECD equivalence scale. Hence, we measure both income and wealth at the household level, whereas the units of analysis are individuals. This choice is consistent with our overarching interest in consumption possibilities since the application of equivalence scales allows for resource sharing among household members.

⁶ Until 1999 wealth information was collected every five years. Since then, it is a regular part of every PSID wave.

3 Multidimensional Equality of Opportunity in the United States

Wealth data in the PSID is often considered inferior to wealth data in the SCF. Therefore, we compare PSID and SCF concerning time trends in household net worth in Appendix Figure C.1. Due to oversampling of wealthy households, the SCF assigns a larger share of total net worth to the top 10% of the wealth distribution. Yet, level differences at the top are the only notable difference between PSID and SCF. Importantly, time trends in household net worth are consistent across both data sources.⁷

Family Background Characteristics and Types. We consider two alternative ways to measure family background. First, we use parental income ranks for the total incomes of mothers and fathers averaged over the years 1983-1988. We residualize parental income from the first and second-order polynomials of parental age to account for life-cycle effects in parental earnings profiles. We then partition the population into 36 types by ranking total parental income. Second, we use a vector of alternative socioeconomic background variables. This vector includes parental education (3 categories), parental occupation (3 categories), race (2 categories), and Census region of upbringing (2 categories).⁸ We partition the population into 36 types based on the combination of these family background variables.

Estimation Samples. We base our estimates on two different samples. First, we construct an *intergenerational sample* of 1,366 individuals. To obtain this sample, we drop all individuals with (i) missing or negative income and wealth, and (ii) missing information on parental education, parental occupation, race and region of upbringing. Then we match all respondents to both of their parents and drop observations with (iii) missing parental income. Lastly, we restrict observations to children (parents) aged 25-60 in the period 2010-2016 (1983-1988).⁹ This sample allows us to proxy ω with parental income rank, which is common practice in the literature on intergenerational mobility. However, it imposes restrictions on the analysis of time trends since one requires information on both parental and child outcomes while allowing for sufficient time between these observations.

⁷ See also Pfeffer et al. (2016) for a detailed comparison of wealth definitions in PSID and SCF.

⁸ Parental education: low (less than high school), intermediate (high school), high (some college and more); parental occupation based on 1-digit 1990 Census codes: low (4,8,9), intermediate (3,5,6,7), high (1,2); race: white (non-Hispanic), other; Census region of upbringing: South, other.

⁹ Appendix Table C.2 details how the different restrictions affect the final sample size.

Second, to investigate time trends, we construct an *individual sample*.¹⁰ In contrast to the previous sample, we drop requirement (iii). Again, we limit the sample to individuals aged between 25-60. We obtain a sample of at least 4,000 observations in every year of the period 1983-2016 which allows us to monitor the development of equality of opportunity in the US over 33 years.

We are conscious that the PSID is subject to selective survey attrition across waves and that our data restrictions may distort our sample through selective item non-response. Therefore, we follow Meyer et al. (2015) and perform all calculations using sampling weights that match the Current Population Survey (CPS). Results remain unchanged when using standard survey weights provided by the PSID instead. Descriptive statistics for all estimation samples are disclosed in Table C.1.

3.4 Results


Our analysis proceeds in two steps. First, we measure equality of opportunity in the *intergenerational sample*. Thereby, we either use parental income ranks or the vector of alternative socioeconomic background variables to proxy for family background. We will show that both approaches yield very similar results. Second, having validated the use of alternative socioeconomic background variables, we use the *individual sample* to analyze trends in equality of opportunity in the period 1983–2016.

Intergenerational Estimates. Figure 3.2 shows estimates for inequality of opportunity in the *intergenerational sample* for different combinations of outcomes and family background variables.

First, we focus on the dark-blue bars that show estimates based on parental income ranks. In Panel (A), we measure monetary resources by income only and inequality of opportunity amounts to 0.19. In Panel (B), we measure monetary resources by wealth and inequality of opportunity doubles to a level of 0.39. Finally, in Panel (C) we account for the multidimensionality of monetary resources by considering both income and wealth. Then, inequality of opportunity amounts to 0.29. These results suggest that we tend to underestimate tilt in the playing field when relying on income as the sole proxy for monetary resources.

¹⁰ We note it is possible to analyze time trends in an intergenerational sample when focusing on income and wealth measures in early adulthood (e.g. Hartley et al., 2022). However, such an age restriction would not be adequate in our setting due to lifecycle gradients in both income and wealth.

Figure 3.2: Inequality of Opportunity in the US - Intergenerational Sample



Data: PSID.

Note: This figure shows estimates of inequality of opportunity in the US for the *intergenerational sample*. Panel (A) (Panel [B]) shows results for a unidimensional definition of monetary resources based on income (wealth). Panel (C) shows results for a multidimensional definition of monetary resources based on income and wealth. In each panel, inequality of opportunity estimates are based on 36 types according to alternative definitions: parental income rank or self-reported family background. Estimates are computed based on Equation (3.3) with dimension weights $r_{Income} = r_{Wealth} = -0.2$. Δ indicates the difference in inequality estimates across type definitions. p -values for the null hypothesis that $\Delta = 0$ are bootstrapped using 1,000 draws.

Second, we focus on a comparison between dark-blue bars and light-blue bars. To estimate the latter, we replace parental income ranks with a vector of alternative socioeconomic background characteristics. Point estimates remain virtually unchanged by this alternation and we cannot reject the equality of estimates at conventional levels of significance. This result suggests that parental income ranks and alternative socioeconomic background characteristics contain similar information about family background. In general, this is an encouraging message as data sets including intergenerational links are much scarcer than data sets including retrospective information on various socioeconomic background variables.¹¹

¹¹ Jácome et al. (2021) use a similar strategy and approximate parental income with self-reported background characteristics, i.e., retrospective information collected from the respondents of interest and not their parents. In Appendix Figure C.2, we show that income distributions within family background types are broadly comparable regardless of whether we use types based on parental income ranks or self-reported background characteristics.

This conclusion is robust to a variety of checks. First, it is well-known that PSID subsamples with intergenerational links are positively selected on their socioeconomic status (Ward, 2021). Therefore, we re-weight the *intergenerational sample* to match the broader population characteristics concerning parental education, parental occupation, race, Census region of upbringing, and age. The re-weighting has little effect on inequality of opportunity estimates (Appendix Figure C.3). Second, existing literature documents life-cycle bias in intergenerational mobility estimates. Estimates of both income and wealth mobility tend to be downward (upward) biased when children are young (old) (Haider and Solon, 2006; Nybom and Stuhler, 2016; Mazumder, 2018). This bias is usually addressed by measuring income in midlife. There are slight upward corrections of inequality of opportunity in wealth when we restrict our sample to the age range 40-45 (Appendix Figure C.4). Importantly, however, differences in results based on income ranks and alternative socioeconomic background characteristics remain small for all considered age ranges. Third, we compare estimates based on the alternative background characteristics to expanded sets of family backgrounds where we add parental income and parental wealth ranks. The resulting estimates are very close to our baseline estimates suggesting that the vector of alternative socioeconomic background characteristics captures most of the relevant cross-family variation in socioeconomic status (Appendix Figure C.5).


We conclude that the vector of alternative socioeconomic background characteristics provides suitable information to capture intergenerational disadvantage. As these data are available on an annual basis, we can use them to assess time trends in inequality of opportunity.

Time Trend (1983-2016). Figure 3.3 shows the development of inequality of opportunity in the US over the period 1983-2016. The following patterns emerge.

First, inequality of opportunity in income increased from 0.10 to 0.26 over time. We can distinguish two distinct periods. On the one hand, we observe marked increases from 1983 to 1998. On the other hand, there are only moderate increases after the year 2000. This two-partite pattern is consistent with findings from the literature on intergenerational income mobility. For example, Davis and Mazumder (2017) show that equality of opportunity decreased for cohorts born in the 1960s and that entered the labor market after 1980. Chetty et al. (2014b) show that this trend flattens for cohorts born in the 1970s that enter labor markets in the 1990s and 2000s. Likewise, Hartley et al. (2022) document a flat time trend in the intergenerational income correlation of mothers and daughters after 2000.

3 Multidimensional Equality of Opportunity in the United States

Figure 3.3: Inequality of Opportunity in the US, 1983-2016
Baseline Estimates



Data: PSID.

Note: This figure shows estimates of inequality of opportunity in the US for the *individual sample* over the period 1983-2016. Inequality of opportunity estimates are based on 36 types according to the following socioeconomic background characteristics: parental education, parental occupation, race, and region of upbringing. Estimates are computed based on Equation (3.3) with dimension weights $r_{Income} = r_{Wealth} = -0.2$. 95% confidence intervals (shaded areas) are bootstrapped using 1,000 draws.

Second, inequality of opportunity in wealth increased from 0.34 to 0.50 over time. Again, we can distinguish two distinct periods. On the one hand, we observe moderate increases from 1983 to 2006. In these years, increases in the stock market were accompanied by a robust housing market (Kuhn et al., 2020; Wolff, 2017). Since owner-occupied housing has a higher weight in the portfolios of individuals from lower socioeconomic backgrounds, increasing house prices attenuated the tendency towards a less opportunity-egalitarian distribution of wealth. On the other hand, differences in portfolio compositions across socioeconomic backgrounds started working in the opposite direction with the financial crisis in 2008. While

the stock market experienced a quick recovery, house prices did not catch up to their pre-crisis level. As a consequence, the wealth distribution has become less opportunity-egalitarian with the crisis—a trend that has not reverted ever since.

Taken together, the playing field for the acquisition of monetary resources has become more tilted over time. Starting at a level of 0.25 in 1983, inequality of opportunity in the joint distribution of income and wealth reached a level of 0.39 in the latest period of observation. This shift corresponds to an increase of 56%. Importantly, the trend towards decreasing opportunities to acquire monetary resources continues after the year 2000. This finding can be related to extant literature invoking intergenerational income mobility estimates to conclude relative stability in equality of opportunity in recent years (Hartley et al., 2022; Chetty et al., 2014b). To the extent that these works aim to proxy financial opportunities more generally, they miss important information by focusing on income only. When accounting for the multidimensionality of monetary resources, one cannot reject the claim that opportunities in the US have further declined after the year 2000.

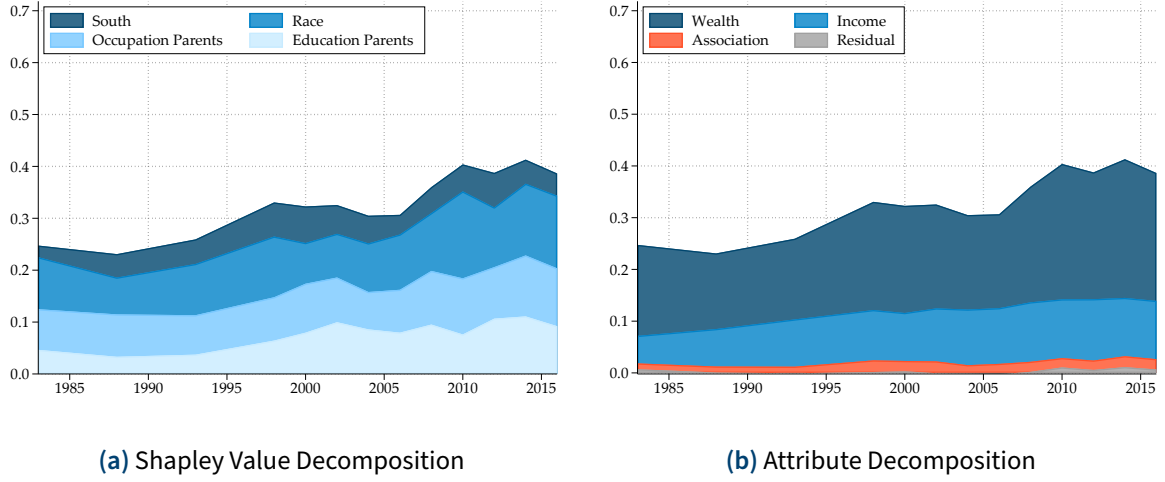
Decomposition. To develop a better understanding of these trends, we conduct a Shapley value decomposition, i.e., we decompose the trend in equality of opportunity into the contributions from different socioeconomic background characteristics: parental education, parental occupation, race, and the region of upbringing (Figure 3.4, Panel [A]).

First, 57% of the overall increase in inequality of opportunity is accounted for by parental education and occupation. This finding is consistent with Hufe et al. (2022b) who identify these components as the strongest drivers of increasing inequality of opportunity in incomes in the US. Second, 29% of the overall increase is accounted for by race. At first glance, this finding appears at odds with the stagnation of racial income gaps since the civil rights era (Derenoncourt and Montialoux, 2021; Bayer and Charles, 2018). However, the importance of race increases only after the 2008 financial crisis. Therefore, decreased opportunities to acquire monetary resources are most likely driven by the sustained effect of the financial crisis on the housing wealth of Black Americans (Wolff, 2017; Kuhn et al., 2020). Lastly, the contribution of the region of upbringing remains stable over time.

We also conduct an attribute decomposition, i.e., we decompose the time trend into the contributions of (i) inequality of opportunity in income, (ii) inequality of opportunity in wealth,

3 Multidimensional Equality of Opportunity in the United States

Figure 3.4: Inequality of Opportunity in the US, 1983-2016
Decomposition by Background Characteristic and Outcome Dimension



Data: PSID.

Note: This figure shows a decomposition of inequality of opportunity in the US for the *individual sample* over the period 1983-2016. Inequality of opportunity estimates are based on 36 types according to self-reported socio-economic background characteristics. Estimates are computed based on Equation (3.3) with dimension weights $r_{Income} = r_{Wealth} = -0.2$. The decomposition in Panel (A) is based on the Shapley value procedure proposed in Shorrocks (2013). The decomposition in Panel (B) is based on the attribute decomposition derived in Appendix C.2.

as well as (iii) the cross-type association in both outcomes. The last dimension is of particular interest as it cannot be analyzed in unidimensional measures of inequality of opportunity. In Appendix C.2, we show that $I(X)$ can be decomposed as follows:

$$\begin{aligned}
 I(X) &= \frac{r_1}{r_1+r_2} \underbrace{\left(1 - \left(\sum_{t=1}^M \frac{N_t a_t}{\sum_{t=1}^M N_t a_t} \left(\frac{\mu_1^t}{\mu_1} \right)^{r_1} \right)^{\frac{1}{r_1}} \right)}_{=I_1(\text{Income})} \\
 &+ \frac{r_2}{r_1+r_2} \underbrace{\left(1 - \left(\sum_{t=1}^M \frac{N_t a_t}{\sum_{t=1}^M N_t a_t} \left(\frac{\mu_2^t}{\mu_2} \right)^{r_2} \right)^{\frac{1}{r_2}} \right)}_{=I_2(\text{Wealth})} \\
 &+ \frac{1}{r_1+r_2} \underbrace{\left(1 - \frac{\sum_{t=1}^M N_t a_t \sum_{t=1}^M N_t a_t (\mu_1^t)^{r_1} (\mu_2^t)^{r_2}}{\sum_{t=1}^M N_t a_t (\mu_1^t)^{r_1} \sum_{t=1}^M N_t a_t (\mu_2^t)^{r_2}} \right)}_{=\kappa_I(\text{Association})} \\
 &+ R,
 \end{aligned} \tag{3.4}$$

where I_q is a unidimensional index of inequality of opportunity in outcome dimension q , κ_I is a measure of cross-type association in outcomes, and R is a residual resulting from linear approximation.¹²

The results of this decomposition are shown in Panel (B) of Figure 3.4. 43% and 51% of the overall increase in inequality of opportunity can be explained by trends in unidimensional inequality of opportunity in income and wealth, respectively. The cross-type association of income and wealth explains only 6% of the overall increase in unequal opportunities. This finding is somewhat surprising since recent research points to an increased correlation between income and wealth in the US (Berman and Milanovic, 2023; Kuhn and Ríos-Rull, 2016). Our results suggest that these increases at the individual level are mostly driven by increased correlation within family background types while the association of these outcomes across family background types remains rather stable. However, we note that the relative stability of cross-type association κ_I depends on the parameter choices for a_t and r_q . In Appendix Table C.3, we show the decomposition of time trends under different plausible assumptions for r_q . For example, if we allow for a higher degree of inequality aversion by choosing $r_{Income} = -0.4$ and $r_{Wealth} = -0.4$, κ_I explains up to 17% of the overall increase in unequal opportunities. This finding indicates that family background types in the lower tail of the distribution have become more resource-constrained than the rest of the population in both income and wealth simultaneously.

3.5 Sensitivity Analysis

Parameter Choices. We assess the sensitivity of our main conclusions to changes in the measurement parameters, i.e., dimension weights r_q and type weights a_t . Alternative parameter choices correspond to different normative assumptions about inequality aversion. Therefore, they will lead to level shifts in the extent of inequality of opportunity—a property that is well-known in the literature (Atkinson, 1970). However, we are especially concerned with the development of inequality of opportunity over time. In the following discussion, we will therefore abstract from levels and focus on whether changes in unequal opportunities are sensitive to different assumptions about these parameters.

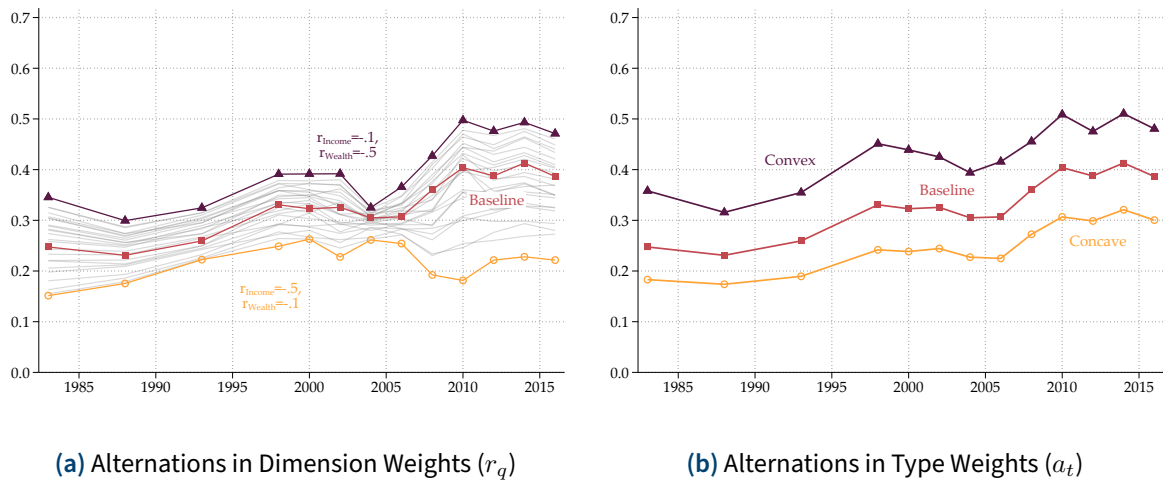
First, dimension weights r_q determine inequality aversion in income and wealth, respectively. In our baseline estimates, we give both dimensions equal weight and choose $r_{Income} =$

¹² I_q are unidimensional inequality of opportunity measures based on the Atkinson (1970) index of inequality—see our discussion in section 3.2.

3 Multidimensional Equality of Opportunity in the United States

$r_{Wealth} = -0.2$. However, there may be good reasons to give different weights to different dimensions of monetary resources. For example, one could argue that wealth should receive a higher weight due to its insurance value. Reversely, one could argue that wealth should receive a lower weight since it is less liquid and might not be available for instantaneous consumption. Panel (A) of Figure 3.5 shows alternative results for all pairwise combinations over the parameter grid $r_q \in (-0.1, -0.2, -0.3, -0.4, -0.5)$. Lowest estimates of inequality of opportunity are obtained for $r_{Income} = -0.5$ and $r_{Wealth} = -0.1$; that is, in the case where we place little weight on the wealth dimension, and more weight on the income dimension. We note that such income-focused parameterization yields a flat trend after the year 2000. This result is expected and consistent with existing work on intergenerational income mobility (Hartley et al., 2022; Chetty et al., 2014b). However, even small increases in the wealth focus lead to upward corrections in inequality of opportunity estimates and overturn the conclusion of flat time trends after 2000. The highest estimates of inequality of opportunity are obtained for $r_{Income} = -0.1$ and $r_{Wealth} = -0.5$; that is, in the case where we place more weight on the wealth dimension, and little weight on the income dimension.

Figure 3.5: Inequality of Opportunity in the US, 1983-2016
Sensitivity to Parameter Choices



Data: PSID.

Note: This figure shows the sensitivity of inequality of opportunity in the US for the *individual sample* over the period 1983-2016 under different parameter choices. Panel (A) shows the sensitivity to alternations in r_q . We display are all pairwise combinations of $r_{Income} \in (-0.1, -0.2, -0.3, -0.4, -0.5)$ and $r_{Wealth} \in (-0.1, -0.2, -0.3, -0.4, -0.5)$. The central line replicates our baseline estimates from Figure 3.3 where we use linear $r_{Income} = r_{Wealth} = -0.2$. Panel (B) shows the sensitivity to alternations in a_t . We construct convex (concave) weights as a_t^2 ($a_t^{0.5}$). The central line replicates our baseline estimates from Figure 3.3 where we use linear a_t .

Second, type weights a_t determine the degree of inequality aversion between types. In our baseline estimates, we choose linear a_t that are inversely related to type ranks in monetary resources. Panel (B) of Figure 3.5 shows alternative results for convex (a_t^2) and concave type weights ($a_t^{0.5}$). The lowest estimates of inequality of opportunity are obtained for concave type weights, where we place relatively less weight on inequality in the lower tail of the type distribution. Conversely, the highest estimates are obtained for convex type weights, where we place relatively more weight on inequality in the upper tail of the type distribution. Despite changes in levels, our conclusions concerning time trends are insensitive to parameter choices in a_t .

Data Choices. In Appendix Figure C.6, we furthermore document that our main conclusions are robust to different data choices.

First, we recompute inequality of opportunity while smoothing transitory changes in income and wealth, i.e., we replace annual values of income and wealth with 5-year averages. As a consequence, outcome variables provide better proxies for the long-term income and wealth potential of individuals (Solon, 1992). Time trends are very close to our baseline estimates.

Second, we recompute inequality of opportunity for different type partitions. To this end, we code three additional variables and add them to the vector of socioeconomic background characteristics: the number of siblings (11 categories), a dummy for foreign-born parents, and a dummy for single-parent families. In turn, we follow Brunori et al. (2023) and let a regression tree algorithm decide on the optimal type partition in each year of our analysis. Again, time trends are very similar to our baseline estimates.

Third, we recompute inequality of opportunity for different ways of dealing with non-positive income and wealth. For our baseline, we drop observations with negative income/wealth and set observations with zero income/wealth to 1 USD, respectively. Alternatively, we (i) drop all observations with negative and zero income/wealth, or (ii) retain all observations with negative and zero income/wealth in the sample. Time trends are again very similar, regardless of the chosen specification.

Fourth, we recompute inequality of opportunity using alternative definitions of income and wealth. Our baseline definitions may contain mechanical relationships between income and wealth. Wealth enters household income through capital returns; reversely, savings from household income increase wealth in a given period. Therefore, we divorce both concepts as

follows: first, we replace household disposable income with household labor market earnings, i.e., we use an income concept that is not mechanically related to asset returns. Second, we adjust household net worth by deducting active savings in a given year, i.e., we use a wealth concept that is not mechanically related to contemporaneous saving decisions. Our time series are not sensitive to these adjustments, suggesting that mechanical relationships between income and wealth are not the main driver of our results.

We conclude: while the level of inequality of opportunity and the magnitude of its increase varies with different measurement choices, all main conclusions from our baseline estimates remain in place. The only exception arises if we parameterize our index in ways that give little weight to the wealth dimension. In this case, we replicate analyses that focus on the income dimension only and we obtain a flat time trend after the year 2000.

3.6 Conclusion

In this paper, we study inequality of opportunity for the acquisition of monetary resources in the US over the period 1983-2016. In contrast to existing work, we account for the multidimensionality of monetary resources by targeting the joint distribution of income and wealth. Our results show that unidimensional analyses may miss important information when analyzing the playing field in the US: first, we document a more unequal distribution of opportunities when complementing income with the wealth dimension. Second, there are strong and consistent increases in inequality of opportunity over time. This trend is driven by a less opportunity-egalitarian distribution of income until 2000, and a less opportunity-egalitarian distribution of wealth after the financial crisis in 2008.

4 Career Preferences and Socio-Economic Background

This chapter is single authored. A previous version has been published as an ifo Working Paper. See Schüle (2023) for the full reference.

4.1 Introduction

Earnings in the labor market crucially depend on parental background, as documented in the extensive literature on intergenerational income mobility (e.g. Black and Devereux, 2011). These gradients between earnings and parents' socio-economic status (SES) could arise for two reasons. On the one hand, children from different SES face different economic opportunities, for example due to credit constraints or access to high quality schools. On the other hand, the preferences about what is valuable in a career, for example a high income or rather an interesting job, may differ with respect to SES. While economic research has long explored the differences in opportunities, we still know little about the relationship between SES and career preferences.

These preferences are important, as recent work has emphasized that education and career decisions are not only taken on the basis of expected financial returns, but also critically depend on the non-monetary aspects associated with these choices (Heckman et al., 2006). For example, adolescents with high fertility desires sort into occupations where work and family commitments are more compatible (Adda et al., 2017; Keane and Wolpin, 2010; Erosa et al., 2022). Important career choices such as whether to attend college, whether to relocate to other labor markets, or which occupation to choose thus depend on career preferences that entail more than just monetary aspects.

In this paper, I contribute to our understanding of the role played by career preferences in the intergenerational persistence of socio-economic status. Using panel data from the German Socio-Economic Panel (SOEP) and the British Cohort Study (BCS), I combine information on parental background with an array of questions on both general preferences and goals in life over various outcome domains when the respondents are 16 to 17 years old. These preferences include diverse career aspects, comprising among others the desire to achieve a high income, the desire to help others, health and safety conditions, job security, chances

4 Career Preferences and Socio-Economic Background

of promotion, family time, and the perceived importance of a job for society. I subsequently demonstrate that these preferences predict future outcomes in the labor market in both Germany and the UK, and document their correlation with socio-economic status.

Since identifying career preferences from realized job choices is difficult, researchers typically have to rely on stated preferences.¹ In recent work, hypothetical choice questions have been used to elicit stated preferences, in particular to quantify to what extent gender differences in preferences can explain the gender wage gap (e.g. Wiswall and Zafar, 2018; Burbano et al., 2020; Valet et al., 2021). An advantage of hypothetical choice surveys is the possibility of asking targeted questions about both preferences and beliefs, and estimating the willingness to pay for certain job attributes. However, studies relying on this methodology cannot always link preferences to actual choices, which would constitute an important limitation in the context of my research questions.² In the SOEP and the BCS, in contrast, I observe individuals for 12 to up to 30 years after having reported their preferences. This enables me to connect career preferences to actual choices in the education system and the labor market.

A second advantage of the household survey data used in this paper is the rich information on individual characteristics and parental background. In both data sets, I observe parental income, education and occupation, measures of ability like school grades or test scores, and beliefs about success in the labor market. In the SOEP, I can additionally measure personality traits, trust, and risk preferences. I can therefore condition on a rich vector of covariates when relating career preferences to both SES and labor market outcomes.

I present three main findings. First, I show that most career preferences exhibit a significant correlation with parental income, education, and occupation. For example, high SES children report that they place less value on income and job security, but more on having an interesting job than low SES children. While I cannot directly observe parental career preferences in my data, I find that related values and goals in life are intergenerationally very persistent. This suggests that career preferences are, at least to some extent, directly transmitted across generations.³

¹ As discussed in Wiswall and Zafar (2018), isolating occupational preferences from job choice is very challenging, because the equilibrium matching of jobs to workers reflects the preferences of both workers and firms. Furthermore, labor market frictions additionally blur the relationship between choices and preferences.

² An exception is Wiswall and Zafar (2018), who observe in a follow-up survey realized earnings for a subset of respondents. Boneva and Rauh (2017), for example, predict college choice using a random forest, whereas Burbano et al. (2020) observe neither occupational choice nor earnings.

³ For the decision whether to attend college, Müller (2021) provides the first evidence that parental preferences causally affect their children's college aspirations. Relatedly, Altmejd (2023) shows that the choice of field of study is causally determined by the parental field of study, presumably because parents serve as role models

Second, I show that career preferences elicited at age 16 or 17 are highly predictive of the labor market outcomes of these children when they are 28 to 46 years old. For example, a one standard deviation increase in the preference for a career with good chances of promotion is associated with a two percentile ranks higher labor income in the UK, and a six percentile ranks higher labor income in Germany. In turn, children who stated that it is important for them to help others in their career have significantly lower earnings in midlife than their peers who placed less value on this domain. These correlations hold both unconditionally and when controlling for an extensive set of personal characteristics, including measures of ability, the Big Five personality traits, and trust and risk preferences. I likewise find that individuals who had placed more value on chances of promotion also work more hours on average, whereas individuals who place more value on family time work fewer hours.

Third, since career preferences correlate with both own and parental income, they are a potential channel for intergenerational income transmission. In my data, I find that statistics of intergenerational income mobility—the most popular measures of equality of opportunity—are reduced by 7 to 28 percent when controlling for career preferences. This is because children from low-income families are on average more likely to value career aspects associated with low earnings. While these findings are descriptive, they nevertheless suggest that intergenerational income mobility could potentially be increased by targeting the career preferences of children from disadvantaged backgrounds. In contrast to cognitive ability, which already manifests in early childhood (Cunha et al., 2006), career preferences are still malleable during adolescence and are thus more easily influenced later in life. At least, this is the hope of policymakers, who in the US for example have mandated career guidance in school for more than a century (Gysbers, 2005), or tried to increase preferences for STEM enrollment for women, ethnic minorities and low SES children (Seymour and Hewitt, 1997; Best et al., 2013; Arcidiacono et al., 2016; Hill, 2017; De Philippis, 2021).

While the measurement of career preferences, incomes of children and parents, and the observation period differ between the SOEP and the BCS, the relationships between career preferences and own and parental income are surprisingly similar in both samples. The patterns documented in this paper are therefore likely to have external validity beyond the respective country and the cohorts under consideration.

when choosing a field. Likely,—but this has not been shown empirically so far—parental preferences also causally affect the career preferences of their children. Zumbuehl et al. (2020) suggest parental involvement as one potential channel for preference transmission.

4 Career Preferences and Socio-Economic Background

By providing a comprehensive account of the nexus between career preferences, own labor market outcomes, and parental background, this paper makes three contributions. First, to the best of my knowledge, it is the first paper to systematically document how career preferences vary with parental background. A small but growing set of studies shows that preferences such as patience (Kosse and Pfeiffer, 2012, 2013), risk-taking (Dohmen et al., 2012; Alan et al., 2017), or general social preferences (Kosse et al., 2020; Attanasio et al., 2020) are persistent across generations and correlated with socio-economic status. I complement this literature by showing that in addition to these very general preferences, also the far more specific career preferences differ significantly by parental income and other markers of socioeconomic status.

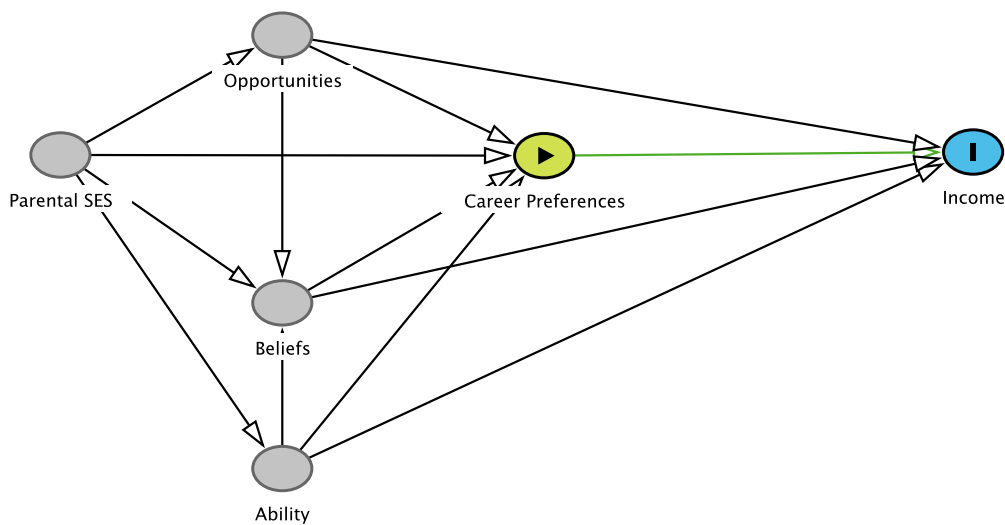
Second, by showing that career preferences predict future earnings, I complement studies demonstrating that basic preferences like trust, patience, and risk aversion are correlated with future economic outcomes (e.g. Barsky et al., 1997; Bonin et al., 2007; Dohmen et al., 2009; Algan and Cahuc, 2010; Fouarge et al., 2014; Golsteyn et al., 2014; Sunde et al., 2022). Career preferences still have substantial additional predictive power for earnings when controlling for these basic preferences. This is encouraging for policymakers wanting to improve the outcomes of children from disadvantaged households, as it may be easier to alter career preferences than changing patience, trust, or risk aversion.

Third, by testing how conditioning on career preferences alters estimates of intergenerational income mobility, I address a recent literature which emphasizes the role of preferences for explaining labor market inequalities. So far, this literature has mainly focused on gender gaps (Mas and Pallais, 2017; Wiswall and Zafar, 2018; Burbano et al., 2020; Valet et al., 2021), and income differences have been predominantly measured using expected as opposed to realized incomes. For differences by parental income and other markers of SES, very little evidence is available, as studies decomposing the intergenerational elasticity (e.g. Bolt et al., 2021) typically do not model or observe preferences explicitly. The most closely related study in this domain is Boar and Lashkari (2021), who show that the children of richer US parents do not only enjoy higher incomes, but also more favorable work conditions. Through the lens of a model of occupational choice, they conclude that at least part of this gradient may be attributed to preference heterogeneity. I complement these insights empirically, by directly observing preferences and occupational choice for the same individuals.

4.2 Conceptual Framework

The main aim of the paper is to provide descriptive evidence on the correlation between career preferences, future labor market outcomes and parental SES. Nevertheless, it is useful to consider under what conditions a regression coefficient of today's income on past career preferences could in theory be interpreted as a causal effect. Figure 4.1 represents a simple causal model of the interaction between parental SES, career preferences and own income.

Figure 4.1: Causal Model



Notes: This figure shows a directed acyclical graph (DAG) of how career preferences are causally connected to own income later in life, and various other personal characteristics. Arrows indicate the direction of causality.

The framework is very stylized in the sense that intermediate outcomes such as educational attainment are not explicitly modelled. Income in this model depends on four factors: ability, opportunities, beliefs and preferences. Ability should be understood as the genetic predisposition for achieving a high income. Opportunities are very broadly defined as all environmental factors influencing income, for example credit constraints, school and peer quality, parental connections to employers, and asymmetric information about the labor market and career pathways. Both ability and opportunities will influence career preferences. For example, only individuals with a sufficiently high IQ would enjoy embarking upon a career as a chess player or software engineer. In terms of opportunities, children will on average be more likely to become physicians if they know that they have the opportunity to take over their parents' doctor's office. As a fourth factor, parental SES, which is a strong predictor of both ability and opportunities, will also directly influence children's career preferences. If the parents

4 Career Preferences and Socio-Economic Background

are physicians, their children may themselves perceive the job of a physician to be more appealing. Parental SES thus comprises the whole socio-economic situation of the parental household, including parental income, education and occupation, but also parental values and preferences. Controlling for opportunities, ability and beliefs, parental SES in this model does not directly influence the income of the children.

As always, the ideal experiment to estimate the effect of career preferences on income would be to randomly assign different preferences to different individuals. This is fundamentally not possible.⁴ Instead, researchers must invoke the conditional independence assumption $Y_i \perp P_i | X_i$, arguing that the distribution of preferences P_i is independent of potential future earnings in the labor market Y_i when controlling for a covariate vector X_i , comprising opportunities, ability, and beliefs. Assuming that the empirical conditional expectation function of income with respect to X_i is well approximated by a linear relationship, Figure 4.1 can be represented by the following equation

$$y_{i,t+1} = \beta_0 + \beta_1 \text{Preferences}_{i,t} + \beta_2 \text{Ability}_{i,t} + \beta_3 \text{Opportunities}_{i,t} + \beta_4 \text{Beliefs}_{i,t} + \varepsilon_{i,t}, \quad (4.1)$$

where β_1 identifies the causal impact of a variation in preferences on realized income $y_{i,t+1}$ if there are no omitted variables that correlate with both preferences and the error term after controlling for ability, opportunities, and beliefs. t subscripts indicate that preferences, ability, and opportunities are measured at the time individuals take their career decisions as they reach adulthood, whereas income is measured later in life. Without this important restriction, we would also have to worry about reverse causality, i.e. income influencing preferences.

In general, it is not plausible that any given set of controls will achieve full conditional independence. Studies relying on specifically designed surveys to elicit career preferences have more scope to control for beliefs, but typically cannot approximate the other two factors very well. In contrast, the panel data used in my analysis have less information on beliefs, but provide an exceptionally in-depth set of covariates to capture both opportunities and ability.

⁴ While interventions can and have been used to affect preferences (e.g. Abeler et al., 2021), it is very difficult to vary a certain preference without altering other preferences or beliefs. Likewise, it is hard to think of a valid IV strategy here, as the instrument would have to be correlated with preferences but uncorrelated with both ability and opportunities.

4.3 Data

My analysis requires data that make it possible to observe both own and parental income, and a reliable measure of career preferences. To the best of my knowledge, the only publicly available data sets that meet all three requirements are the German Socio-Economic Panel (SOEP) and the British Cohort Study (BCS). Although the elicitation of career preferences is not harmonized between both surveys, the respective questions are very similar. Hence, the use of both data sets allows me to test the robustness of my results across two different samples. In addition, both data sets have specific benefits, which I exploit in the analysis. While the measurement of parental background and the coverage of important covariates such as risk and trust preferences are superior in the SOEP, the BCS presents the advantage of a larger sample size and a longer observation period in adulthood.

Below I describe how career preferences and the incomes of children and parents are measured in both surveys. Appendix D.1 provides more details and describes all further variables used in the analysis.


4.3.1 The German Socio-Economic Panel (SOEP)

Established in 1984, the SOEP is a nationally representative household panel survey of the German population. In its more recent waves, it samples around 15,000 German households or 25,000 individuals each year (Goebel et al., 2019). All children from sample households become regular members of the panel themselves once they reach the age of 18. One year in advance, at age 17, these children additionally answer a youth questionnaire, which was introduced in the year 2000. This questionnaire asks them about their current situation in the education or employment system, their values and attitudes, and their aspirations and goals for the future. In my analysis, I combine information from the youth questionnaire with the panel dimension of the SOEP's main questionnaire.

Career Preferences. Career preferences in the SOEP are surveyed with a single question that is posed annually to all participants in the youth questionnaire. In the English translation, this question reads: *Different things may be important to people when choosing a career. Please state how important each of the following is to you—very important, important, not so important, completely unimportant. How important for your career is....*, followed by a list of twelve different sentence endings, which capture the most important aspects in choosing

a career.⁵ In the following, I will refer to the answers to this question as career preferences. If respondents answered the career question but failed to rate up to three items, I impute missing values based on the ratings in all other categories using chained equations.⁶ This affects less than 2% of the respondents who answered any item of the question, and for only 0.2% more than one item is missing. To enable a better comparison of the estimates, I use standardized measures of career preferences and all other measures of traits, preferences, and values which are elicited on a Likert scale.

Figure 4.2: Correlation of Career Preferences



Notes: This figure shows a heatmap of the pairwise Pearson correlation coefficients between all career preferences. To account for the fact that some respondents are more likely to assign high importance to all career aspects, the correlations are adjusted by individual fixed effects computed by averaging over all preferences. The exact point estimates of the unconditional and conditional correlations are disclosed in Appendix Table D.1.


Figure 4.2 displays all twelve preferences and their correlations. For example, respondents who highly value income in their career are also more likely to value chances of promotion and

⁵ Appendix Figure D.2 displays the original question including all answer categories in German, along with the English translation.

⁶ The idea behind multiple imputation by chained equations is to impute multiple variables iteratively via a sequence of univariate imputation models (here: ordered logit), with fully conditional specifications of the prediction equations. See White et al. (2011) for a description of the method and Royston and White (2011) for more details on the Stata implementation.

leisure time. In contrast, they are less likely to value a job where they have the option of helping others, contacting others, or a job they consider to be important for society. Respondents who value having enough time for family commitments also consider leisure time and health and safety conditions to be more important. While the correlations between the different career aspects are meaningful, they are not very high on average and collinearity is low. This highlights that each question captures a distinct and genuine aspect of choosing a career. As such, the answers to the questions cannot be easily reduced via a principal component analysis (PCA), for which I find a Kaiser-Meyer-Olkin measure of sampling adequacy of 0.77.

Figure 4.3: Principal Component Analysis Career Preferences



Notes: This figure shows a scatter plot of the factor loadings of the first and second component of a principal component analysis (PCA) of all twelve career preferences.

Just descriptively, it is still interesting to see the factor loadings of each preference on the first two components of a PCA. As shown in Figure 4.3, the two components that explain most of the joint variance in career preferences are best described as referring to a contrast between extrinsic versus intrinsic career aspects, and to a dichotomy between more public/visible versus private features of a career. Preferences correlating positively, such as income/leisure or helping others/important for society, tend to display similar factor loadings in both components.

4 Career Preferences and Socio-Economic Background

On average, the most desired career aspects are an interesting job, job security, and health and safety conditions (compare Appendix Figure D.1). In contrast, high income and leisure—corresponding to the ingredients of the classical utility function in labor economics—are rated as comparatively less important.

Income, Education and Occupation. My main income measure is individual gross labor earnings. Labor earnings encompass wages and salaries from all employment, including training, primary and secondary jobs, and self-employment, plus income from bonuses, overtime, and profit-sharing. In order to also capture other income sources and the income flows of other household members, I additionally use gross and net household income in some specifications. A further advantage of using household income is that values of zero are the exception, which avoids problems when computing log incomes or income ranks. All income figures are CPI adjusted to the year 2016. Occupational status is measured in two-digit ISCO codes, education either as years in formal education and training or as a three-category variable, indicating whether someone has more, less or the equivalent of a high school degree (*Abitur*).

All variables are observed not only for the respondents answering the career preference question, but also for their parents. In my baseline estimates, I focus on parental income to capture children's socio-economic background, measured as the five-year average when children are between 15 and 19 years old. In addition to the central importance of parental income as a measure of SES, this choice allows me to directly link my results to the large literature on intergenerational income mobility. When I use parental occupation or education to measure socio-economic background instead, the correlation with career preferences is very similar.

Incomes of children are measured in the five-year interval from age 28 to 32. This ensures that (most) children have already left education and entered the labor market. The cut-off value of 28 is chosen so as to balance the tradeoff between sample size and lifecycle bias.⁷

4.3.2 The British Cohort Study

The British Cohort Study (BCS) is a cohort study of all children born in Great Britain in a week in April 1970. Information was obtained about the sample members and their families at birth

⁷ As the youth questionnaire was only introduced in 2000, the oldest respondents are 38 years old in the last survey wave of the SOEP.

and at ages 5, 10, 16, 26, 30, 34, 38, 42, 46 and 51. The cohort began with more than 17,000 children. At age 16, when career preferences are elicited, it was still possible to reach 11,620 individuals. Since then, attrition has been low, with 8,581 study members having participated in the age 46 survey.


Career Preferences. In the British Cohort Study, career preferences are elicited when respondents are 16 years old. Similar to the SOEP, children were asked to indicate which aspects would matter for them when choosing a career.⁸ While some answer categories directly correspond to the respective aspect in the SOEP (e.g. helping others), other important categories such as family commitments or importance for society are lacking in the BCS questionnaire. On the other hand, the BCS asks about some additional aspects such as “working with figures” or “working outside”. As shown in Appendix Figure D.3, respondents rated these additional categories to be of mostly minor importance in choosing a career. Instead, the most desired career aspects are an interesting job, an understanding boss, and long-term security, similar to the most highly rated preferences in the SOEP. Career preferences were again imputed using chained equations if respondents answered the career question but failed to rate up to three items, and standardized to have mean zero and standard deviation one.

Figure 4.4, Panel A, displays all 16 preferences and their correlation among each other. The patterns closely reflect those in the SOEP. For example, respondents that highly value income in their career are also more likely to value chances of promotion, but less likely to care about a job where they have the possibility to help others. Respondents who want a quiet life also tend to favor regular hours and not having to work too hard, but dislike a job with a real challenge.

In Panel B, I correlate these preferences with a second measure of career preferences available at age 16 in the BCS, the JIIG-Cal scores. JIIG-Cal (“Job Ideas and Information Generator - Computer-Assisted Learning”) scores are derived from a separate questionnaire where respondents make preference choices between 30 pairs of occupational activities, while at the same time they are asked whether they like each activity or not. The occupational activities presented are designed to fall into one of six types, which are not made explicit to the participants: scientific/practical, work with living things, business/commerce, artistic leanings, caring for people, and communication activities. An algorithm is then used to produce 0 to 100 scores for each of the types, where a higher score indicates a stronger preference for a

⁸ Appendix Figure D.5 provides a screenshot of the original questionnaire.

Figure 4.4: Correlation of Career Preferences in the BCS



Notes: This figure shows a heatmap of the pairwise Pearson correlation coefficients between all career preferences (Panel A) and between career preferences and JIIG-Cal scores (Panel B, see text for details).

respective job type. Career preferences correlate with these scores in a meaningful way. For example, the desire to find a job where one is able to help others is a strong predictor for the “caring for people” type, and individuals who desire a high income are more likely to be categorized into enjoying jobs of the “business/commerce” type. As the directly elicited career preferences, JIIG-Cal scores are highly predictive for future labor market outcomes.

Income, Education and Occupation. Income among children is defined as gross pay from both employment and self-employment. I use observations at ages 30, 34, 38, 42 and 46.⁹ For individuals with missing wages but known employment status, I assign a wage of zero if the individual was out of work at that time. For individuals with missing information in some of the five observation ages and unknown employment status, I use the row and column imputation procedure also employed to impute incomes in the SOEP (Frick and Grabka, 2014). Incomes at different ages are CPI adjusted and averaged over all years. I finally assign each child its percentile ranks among all other children.

Parental income in the BCS is elicited in the third and fourth wave when children are ten and 16 years old. Parents report gross weekly family income in seven bands when children are ten, and in eleven bands when children are 16. Following Gregg et al. (2017), I fit a Singh-Maddala distribution in each wave to derive a continuous measure of income from the banded data. For parents who reported income in both waves (53%), I average incomes over both periods. For parents with missing incomes in one period, I impute income missing values based on income in the other period, age of the mother, and changes in social class, employment status, housing tenure and lone parent status across the two periods, again following Gregg et al. (2017). Based on this continuous measure of income, I assign parents their percentile rank in the income distribution comprising all other parents.

Occupational status is measured in two-digit SOC codes, education either as years in formal education and training or as a four-category variable, indicating whether someone has an educational qualification below/equal to an O-Level degree, or equal to/higher than an A-Level degree.

⁹ At age 26, only net incomes are reported, and the question was asked in terms of hourly wages, which produced a high number of implausible outliers. I also omit the age 51 wave, which was administered in 2021, still during the wake of the COVID-19 pandemic.

4.4 Empirical Evidence

This section presents empirical results in three steps: First, I document how career preferences differ by parental SES, before then demonstrating that children's preferences and those of their parents are positively and significantly correlated. Third, I present evidence on the predictive power of career preferences for future earnings, work hours, and education.

4.4.1 Career Preferences Vary with Parental Background

Starting with the German data from the SOEP, Figure 4.5 shows that most career preferences vary significantly by parental background. The first estimate (circle) in each category is obtained by regressing the respective preference on the parental labor income rank, the second estimate (diamond) by regressing it on the parental net income rank, and for the third estimate (triangle), I additionally control for all other career preferences. In all specifications, the importance children assign to job security, achieving a high income, pursuing a respected and recognized career, working in a job that is important for society, and the possibility of helping others in the job decreases in parental income rank. In contrast, children from high income households are more likely to value pursuing an interesting career and having enough time for family commitments.

Many of these correlations seem intuitive. For example, as children from low-income families have fewer parental resources available, it makes sense for them to place more value on job security and achieving a (relatively) high income. Children from high-income households, who on average are more likely to fall back into the parental safety net in times of crises, can afford to place more value on “softer” career aspects such as how interesting they consider a job, or to have a job which is well compatible with family commitments. Interestingly, though, this line of reasoning does not apply to two other soft factors: the wish to work in jobs that are important for society and where one is able to help others. While parental income is just one possible proxy of socio-economic status, I document in Appendix Figures D.6 and D.7 that the correlations between career preferences and SES are very similar when using parental years of education or the parental average occupational prestige score to measure SES instead.

Figure 4.6 shows how the career preferences reported by the respondents in the BCS differ with parental income. As their German counterparts, British adolescents from high-income families are less likely to value helping others and a high income, and more likely to value an interesting job. Unlike the German results, it also becomes apparent that the willingness

Figure 4.5: Variation in Career Preferences by Parental Income Rank

Notes: This figure shows point estimates and the corresponding 95% confidence intervals (robust standard errors) of separate regressions of one of the twelve career preferences on parental income rank. The first estimate (circle) in each category is obtained by regressing the respective preference on the percentile rank in labor earnings of the father, the second estimate (diamond) by regressing it on the parental gross household income rank, and for the third estimate (triangle), I additionally control for all other eleven career preferences. A coefficient of 0.002 for example implies that a 10 percentile higher parental income rank is associated with an increase of 2 percent of a standard deviation in the respective preference. The sample consists of 8,185 parent-child pairs.

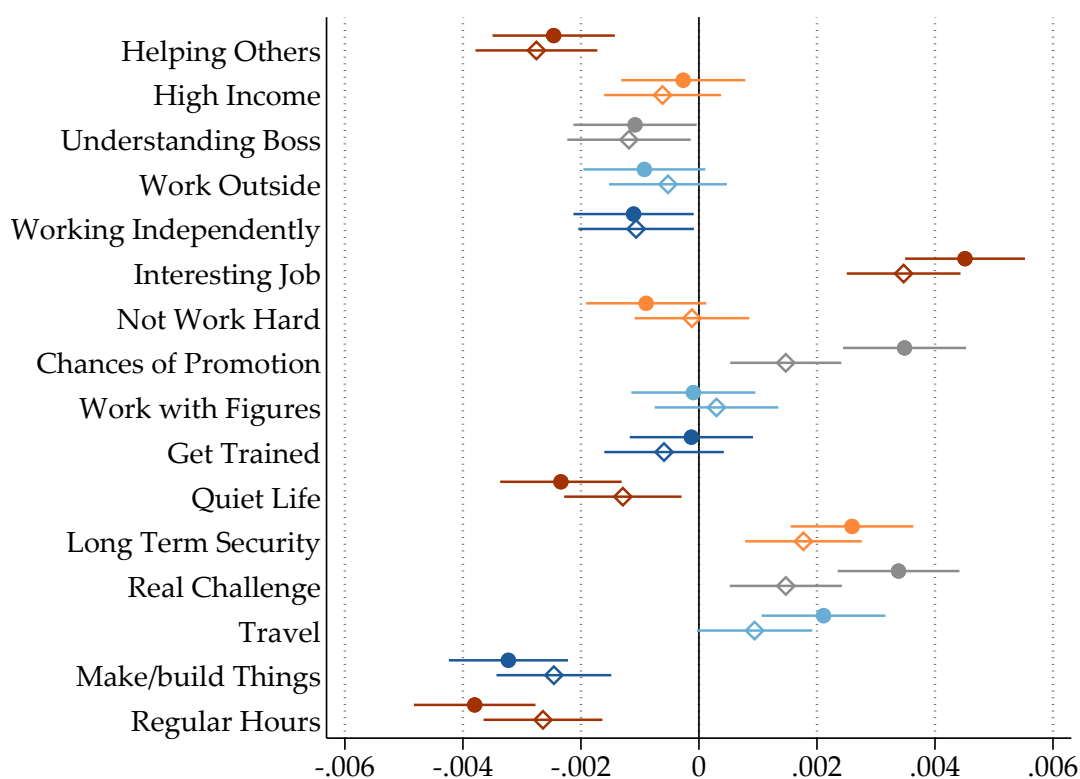
Source: German Socio-Economic Panel (SOEP)

to exert effort in a career increase with parental income: children from high-income families place more emphasis on chances of promotion and taking a real challenge, but less emphasis on not working hard, working regular hours, or leading a quiet life.

4.4.2 Preferences Are Intergenerationally Persistent

I also find that—not surprisingly—preferences and values are very persistent across generations. While I cannot test this for career preferences directly, as this would require an even

Figure 4.6: Variation in Career Preferences by Parental Income Rank in the British Cohort Study



Notes: This figure shows point estimates and the corresponding 95% confidence intervals (robust standard errors) of separate regressions of one of the 16 career preferences on parental income rank. The first estimate (circle) in each category is obtained by regressing the respective preference on the mean parental income rank, the second estimate (diamond) by additionally controlling for all other eleven career preferences. The sample consists of 4,365 parent-child pairs.

Source: British Cohort Study (BCS)

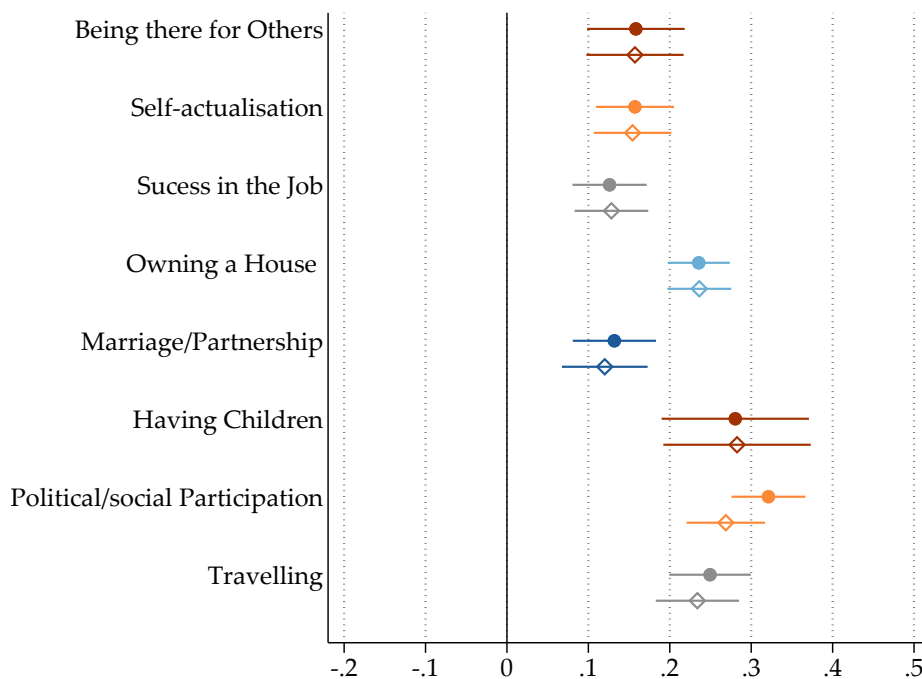
longer panel dimension, I can evaluate the answers to a closely related question in the SOEP, asking about the importance of several life aspects in a related manner, which respondents assess on a four-point Likert scale.¹⁰

Figure 4.7 shows that preferences across all of those life domains are strongly correlated between children and their parents. While the first estimate (circle) displays the unconditional intergenerational correlations, the second estimate additionally controls for parental income rank and parental years of education. For example, individuals whose parents value caring for

¹⁰ As career preferences are only elicited in the SOEP youth questionnaire which became available in 2000, I do not observe them for the majority of parents. This is different for the importance question, which is asked in the main questionnaire of the SOEP. In the BCS, parents never report career preferences or similar values.

others are significantly more likely to hold this aspect in high regard themselves. Likewise, a one standard deviation increase in the parental assessment of success in the job is associated with a 0.12 higher standard deviation of this measure among their children. These correlations are essentially unaffected by controlling for parental income and education. This indicates that they reflect a direct transmission of preferences and values from parents to their children, as opposed to a joint association with omitted variables. For example, if having a low income were to induce individuals to adjust their preferences away from “expensive” values such as owning a house or travelling, these values would be intergenerationally persistent to the extent that income is persistent across generations, too. However, as accounting for parental income barely changes the estimates, I can rule out that it constitutes a relevant omitted variable in these regressions. Appendix Figure D.8 provides additional evidence, by plotting the correlations separately for each quintile of parental income.

Figure 4.7: Intergenerational Persistence in Preferences and Values



Notes: This figure shows point estimates and corresponding 95% confidence intervals (robust standard errors) of separate regressions of eight different preferences and values on the respective measure of their parents, based on 4,474 parent-child pairs. The first estimate (circle) in each category is obtained by regressing the respective value of the child on the average value of the parents, the second estimate (diamond) is obtained by additionally controlling for the percentile rank in the parental gross household income distribution and parental years of education of the more educated parent. The values of the children are measured when children are between 16-25 years old, the values of the parents when children are between 15-19 years old.

Source: German Socio-Economic Panel (SOEP)

Overall, the positive and significant correlations in Figure 4.7 corroborate the conjecture that the related career preferences are at least partially directly transmitted from parents to their children.

4.4.3 Career Preferences Predict Future Earnings

To demonstrate that career preferences elicited in adolescence convey meaningful information about the economic outcomes of these children, I next document their relation to realized earnings. Here, the timing structure is crucial. Contemporaneous correlations between preferences and outcomes are difficult to interpret, as economic conditions are expected to influence preferences (e.g. Doepke and Zilibotti, 2017), and causality can run in both directions. To study the role of career preferences, we would ideally like to measure preferences at the time when career decisions are taken, and economic outcomes in midlife around the ages 30 to 50, where annual incomes provide the best approximation of lifetime incomes.¹¹ The panel structure of the SOEP and the BCS enable me to closely adhere to this ideal setting. Career preferences are elicited at ages 16 and 17, just before children legally become adults and start moving out of the parental household (Dodin et al., 2024).¹² Furthermore, most children at this age have not yet completed secondary education and are about to decide on their career paths. Data on earnings are observed in the age range 30 to 46 in the BCS and in the age 28 to 32 in the SOEP. In the latter case, this is younger than the ideal time span, but still an age range where incomes well approximate lifetime income.¹³ As such, both data sets are highly suited to studying the association between career preferences and future earnings.

Table 4.1 shows the results in the SOEP. In the first column, I regress the own income rank on all twelve career preferences jointly.¹⁴ The first coefficient of -0.63 for job security means that,

¹¹ Due to heterogeneity in life cycle earnings profiles, estimates obtained when children are young (old) tend to be downward (upward) biased (Haider and Solon, 2006; Nybom and Stuhler, 2016). This lifecycle bias is smallest when measuring income in midlife. Nybom and Stuhler (2016) also show that lifecycle bias is less of an issue if incomes are measured in ranks, as done in this paper.

¹² Furlong and Biggart (1999) show that in the UK, occupational aspirations are quite stable between the ages of 13 to 16, indicating that the exact time of measurement is of minor importance.

¹³ Without limiting my sample to individuals who answered the youth questionnaire, I can directly test the correlation between incomes measured at different points in the lifecycle. For the age range 28 to 32, I find a rank correlation of 0.79 with lifetime income (defined as average income in the age range 18 to 65 for all respondents with at least 30 years of data in this age range). While the main reason for not choosing an older age range is data availability, focusing on the age range 28 to 32 also affords the advantage of comparability with other estimates in the literature. Following the influential study of Chetty et al. (2014a), further empirical estimates of intergenerational income mobility have been obtained for children in their early 30s (e.g. Chuard and Grassi, 2020; Helsø, 2021).

¹⁴ Appendix Figure D.9 shows that the unconditional correlations tend to be similar.

4 Career Preferences and Socio-Economic Background

Table 4.1: Labor Earnings and Career Preferences in the SOEP

	Percentile Rank in Individual Labor Earnings					
	(1)	(2)	(3)	(4)	(5)	(6)
Job Security	-0.63 (1.70)	-0.58 (1.64)	-0.86 (1.63)	-0.32 (1.56)	-0.06 (1.52)	-0.76 (1.47)
High Income	-3.18** (1.60)	-2.29 (1.60)	-2.19 (1.60)	-1.19 (1.55)	-1.08 (1.57)	-0.71 (1.49)
Chances of Promotion	6.86*** (1.51)	6.11*** (1.48)	6.01*** (1.47)	6.03*** (1.50)	5.56*** (1.49)	4.89*** (1.48)
Respect and Recognition	-0.48 (1.52)	0.15 (1.46)	0.05 (1.46)	0.09 (1.47)	0.51 (1.48)	0.84 (1.45)
Leisure	1.08 (1.26)	0.44 (1.25)	0.43 (1.25)	0.70 (1.27)	1.00 (1.25)	1.24 (1.26)
Interesting Job	0.83 (1.60)	0.44 (1.58)	0.55 (1.57)	0.47 (1.39)	0.47 (1.41)	0.01 (1.38)
Working Independently	1.00 (1.58)	1.04 (1.57)	0.86 (1.57)	0.06 (1.47)	-0.39 (1.49)	-0.74 (1.41)
Contact to Others	-2.10 (1.83)	-1.99 (1.78)	-1.79 (1.77)	-1.97 (1.63)	-2.46 (1.61)	-2.13 (1.59)
Important for Society	2.31 (1.87)	2.47 (1.83)	2.33 (1.83)	2.02 (1.68)	1.40 (1.66)	1.32 (1.55)
Health and Safety Conditions	0.35 (1.64)	0.47 (1.53)	0.52 (1.53)	-0.33 (1.53)	-0.60 (1.52)	-0.55 (1.46)
Family Time	2.26 (1.49)	1.62 (1.47)	1.71 (1.46)	1.83 (1.46)	1.81 (1.38)	2.01 (1.33)
Helping Others	-5.18*** (1.63)	-4.90*** (1.55)	-5.03*** (1.54)	-5.82*** (1.60)	-5.01*** (1.61)	-4.96*** (1.59)
Parental Income Rank		0.20*** (0.04)	0.23*** (0.05)	0.21*** (0.05)	0.18*** (0.05)	0.17*** (0.05)
Parental Years of Education			-0.84 (0.58)	-0.77 (0.57)	-1.02* (0.60)	-0.77 (0.60)
Probability Desired Career				1.55** (0.70)	1.31* (0.68)	1.17* (0.68)
Gender, State of Birth	-	-	-	✓	✓	✓
Grades, Tracking Recommendation	-	-	-	-	✓	✓
Trust/Risk Preferences, Big Five	-	-	-	-	-	✓
Observations	787	787	787	787	787	787
R^2_{adj}	0.089	0.128	0.131	0.184	0.214	0.231

Notes: This table shows estimates of six separate regressions of the percentile rank in own labor income for children aged 28 to 32 on past career preferences reported at age 17. Parental income rank refers to gross household income, parental education to years of education of the more educated parent. Column (4) additionally includes gender and dummies for the state of birth, Column (5) dummies for the recommended school track after primary school, the grade average, and interactions between the grade average and the track recommendation. In Column (6), I further control for trust and risk preferences, and measures of the Big Five personality traits. Robust standard errors in parentheses, * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. *Source:* German Socio-Economic Panel (SOEP)

4 Career Preferences and Socio-Economic Background

conditional on all other career preferences, a one standard deviation increase in the preference for job security in one's career is associated with a 0.6 percentile lower income rank when individuals are around thirty years old. Below average earnings are also obtained by children who greatly valued achieving a high income, contact to others, and helping others when choosing a career. Conversely, a one standard deviation increase in the desire to encounter good chances of promotion corresponds to a six percentile ranks higher income.

While only the coefficients for chances of promotion and helping others are statistically significantly different from zero throughout all specifications, most of the aforementioned associations are robust to including a comprehensive set of control variables and economically large. From Column (2) for example, where I add the parental income rank as control, it can be deduced that a one standard deviation increase in the desire for helping others has the same effect as a 25 percentile drop in parental income rank. Furthermore, the associations between career preferences and future earnings rank remain quantitatively similar. One main exception is the coefficient for the high-income preference, which is substantially reduced when conditioning on parental income. Apparently, the initially negative and non-intuitive correlation between the desire to earn a high income and realized income later on is partially driven by the fact that children from low-income households are more likely to desire high earnings, yet less likely to achieve this later in life. Furthermore, Appendix Figure D.9 shows that the unconditional correlation between the income preference and future earnings is slightly positive, with the difference being mostly driven by not conditioning on chances of promotion.

In Columns (3)-(5), I add the highest number of years of schooling that one of the parents has achieved, dummies for gender and state of birth, and school grades as a measure of ability. Furthermore, I control for respondent's subjective beliefs of how likely they are to achieve their desired career. While the single coefficient estimates of course vary by specification, the general patterns remain remarkably stable. Overall, the results indicate a persistent link between career preferences elicited at age 17 and earnings ten to 15 years later. In the last column, I additionally control for trust, risk aversion and measures of the Big Five personality traits. The results demonstrate that career preferences continue to predict future labor market earnings, even conditional on these basic economic preferences often emphasized in the literature.¹⁵

¹⁵ Note also that in my data, career preferences explain a substantially larger share of the variation in earnings than the measures of trust, risk and the Big Five personality traits combined ($R^2 = 0.103$, $R^2_{\text{adj}} = 0.089$ versus $R^2 = 0.046$, $R^2_{\text{adj}} = 0.037$).

On top of predicting earnings, career preferences contain information about future work hours (Appendix Table D.3). For example, a one standard deviation increase in the job security or family time preference corresponds to around 20 fewer annual hours worked at ages 28 to 32. The strongest predictors for working more are the preferences for chances of promotion and working independently, where a one standard deviation increase is associated with up to 80 additional annual work hours. This corresponds to a three percentile increase in the annual work hours rank. Overall though, compared to earnings, variation in work hours is higher and more idiosyncratic, and few of these correlations are statistically different from zero. Likewise, the predictive power of career preferences for educational attainment is weaker compared to earnings (Appendix Table D.4). The most robust and significant correlations which arise are that a high preference for an interesting job is associated with more years of education, whereas individuals who valued health and safety conditions and helping others at age 17 had obtained two to four months less of education or training in their early 30s.

In this context, it is interesting to ask which of these factors mediate the impact of preferences on income. To touch on this, I run a horse-race regression, where I re-estimate Column (6) of Table 4.1, additionally controlling one-by-one for four levels of education, occupation fixed effect at the two-digit ISCO level, and annual work hours.¹⁶ The results are shown in Table 4.2. For the purpose of presentation, only the estimates for the two most predictive preferences are displayed. Comparing Columns (1) and (2), educational attainment seems to matter little in the relationship between chances of promotion and future labor market outcomes. In contrast, the coefficient drops from 4.59 to 2.99 when including occupational fixed effects. This suggests that occupational choice is one important channel through which career preferences translate into future earnings. At least equally important though are annual work hours (Column 4): presumably, one reason why individuals who place great value on chances of promotion ultimately earn comparatively more is that they will also work more hours. For helping others, the education channel seems to be of greater relevance, but again not as important as occupational choice or hours worked. Jointly, education, occupation, and work hours explain half of the size of each coefficient. However, it is likely that I would be able to explain even more in a larger sample that would facilitate the use of a more fine-grained occupation classification or to account for non-linearities in work hours due to part-time penalties or cutoff thresholds in the tax code. However, additional factors will play a role as well. For example, within an occupation, individuals with a strong desire for promotion might be more likely to aim for better-paying management positions.

¹⁶ Due to the categorical nature of the occupation variable, I refrain here from a more formal mediation analysis, as is presented later in Section 4.5.

Table 4.2: Controlling for Mediating Factors

	Percentile Rank in Individual Labor Earnings				
	(1)	(2)	(3)	(4)	(5)
Chances of Promotion	4.59*** (1.54)	4.68*** (1.53)	2.99** (1.39)	2.87** (1.11)	2.37** (1.05)
Helping Others	-5.68*** (1.71)	-4.83*** (1.65)	-4.56*** (1.64)	-4.50*** (1.21)	-3.28*** (1.09)
Full Set of Table 4.1 Controls	✓	✓	✓	✓	✓
Education FE	-	✓	-	-	✓
Occupation FE	-	-	✓	-	✓
Annual Work Hours	-	-	-	✓	✓
Observations	702	702	702	702	702
R^2_{adj}	0.243	0.279	0.465	0.620	0.672

Notes: This table shows estimates of six separate regressions of the percentile rank in own labor income for children aged 28 to 32 on past career preferences reported at age 17. In all columns, I control for parental gross household income rank, years of education of the more educated parent, gender, dummies for the state of birth, dummies for the recommended school track after primary school, the grade average, interactions between the grade average and the track recommendation, trust and risk preferences, and measures of the Big Five personality traits. Furthermore, I control for all ten other career preferences not displayed in the table. Education fixed effects refer to four categories (below/equal to an O-Level degree, or equal to/higher than an A-Level degree), occupation fixed effect are at the 2-digit ISCO level. Robust standard errors in parentheses, * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. *Source:* German Socio-Economic Panel (SOEP)

I now turn to the evidence from the UK. Table 4.3 replicates Table 4.1 by showing how career preferences as elicited in the BCS at age 16 predict the own earnings rank during adulthood, measured between the ages 30 to 46. The results are very similar: individuals who had reported placing more value on chances of promotion earn significantly more during adulthood, and the desire to help others is associated with a lower earnings rank. Further significant predictors of future earnings in Column (1) are the preferences for a high income, working with figures, long term security, and facing a real challenge (positive), as well as working independently and working regular hours (negative). The fact that the preference for a high income now clearly predicts a higher income during adulthood likely reflects that in the BCS, there exists no association between *parental* income and this preference. Similar to the SOEP, career preferences explain around 9% of the variation in future earnings rank.

Table 4.3: Labor Earnings and Career Preferences in the BCS

	Percentile Rank in Individual Labor Earnings					
	(1)	(2)	(3)	(4)	(5)	(6)
Helping Others	-2.83*** (0.50)	-2.38*** (0.50)	-2.32*** (0.49)	-0.58 (0.47)	-0.01 (0.47)	-0.06 (0.47)
High Income	1.25** (0.53)	1.38*** (0.52)	1.44*** (0.52)	0.48 (0.48)	0.73 (0.48)	0.83* (0.48)
Understanding Boss	-0.59 (0.51)	-0.40 (0.50)	-0.36 (0.50)	0.03 (0.47)	0.14 (0.47)	0.22 (0.46)
Working Outside	-0.18 (0.54)	-0.10 (0.53)	-0.09 (0.53)	-1.58*** (0.51)	-1.18** (0.50)	-1.03** (0.50)
Working Independently	-1.45*** (0.56)	-1.29** (0.55)	-1.23** (0.54)	-1.04** (0.52)	-0.74 (0.51)	-0.61 (0.51)
Interesting Job	0.46 (0.54)	-0.16 (0.53)	-0.31 (0.53)	0.51 (0.50)	-0.12 (0.49)	-0.32 (0.49)
Not Work Hard	-0.39 (0.53)	-0.44 (0.52)	-0.53 (0.52)	-0.99** (0.49)	-0.78 (0.48)	-0.71 (0.48)
Chances of Promotion	2.56*** (0.57)	2.37*** (0.56)	2.47*** (0.56)	2.22*** (0.52)	1.86*** (0.51)	1.63*** (0.51)
Work with Figures	3.80*** (0.51)	3.70*** (0.50)	3.69*** (0.50)	2.77*** (0.48)	2.49*** (0.47)	2.40*** (0.47)
Get Trained	0.16 (0.51)	0.21 (0.50)	0.23 (0.49)	0.24 (0.46)	0.26 (0.46)	0.16 (0.46)
Quiet Life	-0.17 (0.52)	-0.03 (0.51)	-0.03 (0.51)	-1.15** (0.48)	-0.90* (0.47)	-0.76 (0.47)
Long Term Security	2.50*** (0.53)	2.21*** (0.52)	2.20*** (0.52)	1.07** (0.50)	0.70 (0.49)	0.56 (0.49)
Real Challenge	1.91*** (0.56)	1.58*** (0.55)	1.46*** (0.55)	1.55*** (0.51)	1.31*** (0.51)	0.96* (0.51)
Travel	-0.61 (0.53)	-0.67 (0.52)	-0.67 (0.52)	0.41 (0.49)	0.44 (0.48)	0.39 (0.48)
Make/build Things	-1.02* (0.55)	-0.60 (0.54)	-0.59 (0.53)	-2.67*** (0.51)	-2.36*** (0.51)	-1.94*** (0.50)
Regular Hours	-4.39*** (0.52)	-3.99*** (0.51)	-3.78*** (0.51)	-3.25*** (0.48)	-2.40*** (0.48)	-2.01*** (0.49)
Parental Income	-	✓	✓	✓	✓	✓
Parental Education	-	-	✓	✓	✓	✓
Gender, Country of Birth	-	-	-	✓	✓	✓
Test Scores Age 10	-	-	-	-	✓	✓
No. Career Talks, Beliefs	-	-	-	-	-	✓
Observations	3165	3165	3165	3165	3165	3165
R^2_{adj}	0.082	0.111	0.116	0.225	0.250	0.259

Notes: This table shows estimates of six separate regressions of the percentile rank in own labor income for children aged 30-46 on past career preferences reported at age 16. Parental income rank refers to gross household income, parental education indicates if father and/or mother have an A-level degree. Column (5) adds age 10 test scores for language, reading, math and matrices. In Column (6), I further control for the number of attended career talks by age 16, and a set of further beliefs on what helps in advancing careers. Robust SEs in parentheses, * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Source: British Cohort Study.

Whereas most estimates remain fairly robust to the inclusion of a comprehensive set of controls in Columns (2)-(6), some coefficients change with the inclusion of a gender dummy in Column (4). In particular, the negative correlation between helping others and earnings vanishes, as females are both more likely to highly rate this aspect and earn less on average. Other estimates only become significant when controlling for gender: working outside and making/building things are now aspects that negatively predict future earnings. As in the German data, I find suggestive evidence that an important channel through which career preferences translate into differences in future labor market outcomes is occupational choice. In contrast, hours worked are comparatively less important.

4.5 Career Preferences and Income Mobility

The previous section documented that career preferences predict future earnings and are correlated with parental income. As such, they are a potential channel for intergenerational income transmission. To test this hypothesis, Table 4.4 presents estimates of intergenerational income mobility in the SOEP, once as the raw “regression correlation”, and once additionally controlling for the full set of career preferences.

In Panel A, the incomes of children and parents are measured in percentile ranks. As a consequence, the first entry in Column (1) represents the rank-rank slope of individual labor earnings. The estimate of 0.160 implies that a ten percentile increase in fathers’ earnings rank is associated with a 1.6 percentiles higher earnings rank of the child. In Column (2), I additionally control for the full set of career preferences. The estimate drops to 0.139, which means the association between parent and child incomes is weaker when conditioning on career preferences. This suggests that income transmission is indeed partially mediated by the transmission of preferences. In the remaining columns, I repeat this exercise for different income concepts. Column (3) shows that the intergenerational rank-rank correlation is much stronger when using parental gross household income as explanatory variable instead. Again, however, this correlation is reduced when controlling for career preferences. The same applies to the intergenerational correlations in net and gross household income.

To investigate which preferences account for the decrease in income persistence, I conduct a simple mediation analysis as in Heckman and Pinto (2015), based on the following three equations:

$$y_i = \alpha_0 + \alpha_1 y_i^{Parent} + \varepsilon_i \quad (4.2)$$

$$y_i = \beta_0 + \beta_1 y_i^{Parent} + \sum_k \beta_2^k \text{Pref}_i^k + \varepsilon_i \quad (4.3)$$

$$\text{Pref}_i^k = \gamma_0 + \gamma_1^k y_i^{Parent} + \varepsilon_i \quad \forall k \quad (4.4)$$

Equations (2) and (3) just replicate the estimates from Table 4.4 by regressing child on parent incomes, once controlling for career preferences. In Equation (4), I regress each career preference separately on parental income. The direct effect of parent on child income is then given by β_1 , the indirect effect via preference k by $\gamma_1^k \times \beta_2^k$. The share of the total effect mediated via mediator k can be computed as the fraction between indirect effect and total effect α_1 . The preferences that explain most of the decrease in income persistence are chances of promotion, family commitment, and helping others (Appendix Table D.5). In contrast, the preferences for respect and recognition and pursuing a career that is important for society tend to push in the other direction, increasing persistence.

Panel B of Table 4.4 replicates these results if the intergenerational income association is measured in logarithmic form. In consequence, the respective entry in Column (1) represents the intergenerational elasticity (IGE) of individual labor earnings. The coefficient of 0.224 implies that a 10% increase in parental income is associated with a 2.24% higher child income. All IGE estimates in Panel B are once again reduced when controlling for career preferences. Note though that due to the small sample size, none of the differences in Panel A and B are statistically significant. Nevertheless, there is overall clear evidence that point estimates of the intergenerational association between child and parent incomes are lower when accounting for career preferences.¹⁷ Depending on the specification, 7 to 28% of intergenerational income persistence can be explained by career preferences. This finding is also robust to including the full set of controls used in Column (6) of Table 4.1: I still find that controlling for career preferences reduces income persistence even conditional on the full control vector (Appendix Table D.7).

I find very similar patterns in the BCS. Table 4.5 shows that in the UK, too, the intergenerational correlation of earnings is reduced when controlling for career preferences. For the IGE, income

¹⁷ This conclusion is robust to different weighting schemes. Subsamples with intergenerational links are typically positively selected with respect to their socio-economic status (see, e.g., Ward, 2021). As a robustness test, I therefore re-weight the sample to match the broader population characteristics with respect to gender, state of upbringing, migration background and parental education (in three categories). As documented in Appendix Table D.6, results are very similar. The same applies when omitting weights altogether.

4 Career Preferences and Socio-Economic Background

Table 4.4: Career Preferences and Intergenerational Income Mobility

	Individual Labor Earnings				Gross HH Income		Net HH Income	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<i>Panel (A): in Ranks</i>								
Ind. Labor Earnings Father	0.160*** (0.054)	0.139*** (0.051)						
Gross HH Income			0.221*** (0.045)	0.188*** (0.044)	0.336*** (0.048)	0.311*** (0.047)		
Net HH Income							0.319*** (0.051)	0.295*** (0.048)
Preferences Observations	- 700	✓ 700	- 825	✓ 825	- 825	✓ 825	- 825	✓ 825
<i>Panel (B): in Logs</i>								
Log Ind. Labor Earnings Father	0.224** (0.089)	0.187** (0.073)						
Log Gross HH Income			0.218*** (0.070)	0.157** (0.066)	0.357*** (0.081)	0.300*** (0.084)		
Log Net HH Income							0.344*** (0.121)	0.297** (0.127)
Preferences Observations	- 652	✓ 652	- 784	✓ 784	- 818	✓ 818	- 825	✓ 825

Notes: This table shows estimates of separate regressions of child on parental income. Three different income concepts are used: gross individual labor earnings, gross household income and net household income. For individual labor earnings of the parents, I focus only on earnings of the father, as mothers display large variation at the extensive margin of labor supply. In Panel A, both child and parent incomes are measured in 100 percentile ranks and the estimates represent rank-rank slopes. In Panel B, incomes are measured in logarithmic form and the estimates represent the intergenerational elasticity (IGE). In every second column, I additionally control for all 12 career preferences. Robust standard errors in parentheses, * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Source: German Socio-Economic Panel (SOEP)

persistence drops by 18% from 0.372 to 0.304. The rank-rank slope declines by 20% from 0.223 to 0.179. Both reductions are statistically significant at the one percent level. Furthermore, considering that in a very comprehensive mediation analysis in the BCS looking at a multitude of channels, Bolt et al. (2021) can only explain 54-62% of the IGE in total, the magnitude of the reduction is sizable. As in the SOEP, important preferences mediating this reduction in income persistence are the desire to help others and chances of promotion. Further mediators are

Table 4.5: Career Preferences and Intergenerational Income Mobility in the British Cohort Study

	Log Gross Earnings		Gross Earnings Rank	
	(1)	(2)	(3)	(4)
Log Gross Income Parents	0.372*** (0.030)	0.304*** (0.029)		
Gross Income Rank Parents			0.223*** (0.015)	0.179*** (0.015)
Preferences	-	✓	-	✓
Observations	4047	4047	4365	4365

Notes: This table shows estimates of separate regressions of child on parental income. Income of children is measured as gross weakly individual labor earnings between ages 30-46, parental income as gross weakly household income when children are 10-16 years old. In the first two columns, both child and parent incomes are measured in logs, whereas in the last two columns incomes are measured in 100 percentile ranks. In every second column, I additionally control for all 16 career preferences. Robust standard errors in parentheses, * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Source: British Cohort Study (BCS)

the preferences for long term security, facing a real challenge, and working regular hours (Appendix Table D.8). Again, I find that controlling for career preferences continues to reduce income persistence even conditional on the full control vector (Appendix Table D.9).

What are the implications of these findings? Statistics of intergenerational income mobility are of interest because they are interpreted as a measure of equality of opportunity: if children from high-income households have higher earnings, this is because they experience better opportunities to achieve a high income. For example, it has been shown that financing constraints can partially explain the parental income gradient in college enrollment (Solis, 2017; Manoli and Turner, 2018), parental social networks foster own earnings (Corak and Piraino, 2011; Kramarz and Skans, 2014; San, 2020), and that children from high-income households receive more parental time investments (Guryan et al., 2008), affecting skill acquisition and resulting in more favorable labor market outcomes later on.

The often implicit assumption behind this line of reasoning is that conditional on opportunities, children from low- and high-income households would achieve the same incomes on average. However, this neglects the fact that maximizing utility is not necessarily the same as maximizing expected incomes. As shown in this paper, non-monetary aspects such as the (dis)utility from work and leisure, interest in a particular job, or health and safety condi-

4 Career Preferences and Socio-Economic Background

tions are all examples of factors why individuals may not choose the highest paying career. And in Germany and the UK, these factors indeed explain part of the earnings gap between individuals with different SES.

Is this conditional estimate of intergenerational income mobility a more accurate measure of equality of opportunity? The answer to this question is yes if one believes that individuals should be held responsible for their preferences, as argued by Rawls (1971) and Dworkin (1981). According to this view, my estimates are encouraging news, suggesting that intergenerational income persistence only partially reflects lacking opportunities for low SES children. Germany and the UK would then have a more opportunity-egalitarian society than previously thought. However, this position is increasingly difficult to defend on ethical grounds, as we come to understand that career preferences are themselves influenced by parental background (Bowles, 1998; Müller, 2021). Taking this into account, also an opportunity egalitarian social planner could therefore want to compensate for preferences about career and occupational choice.¹⁸

On these grounds, many real world policies that are actively designed to guide and shape the career preferences of adolescents enjoy widespread support. For example, governments foster STEM participation for women, ethnic minorities and low SES children (Seymour and Hewitt, 1997; Best et al., 2013; Arcidiacono et al., 2016; Hill, 2017; De Philippis, 2021). In Germany, the nationwide Girls' Day, initiated in 2001, aims to familiarize teenagers with fields of work in which women are underrepresented, and reaches around 100,000 girls each year. At a more general level, secondary schools in many countries, including Germany and the US, have the mandate to provide career guidance in school (see, e.g., Gysbers, 2005). The results of this paper suggest that further targeting such policies further to reach children from disadvantaged backgrounds in particular could be one step forward in improving equality of opportunity.

4.6 Conclusion

This paper has documented novel stylized facts about career preferences, that is adolescents' goals and desires regarding their career paths. Career preferences are correlated with parental

¹⁸ In this light, Arneson (1989), Cohen (1989), and Roemer (1998) reject Dworkin's view on holding individuals responsible for their preferences and advocate for a control approach to responsibility. In their conception, people should only be held responsible for outcomes resulting from genuine choices, which means that one should correct for the influence of circumstances like parental income on preferences.

background, and highly informative of future labor market outcomes. As such, they constitute a potential mechanism of intergenerational income persistence. In both Germany and the United Kingdom, I find that statistics of intergenerational income mobility are reduced by 7 to 28 percent when controlling for career preferences.

The latter result suggests that altering career preferences opens up a potential pathway for reducing the intergenerational persistence of socio-economic status. For example, schools could offer more advanced career preparation or enhance the attractiveness of enrolling in STEM majors. As the results in this paper remain descriptive, a causal impact evaluation of such policies would constitute a fruitful avenue for future research.

A Appendix to Chapter 1

A.1 Business Taxation in Germany

In Germany, business profits are subject to two different taxes. At the national level, profits are either taxed under the personal income tax or under the corporate income tax, depending on the legal form of the firm. In addition, both corporate and non-corporate firms are subject to the local business tax (LBT) at the municipality level.

Corporate Income Tax Profits of incorporated firms are subject to the national corporate income tax (*Körperschaftsteuer*). The rate of the corporate income tax is currently 15 percent. Until 2000, a split rate imputation system existed in Germany, where retained profits were subject to a tax rate of 40-45 percent, whereas distributed profits were taxed at a rate of 30 percent. From 2001 to 2007, all profits were equally taxed at 25 percent. In all years since 1991, a so-called solidarity surcharge (*Solidaritätszuschlag*) of 5.5 percent of the corporate tax rate was added, dedicated to financing the costs of the German reunification.

Personal Income Tax Profits of non-corporated firms are subject to the progressive income tax (*Einkommensteuer*). The top marginal tax rate of the personal income tax is currently 45 percent but has been higher in the past, with a maximum of 56 percent in the 1980s. Since 2001, sole proprietors and partners in a partnership have been able to partially offset LBT payments tax against their income tax. This regulation, limiting the bite of the LBT, is however not relevant in our setting, as it only applies to unincorporated businesses, whereas we focus exclusively on the corporate sector.


Local Business Tax In addition, both corporate and non-corporate firms are subject to the LBT (*Gewerbesteuer*). As the corporate tax and the personal income tax, the LBT is a federal tax. For this reason, tax base and liability criteria of the LBT are set at the federal level. The tax rate, in turn, falls under the discretion of the municipalities. More precisely, municipalities decide autonomously on a scaling factor that is then multiplied with a uniform basic tax rate. This results in the following formula:

$$\text{Local Business Tax Rate} = \text{Basic Federal Tax Rate} \times \text{Municipal Scaling Factor}$$

The basic rate, which is fixed at the national level, has been constant with exception to a change in 2008, when it was decreased from 5.0 to 3.5 percent. This means that for the median municipal scaling factor of 3.2, the resulting LBT rate was 16 percent before 2008. After 2008, the tax rate for the median scaling factor of 3.5 was 12.25 percent.

Each year, the municipal council has to vote on next year's municipal scaling factor, even if it remains unchanged. The decision on next year's local scaling factor is taken jointly with the adoption of the budget in the year's last meeting of the municipal council. For this reason, tax hikes are typically announced in December. In Figure A.1, we substantiate this empirically, showing that newspaper coverage of municipal tax hikes in a given year indeed peaks in December. This holds for both a narrower definition (in red) and a broader definition (in blue) of newspaper coverage of a hike in the LBT. As documented in Appendix A.2.1, a decision to increase the LBT sends no clear signal about the likelihood of future tax changes.

Figure A.1: Timing of Tax Hike News



Notes: This figure provides evidence on the point in time when firms typically learn about a tax hike by displaying the number of monthly newspaper articles covering increases in the LBT, obtained from the German press database Genios. Under the broad definition, we counted search matches for “gewerbesteuer erhöh*”, under the narrow definition for “gewerbesteuer (erhöht* || angehob* || erhöhung) (beschl* || entschei*)”.

Around three quarters of the revenues of the LBT accrue directly to the municipalities, whereas one quarter is transferred to the federal government. Taxable profits of firms with establishments in more than one municipality are divided between municipalities according to formula apportionment based on the payroll share. As a consequence, profit shifting between municipalities requires the actual re-allocation of the employees (or wages) of a firm, and is thus associated with relatively high costs. The revenues from the LBT are of key importance for municipal budgets, as the LBT constitutes the most important original source of revenue for municipalities in Germany. Besides own tax revenues, municipal budgets are strongly dependent on fiscal transfers from the federal government or the federal states. As the municipalities cannot directly influence these fiscal transfers, the rate of the LBT is the central budget parameter under their control.

A.2 Data Appendix

This appendix provides comprehensive information on the data sets used in the empirical analysis (including the translated wording of the relevant survey questions from the ifo Investment Survey), explains how we obtain our analysis sample, and reports summary statistics and aggregate time series of our final sample.

A.2.1 Administrative Data at the Municipality Level

The administrative data on tax rates and municipality revenues and expenditures used in this paper cover the period from 1980 to 2018. The data largely correspond to the municipality data underlying the analysis in Fuest et al. (2018), comprising the period 1993 to 2018. Data for the period from 1980 to 1992 were obtained by filing individual requests to the respective Statistical Offices of the German Federal States. For the state of Schleswig-Holstein, data were not available in the year 1980. For Bremen and Saarland, data are only available since 1990. As these are the two smallest states of Germany in terms of GDP and population, jointly comprising less than 2% of the German population, this does not substantially change the composition of our sample. For all years, the data contain information on scaling factors of the LBT. In addition, we know the full municipality budget, that is all categories of expenditures and revenues, for most years. For a more detailed description of the data, we refer to Fuest et al. (2018) and Isphording et al. (2021).

There is substantial variation in LBT rates across municipalities and over time. To document this variation, we use the subset of municipalities, where we observe at least one firm during our sample period in the ifo Investment Survey. Figure A.2 plots the raw data of the local scaling factors for each municipality in Western Germany (excl. Berlin) over time, demonstrating that there is a lot of variation in local business taxes in any given year. Municipalities tend to increase the LBT approximately ten times as frequently as decreasing it. In consequence, the statistical power of this variation is too low to investigate the effect of tax drops in our data, and the analysis is thus restricted to tax hikes.¹ Accordingly, Figure A.3 shows that the share of municipalities that increased the LBT in a given year is relatively stable over time and does not differ between recessions and expansions. Moreover, Panel (A) of Figure A.4 plots the fraction of municipalities that underwent a given number of tax hikes in the period between

¹ The number of tax decreases that could in principle be used in the analysis is very low. If we followed the protocol in Appendix A.2.3 to combine the municipality-level data on LBT rates and the firm-level data from the IVS, our analysis could only exploit 236 firm observations (0.7% of all firm-year observations) that face a tax drop in a given year despite spanning a time frame of almost four decades.

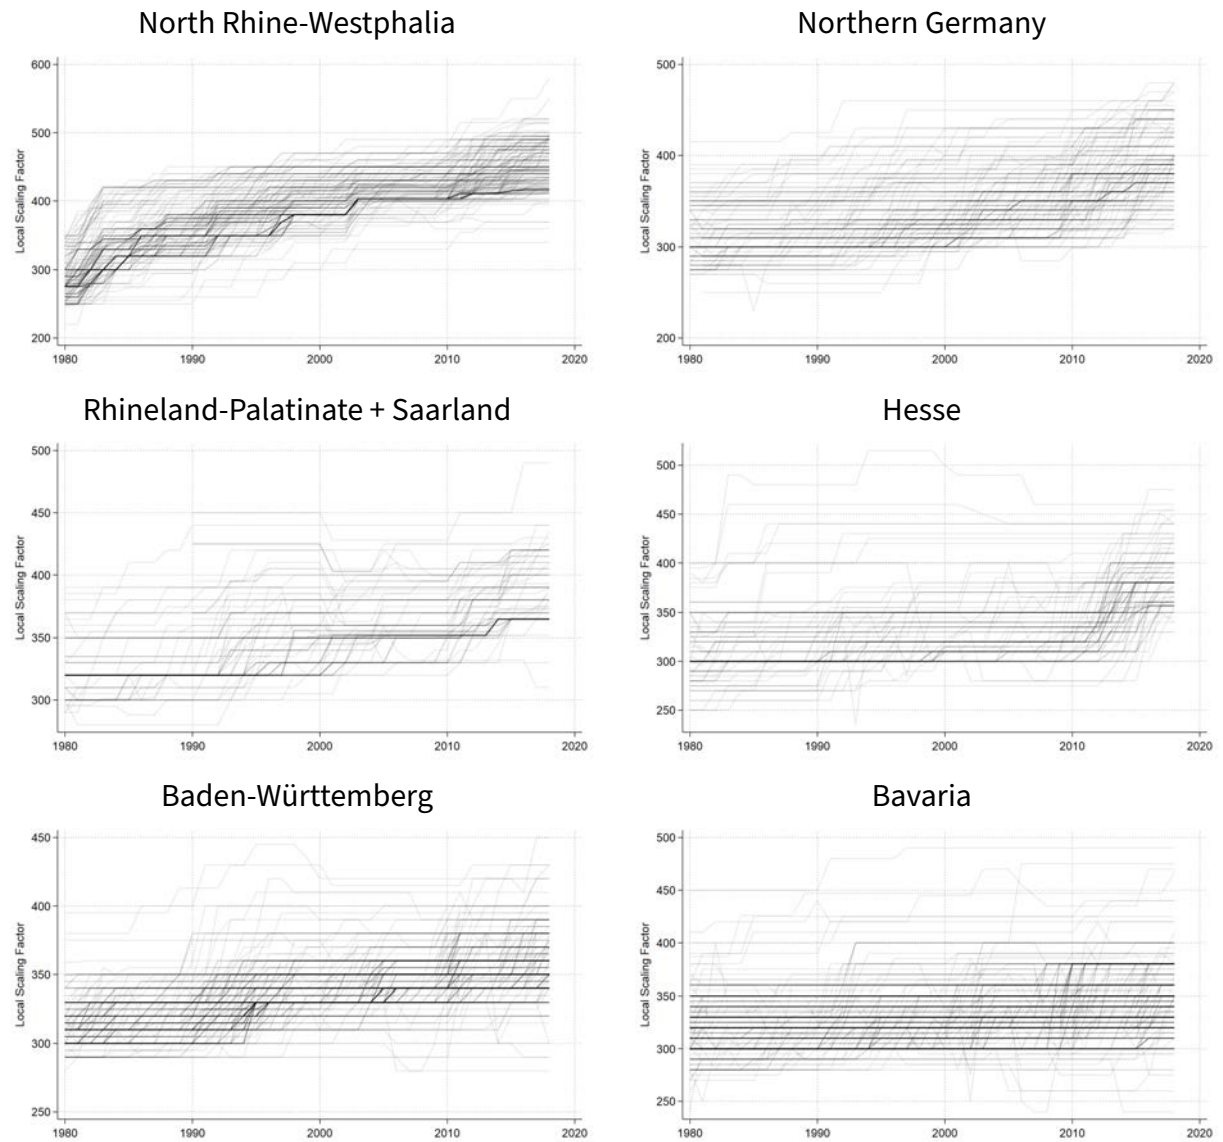
1980 and 2018. The median municipality experienced three tax hikes, while taxes were never increased in only 7% of municipalities. The average duration between two tax hikes in our sample is 14.6 years, the median duration 13 years. Panel (B) displays the mean and various percentiles of the size of tax hikes over time. The distribution of tax hikes is rather stable over time in terms of average size and dispersion. If anything, tax hikes were slightly larger in the early 1980s and slightly lower in the 2010s.

To shed light on the dynamic aspects of tax hikes, Figure A.5 documents how a tax hike in year t_0 influences the probability for future tax hikes in the same municipality. Specifically, the figure displays the coefficients of separate regressions of the following form


$$TaxHike_{m,t+x} = \beta TaxHike_{m,t} + \mu_m + \epsilon_{m,t} \quad \forall x = \{1, 20\},$$

where $TaxHike_{m,t+x}$ is an indicator for a tax hike occurring x years after a tax hike in the same municipality m in year t that is estimated separately for each year in the future $x \in \{1, 20\}$. In the right panel, we include municipality fixed effects. The results show that tax hikes contain little predictive power for future tax hikes. While the unconditional probability for future tax hikes is slightly elevated if a tax hike has recently been enacted, the association is very weak and completely vanishes when including municipality fixed effects, which corresponds to the tax rate variation exploited in our main analyses (that applies firm fixed effects which are themselves nested within municipalities).



Figure A.2: Time Series of Local Scaling Factors by Municipality



Notes: This figure shows the local scaling factors underlying the LBT for each municipality in West Germany (excl. West-Berlin) over the period between 1980 and 2018. “Northern Germany” summarizes the states of Schleswig-Holstein, Hamburg, Bremen, and Lower Saxony.

Figure A.3: Share of Municipalities Increasing the LBT over Time

Notes: This figure shows the share of municipalities that increased the LBT in a given year. Gray shaded areas indicate recessions as defined by the German Council of Economic Experts.

Figure A.4: Number of Tax Hikes and Distribution of Tax Changes**(A) Number of Tax Hikes per Municipality****(B) Distribution of Tax Hikes over Time**

Notes: Panel (A) plots the fraction of municipalities that underwent a given number of tax hikes in the period between 1980 and 2018. Panel (B) displays the average size of tax hikes (in percentage points) along with various distributional parameters, i.e., the median, the interquartile range, and the range between the 10th and 90th percentile of tax hikes in a given year.




Figure A.5: Predictability of Tax Hikes as a Function of Past Tax Hikes in the Same Municipality

Notes: This figure reports how a tax hike in year t_0 influences the probability for future tax hikes in the same municipality, by showing the estimates of separate regressions with tax hike indicators x years in the future as dependent variable and a tax hike indicator for the current year as explanatory variable: $TaxHike_{m,t+x} = \beta TaxHike_{m,t} + \mu_m + \epsilon_{m,t} \forall x = \{1, 20\}$. In the right panel, we include municipality fixed effects, so that the graph shows the probability of future tax hikes conditional on knowing the institutional and political economy patterns of the own municipality.

A.2.2 The ifo Investment Survey

General Information The ifo Investment Survey (IVS, 2019) is a firm-level survey of the German manufacturing sector. Since its inception in 1955, it is conducted biannually by the ifo Institute, with survey waves in spring and fall of each year. The aim of the IVS is to supplement investment data collected by the German Statistical Office, which is only available with a time lag of two years, with more recent data by means of extrapolations at the industry level. The survey is part of the European Commission's sponsored investment surveys in its member countries and participation in the survey is supported by several industry associations. All of this background information is contained in the cover letter of each survey. The aggregated investment volume of the participants of the IVS represents approximately 56% of overall investment in the manufacturing sector (see Sauer and Wohlrabe (2020), p. 145).

The repeated panel structure of the ifo Investment Survey allows tracking approximately 1,500 firms over time. As outlined in greater detail below, the questionnaire elicits three types of questions, covering (i) the planned volume of investment, (ii) the realized volume of investment, and (iii) investment objectives. Realized investment is always reported for the previous year. Next to these investment-related variables, firms also report annual revenues and the number of employees. For all model specifications which include year fixed effects at the industry level, we rely on the ifo industry classification that maps firms into 34 industries over the entire sample period. The ifo industry classification is slightly more granular than, but largely comparable to two-digit NACE industries. All items of the questionnaire refer to the firms' plants located in Germany. Sauer and Wohlrabe (2020) provide a comprehensive overview and detailed description of this data source. After a protection period of one and a half years, the anonymized data can be accessed via the LMU-ifo Economics & Business Data Center under strict non-disclosure regulations (<https://www.ifo.de/en/ebdc>).

The survey is usually completed by high-level management personnel at the firms' controlling departments (Sauer and Wohlrabe, 2020). The ifo Institute incentivizes the participation to the survey by automatically providing the participants with the survey results free of charge as a thank-you for their cooperation. In order to create an additional incentive for participation in the investment survey, this reporting includes more detailed information, e.g., at more disaggregate sectoral levels, compared to the results that are reported publicly.

Representativeness and Accuracy In Table A.1, we demonstrate the representativeness of the ifo Investment Survey by comparing it to the distribution of firms in administrative data

Table A.1: Distribution of Firms in the IVS by Industry and Size

WZ08	Industry	ifo Investment Survey				Actual Germany by			
		Small	Medium	Large	Total	Count	Employees	GVA	Payroll
10-12	Food, beverages, and tobacco	1.1	3.6	3.6	8.2	14.0	12.4	7.8	7.0
13-15	Textiles, apparel, and leather	1.2	1.8	1.0	4.1	4.2	1.8	1.1	1.1
16-18	Wood/paper products and printing	3.0	5.7	3.5	12.2	11.8	5.5	4.3	4.0
19	Coke and refined petroleum	-	-	0.3	0.6	0.0	0.3	0.8	0.5
20	Chemicals	-	1.1	3.4	4.7	1.5	4.6	6.9	6.0
21	Pharmaceuticals, medicinal chemical, and botanical	-	0.4	1.0	1.4	0.3	1.9	3.1	3.0
22+23	Rubber/plastic products, and other non-metallic	1.4	6.4	6.4	14.2	8.1	9.0	7.6	7.7
24+25	Basic and fabricated metal products	2.1	6.8	8.4	17.3	21.9	15.7	13.2	13.5
26	Computers, electronics, and optical products	-	1.0	2.4	3.6	3.7	4.8	5.4	5.5
27	Electrical equipment	-	1.3	3.8	5.3	2.9	6.4	7.0	7.5
28	Machinery and equipment	0.5	5.3	11.1	17.0	7.7	15.7	16.7	18.0
29+30	Transport equipment	-	0.8	4.0	4.9	1.9	13.2	19.0	19.1
31-33	Other, and installation of machinery and equipment	1.2	2.1	3.2	6.4	21.9	8.6	7.0	7.0
Total		11.4	36.5	52.2	100	100	100	100	100
Actual GER by Count		89.7	7.7	2.6	100				
Actual GER by Employees		19.1	18.6	62.3	100				
Actual GER by Gross Value Added (GVA)		10.6	13.2	76.1	100				
Actual GER by Payroll		10.0	13.9	76.1	100				

Notes: This table compares the distribution of firms in the ifo Investment Survey to the distribution of firms in administrative data by industry and firm size. The ifo Investment Survey data is based on the year 2018. The administrative data is based on the 2018 Statistics on Small and Medium-sized Enterprises (*“Statistik für kleine und mittlere Unternehmen”*) provided by the Federal Statistical Office (EVAS Code 48121). Definition of size classes: small: 0-49 employees; medium: 50-249 employees; large: 250+ employees. Cells are empty if there are less than 4 observations due to data protection.

by industry and firm size. The numbers depicted in the table display the percentage share of firms in the respective cells. For instance, 17.3% of firms in the 2018 ifo Investment Survey are in the basic metals and fabricated metal products industry (2-digit WZ08: 24 and 25). This is in between the share of firms by count (21.9%) and weighted by employees (15.7%) in the administrative data. The share of firms by gross value added and payroll in this industry is around 13% in population. Overall, the industry-composition of the ifo Investment Survey is very close to the distribution in administrative data. Regarding the distribution across firm size, the ifo Investment Survey covers a substantial share in each size category. Around a third of firms have between 50 and 249 employees. Thereby, the survey slightly oversamples medium-sized firms while still being representative for small and large firms, since the share of firms is in between the population share of firms by count on the one hand, and by employees, gross value added, or payroll on the other hand.

In general, the accuracy of the IVS data appears to be quite high, as the average deviations of the survey results from the data of the Federal Statistical Office for the manufacturing sector as a whole are only relatively minor. For instance, Bachmann and Zorn (2020) show

that aggregate investment growth calculated from the microdata of the ifo Investment Survey is highly correlated with manufacturing investment growth reported by the Federal Statistical Office. Similarly, benchmarking the investment growth rates calculated from the survey against official statistics from the German Statistical Office for the period 1980 to 2016, Sauer and Wohlrabe (2020) report an average absolute estimation error of less than two percentage points. Sauer and Wohlrabe (2020) stress that it should be borne in mind that, at the time investments were recorded in the survey, the balance sheets of some of the companies may not yet be final, while the official results, on the other hand, are based on the final balance sheet figures.

Lastly, and in line with evidence presented in Appendix A.2.3, Bachmann et al. (2017) present a series of stylized facts on the cross-sectional and time-series properties of revisions of investment plans, i.e., the difference between ex ante planned and ex post realized investment volumes, showing that these deviations are meaningful along many dimensions. For example, they document that the overall distribution of revisions is not systematically skewed, while their cross-sectional average is procyclical. This indicates that participants provide accurate investment plans given their current level of knowledge at the time of the survey.

Wording of Questions in the IVS Used in the Paper In the following, we present the translated wording of the questions of the IVS that we use in the paper.

Fall Questionnaire

1. General company information on the current financial year

Employees (as of Sept. 30th): _____ Total revenue (TEUR): _____

2. Gross fixed capital formation (equipment and buildings) in TEUR

	last year	this year	next year
Total (equipment + buildings):	_____	_____	_____

3. Investment targets this year and next year

A Appendix to Chapter 1

Our domestic investment activity is influenced positively/negatively by the following factors:

		inducement		no	hampering	
		strong	little	influence	little	strong
This year:	a) Financing situation	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
	b) ...					
Next year:	a) Financing situation	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
	b) ...					

Spring Questionnaire

1. General company information on the last financial year

Financial year from: ____ to: ____	Focus of production: _____
Employees (as of Sept. 30th): _____	Total revenue (TEUR): _____

2. Gross fixed capital formation (equipment and buildings) in TEUR

	two years ago	last year	this year
Buildings:	_____	_____	_____
Equipment:	_____	_____	_____
Total (buildings + equipment):	_____	_____	_____

A.2.3 Construction and Descriptive Statistics of the Merged Dataset

Protocol for Construction of Merged Dataset In constructing the final sample used for analysis, we have aimed at establishing a valid control group to analyze corporate tax hikes over time, and at cleaning the data to ensure that the results are not driven by outliers. To obtain our final sample, we follow the protocol outlined below:

- We restrict our sample to West Germany and, as Fuest et al. (2018), drop all municipalities which underwent municipal mergers in the observation period. As most of these municipalities were located in East Germany anyway, this does not substantially restrict our sample further (less than 1% of municipalities affected).
- We drop observations in a window of two years before and after a tax hike, if another tax hike occurred in that window.
- We drop all observations for which a tax decrease was enacted, as well as the two years before and after the tax decrease. Fuest et al. (2018) find that while tax hikes are arguably exogenous to shocks to economic variables, a potential endogeneity to economic conditions cannot be ruled out for tax decreases. In addition, only 13.5% of tax changes in the sample are tax decreases. In our setting, we do not have enough statistical power to separately analyze tax decreases.
- In total, the outlined sample selection above reduces the sample size from 8,522 municipalities and 326,274 municipality \times year observations to 8,266 municipalities and 283,846 municipality \times year observations.
- In the firm survey, for variables that are elicited both in the spring and the fall (last year's number of employees, revenues, and total investment volume), we follow Bachmann et al. (2017) and compute a yearly value by taking the average. We drop the observation if both values deviate more than 20% from the mean.
- As Fuest et al. (2018), we drop firms with legal forms which are exempt from paying the LBT (this affects only 6.2% of the observations).
- We drop firms for which we observe revisions in investment plans in less than 5 years.
- To construct the Log Revision Ratio, we calculate the ratio of realized investments over planned investments, take the natural logarithm, and drop outliers (all values smaller/larger than $p1/p99$ in each year).

- Matching the municipal and firm-level samples, the final sample consists of 35,310 observations that are spread across 1,192 municipalities.
- We express all nominal variables, i.e., the amounts of revenues and investments, in real terms of constant 2015 Euro by converting German Mark to Euro and adjusting for inflation using the German Consumer Price Index (CPI).

Firms in the Merged Dataset: Descriptive Statistics Table A.2 displays summary statistics for the firms in our sample. For each firm, we can rely on information on reported planned and realized investment volumes in 17 years on average. The median firm is a typical representative of the “German Mittelstand” employing 264 workers, generating annual revenues of 45 million Euro (CPI inflation-adjusted and—if denominated in German Mark—converted to 2015 Euros), and investing 1.4 million Euro each year. As described in Appendix A.2.2, the IVS covers firms of all sizes. While slightly oversampling medium-sized firms, it is still representative for small and large firms. Accordingly, 10% of firms in our sample have at most 38 employees, annual revenues of 5.2 million Euro and invest as little as 88,000 Euro per annum. In contrast, the 10% largest firms employ at least 1,950 workers and have annual revenues of almost half a billion Euro and total annual investment of at least 19 million Euro. As shown in Figure A.6, the firm size is consequently highly skewed according to the number of employees (Panel A), while the distribution of its logarithm displays a bell-shape (Panel B). As demonstrated in Panel A of Table A.3, firms experiencing a tax hike in year t are largely comparable to firms in the control group according to key firm characteristics measured in year $t - 1$. Further, the pre-treatment averages of both main outcome variables (Log Revision Ratio and Downward Revision Indicator measured in year t_{-1}) are not statistically different for firms that eventually are affected by a tax hike in year y_0 and firms ending up in the control group, see Panel B.

Documenting variation in investment over time, Figure A.7 displays a calendar time graph of the investment plans and investment realizations. Relatedly, Figure A.8 presents the share of downward revisions of investment (blue, solid) and the average log revision ratio (red, dashed) over time. The gray shaded areas indicate recession periods. During recessions, the share of downward revisions increases and the log revision ratio decreases. In addition, there might be a slight time trend towards a higher share of firms that revise their investments downwards. Note, however, that this potential trend does not affect our analysis since we include year fixed effects in the regressions and thus rely on differences between firms in a given year for identification.

Figure A.9 shows the share of firms that report a decline in revenues by more than 10% compared to the previous year. In normal times, we observe that around 10% of firms experience such a revenue drop. In recessions, this share spikes up to 60%. This variable is used in Section 1.4.3, where we discuss potential channels of state-dependence in the effect of tax hikes on investment revisions.

Table A.2: Summary Statistics of Firms in the Sample

	p10	p50	p90	Mean
Employees	38	264	1,950	1,361
Revenues	5,194	44,901	451,899	418,842
Investment	88	1,435	19,163	17,751
Obs. per Firm	7	16	29	17


Notes: This table shows the 10th, 50th, and 90th percentile, and the mean of employees, revenues, and realized investment for the firms in our sample. “Obs. per Firm” refers to the number of years a firm is observed in our sample, i.e., the number of years for which firms report both ex ante planned and ex post realized volumes of investment. Revenues and investment are displayed in thousands of Euro.

Table A.3: Balance Statistics of Firms in the Treatment and Control Group

	Treated	Control	p-value
	(1)	(2)	(3)
	Mean _{<i>t</i>-1}	Mean _{<i>t</i>-1}	(1)=(2)
<i>Panel (A): Firm Characteristics</i>			
Employees	1,283	1,376	0.54
Revenues	337,895	434,071	0.14
Realized Investment	14,401	18,372	0.15
<i>Panel (B): Main Outcome Variables</i>			
Downward Revision of Investment Plans (Share)	0.54	0.53	0.34
Log Revision Ratio	-0.04	-0.03	0.59


Notes: This table presents balance statistics of firms experiencing a tax hike in year t_0 and untreated firms, respectively (both measured in the pre-event year t_{-1}). Panel (A) covers firm characteristics including the number of employees, annual revenues (in thousands of Euro), and realized investment (in thousands of Euro). Panel (B) refers to the main outcome variables, i.e., the Downward Revision Indicator and the Log Revision Ratio. Column (3) presents the p-values of a t-test on the equality of the means depicted in Columns (1) and (2).

Figure A.6: Distribution of Firms by Number of Employees




Notes: Panel (A) shows a histogram of the number of employees for the firms in our sample. The distribution is winsorized at a value of 4,000 employees. The vertical line denotes the median number of employees, which is 264. Panel (B) shows a histogram of the natural logarithm of the number of employees for the firms in our sample. The labels on the x-axis refer to the absolute number of employees.

Figure A.7: Time Series of Investment Plans and Realizations




Notes: This figure shows time trends of log planned investment and log realized investment in the period 1980 to 2018, for all firms with a non-missing log revision ratio. The shaded areas indicate 95% point-wise confidence intervals. Gray shaded areas indicate recessions as defined by the German Council of Economic Experts.

Figure A.8: Time-Series of Investment Revisions



Notes: This figure shows time series of the Log Revision Ratio (right axis), defined as the logarithm of the ratio between realized and planned investment, and the downward revision dummy (left axis), indicating whether a firm has invested less than planned, for the period 1980 to 2018 in our sample. Blue and red shaded areas indicate 95% point-wise confidence intervals. Gray shaded areas indicate recessions as defined by the German Council of Economic Experts.

Figure A.9: Time-Series of Share of Large Revenue Drops




Notes: This figure depicts the time series of the share of firms with large revenue drops, defined as a year-to-year decline in revenues of more than 10%, over the period 1980 to 2018. Blue shaded areas indicate 95% point-wise confidence intervals, while gray shaded areas indicate recessions as defined by the German Council of Economic Experts.

Relationship between Planned and Realized Investment Our identification approach relies on the investment plans of firms. In the following, we display the distribution of the investment revision ratio and illustrate the strong explanatory power of investment plans for actual investments.

Figure A.10 shows the distribution of the log revision ratio, trimmed at the first and 99th percentile. The log revision ratio is centered around zero, which means that on average, firms invest as much as they have previously planned. Overall, the approximately normal distribution in Figure A.10 indicates that firms revise investments frequently and similarly upwards and downwards.

Next, we provide further evidence that investment plans are highly informative for subsequently realized investment volumes. As shown in Figure 1.2 in the main part of the paper, the relationship between planned and realized investment volumes is highly linear and virtually corresponding to the 45 degree line. According to the corresponding regression output presented in Column (1) of Table A.4, 84% of the unconditional variation in (log) realized investment is explained by the investment plans for the respective year ($R^2 = 0.84$). The estimated slope is 0.91 and thus close to one. Moreover, the horse-race regression depicted in Column (2) demonstrates that planned investment regarding year t is much more strongly correlated with the ex post realizations in t than with realized levels in the previous year. As shown in Columns (3) and (4), these patterns even hold when controlling for firm fixed effects and investment plans are still strongly positively associated with ex post realized investment. Taken together, investment plans appear to contain accurate information on subsequent year's investment that goes beyond the extrapolation of the level of investment that was realized in the year these plans are reported to the IVS.

Figure A.10: Distribution of the Log Revision Ratio

Notes: This figure shows a histogram of the Log Revision Ratio in our sample. The Log Revision Ratio is defined as the logarithm of the ratio between realized and planned investment and constitutes one of the two main variables used in the analysis. For exhibitional reasons, the outliers below p1 and above p99 are not depicted here.

Table A.4: Information Content of Investment Plans for Realized Investment


	Log(Realized Investment)			
	(1)	(2)	(3)	(4)
Log(Planned Investment)	0.908 (0.005)	0.552 (0.011)	0.574 (0.012)	0.462 (0.012)
L.Log(Realized Investment)		0.395 (0.011)		0.195 (0.011)
Constant	1.276 (0.067)	0.731 (0.047)	6.064 (0.165)	4.886 (0.164)
Observations	25282	25282	25282	25282
R^2	0.84	0.87	0.89	0.89
R^2 (within)	-	-	0.27	0.30
Firm FE	-	-	✓	✓

Notes: This table reports estimates from linear regressions of log realized investment in year t_0 (I_{t0}) on log planned investment ($E_{t-1}(I_{t0})$) and log realized investment in the previous year (I_{t-1}). Columns (3) and (4) in addition purge for fixed effects at the firm-level. The sample is restricted to observations in years without changes in the LBT. Robust standard errors in parentheses.


A.3 Supplementary Figures and Tables

**Figure A.11: Long Event Study:
Effect of Tax Hike on Investment Plans, Realizations, and Revisions**

Downward Revision and Log Revision Ratio




Log Planned and Realized Investment



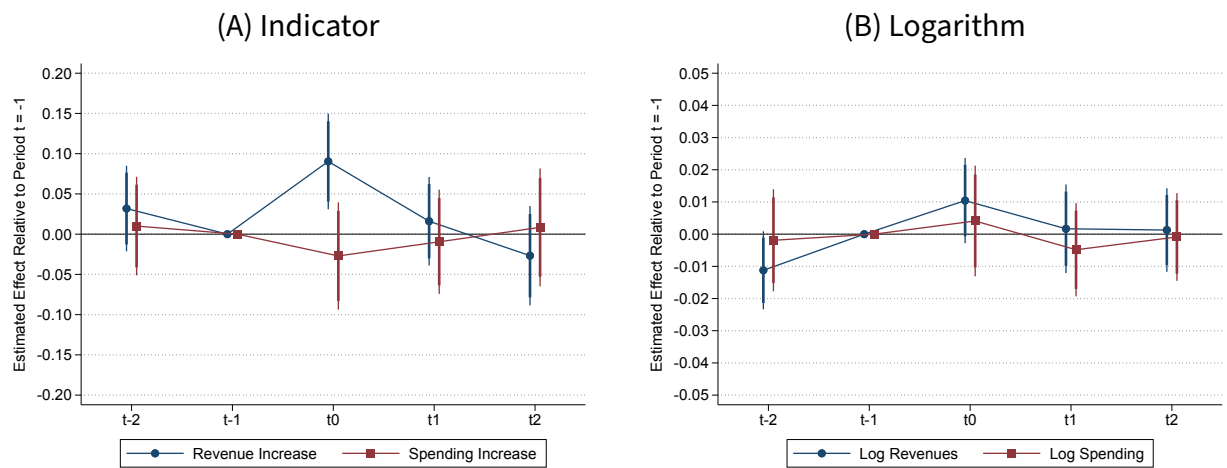
Notes: This figure shows event-study estimates of the downward revision (blue, solid line) and the log revision ratio (red, solid lines) in the left panel and log planned investment (green, short dashed lines) and log realized investment (orange, long dashed lines) in the right panel on the tax hike indicator and fixed effects at the levels of firm identifiers and years with longer event windows. The reference period is t_{-1} . “Log Revision Ratio” is the natural logarithm of the ratio between the ex post realized and ex ante planned volume of investment. In the right panel, the sample is trimmed outside the event window. Inspired by Dube et al. (2023), when estimating the effects with respect to log planned and realized investment, firms are assigned to another firm identifier after the year that is in the middle between two tax hikes in order to ensure that there is only one treatment for each unit and to allow for different long-run trends. In addition, end-periods $t-4$ and $t+4$ are binned in the right panel. The confidence intervals refer to the significance levels of 90% (thick lines) and 95% (thin lines).

**Figure A.12: Investment Revision Effect after a Tax Hike:
Alternative Estimators**




Notes: This figure shows the estimates of the imputation estimator introduced by Borusyak et al. (forthcoming) (solid lines) and the interaction-weighted estimator by Sun and Abraham (2021) (dashed lines). The dependent variable is based on the ratio of realized investments over planned investments (elicited in fall the year before). “Downward Revision” is an indicator that is one if the ratio is below one (blue/circle). “Log Revision Ratio” is the log of the ratio (red/square). The treatment “Tax Hike Indicator” is an indicator that is one if the local corporate tax rate is higher than in the previous year. Time fixed effects and firm fixed effects are absorbed in the estimation. Confidence bands refer to the 95% level.

Figure A.13: Event Study: Expenditures and Revenues of Municipalities




Notes: This figure shows the estimates of the following event-study regression: $Y_{m,t} = \sum_{j=-2}^2 \gamma_j \text{TaxHike}_{m,t}^j + \mu_i + \phi_{l,t} + \psi_{s,t} + \varepsilon_{i,t}$, where μ_i are firm fixed effects, $\psi_{s,t}$ year fixed effects at the industry level, and $\phi_{l,t}$ state-year fixed effects. In the left panel, $Y_{m,t}$ represents an indicator that is one when municipal revenues/spending increases compared to the previous year. In the right panel, $Y_{m,t}$ represents log municipal revenues/spending. The reference period is $t - 1$. Industry fixed effects are at the ifo industry classification level that is comparable to the level of two-digit NACE industries. Standard errors are clustered at the municipality level. The thick and thin confidence bands refer to the levels of 90% and 95%.

Figure A.14: Investment Revisions after a Tax Hike: Permutation Test

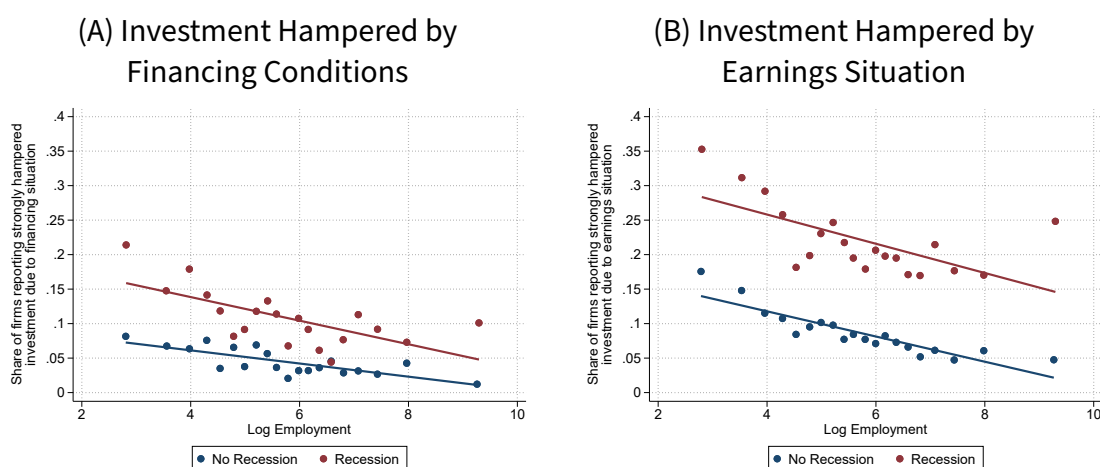
Notes: This figure reports the empirical cumulative distribution functions of estimates from 2000 placebo tests. In a Monte Carlo exercise, tax hikes ($\mathbb{1}(\Delta tax_{m,t} > 0)$) are randomly allocated to municipalities by holding the share of treated municipalities constant. Then, Model (1) is estimated with the full set of fixed effects. In Panel (A), the dependent variable is $\mathbb{1}\left(\frac{I_{i,t}}{E_{i,t-1}(I_{i,t})} < 1\right)$, i.e., an indicator that is one if the fraction of realized investment over planned investment is below one. In Panel (B), the dependent variable is $\ln\left(\frac{I_{i,t}}{E_{i,t-1}(I_{i,t})}\right)$, i.e., the natural logarithm of the investment revision ratio. The vertical lines correspond to the baseline estimates from Column 5 in Panels A1 and B1 of Table 1.2. In Panel (A), 0.05% of the estimates are equal or larger than the baseline estimate. In Panel (B), 1.15% of the estimates are equal or smaller than the baseline estimate.

Figure A.15: Collectively Bargained Wage Growth in Manufacturing



Notes: This figure shows year-on-year changes of the index of hourly earnings in the manufacturing sector without special payments obtained from the German Statistical Office. Grey shaded areas indicate recessions as defined by the German Council of Economic Experts.

Figure A.16: Obstacles to Investment by Firm Size



Notes: This binscatter plot depicts the share of firms reporting that their investment activity is strongly negatively affected by adverse financing conditions (Panel A) and the earnings situation (Panel B) by firm size, separately for recession and non-recession years. Recession years are defined as 1980-1982, 1992-1993, 2001-2003, and 2008-2009, as classified by the German Council of Economic Experts. Panel A uses the same survey question as Table A.11 in which firms report how financial constraints influence their current year's investment activity on a scale between 1 (strongly induced) and 5 (strongly hampered), see Appendix B for the exact wording. Panel B uses the answer category "role of earnings situation" of the same survey question.

Table A.5: Robustness: Baseline Estimates Excl. Reunification Period

	(1)	(2)	(3)	(4)	(5)
<i>Panel (A): Downward Revision</i>					
A1: Tax Hike Indicator: $\mathbb{1}(\Delta tax_{m,t} > 0)$					
	0.026	0.031	0.033	0.039	0.049
	(0.013)	(0.013)	(0.014)	(0.014)	(0.014)
Constant	0.540	0.539	0.539	0.539	0.538
	(0.005)	(0.005)	(0.001)	(0.001)	(0.001)
A2: Tax Hike in Percentage Points: $\Delta tax_{m,t}$					
	0.008	0.019	0.022	0.030	0.038
	(0.012)	(0.012)	(0.012)	(0.012)	(0.013)
Constant	0.541	0.540	0.540	0.540	0.539
	(0.005)	(0.005)	(0.001)	(0.001)	(0.001)
Observations	25960	25960	25911	25911	25911
<i>Panel (B): Log Revision Ratio</i>					
B1: Tax Hike Indicator: $\mathbb{1}(\Delta tax_{m,t} > 0)$					
	-0.039	-0.049	-0.035	-0.046	-0.062
	(0.020)	(0.019)	(0.020)	(0.020)	(0.022)
Constant	-0.039	-0.038	-0.039	-0.038	-0.037
	(0.007)	(0.007)	(0.001)	(0.001)	(0.001)
B2: Tax Hike in Percentage Points: $\Delta tax_{m,t}$					
	-0.031	-0.051	-0.047	-0.062	-0.073
	(0.017)	(0.016)	(0.018)	(0.018)	(0.021)
Constant	-0.039	-0.038	-0.038	-0.037	-0.037
	(0.007)	(0.007)	(0.001)	(0.001)	(0.001)
Observations	25310	25310	25255	25255	25255
Firm FE	-	-	✓	✓	✓
Year FE	-	✓	-	✓	-
Year × State FE	-	-	-	-	✓
Year × Industry FE	-	-	-	-	✓

Notes: This table re-estimates our baseline results from Table 1.2, excluding the years between the reunification of Germany in 1990 and the end of the government of Helmut Kohl in 1998, i.e., a period when many subsidy programs for investment, especially in East Germany, were in place that might have influenced investment decisions of West German firms. “Downward Revision” is an indicator that is one if the fraction of realized investment over planned investment is below one. “Log Revision Ratio” is the natural logarithm of this ratio. “Tax Hike Indicator” is an indicator that is one if the local corporate tax rate is higher than in the year before. “Tax Hike” is the change in the local corporate tax rate in percentage points compared to the previous year. Industry fixed effects refer to the ifo industry classification, comparable to two-digit NACE industries. Standard errors in parentheses are clustered at the municipality level.

Table A.6: Treatment Effect Heterogeneity: State Dependence

	Downward Revision				Log Revision Ratio			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<i>Panel (A): Baseline Recession Definition by the German Council of Economic Experts</i>								
Tax Hike Indicator ×								
No Recession	0.018 (0.012)	0.021 (0.013)			-0.011 (0.018)	-0.019 (0.020)		
Recession	0.062 (0.022)	0.069 (0.024)			-0.084 (0.034)	-0.086 (0.036)		
Tax Hike ×								
No Recession			0.013 (0.011)	0.015 (0.012)			-0.019 (0.016)	-0.024 (0.019)
Recession			0.037 (0.016)	0.043 (0.018)			-0.064 (0.027)	-0.063 (0.028)
Constant	0.536 (0.001)	0.535 (0.001)	0.536 (0.001)	0.536 (0.001)	-0.033 (0.001)	-0.032 (0.001)	-0.033 (0.001)	-0.032 (0.001)
H0: Coefficients Equal	0.069	0.074	0.201	0.195	0.059	0.105	0.16	0.266
<i>Panel (B): Alternative Recession Definition by Negative Year-on-Year Real GDP Growth</i>								
Tax Hike Indicator ×								
No Recession	0.019 (0.011)	0.020 (0.012)			-0.021 (0.017)	-0.024 (0.018)		
Recession	0.089 (0.028)	0.112 (0.030)			-0.079 (0.047)	-0.107 (0.051)		
Tax Hike ×								
No Recession			0.013 (0.010)	0.013 (0.011)			-0.024 (0.015)	-0.024 (0.017)
Recession			0.066 (0.024)	0.084 (0.026)			-0.089 (0.039)	-0.107 (0.042)
Constant	0.536 (0.001)	0.535 (0.001)	0.536 (0.001)	0.536 (0.001)	-0.033 (0.001)	-0.032 (0.001)	-0.033 (0.001)	-0.032 (0.001)
H0: Coefficients Equal	0.017	0.004	0.038	0.012	0.243	0.126	0.12	0.065
Observations	35310	35310	35310	35310	34421	34421	34421	34421
Firm FE	✓	✓	✓	✓	✓	✓	✓	✓
Year FE	✓	-	✓	-	✓	-	✓	-
Year × State FE	-	✓	-	✓	-	✓	-	✓
Year × Industry FE	-	✓	-	✓	-	✓	-	✓

Notes: This table reports estimates from linear regressions based on Equation (1.1), where the tax hike treatment is split into recession and non-recession years. In Panel (A), 1980-1982, 1992-1993, 2001-2003, and 2008-2009 are classified as recession years as defined by the German Council of Economic Experts. In Panel (B), 1982, 1993, 2002, 2003, and 2009 are classified as recession years as these years showed negative real GDP growth according to World Bank data: <https://data.worldbank.org/indicator/NY.GDP.MKTP.KD.ZG?locations=DE>. The dependent variable is based on the ratio of realized investments over planned investments (elicited in fall the year before). “Downward Revision” is an indicator that is one if the ratio is below one. “Log Revision Ratio” is the log of the ratio. “Tax Hike Indicator” is an indicator that is one if the local corporate tax rate is higher than in the year before. “Tax Hike” is the change in the local corporate tax rate in percentage points compared to the previous year. Industry fixed effects are at the ifo industry classification level that is comparable to the level of two-digit NACE industries. The p-values at the bottom of each panel indicate whether the coefficients are statistically different from each other. Standard errors in parentheses are clustered at the municipality level.

Table A.7: Difference-in-Differences: Investment Revisions after a Tax Hike

	Log Revision Ratio				
	(1)	(2)	(3)	(4)	(5)
Net-of-Tax Change	1.849 (1.119)	2.617 (1.070)	2.217 (1.177)	2.764 (1.138)	3.032 (1.312)
Constant	-0.033 (0.006)	-0.033 (0.006)	-0.033 (0.001)	-0.033 (0.001)	-0.032 (0.001)
Observations	34421	34421	34421	34421	34421
Firm FE	-	-	✓	✓	✓
Year FE	-	✓	-	✓	-
Year × State FE	-	-	-	-	✓
Year × Industry FE	-	-	-	-	✓

Notes: This table reports estimates from linear regressions of the log revision ratio on the percent change in the net-of-tax rate, defined as $\log(1 - \tau_t) - \log(1 - \tau_{t-1})$, as main explanatory variable. Industry fixed effects refer to the ifo industry classification, comparable to two-digit NACE industries. Standard errors in parentheses are clustered at the municipality level.

Table A.8: Treatment Effect Heterogeneity: Volatility of Revenue Growth

	Downward Revision				Log Revision Ratio			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Tax Hike Indicator ×								
Low Revenue Growth Volatility	0.029 (0.015)	0.034 (0.015)			-0.012 (0.021)	-0.020 (0.022)		
High Revenue Growth Volatility	0.029 (0.015)	0.032 (0.016)			-0.045 (0.023)	-0.050 (0.024)		
Tax Hike Indicator ×								
Low Revenue Growth Volatility			0.022 (0.013)	0.024 (0.014)			-0.016 (0.018)	-0.019 (0.020)
High Revenue Growth Volatility			0.022 (0.013)	0.025 (0.014)			-0.052 (0.020)	-0.054 (0.023)
Constant	0.536 (0.001)	0.535 (0.001)	0.536 (0.001)	0.536 (0.001)	-0.033 (0.001)	-0.033 (0.001)	-0.033 (0.001)	-0.033 (0.001)
H0: Coefficients Equal: p-value	0.993	0.912	0.987	0.967	0.275	0.331	0.18	0.227
Observations	35155	35151	35155	35151	34281	34277	34281	34277
Firm FE	✓	✓	✓	✓	✓	✓	✓	✓
Year FE	✓	-	✓	-	✓	-	✓	-
Year × State FE	-	✓	-	✓	-	✓	-	✓
Year × Industry FE	-	✓	-	✓	-	✓	-	✓

Notes: This table reports estimates from linear regressions based on Equation (1.1), where the tax hike treatment is split into firms with low and high revenue growth volatility (split at median of firm-level standard deviation in revenue growth elicited in the ifo Investment Survey). The dependent variable is based on the ratio of realized investments over planned investments (elicited in fall the year before). “Downward Revision” is an indicator that is one if the ratio is below one. “Log Revision Ratio” is the log of the ratio. “Tax Hike Indicator” is an indicator that is one if the local corporate tax rate is higher than in the year before. “Tax Hike” is the change in the local corporate tax rate in percentage points compared to the previous year. Industry fixed effects are at the ifo industry classification level that is comparable to the level of two-digit NACE industries. The p-values at the bottom of each panel indicate whether the coefficients are statistically different from each other. Standard errors in parentheses are clustered at the municipality level.

Table A.9: Treatment Effect Heterogeneity: Current Revenue Growth I

	Downward Revision				Log Inv. Revision			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Tax Hike Indicator ×								
No Recession ×								
No Strong Revenue Drop	0.018 (0.013)	0.024 (0.013)			-0.013 (0.019)	-0.022 (0.021)		
Strong Revenue Drop	0.009 (0.034)	-0.002 (0.034)			0.010 (0.053)	0.009 (0.054)		
Recession ×								
No Strong Revenue Drop	0.072 (0.026)	0.080 (0.029)			-0.085 (0.040)	-0.088 (0.042)		
Strong Revenue Drop	0.032 (0.037)	0.037 (0.038)			-0.067 (0.064)	-0.066 (0.068)		
Strong Revenue Drop	0.105 (0.008)	0.094 (0.009)	0.105 (0.008)	0.094 (0.008)	-0.190 (0.014)	-0.172 (0.014)	-0.191 (0.014)	-0.173 (0.014)
Tax Hike ×								
No Recession ×								
No Strong Revenue Drop			0.014 (0.011)	0.017 (0.013)			-0.021 (0.017)	-0.027 (0.020)
Strong Revenue Drop			0.007 (0.030)	-0.006 (0.031)			-0.004 (0.046)	0.002 (0.048)
Recession ×								
No Strong Revenue Drop			0.047 (0.019)	0.054 (0.022)			-0.077 (0.031)	-0.077 (0.034)
Strong Revenue Drop			0.004 (0.031)	0.008 (0.033)			-0.019 (0.052)	-0.017 (0.056)
Constant	0.518 (0.002)	0.520 (0.002)	0.519 (0.002)	0.520 (0.002)	-0.002 (0.003)	-0.004 (0.003)	-0.002 (0.002)	-0.004 (0.002)
H0: Coefficients Equal: p-value	0.379	0.354	0.243	0.256	0.817	0.787	0.347	0.376
Observations	35138	35138	35138	35138	34257	34257	34257	34257
Firm FE	✓	✓	✓	✓	✓	✓	✓	✓
Year FE	✓	-	✓	-	✓	-	✓	-
Year × State FE	-	✓	-	✓	-	✓	-	✓
Year × Industry FE	-	✓	-	✓	-	✓	-	✓

Notes: This table reports estimates from linear regressions based on Equation (1.1), where the tax hike treatment effect is estimated separately for each combination of recession and non-recession years and indicators of strong and weak revenue drops. A strong revenue drop is defined as a decline in revenues by more than 10% compared to the previous year. The dependent variable is based on the ratio of realized investments over planned investments (elicited in fall the year before). “Downward Revision” is an indicator that is one if the ratio is below one. “Log Revision Ratio” is the log of the ratio. “Tax Hike Indicator” is an indicator that is one if the local corporate tax rate is higher than in the year before. “Tax Hike” is the change in the local corporate tax rate in percentage points compared to the previous year. 1980-1982, 1992-1993, 2001-2003, and 2008-2009 are classified as recession years as defined by the German Council of Economic Experts. Industry fixed effects are at the ifo industry classification level that is comparable to the level of two-digit NACE industries. The p-values at the bottom of each panel indicate whether the coefficients are statistically different from each other. Standard errors in parentheses are clustered at the municipality level.

Table A.10: Treatment Effect Heterogeneity: Current Revenue Growth II

	Downward Revision				
	(1)	(2)	(3)	(4)	(5)
Tax Hike Indicator ×					
No Recession ×					
No Strong Revenue Drop	0.024 (0.012)	0.018 (0.013)	0.024 (0.013)	0.027 (0.017)	0.034 (0.018)
Strong Revenue Drop	0.001 (0.033)	0.009 (0.034)	-0.002 (0.034)	-0.035 (0.057)	-0.009 (0.061)
Recession ×					
No Strong Revenue Drop	0.058 (0.026)	0.072 (0.026)	0.080 (0.029)	0.078 (0.035)	0.078 (0.039)
Strong Revenue Drop	0.011 (0.034)	0.032 (0.037)	0.037 (0.038)	-0.036 (0.073)	-0.020 (0.076)
Strong Revenue Drop	0.122 (0.008)	0.105 (0.008)	0.094 (0.009)	0.087 (0.013)	0.072 (0.014)
Constant	0.515 (0.005)	0.518 (0.002)	0.520 (0.002)	0.497 (0.002)	0.498 (0.002)
<i>N</i>	35139	35138	35138	21255	21193
Firm FE	-	✓	✓	✓	✓
Year FE	✓	✓	-	✓	-
Year × State FE	-	-	✓	-	✓
Year × Industry FE	-	-	✓	-	✓
Exclude Labor Drop	-	-	-	Yes, > 5%	Yes, > 5%

Notes: This table reports estimates from linear regressions based on Equation (1.1), where the tax hike treatment effect is estimated separately for each combination of recession and non-recession years, as well as indicators of strong and weak revenue drop observations. A strong revenue drop is defined as a decline in revenue by more than 10% compared to the previous year. In Columns (4) and (5), we drop firm observations that have a decrease in employees by more than 5% compared to the previous year. “Downward Revision” is an indicator that is one if the ratio of realized investments over planned investments (elicited in fall the year before) is below one. “Tax Hike Indicator” is an indicator that is one if the local corporate tax rate is higher than in the year before. 1980-1982, 1992-1993, 2001-2003, and 2008-2009 are classified as recession years as defined by the German Council of Economic Experts. Industry fixed effects are at the ifo industry classification level that is comparable to the level of two-digit NACE industries. Standard errors in parentheses are clustered at the municipality level.

Table A.11: Treatment Effect Heterogeneity: Financial Constraints

	Downward Revision				Log Inv. Revision			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Tax Hike Indicator ×								
No Recession ×								
No Fin. Constr.	0.025 (0.015)	0.024 (0.016)			-0.009 (0.022)	-0.011 (0.024)		
Fin. Constr.	-0.004 (0.055)	-0.017 (0.057)			-0.051 (0.116)	-0.045 (0.121)		
Recession ×								
No Fin. Constr.	0.024 (0.029)	0.039 (0.030)			-0.040 (0.044)	-0.061 (0.047)		
Fin. Constr.	0.126 (0.065)	0.153 (0.070)			-0.066 (0.125)	-0.095 (0.127)		
Fin. Constr.	0.113 (0.017)	0.111 (0.017)	0.115 (0.017)	0.113 (0.017)	-0.225 (0.031)	-0.214 (0.032)	-0.226 (0.031)	-0.214 (0.031)
Tax Hike ×								
No Recession ×								
No Fin. Constr.			0.026 (0.014)	0.021 (0.016)			-0.023 (0.023)	-0.018 (0.025)
Fin. Constr.			-0.025 (0.050)	-0.039 (0.053)			-0.048 (0.101)	-0.045 (0.108)
Recession ×								
No Fin. Constr.			0.013 (0.023)	0.024 (0.025)			-0.019 (0.036)	-0.030 (0.040)
Fin. Constr.			0.066 (0.064)	0.089 (0.077)			-0.056 (0.107)	-0.066 (0.116)
Constant	0.550 (0.001)	0.550 (0.001)	0.551 (0.001)	0.551 (0.001)	-0.054 (0.002)	-0.054 (0.002)	-0.054 (0.002)	-0.054 (0.002)
H0: Coefficients Equal: p-value	0.168	0.138	0.449	0.426	0.849	0.802	0.746	0.767
Observations	23661	23640	23661	23640	23123	23101	23123	23101
Firm FE	✓	✓	✓	✓	✓	✓	✓	✓
Year FE	✓	-	✓	-	✓	-	✓	-
Year × State FE	-	✓	-	✓	-	✓	-	✓
Year × Industry FE	-	✓	-	✓	-	✓	-	✓

Notes: This table reports estimates from linear regressions based on Equation (1.1), where the tax hike treatment effect is estimated separately for each combination of recession and non-recession years, as well as indicators on whether the financing situation is reported to be a factor for a strong slowdown in investment volumes or not. To construct the financing indicator, we use a question from the fall survey (available since 1989), where firms rate on a scale from 1 (strong stimulus) to 5 (strong slowdown) different factors that influence investments in the current year, see Appendix A.2 for the exact wording. We construct an indicator that is one if a firm reports the highest category (5). The dependent variable is based on the ratio of realized investments over planned investments (elicited in fall the year before). “Downward Revision” is an indicator that is one if the ratio is below one. “Log Revision Ratio” is the log of the ratio. “Tax Hike Indicator” is an indicator that is one if the local corporate tax rate is higher than in the year before. “Tax Hike” is the change in the local corporate tax rate in percentage points compared to the previous year. 1980-1982, 1992-1993, 2001-2003, and 2008-2009 are classified as recession years as defined by the German Council of Economic Experts. Industry fixed effects are at the ifo industry classification level that is comparable to the level of two-digit NACE industries. The p-values at the bottom of each panel indicate whether the coefficients are statistically different from each other. Standard errors in parentheses are clustered at the municipality level.

Table A.12: Treatment Effect Heterogeneity: Firm Size and Settlement Structure

	Downward Revision				Log Revision Ratio			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<i>Panel (A): Heterogeneity by Firm Size</i>								
Tax Hike Indicator ×								
Small Firms	0.029 (0.017)	0.031 (0.017)			-0.031 (0.027)	-0.035 (0.028)		
Large Firms	0.028 (0.014)	0.034 (0.014)			-0.027 (0.017)	-0.036 (0.019)		
Tax Hike ×								
Small Firms			0.021 (0.015)	0.024 (0.016)			-0.040 (0.023)	-0.043 (0.025)
Large Firms			0.022 (0.012)	0.024 (0.013)			-0.030 (0.016)	-0.034 (0.018)
Constant	0.536 (0.001)	0.535 (0.001)	0.536 (0.001)	0.536 (0.001)	-0.033 (0.001)	-0.032 (0.001)	-0.033 (0.001)	-0.032 (0.001)
H0: Coefficients Equal	0.951	0.869	0.972	0.989	0.899	0.96	0.73	0.756
<i>Panel (B): Heterogeneity by Settlement Structure</i>								
Tax Hike Indicator ×								
Urban Area	0.027 (0.012)	0.030 (0.013)			-0.020 (0.018)	-0.028 (0.019)		
Rural Area	0.037 (0.023)	0.043 (0.025)			-0.069 (0.034)	-0.070 (0.037)		
Tax Hike ×								
Urban Area			0.019 (0.011)	0.022 (0.012)			-0.023 (0.016)	-0.027 (0.019)
Rural Area			0.029 (0.017)	0.032 (0.019)			-0.072 (0.026)	-0.072 (0.028)
Constant	0.536 (0.001)	0.535 (0.001)	0.536 (0.001)	0.536 (0.001)	-0.033 (0.001)	-0.032 (0.001)	-0.033 (0.001)	-0.032 (0.001)
H0: Coefficients Equal	0.688	0.641	0.64	0.649	0.199	0.314	0.106	0.184
Observations	35310	35310	35310	35310	34421	34421	34421	34421
Firm FE	✓	✓	✓	✓	✓	✓	✓	✓
Year FE	✓	-	✓	-	✓	-	✓	-
Year × State FE	-	✓	-	✓	-	✓	-	✓
Year × Industry FE	-	✓	-	✓	-	✓	-	✓

Notes: This table reports estimates from linear regressions based on Equation (1.1). In Panel (A), the tax hike treatment is split into small (< 250 employees) and large (\geq 250 employees) firms. In Panel (B) the treatment variables are interacted with indicators of urban and rural areas following the classification of the Federal Institute for Research on Building, Urban Affairs and Spatial Development (BBSR) that is mainly based on population density. The dependent variable is based on the ratio of realized investments over planned investments (elicited in fall the year before). “Downward Revision” is an indicator that is one if the ratio is below one. “Log Revision Ratio” is the log of the ratio. “Tax Hike Indicator” is an indicator that is one if the local corporate tax rate is higher than in the year before. “Tax Hike” is the change in the local corporate tax rate in percentage points compared to the previous year. Industry fixed effects are at the ifo industry classification level that is comparable to the level of two-digit NACE industries. The p-values at the bottom of each panel indicate whether the coefficients are statistically different from each other. Standard errors in parentheses are clustered at the municipality level.

Table A.13: Treatment Effect Heterogeneity: Tax Hike Dynamics

	Downward Revision				Log Revision Ratio			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<i>Panel (A): Heterogeneity by the Frequency of Tax Hikes</i>								
Tax Hike Indicator ×								
Few Tax Hikes	0.024 (0.015)	0.028 (0.016)			-0.014 (0.026)	-0.020 (0.029)		
Many Tax Hikes	0.031 (0.015)	0.036 (0.015)			-0.038 (0.020)	-0.046 (0.021)		
Tax Hike ×								
Few Tax Hikes			0.021 (0.013)	0.021 (0.014)			-0.030 (0.019)	-0.029 (0.022)
Many Tax Hikes			0.021 (0.013)	0.027 (0.014)			-0.038 (0.019)	-0.046 (0.022)
Constant	0.536 (0.001)	0.535 (0.001)	0.536 (0.001)	0.536 (0.001)	-0.033 (0.001)	-0.032 (0.001)	-0.033 (0.001)	-0.032 (0.001)
H0: Coefficients Equal:								
p-value	0.721	0.733	0.998	0.775	0.464	0.468	0.762	0.584
Observations	35310	35310	35310	35310	34421	34421	34421	34421
<i>Panel (B): Heterogeneity by Occurrence of a Tax Hike in the Last 5 Years</i>								
Tax Hike Indicator ×								
≥ 1 Hike in Last 5 Years	0.039 (0.019)	0.052 (0.021)			-0.050 (0.026)	-0.067 (0.029)		
No Hike in Last 5 Years	0.019 (0.013)	0.020 (0.013)			-0.012 (0.021)	-0.014 (0.022)		
Tax Hike ×								
≥ 1 Hike in Last 5 Years			0.030 (0.016)	0.044 (0.019)			-0.043 (0.022)	-0.062 (0.027)
No Hike in Last 5 Years			0.013 (0.011)	0.013 (0.012)			-0.023 (0.017)	-0.021 (0.019)
Constant	0.540 (0.001)	0.540 (0.001)	0.541 (0.001)	0.541 (0.001)	-0.040 (0.001)	-0.039 (0.001)	-0.039 (0.001)	-0.039 (0.001)
H0: Coefficients Equal:								
p-value	0.358	0.195	0.386	0.155	0.257	0.145	0.489	0.204
Observations	33220	33201	33220	33201	32375	32356	32375	32356
Firm FE	✓	✓	✓	✓	✓	✓	✓	✓
Year FE	✓	-	✓	-	✓	-	✓	-
Year × State FE	-	✓	-	✓	-	✓	-	✓
Year × Industry FE	-	✓	-	✓	-	✓	-	✓

Notes: This table reports estimates from linear regressions based on Equation 1.1. In Panel (A), the tax hike treatment variable is interacted with dummies splitting the sample into municipalities with few (≤ 3) and many (> 3) tax hikes over the entire sample period. In Panel (B), the tax hike treatment is split into cases where at least one tax hike has already occurred in the previous five years and where no tax hike occurred in the previous five years in the respective municipality. The dependent variable is based on the ratio of realized investments over planned investments (elicited in fall the year before). “Downward Revision” is an indicator that is one if the ratio is below one. “Log Revision Ratio” is the log of the ratio. “Tax Hike Indicator” is an indicator that is one if the local corporate tax rate is higher than in the year before. “Tax Hike” is the change in the local corporate tax rate in percentage points compared to the previous year. Industry fixed effects are at the ifo industry classification level that is comparable to the level of two-digit NACE industries. Standard errors in parentheses are clustered at the municipality level.

A.4 Back-of-the-Envelope Calculation

In the following, we present the assumptions underlying the back-of-the-envelope calculation used to approximate the investment loss for each additional Euro of tax revenue.

The median firm in our sample generates yearly revenues of 45 million Euro. Among the subsample of firms that can be linked to information on the cashflow/revenue ratio balance sheet data, the median profit margin is 4.4%. Assuming that this figure corresponds to all firms in the sample, this translates into 1.98 million Euro of aggregate profits. A one percentage point increase in the LBT increases the tax burden of the median firm—and thus overall tax revenues—by 19,800 Euro. Moreover, the median investment-revenue ratio amounts to 3% in the microdata of the ifo Investment Survey. Hence, the median firm invests approximately 1.4 million Euro each year. Given the estimated semi-elasticity of 3 (see Section 3.4), a one percentage point increase in the LBT is associated with decreased investment of the median firm by roughly 42,000 Euro. Finally, dividing 42,000 by 19,800 gives that 2.12 Euro of investment volume is lost for each additional Euro of tax revenue. In crisis years, we estimate a semi-elasticity of investments with respect to the LBT rate of 6. Assuming that the relation between the profit margin and investment-revenue ratio is the same in a recession, investments even decrease by 4.24 Euro for each additional Euro of tax revenue.²

It is furthermore necessary to also take the (long-term) behavioral response of firms into account: as tax increases decrease firm investment, future firm profits should also be reduced, resulting in lower tax revenues of the municipalities. Unfortunately, we cannot estimate the elasticity of firms' profits with respect to changes in investment based on our data. We circumvent this constraint by separately calculating the behavioral response for reasonable lower and upper bounds of this elasticity, i.e., assuming that foregone investment maps into foregone future profits with half of the median profit margin (2.2%) or with five times the median profit margin (22%). For the median firm, which lowered investment by 42,000 Euro, this translates into lower profits between 924 Euro and 9,240 Euro. As the average LBT rate is approximately 15%, this leads to an additional reduction in tax revenues between 139 Euro and 1,386 Euro. Taken together, we approximate that incorporating the behavioral response increases the investment loss for each additional Euro of tax revenue from 2.12 Euro to an estimate in the range between 2.14 Euro ($42000/(19800-139)$) and 2.28 Euro ($42000/(19800-1386)$).

² In fact, the profit margin decreases slightly more than the investment-revenue ratio in recessions. Incorporating this relation in the calculation would lead to an even higher loss of investments for each additional Euro of tax revenue in recessions.

From this approximation of the behavioral response, we can also derive the marginal value of public funds (MVPF) in the spirit of Hendren and Sprung-Keyser (2020), given as:

$$MVPF = \frac{\text{Beneficiaries' Willingness to Pay}}{\text{Net Cost to Government}}$$

In our setting, firms are the beneficiaries and their willingness to pay is equal to the change of the tax burden. The net cost of the government equals the change of tax revenues plus the additional revenue changes via the behavioral response. According to this, our estimates point at a MVPF in the range between $(19,800)/(19,800 - 139) = 1.01$ and $1.08 = (19,800)/(19,800 - 1,386)$, i.e., slightly above one.

A.5 Calculation of Effective Tax Rates

This appendix describes how we calculate effective tax rates used in the alternative specification presented in Section 1.4.2 based on the statutory LBT rates used in the main specification. When talking about effective tax rates, we always refer to marginal—rather than average—tax rates.³ Nicodème (2001) provides a helpful overview on different approaches of computing effective tax rates. Our procedure is guided by the classic framework of Hall and Jorgenson (1967), as, e.g., recently applied by Furno (2022). The key difference between the two concepts is that while the statutory tax rate is the one imposed by law on taxable profits, the effective tax rate is the percentage of profits actually paid by a company after taking into account deductions including depreciation of assets, exemptions, tax credits, and preferential rates. Note that in the case of the German case, the most important feature are depreciation rules (denoted with z below) while the other components play a negligible role. To compute effective tax rates in the setting of the LBT, we proceed as follows:

- We first obtain depreciation schedules separately for machinery m and buildings b , the two main types of investment for which different depreciation rules apply. To do so, we use information from the Oxford Corporate Tax Database (<https://oxfordtax.sbs.ox.ac.uk/cbt-tax-database>). Indeed, over our sample period of almost 40 years, the depreciation rules have changed repeatedly.
- To illustrate this change over time, we calculate the present discounted value (PDV) of a depreciation, denoted by z_m (for machinery) and z_b (for buildings), respectively. As the choice of the adequate discount rate is not innocuous in our setting, we employ the following two different specifications, whose resulting z_m and z_b are depicted in Figure A.18:
 - I. We follow Zwick and Mahon (2017) and set the discount rate to 7%.
 - II. We use time-varying interest rates for discounting to accommodate for the fact that over the sample period the interest rates on firm loans have been declining substantially, from close to 10% in 1980 to less than 2% in recent years, with considerable

³ Note that in the German context marginal and average tax rates are approximately the same, since the tax rate is flat and there are only a few exceptions, e.g., no tax credits.

variation in between.⁴ In contrast to most other studies in the literature that rely on a single or few tax reforms within shorter time periods, time-variation in interest rates may have large implications for the PDV of a depreciation in our analysis that covers a period of almost 40 years.

- Next, we calculate the combined depreciation schedule, z , for each firm, i.e., a weighted average of z_m and z_b based on firms' respective share of investment in machinery and buildings. However, as we do not observe these investment shares in machinery and buildings for each firm in all years of the survey, we have to impute these values (in some years). We consider two distinct specifications for this imputation:

- I. In each year, we assign the average share of investment in machinery and buildings based on aggregate data from the Federal Statistical Office of Germany (only available since 1990, imputed for the years before). This way, the investment shares vary over time, but are the same for all firms in our sample in a given year. Across years, the average share of investment in machinery amounts to 88%.
- II. We use the firm-specific share of investment in machinery and buildings reported to the ifo Investment Survey. As this information is provided less frequently to the IVS compared to the overall level of investment, we use the firm-specific mean across all years if firms reported machinery and building investments at least three times. To retain the sample size, we replace missing values by the values obtained from method I.

- Having obtained the depreciation schedule z , the effective tax rate is then given by

$$\tau_{eff} = 1 - \frac{1 - \tau}{1 - z * \tau},$$

which only depends on z and the statutory LBT rate τ in the German case as there are no relevant tax credits in the LBT that would complicate the calculation.

- Finally, we calculate the change in the effective tax rate if a tax hike takes place in a given year. Here, we set z in both years (t_0 and t_{-1}) equal to the value of the tax hike

⁴ The time series of average interest rates on firm loans displayed in Figure A.17 builds on three different charts provided by the Deutsche Bundesbank (German Central Bank), as the effective interest rate for non-financial corporations is only available since 2003. The breaks are indicated by the dashed vertical lines. Over our entire sample period, the average interest rate according to this graph has been 5.1%.

year. Thereby, we isolate the effect of tax changes by making the arguably reasonable assumption that firms know the z value of the next year when forming their investment plans.⁵

The results show that, across all specifications, the variation captured by changes in effective tax rates is strongly associated with the underlying changes in the LBT rate as plotted in Figure A.19.

In alternative specifications, we also express changes in the costs for investment in terms of the user cost of capital instead of effective tax rates. This only requires a simple transformation:


$$UserCost = \frac{1 - z * \tau}{1 - \tau}$$

Hence,


$$\tau_{eff} = 1 - UserCost^{-1},$$

which means that switching from effective tax rates to a user cost approach will not impact our results apart from rescaling the magnitude of the coefficients. That the user cost of capital yields virtually the same results as using effective tax rates is also visible in Figure A.20, which plots the change in the user cost of capital against the change in effective tax rates for all tax hikes in our sample.

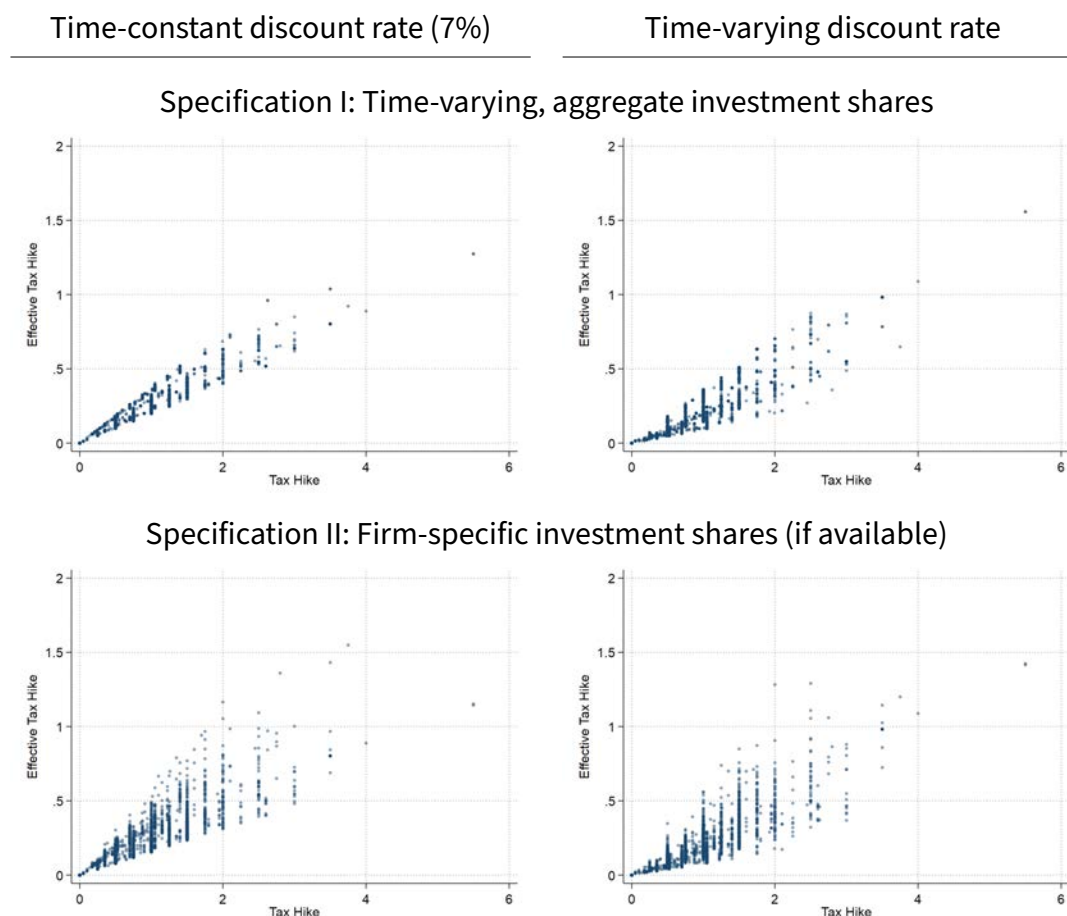
⁵ Note that while German municipalities have the power and discretion to change τ , i.e. the LBT rate, they cannot change the depreciation rules z that are determined at the Federal level.

Figure A.17: Time Series of Average Interest Rate on Loans for Firms

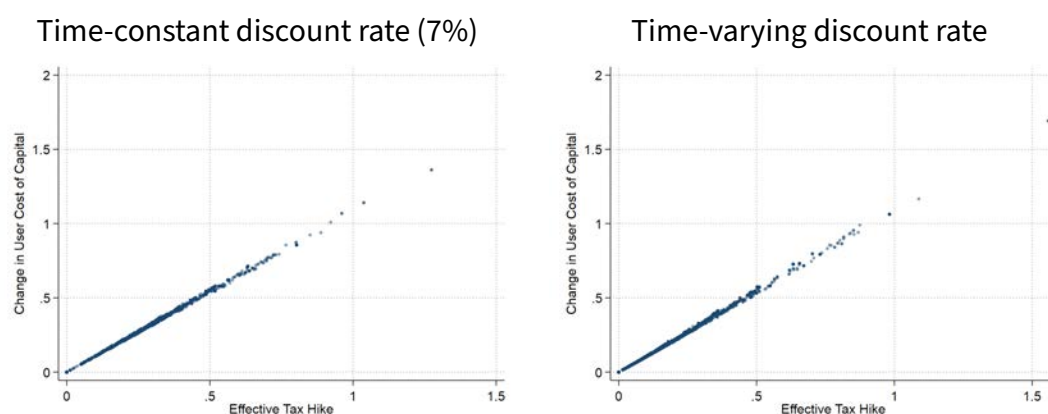
Notes: This figure shows a time series of the average lending rate for firms. From 2003 onward, the effective interest rate for non-financial corporations is used. For the year 1997 to 2002, the effective interest rate to firms for loans between 500,000 and 5 million Euro is used and adjusted upwards (roughly 1 p.p.) to ensure a smooth transition in 2003. For the years 1980 to 1996, the discount rate of the Deutsche Bundesbank is used and adjusted upwards (roughly 4 p.p.) to ensure a smooth transition in 1997. The two dashed vertical lines indicate the breaks in the time series. Source: Deutsche Bundesbank.

Figure A.18: Present Discounted Value of Depreciation: 7% vs Time-Varying Interest Rate

Notes: This figure shows values of the present discounted value of depreciation for machinery (z_m) and buildings (z_b) in the period 1980 to 2018. Depreciation schedules are obtained from the Oxford Corporate Tax Database. The solid line assumes a time-constant discount rate of 7% following Zwick and Mahon (2017), the dashed line calculates the PDV based on the time-varying interest rate on firm loans as displayed in Figure A.17.

Figure A.19: Relation of Changes in LBT Rate and Changes in Effective Tax Rates

Notes: For each tax hike in our sample, this figure plots its size in terms of an effective tax hike (τ_{eff} ; y-axis) against its size as a statutory tax hike (x-axis). As we do not observe the investment shares in machinery and buildings for each firm in all years of the survey, we must impute these values. We consider two distinct specifications in which τ_{eff} is either calculated based on the average share of investment in machinery and buildings based on aggregate data from the Federal Statistical Office of Germany (Specification “I”) or on the firm-specific share of investment in machinery and buildings reported to the ifo Investment Survey whenever available (Specification “II”). In the left panel, we assume a time-constant discount rate of 7% following Zwick and Mahon (2017), in the right panel we calculate the PDV based on the time-varying interest rate on firm loans as displayed in Figure A.17.

Figure A.20: Relation of Changes in Effective Tax Rates and Changes in User Cost of Capital

Notes: This figure plots the change in the user cost of capital (multiplied with 100) against the change in effective tax rates for all tax hikes in our sample, assuming that the share of investment allocated to machinery and buildings is constant across firms, but varying over time (Specification I). In the left panel, we assume a time-constant discount rate of 7% following Zwick and Mahon (2017), in the right panel we calculate the PDV based on the time-varying interest rate on firm loans as displayed in Figure A.17.

B Appendix to Chapter 2

B.1 Additional Information on the Mikrozensus

The Microcensus (Mikrozensus, MZ) is the largest household survey in Europe. Conducted annually with a sampling fraction of 1% of all individuals who have the right of residence in Germany, it yields representative statistics on the German population. The MZ has been conducted in West Germany since 1957 and in the new federal states (East Germany) since 1991. It is planned and prepared by the Federal Statistical Office of Germany and carried out by the statistical offices of the 16 German states. The legal basis of the MZ is the Microcensus Law, which makes it compulsory for households to provide answers to the core items of the survey. The non-response rate is further minimized by repeated visits of interviewers to non-responding households and multiple possible ways for the sampled households to submit information.

Figure B.1: Illustration of the Microcensus Survey Design

Survey Wave	Rotation Quarter							...
	1	2	3	4	5	6	7	
1	✓	✓	✓	✓	X	X	X	...
2	X	✓	✓	✓	✓	X	X	...
3	X	X	✓	✓	✓	✓	X	...
4	X	X	X	✓	✓	✓	✓	...
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮


Since 1972, the MZ uses a single-stage stratified cluster sampling design. The primary sampling units typically consist of neighboring buildings (larger buildings are divided into smaller partitions). For the survey waves utilized in this paper, the target size for a cluster is 7–15 households. All households and residents in the sampled clusters are interviewed. The database used to assign households to clusters is created based on the most recent full census and updated annually using information on new construction activities. Since 1977, each cluster is assigned to a "rotation quarter" that remains in the survey for four years. Each year,

a quarter is replaced by new clusters. The survey does not follow individuals who leave their cluster, but replaces them by the new residents. The MZ survey design results in data best described as a repeated survey with partial overlap of units, as sketched in Figure B.1.

Due to data protection laws, we do only observe this panel structure in our data following wave 2011. In Section 2.4.2, we therefore cannot cluster standard errors at the level of time-constant primary sampling units. We instead cluster standard errors at the household level. As the number of households per cluster is low, the consequences for standard errors are negligible.

Sample Income Distribution and Ranks. Figure B.2 displays the sample distribution of equivalized monthly net household income and the corresponding percentile ranks in the 2011-2018 MZ data. We CPI adjust all household incomes in order to allow for meaningful aggregation of survey-years before computing ranks. Ties are broken by allocating households to the lower quantile. Our findings are insensitive to the choice of tie-breakers. Ranks are computed separately for each year within the sample of all households that have at least one co-resident child in the age range 17-21.

Figure B.2: Household Income by Percentile Rank



Notes: This figure plots equivalized net monthly household income (net of income of dependent children) by parental income rank in the 2011-2018 MZ data. Equivalization is based on the modified OECD scale. For comparison, the non-equivalized values are plotted as well. Both income measures are expressed in constant 2015 Euro.


B.2 Additional Figures and Tables

Table B.1: National Estimates for Different Age-Restrictions

Age	Gradient	Q1	Q5	Q5/Q1	A-Level Share	N
17-21	0.52 (0.004)	0.34 (0.003)	0.76 (0.003)	2.25 (0.021)	0.52	230,972
17-21 (Averaged)	0.52 (0.004)	0.34 (0.003)	0.77 (0.003)	2.26 (0.022)	0.52	230,972
17	0.53 (0.007)	0.30 (0.005)	0.73 (0.004)	2.46 (0.042)	0.49	53,324
18	0.51 (0.007)	0.35 (0.005)	0.77 (0.004)	2.18 (0.033)	0.54	51,278
19	0.51 (0.008)	0.35 (0.005)	0.77 (0.005)	2.19 (0.035)	0.53	46,747
20	0.51 (0.008)	0.35 (0.005)	0.77 (0.005)	2.19 (0.036)	0.53	42,396
21	0.52 (0.008)	0.34 (0.006)	0.77 (0.005)	2.24 (0.039)	0.52	37,227


Notes: This table reports national mobility statistics for the MZ waves 2011-2018. The gradient measures the gap in the probability of obtaining an A-Level between children at the top and the bottom of the parental income distribution. Q1 and Q5 denote the share of children obtaining an A-Level in the first and fifth quintile of parental income; Q5/Q1 is the ratio between both measures. The first row corresponds to our primary sample. The second row replicates these estimates using multi-year averages of parental income before assigning ranks. The additional rows report estimates for samples containing only children of a given age at measurement, as indicated in the first column. The standard errors in parentheses are computed as described in Section 2.3.3.

Figure B.3: A-Level Wage Premium, Years 1997-2016




Notes: This figure shows the development of the A-Level wage premium for the years 1997-2016 as computed in the MZ. We compute the A-Level wage premium by regressing the log of net monthly personal income of full-time working employees aged 30-45 on an A-Level dummy. The adjusted A-Level wage premium is computed by additionally conditioning on a set of age indicators to indirectly account for job experience.

Figure B.4: National Estimates under Different Equivalization Schemes



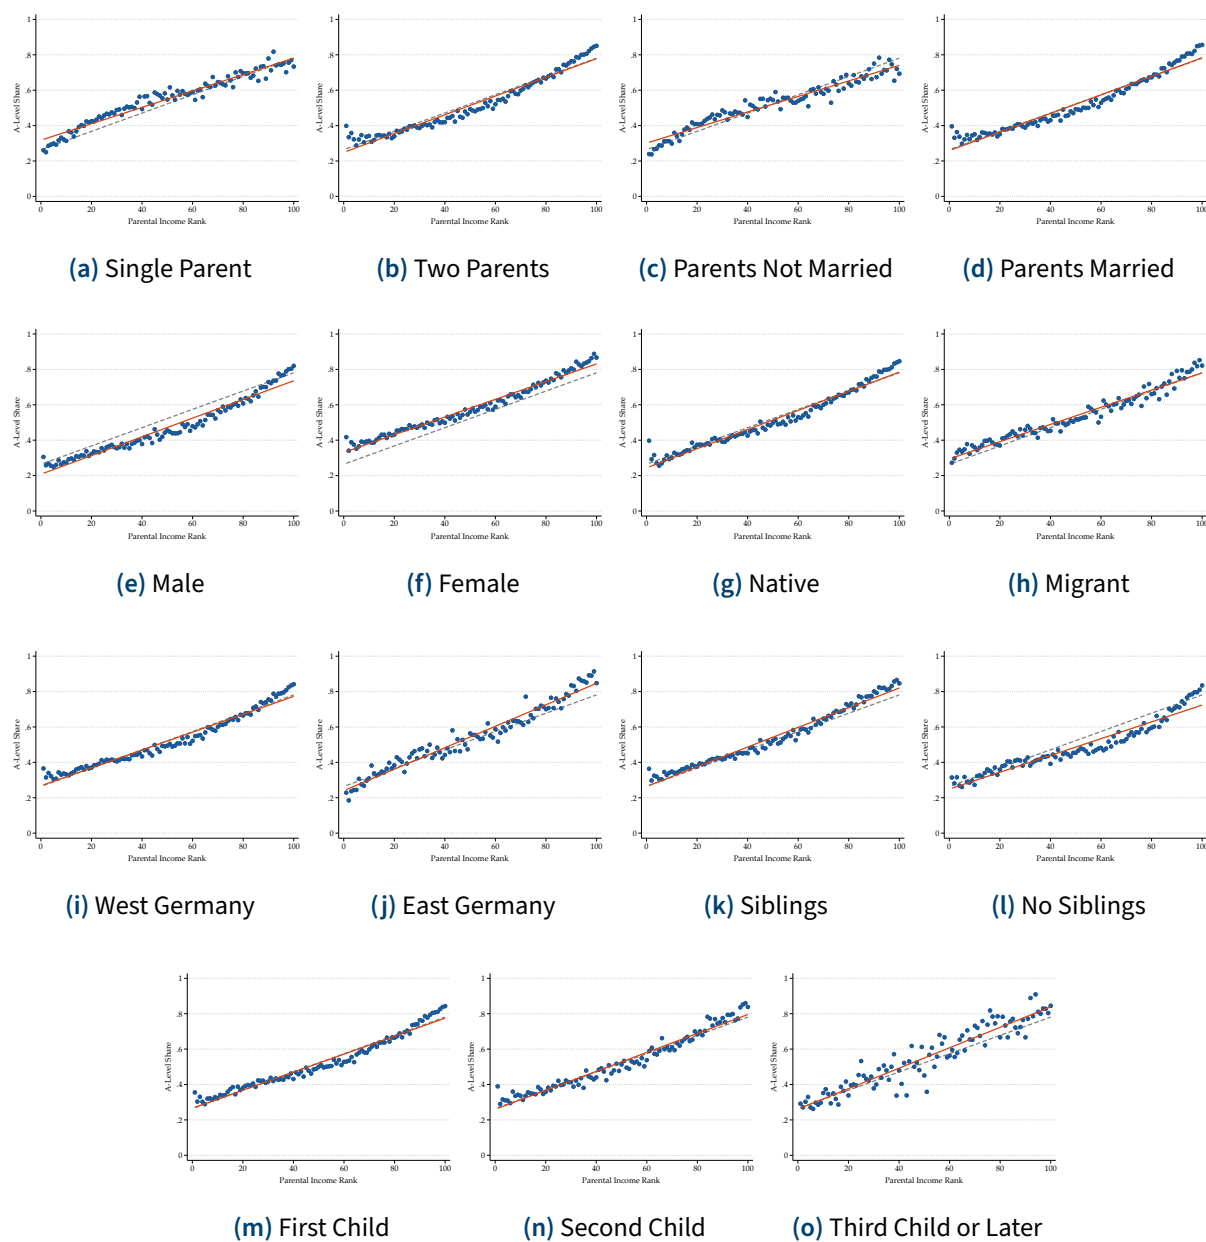
(a) No Adjustment



(b) Per Capita Adjustment


Notes: This figure shows the fraction of children aged 17-21 that are either enrolled in the upper stage of an A-Level track or have already attained an A-Level degree by percentile rank of their parents in the national income distribution based on the MZ waves 2011-2018, as well as the best linear approximation to the empirical CEF. In Panel (A), parental income is not adjusted for household size, whereas in Panel (B) income is divided by the number of household members. The OLS slopes reported in the figure are estimated using the underlying micro data.

Figure B.5: Social Mobility for Subgroups




Notes: This figure shows for different population subgroups the fraction of children aged 17-21 that are either enrolled in the upper stage of an A-Level track or have already attained an A-Level degree by percentile rank of their parents in the national income distribution based on the MZ waves 2011-2018, as well as the best linear approximation to the empirical CEF in orange. The dashed gray line plots the national gradient as a comparison. Migration background subsumes all individuals who immigrated to Germany after 1949, as well as all foreigners born in Germany and all individuals born in Germany with at least one parent who immigrated after 1949 or was born in Germany as a foreigner.


Figure B.6: Time Trend A-Level Share for Subgroups




(a) Gender




(b) Parental Education




(c) Region



(d) Parenting Status



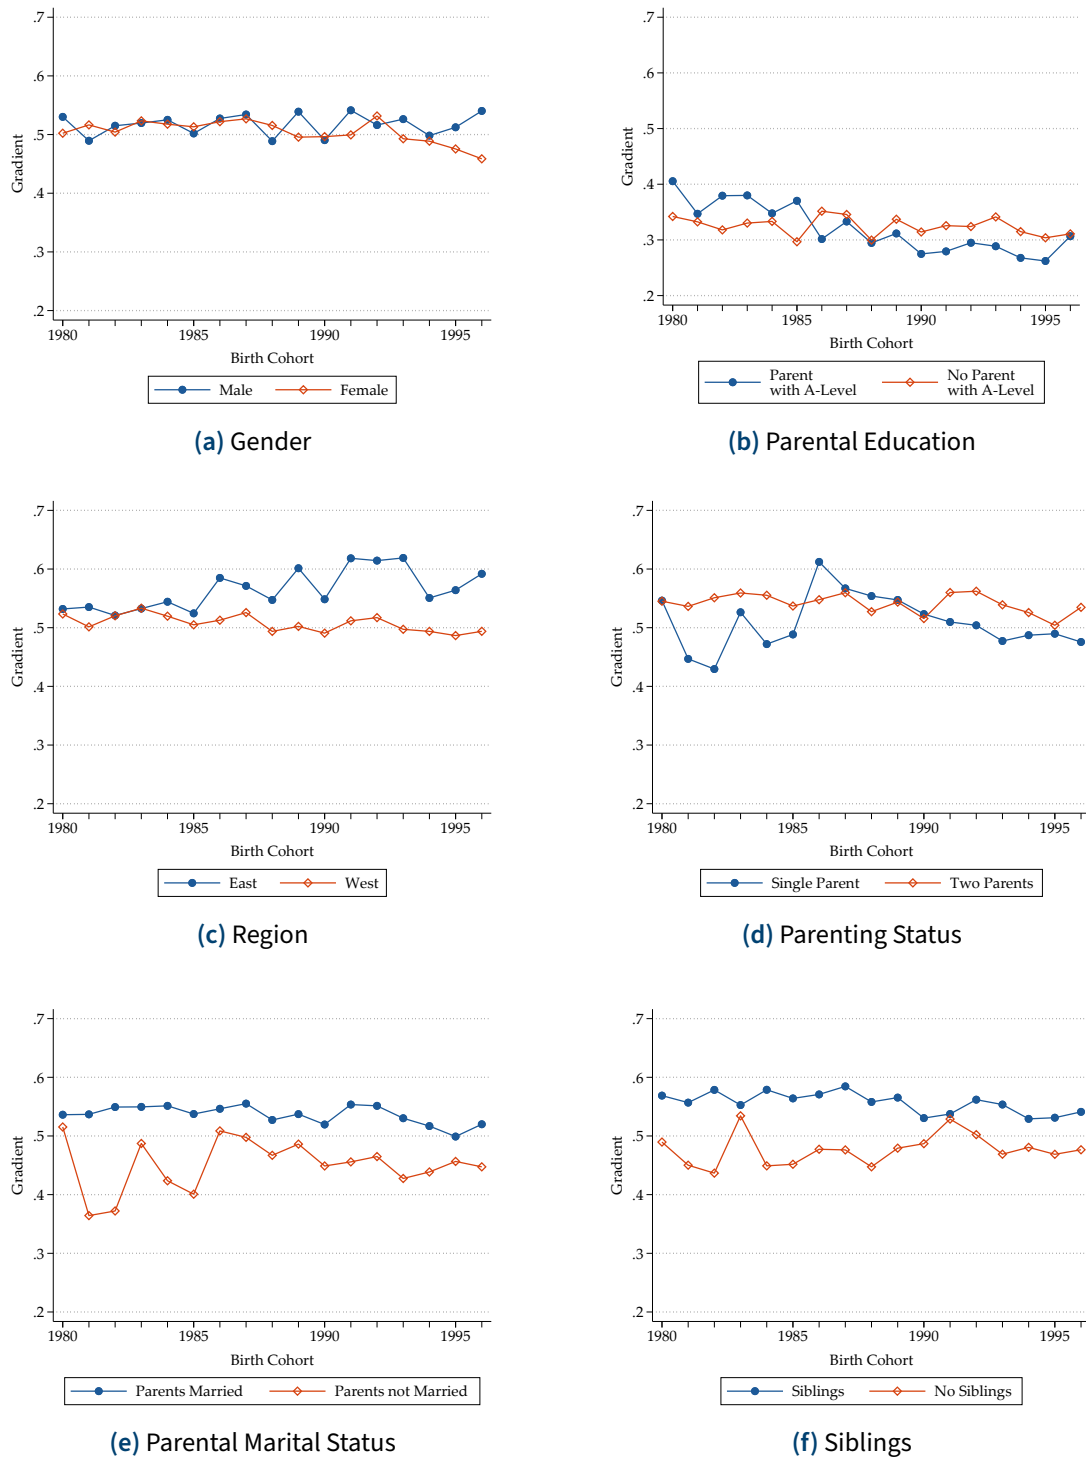
(e) Parental Marital Status



(f) Siblings

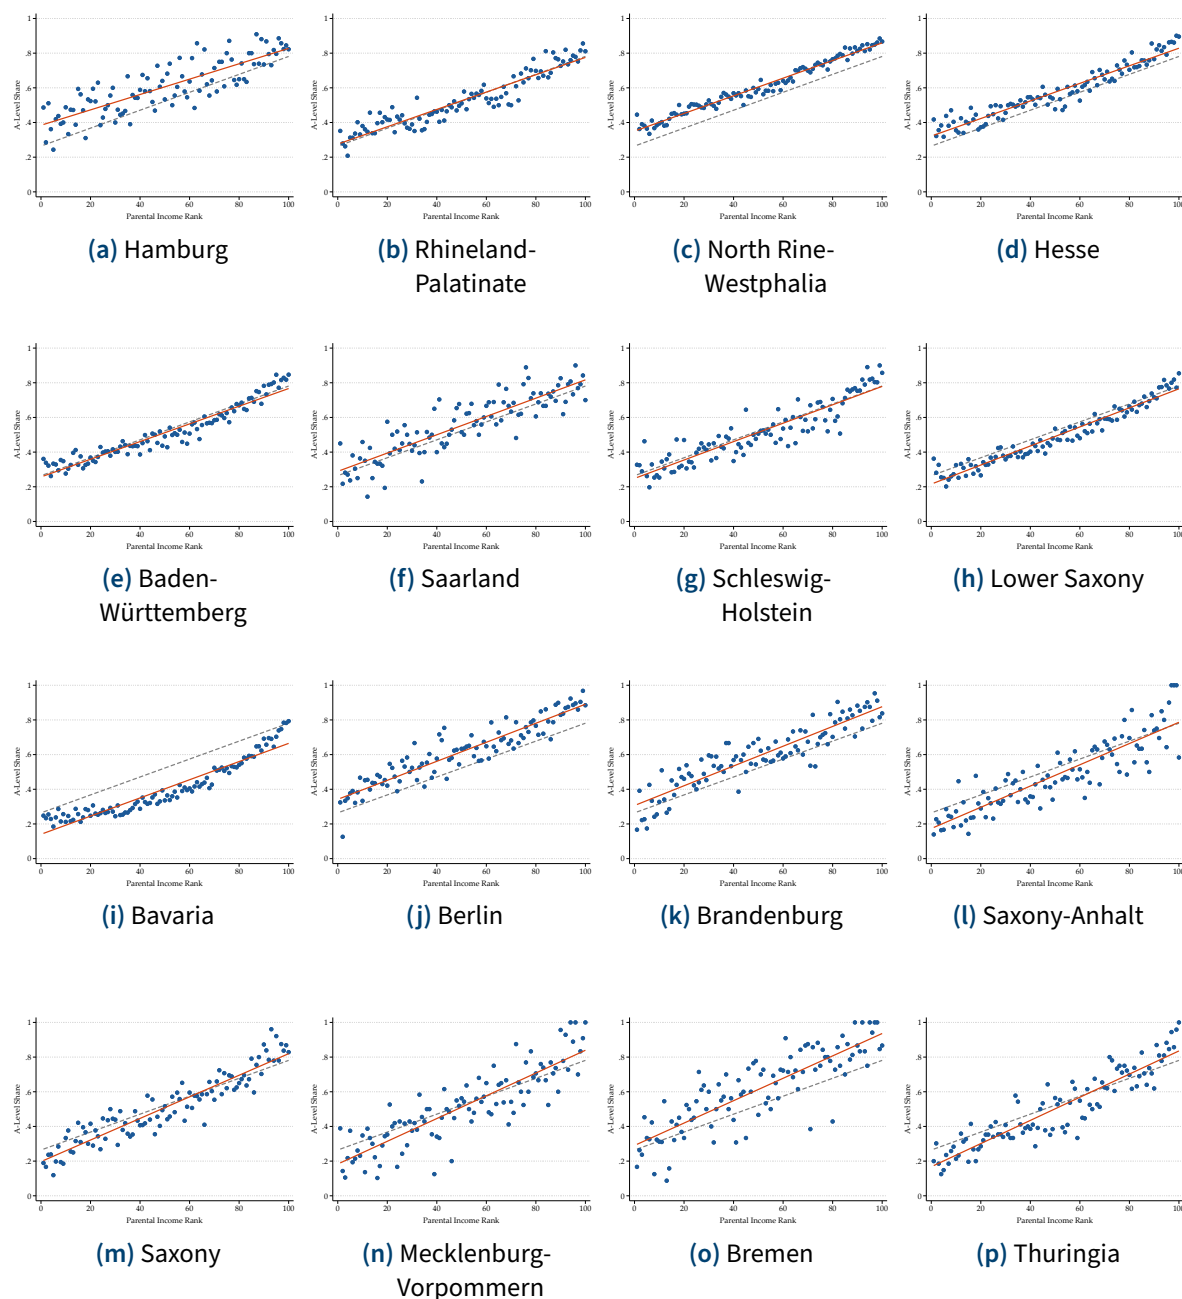
Notes: This figure shows the development of the A-Level share for different population subgroups for birth cohorts 1980-1996 in the MZ. The A-Level share is given as the fraction of children aged 17-21 that are either enrolled in the upper stage of an A-Level track or have already attained an A-Level degree.

Figure B.7: Time Trend Parental Income Gradient for Subgroups



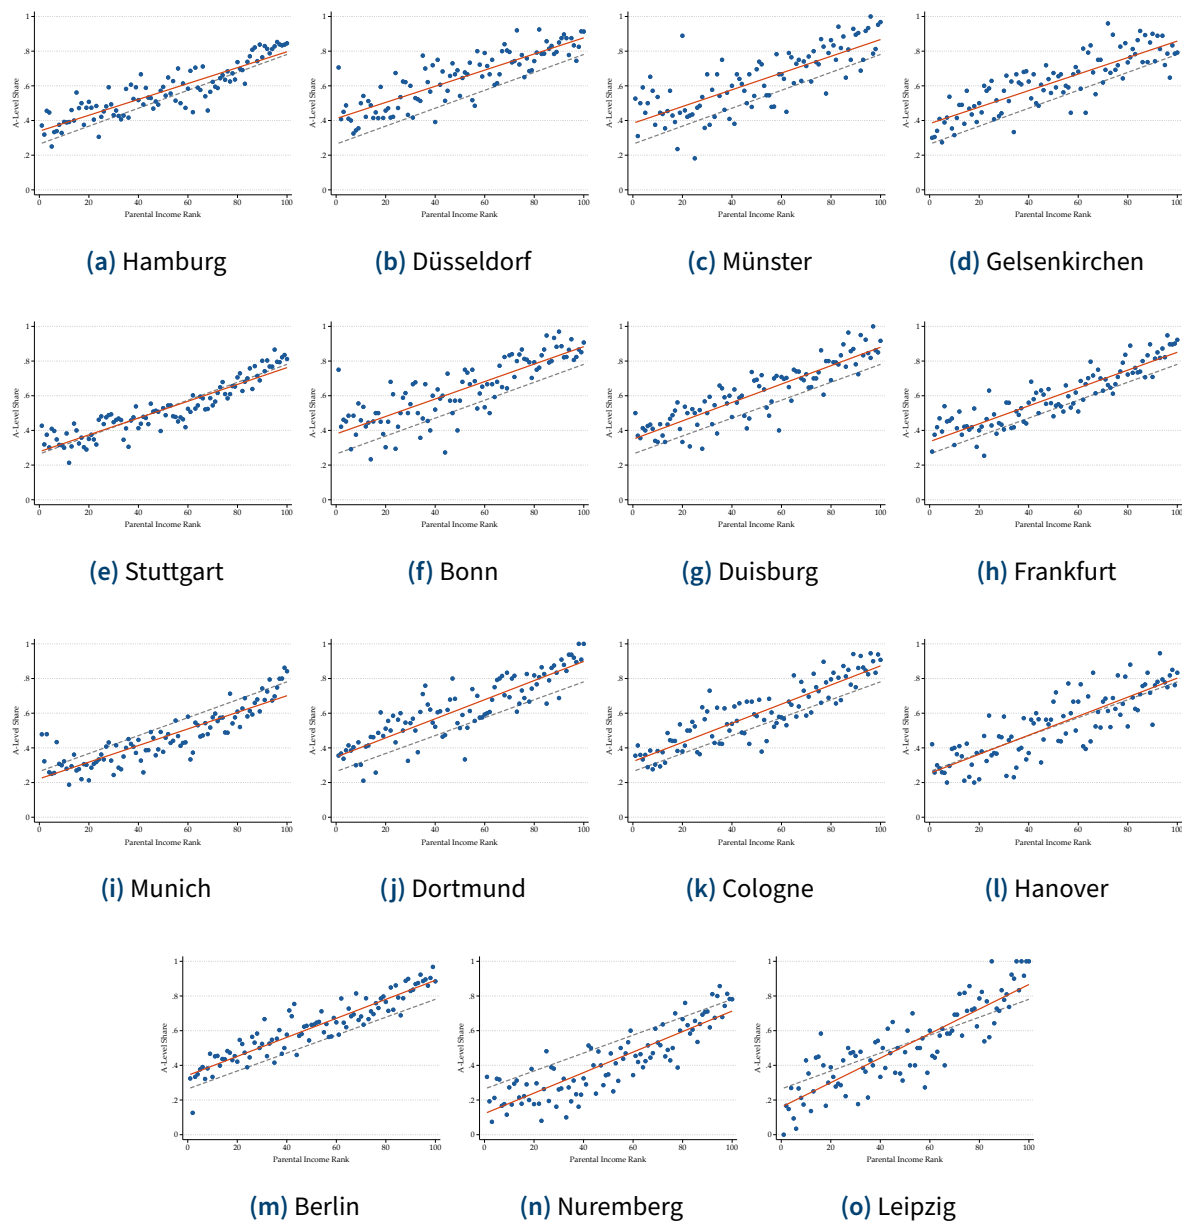
Notes: This figure shows the development of the parental income gradient for different population subgroups for birth cohorts 1980-1996 in the MZ. The parental income gradient per cohort is estimated as $100 \times \gamma_t$ in the following regression: $Y_{i,t} = \alpha + \beta_t C_t + \gamma_t C_t \times R_i + \varepsilon_{i,t}$, where C_t denotes a cohort and $C_t \times R_i$ the interaction between cohort and parental income rank.

Figure B.8: Social Mobility at the State Level



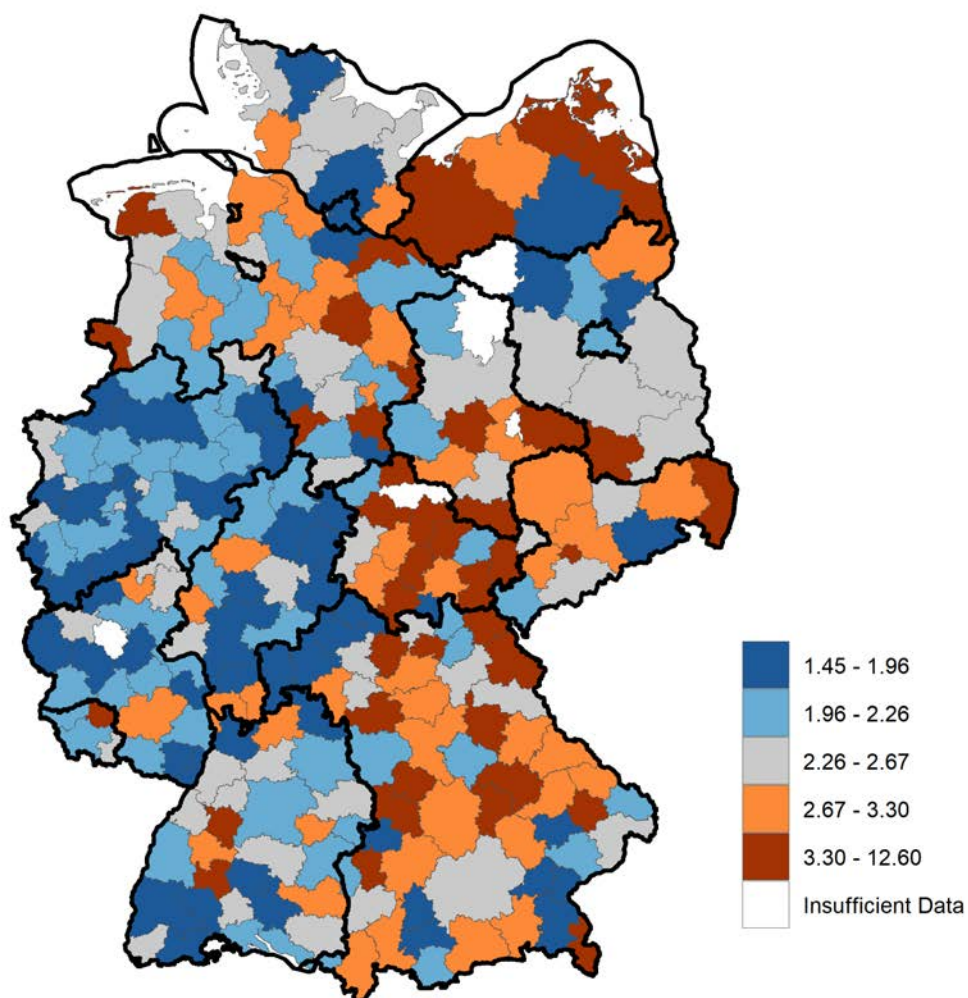
Notes: This figure shows for each German state the fraction of children aged 17-21 that are either enrolled in the upper stage of an A-Level track or have already attained an A-Level degree by percentile rank of their parents in the national income distribution based on the MZ waves 2011-2018, as well as the best linear approximation to the empirical CEF in orange. The dashed gray line plots the national gradient as a comparison.

Figure B.9: Social Mobility for Cities



Notes: This figure shows for the 15 largest (by population size in 2017) local labor markets in Germany the fraction of children aged 17-21 that are either enrolled in the upper stage of an A-Level track or have already attained an A-Level degree by percentile rank of their parents in the national income distribution based on the MZ waves 2011-2018, as well as the best linear approximation to the empirical CEF in orange. The dashed gray line plots the national gradient as a comparison.

Figure B.10: Q5/Q1 Ratio by Local Labor Market



Notes: This figure presents a heat map of the Q5/Q1 ratio by LLM. Children are assigned to LLMs according to their current place of residence. The estimates are based on children aged 17-21 in the years 2011-2018 for which we have non-missing information on educational attainment and parental income. The Q5/Q1 ratio is computed by dividing the share of children with an A-Level degree in the top 20% through the share of children with an A-Level degree in the bottom 20% of the parental income distribution. The colors indicate the quintile of the respective LLM point estimate in the distribution of estimates to account for outliers of the Q5/Q1 ratio induced by small denominators. 6 LLMs with less than three children in the top 20% of the parental income distribution without an A-Level degree are excluded from the analysis.

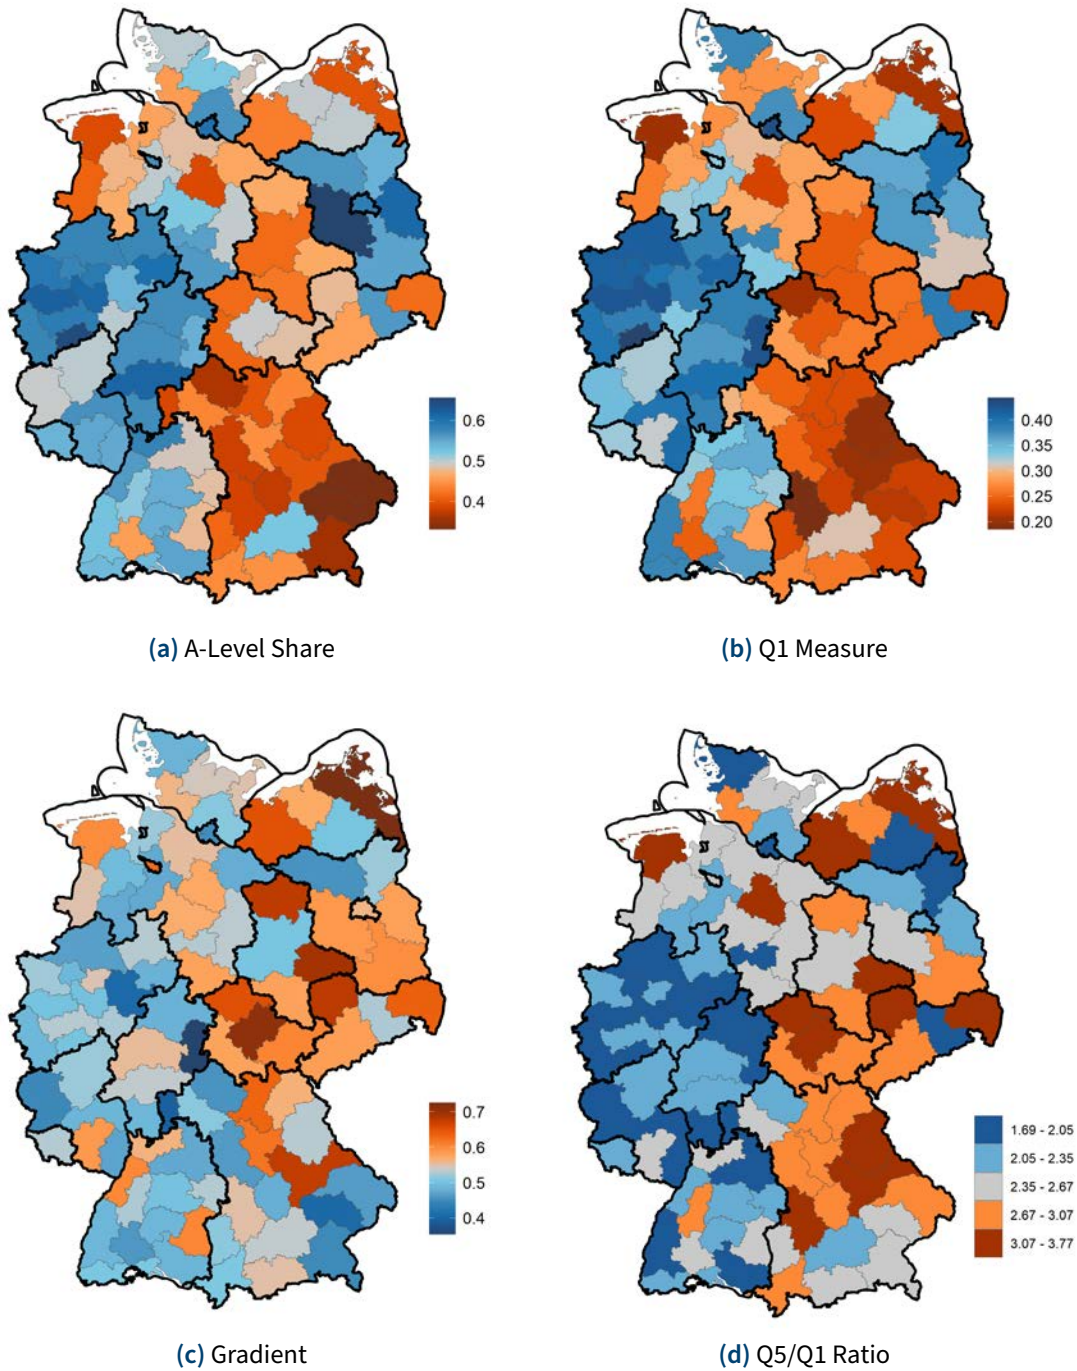
Table B.2: Correlation between Mobility Measures

Measure	Corr.	A-Level	Q1	Q5	Q5/Q1	Gradient
A-Level	ρ	1	-	-	-	-
	r	1	-	-	-	-
Q1	ρ	0.76	1	-	-	-
	r	0.78	1	-	-	-
Q5	ρ	0.70	0.44	1	-	-
	r	0.71	0.48	1	-	-
Q5/Q1	ρ	-0.40	-0.72	0.088	1	-
	r	-0.48	-0.84	-0.04	1	-
Gradient	ρ	-0.01	-0.45	0.45	0.65	1
	r	-0.07	-0.47	0.33	0.76	1

Notes: This table reports the pairwise correlations between estimates of different measures of social mobility across LLMs in Germany. ρ denotes the Pearson correlation coefficient, r denotes the Spearman rank correlation coefficient.

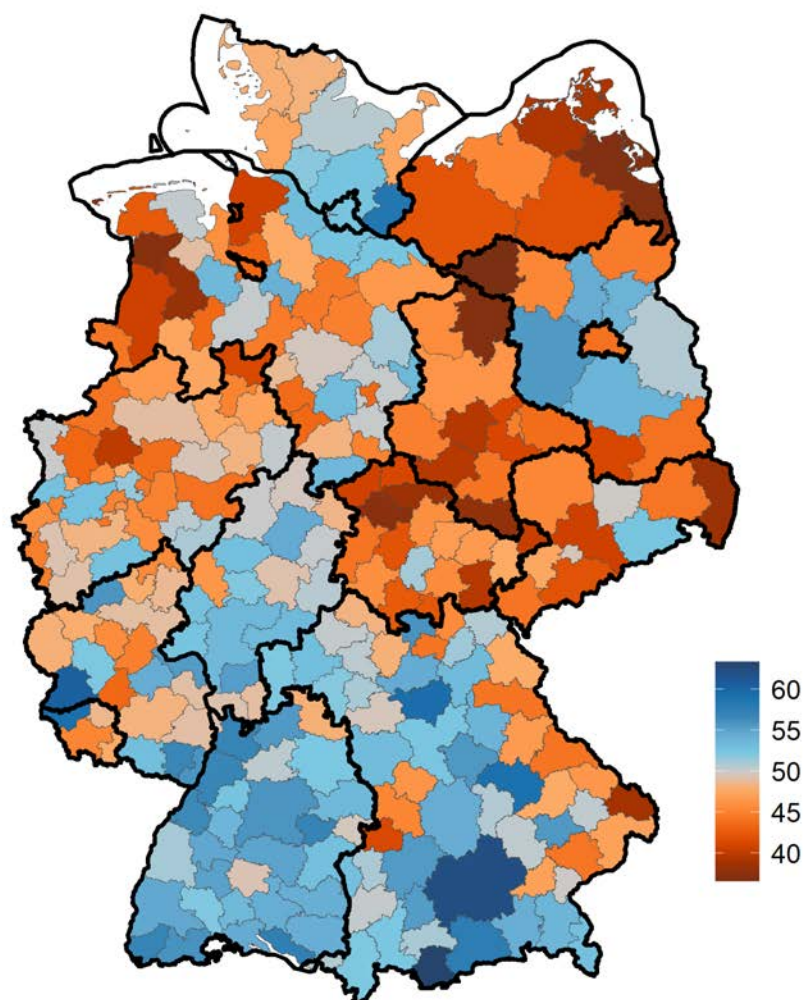
Table B.2 reports the correlations between our mobility measures. While the Q1 measure is well predicted by the unconditional A-Level share, there exists no systematic association between the A-Level share and the parental income gradient, highlighting that the gradient is not sensitive to the baseline probability of obtaining an A-Level degree. Finally, the correlation between the parental income gradient and the Q1 measure ranges below -0.5, demonstrating that a high level of absolute mobility in a given LLM does not always imply a high level of relative mobility.

Figure B.11: Mobility Estimates by Spatial Planning Region







Notes: This figure presents heat maps of the A-Level share (Panel A), the Q1 measure (Panel B), the parental income gradient (Panel C) and the Q5/Q1 ratio (Panel D) for the 96 spatial planning regions of Germany. Spatial planning regions constitute a more comprehensive version of the LLMs discussed in Section 2.5, as they also represent aggregations of counties based on commuting flows. Children are assigned to spatial planning regions according to their current place of residence. The estimates are based on children aged 17-21 in the years 2011-2018 for which we have non-missing information on educational attainment and parental income. In Panel (D), the colors indicate the quintile of the respective point estimate in the distribution of estimates to account for outliers of the Q5/Q1 ratio induced by small denominators.

Figure B.12: Mean Parental Income Rank by Local Labor Market



Notes: This figure presents a heat map of the mean parental income rank by LLM. Children are assigned to LLMs according to their current place of residence. The estimates are based on children aged 17-21 in the years 2011-2018 for which we have non-missing information on educational attainment and parental income. The mean parental income rank is computed as the local labor market specific averages of parental income ranks in the national income distribution.

Figure B.13: Robustness to State and Region Specific Parental Income Ranks**(a) CDFs****(b) Scatter Plot National/State****(c) CDFs****(d) Scatter Plot National/Region Type**

Notes: This figure displays the sensitivity of our LLM-level estimates of the parental income gradient with respect to the reference income distribution. For this aim, the upper two panels compare gradients computed based on the national and the state-specific income distributions: Panel (A) displays the Cumulative Distribution Function (CDF) of both gradients, Panel (B) shows a scatter plot of the point estimates as well as their linear fit. The bottom two panels compare the gradients obtained by computing income ranks based on the national and the region-type-specific income distribution. The region types are defined by the Federal Institute for Building, Urban Affairs and Spatial Research (BBSR) and classify each county into either urban, suburban or rural. For LLMs comprising of counties of different types, we assign the most frequent category. Again, Panel (C) displays the Cumulative Distribution Function (CDF) of both gradients, whereas Panel (D) shows a scatter plot of the point estimates as well as their linear fit. The reported slope parameters of 0.93 and 0.98 correspond to the OLS slope estimates obtained by regressing the gradients computed by using the respective local ranks on the gradients computed by using national income ranks.

B.3 Data Appendix SOEP and PISA

B.3.1 Data Appendix SOEP

The SOEP is a nationally representative household panel survey of the German population, established in 1984. In its more recent waves, it annually samples around 15,000 German households or 25,000 individuals each year (Goebel et al., 2019). Respondents provide information about a broad range of socio-economic variables such as income, education, employment status or biographical characteristics, as well as subjective measures like life satisfaction. Since 2000, participants turning 17 years old answer a youth questionnaire, where they are asked about their current situation in the education system, including school grades, and their aspirations and goals for the future.

Measuring the A-Level degree. In a first step, we replicate our definition of an A-Level degree from the MZ in the SOEP. Because the SOEP follows children even after moving out of the parental household and collects annual information on educational attainment, we choose a cutoff age of 21.¹ At this age, our A-Level dummy turns one if a child has obtained, or is on track to obtain, a degree that is equivalent to an A-Level. Using this definition, Figure B.14 shows that we exactly replicate the parental income gradient of 0.52 from the MZ also in the SOEP, albeit with less precision. For the remaining analyses, we simply assign an A-Level degree to each respondent who reports having obtained such a degree—again including both *Allgemeine Hochschulreife* and *Fachhochschulreife*—but exclude students enrolled in an A-Level track at age 21.

Measuring Ability. Among the educational information in the youth questionnaire, respondents are asked about their last school grade in the subjects of mathematics, German, and the first foreign language. Within the A-Level track, these grades provide a proxy for ability that is broadly comparable among children. We focus on math grades, as they are less likely to be confounded proxies for ability than grades in German, where migration history and command of the German language may result in worse grades even for talented students. As the SOEP does not cover track enrollment in secondary school in sufficient detail, we make the assumption that all students that had obtained an A-Level degree by age 21 were previously enrolled in an A-Level track at age 17.

¹ To the extent that moving out of the parental household leads to panel attrition of some individuals, a small bias is also extant in the SOEP.

Figure B.14: Social Mobility in the SOEP



Notes: This figure shows for the birth cohorts 1980-1996 the fraction of children aged 21 in the German Socio-economic Panel (SOEP) that have already attained an A-Level degree by percentile rank of their parents in the national income distribution. The income ranks are computed with respect to the national distribution of equivalized net household income among households with children in birth cohorts 1980-1996. The reported slope coefficient of 0.0052 is estimated by OLS using the underlying micro data.

Building an Intergenerational Sample. To create an intergenerational sample, we link all children of SOEP respondents that are born during the years 1980-1996 to their parents. In our further sample restrictions, we aim to make our sample as comparable as possible to the data used in Chetty et al. (2014a). For this reason, incomes of children are measured as the average over the five-year interval when children are between 29 and 33 years old. Parental information is measured in the five-year interval when children are between 15 and 19 years old.

B.3.2 Data Appendix PISA

The PISA international student achievement test is conducted by the OECD since the year 2000. PISA assesses achievement in mathematics, science, and reading in a representative cross-section of 15-year-old students, independent of grade level or educational track attended.

To create comprehensive measures of competencies, students complete a broad array of tasks of varying difficulty in assessments that last for up to two hours. PISA achievements in math, science, and reading were standardized to a mean of 500 test-score points and a standard deviation of 100 test-score points for OECD-country students in wave 2000 (and rescaled on the same metric again in 2003 in math and in 2006 in science). PISA test scores are provided as a distribution of five different plausible values. In our analysis, we take the average of all five plausible values. We use the PISA student weights throughout to obtain unbiased parameter estimates.

Measuring Educational Attainment. As PISA test scores are elicited at age 15 before children enter the 2-3 last years of higher secondary schools when different A-Level tracks open up, we focus exclusively on children attending *Gymnasium*, the highest secondary school track and the main avenue to obtaining an A-Level. In our analysis, we thus assume that all children enrolled in *Gymnasium* at age 15 will eventually obtain an A-Level degree. Within *Gymnasium*, PISA test scores provide an excellent proxy for ability and are well comparable among children.

Measuring Parental Income. Since 2006, PISA administers a separate parental questionnaire in selected countries, including Germany. In this questionnaire, parents report gross annual household income in six bands. To derive a continuous measure of income from the banded data, we fit a Singh-Maddala distribution in each wave. Parental income ranks are then computed based on this continuous income measure among all children of each wave.

B.3.3 Adjustment of Ability Trends

In this section, we first describe the calculations behind the results in Table 2.6 and then describe our second approach for predicting changes in grades for inframarginal students.

Assuming Constant Grades/Test Scores for Inframarginal Students. The trends in Figure 2.8 pool the grades of both inframarginal and marginal students. As a first way to obtain an estimate for the grades of marginal students, we assume that grades for inframarginal students have not changed.

Since the A-Level share among children at the top of the income distribution was initially already much higher than among children at the bottom of the distribution, and because the absolute increase in the A-Level share was approximately the same in all parts of the parental income distribution, the share of marginal children among all children obtaining an A-Level degree at the end of the *Bildungsexpansion* is strongly decreasing in parental income rank. For this reason, grade trends among marginal children may look quite different from the patterns in Figure 2.8.

To obtain the grade trend among marginal children in the SOEP, we make use of the following equation

$$\text{Grade}_{1996} = \frac{\text{A-Level}_{1980}}{\text{A-Level}_{1996}} \text{Grade}_{1980} + \frac{\Delta \text{A-Level}}{\text{A-Level}_{1996}} \text{Grade}_{1996}^M, \quad (\text{B.1})$$

where we denote the average grade at the end of the educational expansion in 1996 as Grade_{1996} . It equals the weighted average of the inframarginal students (whose grade we denote by Grade_{1980}) and the marginal students (whose grade we denote by Grade_{1996}^M). Note that we define $\Delta \text{A-Level} = \text{A-Level}_{1996} - \text{A-Level}_{1980}$, i.e. the weights add up to one.

Rewriting, we can express the grade among marginal children as follows:

$$\text{Grade}_{1996}^M = \frac{\text{Grade}_{1996} * \text{A-Level}_{1996} - \text{Grade}_{1980} * \text{A-Level}_{1980}}{\Delta \text{A-Level}} \quad (\text{B.2})$$

The same calculation can be analogously applied to obtain the average PISA test scores among marginal children. We do this analysis separately for above- and below-median-income children and thereby obtain grades and test scores of marginal students from both income groups.

Adjusting for Grade/Test Score Trends among Inframarginal Children. Equation B.2 assumes that the grades of inframarginal children do not change during the educational expansion. While this assumption is a natural starting point, we now explore how our implications about the relative grades of marginal students with high- and low-parental income change if relax this assumption. We therefore incorporate a potential grade trend in the analysis. To achieve this, we predict test scores of inframarginal children based on observables by estimating the following Probit model

$$\begin{aligned} P(\text{A-Level} = 1) = & \beta_0 + \beta_1 \text{Education Mother} + \beta_2 \text{Education Father} \\ & + \beta_3 \text{Occupation Mother} + \beta_4 \text{Occupation Father} \\ & + \beta_5 \text{Migrant} + \beta_6 \text{Gender} + \varepsilon, \end{aligned}$$

among all children of the initial birth cohorts (1982-1984 in the SOEP, 1990 in PISA). Hence, we model the likelihood of attaining an A-level degree as a function of parental education and occupation, migration status and gender.² We then predict among all birth cohorts the probability for each child to graduate with an A-Level degree, and classify all children above the 75th percentile in this probability as “inframarginal”. These are typically children where both parents hold a college degree, or where parents have prestigious occupations. The intuition behind this exercise is that children from these backgrounds would have been very likely to attain an A-Level degree also in absence of the educational expansion.

Table B.3 shows that grades and test scores for children who are inframarginal according to this definition slightly deteriorated. In the SOEP, this decline in our ability proxies happened mainly among low SES children. In the PISA, the decline is more pronounced for children with above-median parental income.

To estimate the grades of marginal students, we now account for these changes in grades among inframarginal children, we can compute the average grade among marginal children according to the following formula:

$$\text{Grade}_{1996} = \frac{\text{A-Level}_{1980}}{\text{A-Level}_{1996}} \text{Grade}_{1991}^I + \frac{\Delta \text{A-Level}}{\text{A-Level}_{1996}} \text{Grade}_{1996}^M, \quad (\text{B.3})$$

which is essentially the same as (B.1) with the only difference that Grade_{1996}^I replaces Grade_{1980} .

² While migration status and gender (2 categories each), and parental education (ISCED 1-6) are defined consistently in both SOEP and PISA, the coding of parental occupation differs slightly between both data sets (6 categories in PISA, 10 categories in the SOEP).

Table B.3: Math Grades and Test Scores of Inframarginal Children

	SOEP		PISA	
	Bottom 50	Top 50	Bottom 50	Top 50
1982	2.3	2.6	-	-
1990	3.2	2.6	585	622
1996	3.0	2.5	583	614

Notes: This table shows the average math grades in the German Socio-Economic Panel (SOEP) and the average PISA math test scores for inframarginal children in birth cohorts 1982, 1990 and 1996, separately for children below and above median parental income. See the text for a definition of inframarginal children. Due to the small sample size of the SOEP, three year averages around the actual birth cohort are used to compute grade averages (1982: 1982-1984, 1990: 1989-1991, 1996: 1994-1996).

Hence, for the inframarginal students we do not assume that their grades equal the grades of 1980 but take the probit model predictions as stated in Table B.3. Rearranging, we get:

$$\text{Grade}_{1996}^M = \frac{\text{Grade}_{1996} * \text{A-Level}_{1996} - \text{Grade}_{1996}^I * \text{A-Level}_{1980}}{\Delta \text{A-Level}}. \quad (\text{B.4})$$

As for the first approach, we do this analysis separately for above- and below-median-income children and thereby obtain grades and test scores of marginal students from both income groups. The results in Table B.4 show that adjusting for grade trends among inframarginal children does slightly alter the conclusions regarding the ability of marginal children. While the grades of marginal children did not differ substantially (or significantly) between children below and above median parental income if we consider cohorts from 1980-1996, there are differences for the time period 1982-1990: among those cohorts, marginal students with lower parental income outperform marginal students with higher parental income.

Table B.4: Math Grades and Test Scores of Marginal Children - Adjusted for Changes among Inframarginal Children

	SOEP			PISA		
	Bottom 50	Top 50	Δ	Bottom 50	Top 50	Δ
1982-1990	2.2	2.5	0.31 SD	-	-	-
1990-1996	4.0	2.9	0.99 SD	554	601	0.68 SD
1982-1996	2.7	2.7	0.04 SD	-	-	-

Notes: This table shows the average math grades in the German Socio-Economic Panel (SOEP) and the average PISA math test scores among “marginal” children in birth cohorts 1982, 1990 and 1996, separately for children below and above median parental income. The grades are computed using Equation B.4 and take into account the differential development in grades among inframarginal children. The third column expresses the differences between both groups in terms of the standard deviations, which is 1.06 for math grades in the SOEP, and 72 points for PISA test scores. Due to the small sample size of the SOEP, three year averages around the actual birth cohort are used to compute grade averages (1982: 1982-1984, 1990: 1989-1991, 1996: 1994-1996).

B.4 Regional Predictors of Mobility

Regional Indicators. We construct a comprehensive database of 73 regional indicators for this analysis, with information on labor market participation, economic conditions, infrastructure, demographics, local educational institutions and social characteristics.

Table B.5 displays all 73 regional indicators we use as predictors in the Random Forest algorithm. In a first step, we retrieve data from the Federal Institute for Building, Urban Affairs and Spatial Research (BBSR), which maintains the INKAR database of regional indicators (<https://www.inkar.de/>). These data are collected from various government bodies in Germany, including the German Statistical Office (Destatis) and the Institute for Employment Research (IAB). We select all indicators which we suppose to be potentially relevant for social mobility and are not collinear: for example, we do not include the general unemployment rate and the unemployment rates among males and females at the same time. In a second step, we add data from Destatis publications with information on the share of Gymnasium students among all secondary school students, the share of A-Level degrees obtained on vocational schools and compute the distance of the geographical center of each LLM to the next college based on data from the website of the Hochschulrektorenkonferenz (HRK; <https://www.hochschulkompass.de/hochschulen/downloads.html>). In a third step, we compute additional regional statistics on the LLM level using the MZ data, like the Gini coefficient in household income, the local A-Level wage premium or the ISEI (an international index of social status). We construct our final variables by averaging the local indicators over the years 2011-2018 at the LLM level.

Prediction Exercise. To study the association between local characteristics and intergenerational mobility, prior literature has typically relied on correlation coefficients or estimated multiple linear models (Chetty et al., 2014a; Corak, 2020). Both approaches have disadvantages. As socio-economic characteristics are highly correlated at the regional level, correlation coefficients are often spurious. While this remedy is overcome in a multiple linear OLS regression, these models are prone to overfitting in high-dimensional data sets (Babyyak, 2004), resulting in diminished external validity. One way to address this it to reduce dimensionality of the covariates via variable selection. Belloni and Chernozhukov (2013) suggest to preselect

covariates via Lasso before estimating a multiple linear model.³ This approach is for example applied by Finkelstein et al. (2016) to explain geographical variation in health care utilization in the US.

We take a similar two-step approach, but preselect variables using a Random Forest variable importance measure instead of a Lasso regression. This is because we find that a linear Lasso model fits our data poorly: To compare the out-of-sample performance of this algorithm against an implementation of a Lasso and an Elastic Net regression with $\alpha = 0.5$, we split our data in a training and test data set (75-25 split). The Random Forest algorithm predicts 39% of the variation in the test sample ($R^2 = 0.39$), whereas the predictive power of Lasso ($R^2 = 0.15$) and Elastic Net ($R^2 = 0.17$) is lower. The results for Lasso and Elastic Net are based on λ chosen by 5-fold cross-validation. For the Random Forest, we fit 1000 trees and randomly select $73/3 \approx 24$ variables for each split.

Before constructing the Random Forest, we standardize all 73 indicators to have mean 0 and standard deviation 1. Once the Random Forest is fitted, we can rank covariates according to their predictive power and thereby obtain a measure of variable importance. We choose the implementation proposed by Strobl et al. (2008), which computes a conditional permutation importance measure that accounts for the dependence structure between the predictors, using the `party` R package (<http://CRAN.R-project.org/package=party>).

Most Informative Predictors. The set of the 15 most informative predictors is displayed in Table B.6, ranked by a measure of variable importance computed by the Random Forest.⁴ The last column displays the sign of the bivariate correlation between each variable and the parental income gradient. A positive sign implies that the indicator predicts low mobility (a high gradient). For example, LLMs with a high prevalence of school dropouts are associated with low relative mobility. Overall, our selection procedure highlights social characteristics, the local organization of the education system and labor market conditions. These correlational findings are consistent with causal studies that emphasize the importance of local characteristics for child and adolescent outcomes (Chetty and Hendren, 2018; Damm and Dustmann, 2014).

³ An alternative approach to deal with model uncertainty is model averaging. See Kourtellis et al. (2016) for an application in the context of social mobility.

⁴ The exact ranking of predictors varies for different implementations of the Random Forest algorithm. We are therefore cautious not to over-interpret the ranking between single predictors.

Regression Estimates. In a second step, we regress the gradient on these 15 indicators selected by the algorithm. All right-hand side variables are standardized so that the coefficients report the association between a one standard deviation change in the covariate and an absolute change in the gradient. The results are reported in Table B.7. The signs of the coefficients mostly match those from the bivariate correlations in Table B.6. For example, a one standard deviation increase in the school dropout rate is associated with a 3.9 percentage point higher parental income gradient.⁵ This association becomes stronger when adding state indicators. A high gradient also aligns with a high number of teenage pregnancies, a high unemployment rate and a large share of households with access to broadband Internet. A negative association with the parental income gradient arises for the share of married individuals, the distance to the next college, the median income for individuals with a recognized vocational qualification, the share of children aged 0-2 in childcare and for the share of children on a vocational A-Level track. Due to the limited sample size of 258 local labor markets, we lack the power to precisely estimate most coefficients. Exceptions are the school dropout rate, broadband availability, the share of married individuals and the share of children on a vocational A-Level track.

⁵ The school dropout rate refers to the share of secondary school students leaving school without the lowest possible certificate (*Hauptschulabschluss*). Although the A-Level share *ceteris paribus* decreases in the school dropout rate, there exists no direct mechanical relationship between the two. For example, any student dropping out of the two higher secondary school tracks (which enroll the vast majority of students) after grade 9 will automatically be awarded a *Hauptschulabschluss*, and thus not fall under the given definition of a school dropout.

Table B.5: List of Regional Indicators

Category	Variable	Source
Labor Market	Unemployment Rate	INKAR
	Share Long Term Unemployed	INKAR
	Share Female Employees	INKAR
	Share Part Time Employees	INKAR
	Share without Vocational Qualification	INKAR
	Share Marginal Employment	INKAR
	Share Employed in Manufacturing Sector	INKAR
	Apprenticeship Positions	INKAR
	Apprentices	INKAR
	Vocational School Students	INKAR
	Employees with Academic Degree	INKAR
	Commuting Balance	INKAR
	Hours Worked	INKAR
	A-Level Wage Premium	MZ
Education	Students (before Tertiary Education)	INKAR
	Students (Tertiary Education)	INKAR
	Students (Universities of Applied Sciences)	INKAR
	School Dropout Rate	INKAR
	Highly Qualified Persons	INKAR
	Share Children 0-2 in Childcare	INKAR
	Share Children 3-5 in Childcare	INKAR
	Share Students Enrolled in Gymnasium	INKAR
	Share Secondary School Students Enrolled in Gymnasium	Destatis
	Distance to Next College	HRK
	Distance to Next Elementary School	INKAR
	Share on Vocational A-Level Track	MZ
	Share A-Level Degree from Vocational Schools	Destatis
	Mean Parental Education	MZ
Income	Median Household Income	INKAR
	Median Household Income with Vocational Qualification	INKAR
	Gender Wage Gap	INKAR
	Child Poverty	INKAR
	Mean Household Income	INKAR
	Gini Household Income	MZ
	Expected Rank Difference Parental Income	MZ
	Mean Parental Income	MZ
	Gini Parental Income	MZ
	Ratio p85/p50 (Household Income)	MZ
	Ratio p50/p15 (Household Income)	MZ

Economy	GDP per Capita	INKAR
	Municipal Tax Revenues per Capita	INKAR
	Municipal Debt per Capita	INKAR
	Business Creation	INKAR
Housing	Construction Land Prices	INKAR
	New Apartments	INKAR
	Building Permits	INKAR
	Living Area	INKAR
	Share Apartment Buildings	INKAR
	Rent Prices	INKAR
Infrastructure	Physician Density	INKAR
	Broad Band Availability	INKAR
	Passenger Car Density	INKAR
	Hospital Beds	INKAR
Demographics	Average Age	INKAR
	Share Female	INKAR
	Share Foreigners	INKAR
	Share Asylum Seekers	INKAR
	Total Net Migration	INKAR
	Births Net of Deaths	INKAR
	Fertility Rate	INKAR
	Teenage Pregnancies	INKAR
	Life Expectancy	INKAR
	Child Mortality	INKAR
	Population Density	INKAR
	Share Single Parents	MZ
	Share Married	MZ
	Share Divorced	MZ
Social	Voter Turnout	INKAR
	Vote Share CDU	INKAR
	Vote Share SPD	INKAR
	Share Social Assistance	INKAR
	Mean ISEI	MZ
	Gini ISEI	MZ

Notes: This table displays all regional indicators considered for our analysis. The third column reports the data source, which is either the INKAR database, the Statistical Office of Germany (Destatis), the Hochschulrektorenkonferenz (HRK) or the Mikrozensus (MZ).

Table B.6: The 15 Most Informative Predictors of Relative Mobility

Variable	Importance Measure	ρ
School Dropout Rate	0.85	+
Share Married	0.60	—
Teenage Pregnancies	0.42	+
Students	0.39	—
Median Income Vocational Qualification	0.18	—
Broadband Availability	0.17	+
Distance to Next College	0.15	—
Unemployment Rate	0.14	+
Gender Wage Gap	0.14	+
Share without Vocational Qualification	0.13	—
Gini Parental Income	0.08	—
Share Marginal Employment	0.07	—
Share Children 0-2 in Childcare	0.07	+
Share Social Assistance	0.07	+
Share on Vocational A-Level Track	0.07	—

Notes: This table lists the optimal predictive set of 15 regional indicators for the local labor market parental income gradient estimates, as chosen by a Random Forest based measure of variable importance (second column, displayed in multiples of 1000). The last column shows the sign of the Pearson correlation coefficient between each variable and the parental income gradient. A positive correlation implies that an indicator is predictive for low relative mobility (a high gradient).

Table B.7: Social Mobility and Regional Characteristics

	(1)	(2)	(3)	(4)	(5)
School Dropout Rate	0.0391 (0.0110)	0.0371 (0.0110)	0.0393 (0.0091)	0.0554 (0.0162)	0.0539 (0.0162)
Share Married	-0.0225 (0.0089)	-0.0286 (0.0089)	-0.0225 (0.0065)	-0.0243 (0.0108)	-0.0278 (0.0109)
Teenage Pregnancies	0.0169 (0.0226)	0.0123 (0.0231)	0.0211 (0.0155)	0.0160 (0.0252)	0.0115 (0.0266)
Students	-0.0143 (0.0131)	-0.0166 (0.0128)	-0.0055 (0.0093)	-0.0214 (0.0164)	-0.0246 (0.0165)
Median Income Vocational Qualification	-0.0179 (0.0129)	-0.0194 (0.0129)	-0.0025 (0.0114)	-0.0234 (0.0167)	-0.0224 (0.0177)
Broadband Availability	0.0260 (0.0100)	0.0274 (0.0100)	0.0194 (0.0085)	0.0231 (0.0105)	0.0261 (0.0109)
Distance to Next College	-0.0048 (0.0072)	-0.0059 (0.0077)	-0.0051 (0.0070)	-0.0025 (0.0072)	-0.0045 (0.0076)
Unemployment Rate	0.0368 (0.0365)	0.0295 (0.0365)	0.0124 (0.0236)	0.0537 (0.0464)	0.0476 (0.0470)
Gender Wage Gap	-0.0029 (0.0142)	-0.0041 (0.0144)	0.0048 (0.0126)	0.0156 (0.0174)	0.0124 (0.0177)
Share without Vocational Qualification	0.0057 (0.0171)	0.0085 (0.0173)	-0.0035 (0.0127)	0.0132 (0.0217)	0.0108 (0.0220)
Gini Parental Income	-0.0171 (0.0147)	-0.0108 (0.0150)	-0.0236 (0.0110)	0.0051 (0.0200)	0.0117 (0.0209)
Share Marginal Employment	-0.0086 (0.0138)	-0.0162 (0.0142)	-0.0183 (0.0121)	-0.0222 (0.0152)	-0.0250 (0.0154)
Share Children 0-2 in Childcare	-0.0398 (0.0192)	-0.0420 (0.0189)	-0.0526 (0.0182)	-0.0246 (0.0234)	-0.0259 (0.0236)
Share Social Assistance	-0.0607 (0.0343)	-0.0429 (0.0361)	-0.0406 (0.0231)	-0.0969 (0.0450)	-0.0782 (0.0498)
Share on Vocational A-Level Track	-0.0165 (0.0092)	-0.0171 (0.0092)	-0.0133 (0.0079)	-0.0213 (0.0100)	-0.0224 (0.0100)
Additional Controls	-	✓	✓	-	✓
State Indicators	-	-	-	✓	✓
Weighted	-	-	✓	-	-
<i>N</i>	258	258	258	252	252
<i>R</i> ²	0.256	0.273	0.253	0.296	0.305

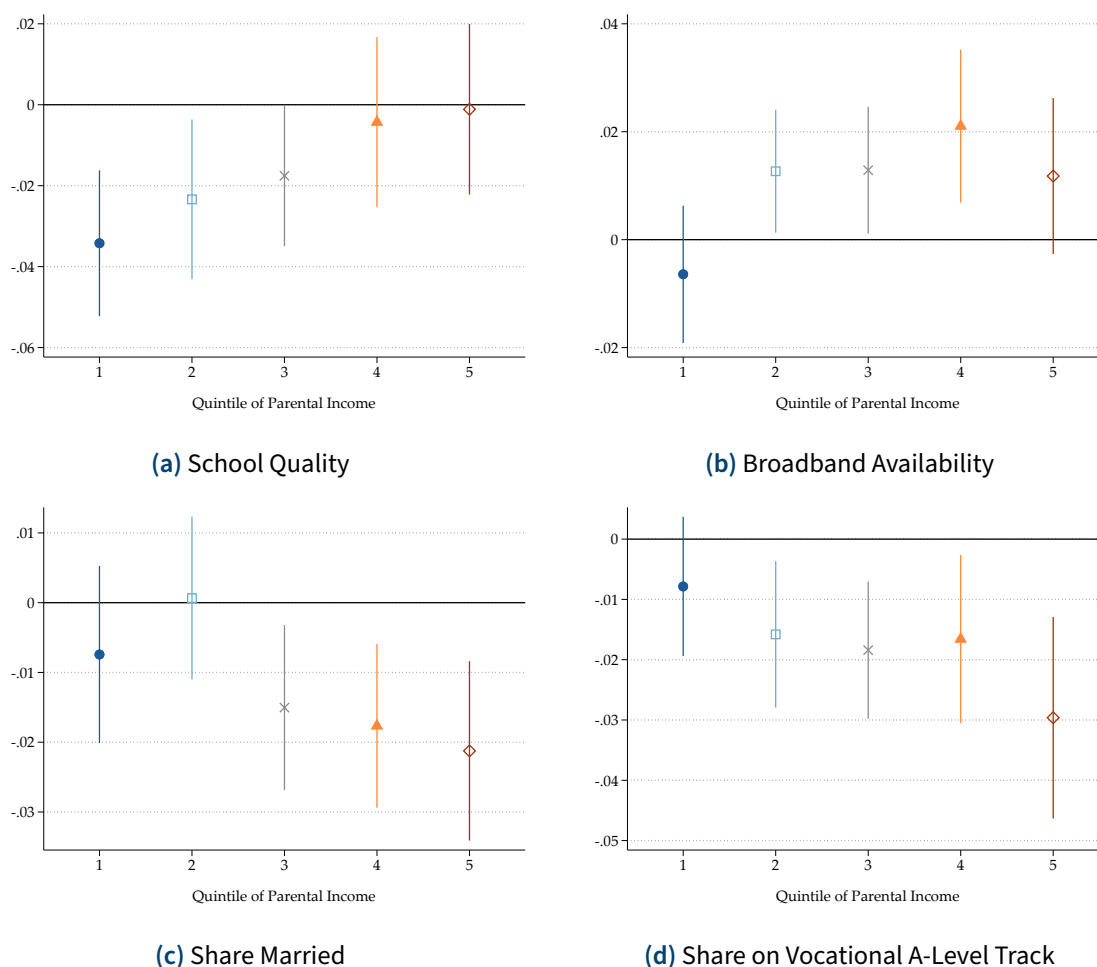
Notes: Each column of this table reports coefficients from a linear regression with robust standard errors reported in parentheses. The dependent variable in all columns is the parental income gradient. The independent variables (as selected by the Random Forest, compare Table B.6) are standardized to have mean 0 and standard deviation 1. Columns (3) and (4) contain state dummies, for which we have to drop five LLMs crossing state borders and the LLM of Berlin. In columns (2) and (4), we additionally control for population, population density and the region type (rural, urban or mixed) to test whether coefficients of the regional indicators are affected by structural differences in mobility between more rural or urban LLMs. In column (3) we weight the regression with the number of observations per LLM.

Graphical Evidence. To understand the relationship between relative mobility and the indicators with the largest t-statistics in more detail, we separately regress the A-Level share in each quintile of the parental income distribution on each indicator and plot the estimates in Figure B.15. These plots reveal whether, for example, a positive relationship between the parental income gradient and an indicator is driven by a lower A-Level share of children from low-income households or by a higher A-Level share of children from high-income households.

We start with the school dropout rate. In the US context, Chetty et al. (2014a) interpret the school dropout rate, adjusted by parental income, as an indicator of school quality and find a strong negative correlation with relative mobility. In close analogy, we regress the dropout rate on mean parental income, the Gini coefficient of parental income, the share of parents holding an A-Level degree and the unemployment rate and take the residuals to obtain a measure of school quality which is adjusted for parental income and labor market conditions. This indicator is still highly correlated with mobility. As depicted in Figure B.15, Panel (A), low school quality (a high value of the indicator) is associated with a lower probability to obtain an A-Level degree for children from low income households but does not seem to affect children in the top two quintiles of the parental income distribution. While this would be consistent with the idea that school quality is crucial for improving opportunities for children from low socio-economic background, further information is needed to test this hypothesis in detail.⁶

Panel (B) sheds light on the negative connection between broadband availability and mobility. While broadband access is associated with a higher A-Level share on average, this is not true for children in Q1, for whom the relationship becomes negative. We can only speculate about the reasons. Broadband access is highly correlated with factors pointing to dynamic and prosperous labor markets, which exhibit above average inequality. For that reason, broadband availability may proxy urban areas in which all but children from the bottom of the income distribution profit from a dynamic and rewarding economic environment. However, broadband availability could also causally influence social mobility. For the US, Dettling et al. (2018) document that increased broadband availability fosters access to college and find

⁶ Most importantly, it remains open if the adjusted school drop out rate is indeed an appropriate proxy for school quality. In the US, Rothstein (2019) studies how closely the transmission of parental income to educational attainment and achievement (test scores) are correlated with income mobility at the commuting zone level. He finds income-income transmission to be closely connected to income-educational attainment transmission but not to income-educational achievement transmission. Rothstein (2019) therefore finds little evidence that differences in the quality of secondary schooling are a key mechanism driving variation in intergenerational mobility. However, the distinct features of the German secondary schooling system could lead to very different patterns in our data. Unfortunately, there exist no comparable data on student test scores in Germany, preventing us from investigating this issue further.

Figure B.15: Predicting the A-Level Share by Parental Income Quintile

Notes: Each panel of this figure reports coefficients from five separate linear OLS regressions with robust standard errors and 95% confidence bands. The dependent variable is the share of children which obtained an A-Level in the respective quintile of the parental income distribution. The independent variable is the adjusted school dropout rate (school quality index) in Panel (A), the share of broadband connections per 100 inhabitants in Panel (B), the share of married individuals in Panel (C) and the share of students on a vocational (rather than general education) A-Level track (Panel D). In addition, all regression include a set of state indicators and control for population, population density and the region type (rural, urban or mixed). We exclude 6 LLMs with insufficient observations for estimating Q5 from the sample. Due to the inclusion of state indicators, we have to further drop five LLMs crossing state borders and the LLM of Berlin from the sample, leaving us with 246 observations. All regressors are standardized to have mean 0 and standard deviation 1.

the effect to be concentrated among students with parents from high socio-economic status. Similarly, Sanchis-Guarner et al. (2021) report a causal (positive) impact of broadband access on student test scores in England but find comparatively lower effects for students eligible for free school meals. Our results would be in line with these findings.

The opposite pattern emerges for the share of married individuals in Panel (C): this statistic is related to higher mobility but a lower A-Level share of children from high-income families. Finally, Panel (D) reports the association between the quintile measures and the share of children on a vocational, rather than general interest, A-Level track. There is reason to believe that the availability of such vocational tracks may dampen the influence of parental background on the opportunities of children. Children in these tracks have typically obtained a degree from the medium track (Realschule) and now attend a specialized vocational school to obtain an A-Level degree on top. In that setting, vocational schools may especially foster the opportunities of children from low-income households initially "misallocated" to the medium instead of the high track. Dustmann et al. (2017) show that vocational schools have the potential to fully offset adverse effects of early age tracking on long-term labor market outcomes, but cannot observe parental background.

Our evidence shows that, relative to children from the top quintile, children from the bottom quintile are more likely to obtain an A-Level in local labor markets with a high prevalence of such schools. In addition, we find that at the national level the parental income rank is more predictive for the probability of attending the general high track (Gymnasium) at the age of 13-14 than of obtaining an A-Level degree later on (gradient of 0.55 versus 0.52), again suggesting that vocational schools may mediate the influence of parental background.

Summary of Results. The key insight from this exercise is that the Random Forest algorithm is able to find meaningful variation in our data at the regional level, corresponding to existing theories of determinants of mobility. For example, as in our data, the school dropout rate is among the most significant negative correlates of relative mobility in the US data analyzed by Chetty et al. (2014a). Similarly, characteristics of the vocational education system, an evergreen in the debate on social mobility in Germany, feature prominently in this list. In light of this evidence, it seems unlikely that the regional variation between LLMs is mainly driven by sampling error. We also repeat the prediction exercise for the 129 largest and 129 smallest LLMs in Table B.8. While this analysis displays some interesting differences between rural and urban areas, the recurring themes are the same.

At the same time, our results do not necessarily imply that mobility differences originate from regional policy-variant parameters like the local school infrastructure, childcare availability or local employment conditions. Some of the predictors in Table B.6, like the school dropout rate or the share of married individuals, could likewise point to the persistence of cultural

norms or the existence of deep-rooted transmission parameters which are hard to capture with a contemporaneous set of regional indicators. For other outcomes of interest, research has shown that regional differences in Germany can reach far back into the past (e.g. Becker et al., 2020; Cantoni et al., 2019). We lack the statistical power for a detailed discrimination of these factors and exogenous variation to identify the causal determinants of mobility at the local level. We hope that future work will be able to build on our analysis and shed more light on these issues.

Table B.8: The 15 Most Informative Predictors by LLM Size

Variable	Importance Measure	ρ
<i>Panel (A): 129 Largest Local Labor Markets</i>		
School Dropout Rate	0.42	+
Gini Parental Income	0.23	—
Share Married	0.16	—
Share without Vocational Qualification	0.10	—
Students	0.09	—
Physician Density	0.09	+
Teenage Pregnancies	0.06	+
Mean Parental Income	0.06	—
Share Marginal Employment	0.06	—
Students (Universities of Applied Sciences)	0.05	+
Median Income Vocational Qualification	0.05	—
Distance to next Elementary School	0.03	—
Share Children 3-5 in Childcare	0.03	+
Child Mortality	0.03	—
Ratio p50/p15	0.03	—
<i>Panel (B): 129 Smallest Local Labor Markets</i>		
School Dropout Rate	0.75	+
Unemployment Rate	0.45	+
Child Poverty	0.40	+
Students	0.40	—
Share Married	0.33	—
Teenage Pregnancies	0.33	+
Median Income Vocational Qualification	0.19	—
Gender Wage Gap	0.19	+
Share Social Assistance	0.18	+
Total Net Migration	0.12	—
Highly Qualified Persons	0.10	+
Broadband Availability	0.10	+
Share on Vocational A-Level Track	0.08	—
Building Permits	0.08	—
Share Apartment Buildings	0.07	+

Notes: This table lists the 15 most predictive indicators for explaining variation in the parental income gradient between local labor markets in Germany, separately for the 129 largest (Panel [A]) and the 129 smallest (Panel [B]) local labor markets. See the text for the details on the implementation via a Random Forest variable importance measure. The second column displays the measure of variable importance (in multiples of 1000). The last column shows the sign of the Pearson correlation coefficient between each variable and the parental income gradient. A positive correlation implies that an indicator is predictive for low relative mobility (a high gradient).

B.5 Robustness of Regional Estimates

A concern with the point estimates that we report in Section 2.5 of the paper is that the heterogeneity we document and depict in the maps could be driven by sampling variation. In order to address this concern in a principled way, we adopt an empirical Bayes (EB) perspective, i.e. we interpret our baseline estimates for each region j as noisy signals of parameters drawn from a distribution in the following hierarchical model:

$$\begin{aligned}\hat{\theta}_j | \theta_j, \sigma_j &\sim N(\theta_j, \sigma_j^2) \\ \theta_j | \sigma_j &\sim G(\theta) \quad j = 1, \dots, J.\end{aligned}$$

The first level of the hierarchy is justified (approximately) by a central limit theorem applying to our estimators of the mobility parameters. The second level of the hierarchy describes the cross-sectional distribution of the respective mobility measure across regions.

Measuring Overdispersion. In this framework, we first ask how much overdispersion we observe in our estimates, i.e. how much excess variation we observe in our estimates beyond what one would expect given the associated sampling uncertainty.⁷ To that end we compute

$$\hat{\sigma}_\theta^2 = \frac{1}{J} \sum_{j=1}^J \left[\left(\hat{\theta}_j - \hat{\mu}_\theta \right)^2 - \hat{s}_j^2 \right],$$

an estimate of the variance of G , where $\hat{\mu}_\theta = J^{-1} \sum_{j=1}^J \hat{\theta}_j$ and \hat{s}_j^2 denotes the estimated variance of $\hat{\theta}_j$. Based on this measure of overdispersion, we report "reliability ratios" (RR_θ ; see also Deutscher and Mazumder (2020)) that capture the share of excess variance in $\hat{\theta}$ via

$$RR_\theta = \frac{\hat{\sigma}_\theta^2}{\hat{s}_\theta^2},$$

where \hat{s}_θ^2 denotes the sample variance of $\hat{\theta}$. Table B.9 reports the results of this exercise. We conclude that, while sampling variation is certainly important, a substantial share of the regional variation that we document does indeed capture regional differences. As expected, our reliability ratios tend to decrease towards more fine-grained regional disaggregations, reflecting the fact that estimation uncertainty increases.

⁷ We thank an anonymous referee for suggesting an exercise along these lines.

Table B.9: Reliability Ratios of Mobility Measures

	A-Level	Q1	Q5	Q5/Q1	Gradient
States	0.92	0.89	0.86	0.94	0.78
Spatial Planning Regions	0.96	0.83	0.82	0.98	0.60
Local Labor Markets	0.93	0.69	0.66	0.71	0.51

Notes: This table reports the "reliability ratios" defined in Appendix B.5 for our mobility measures estimated on different geographical aggregations. There are 16 states, 96 Spatial Planning Regions and 258 Local Labor Markets.

Empirical Bayes Confidence Intervals. In order to provide further transparency regarding the uncertainty associated with the ensemble of our local labor market-level parental income gradient estimates, we report empirical Bayes confidence intervals. Constructed around MSE-optimal linear shrinkage estimates, these intervals allow us to report sets of confidence intervals with coverage guarantees for the ensemble of projection coefficients, yielding visual summaries of the uncertainty associated with our local labor market-level estimates.

Specifically, we linearly shrink the original point estimates of the projection coefficients towards the respective state averages in proportion to the estimated signal-to-noise ratio

$$\hat{\theta}_j^* = \left(\frac{\hat{\sigma}_\theta^2}{\hat{\sigma}_\theta^2 + \hat{s}_j^2} \right) \hat{\theta}_j + \left(\frac{\hat{s}_j^2}{\hat{\sigma}_\theta^2 + \hat{s}_j^2} \right) \hat{\mu}_{s(j)},$$


where $\hat{\mu}_{s(j)}$ denotes the (variance-weighted) state-average gradient estimate, and report intervals with ensemble coverage guarantees under two sets of assumptions on the mixing distribution G .

We first construct parametric empirical Bayes confidence intervals under the assumption that the sampling distribution of θ_j is conditionally normal

$$\begin{aligned} \hat{\theta}_j | \theta_j, X_j, \sigma_j &\sim N(\theta_j, \sigma_j^2) \\ \theta_j | X_j, \sigma_j &\sim N(\mu_\theta, \sigma_\theta^2) \quad j = 1, \dots, J, \end{aligned}$$

where X_j contains state-indicator variables. If G is correctly specified, the resulting parametric empirical Bayes confidence intervals (EBCIs) will cover $(1 - \alpha)$ percent of the true effect parameters under repeated sampling of the data and the parameters (Morris, 1983). The results are depicted in Figure B.16.

Figure B.16: Shrinkage Estimates and Parametric EB Confidence Intervals



Notes: This figure depicts the original point estimates of the projection coefficients (gray dots), as well as the MSE-optimal linear shrinkage estimates (blue triangles) and corresponding 90% parametric empirical Bayes confidence intervals by local labor market. Under repeated sampling of the parameters and data, 90% of the intervals contain the true projection coefficients with high probability.


In order to relax the normality assumption underlying the interval ensemble depicted in B.16, we further report robust empirical Bayes confidence intervals (Armstrong et al. (2022)),⁸ which provide coverage guarantees under the substantially weaker assumption that the conditional second moment and the kurtosis of the projection error $\varepsilon_j = \theta_j - X_j'\delta$ do not depend on (X_j, σ_j) .⁹ Next to relaxing the parametric restriction, these intervals also provide a frequentist

⁸ We implement the procedure using the `ebci` R package of Armstrong et al. (2022) which estimates the hyperparameters of the model (using weights $w_j = \hat{s}_j^{-2}$) via our baseline estimates.

⁹ Conditional moment independence assumptions of this type are common in the literature and were also employed in Chetty and Hendren (2018) (cf. Remark 3.1 in Armstrong et al. (2022)).

coverage guarantee: If the parameters are treated as fixed, at least a fraction $(1 - \alpha)$ of the robust EBCIs will contain their respective parameters with high probability. The results are depicted in Figure B.17.

Figure B.17: Shrinkage Estimates and Robust EB Confidence Intervals



Notes: This figure depicts the original point estimates of the projection coefficients (gray dots), as well as the MSE-optimal linear shrinkage estimates (blue triangles) and corresponding 90% robust empirical Bayes confidence intervals by local labor market. Under mild conditions, at least a fraction $(1 - \alpha)$ of the robust EBCIs will contain their respective parameters with high probability.

The procedures allow us to substantially tighten the confidence intervals relative to those associated with our baseline estimates. At the same time, Figure B.17 shows that the linear shrinkage estimates and confidence sets still display substantial heterogeneity, mitigating concerns that sampling variation is driving our results.

C Appendix to Chapter 3

C.1 Measurement—Some Simple Examples

Consider a society with two types that are of equal size. Both income and wealth are unequally distributed. Since (i) inequality in both dimensions is exactly the same, and (ii) both dimensions are perfectly correlated across types, inequality of opportunity in monetary resources is exactly the same (0.17) regardless of whether we focus on income (I_{Income}) or wealth (I_{Wealth}) in isolation, or whether we focus on the joint distribution of income and wealth ($I_{Income,Wealth}$). We now consider three alternative societies in which unidimensional and multidimensional measures of inequality of opportunity diverge. As in the main part of the paper, estimates are computed based on Equation (3.3) with dimension weights $r_{Income} = r_{Wealth} = -0.2$ and linear a_t that are inversely related to type ranks in monetary resources.

		Income	Wealth
Type 1	50	500	
Type 2	100	1000	
I_{Income}		= 0.17	
I_{Wealth}		= 0.17	
$I_{Income,Wealth}$		= 0.17	
(a)	(b)	(c)	
Income		Wealth	
Type 1	75	500	
Type 2	75	1000	
I_{Income}		= 0.00	
I_{Wealth}		= 0.17	
$I_{Income,Wealth}$		= 0.09	
Income		Wealth	
Type 1	50	1000	
Type 2	100	500	
I_{Income}		= 0.17	
I_{Wealth}		= 0.17	
$I_{Income,Wealth}$		= 0.06	
Income		Wealth	
Type 1	40	1100	
Type 2	110	400	
I_{Income}		= 0.27	
I_{Wealth}		= 0.27	
$I_{Income,Wealth}$		= 0.12	

- (a) We equalize outcomes across types in the income dimension. Therefore, I_{Income} decreases, and I_{Wealth} stays the same. The multidimensional measure $I_{Income,Wealth}$ decreases. This case illustrates the measure's *inequality aversion between types*.
- (b) We maintain inequality across types but reverse the cross-type association of income and wealth. Therefore, I_{Income} stays the same, and I_{Wealth} stays the same. The multidimensional measure $I_{Income,Wealth}$ decreases. This case illustrates the measure's *sensitivity to correlation-increasing transfers*.
- (c) We increase inequality across types in both dimensions and reverse the cross-type association of income and wealth. Therefore, I_{Income} increases, and I_{Wealth} increases. The multidimensional measure $I_{Income,Wealth}$ decreases. This case illustrates the existence of cases where unidimensional and multidimensional measures lead to opposing conclusions. While the former would detect an increase of inequality of opportunity in comparison to the baseline, the latter would detect a decrease in unequal opportunities.

C.2 Attribute Decomposition

In this appendix, we derive and prove the attribute decomposability of $I(X)$ as defined in Equation (3.3). Our derivation is based on results presented in Abul Naga and Geoffard (2006). For the exposition, we focus on the case of two outcome dimensions with $K = 2$.¹ In this case, X consists of two submatrices X_1 and X_2 that denote outcome matrices for dimensions 1 and 2, respectively. Recall that μ_q^t denotes a type mean in outcome dimension q . Given the notation with two submatrices, an element x_{i1} (x_{i2}) of matrix X_1 (X_2) equals μ_1^t (μ_2^t), i.e., the mean value of dimension 1 (dimension 2) in type t to which individual i belongs. Finally, recall that μ_q denotes the population mean of dimension q .

Attribute Decomposability. In general, $I(X) = 1 - \delta(X)$, where $\delta(X) \in [0, 1]$. $I(X)$ is attribute decomposable if and only if

$$\delta(X) = f_1(\gamma_1(X_1)) + f_2(\gamma_2(X_2)) + f_3(\kappa(X)), \quad (\text{C.1})$$

where f_1, f_2, f_3 are increasing functions ($\mathbb{R}_+ \mapsto \mathbb{R}_+$), γ_1 and γ_2 are unidimensional equality indices, and κ is a measure of association between X_1 and X_2 .

Proposition 1. $\delta(X)$ is attribute decomposable as follows:

$$\ln \delta(X) = \frac{r_1}{r_1 + r_2} \ln \gamma_1(X_1) + \frac{r_2}{r_1 + r_2} \ln \gamma_2(X_2) + \frac{1}{r_1 + r_2} \ln \kappa(X), \quad (\text{C.2})$$

where

$$\begin{aligned} \gamma_1(X_1) &= \left(\sum_{t=1}^M \frac{N_t a_t}{\sum_{t=1}^M N_t a_t} \left(\frac{\mu_1^t}{\mu_1} \right)^{r_1} \right)^{\frac{1}{r_1}}, \\ \gamma_2(X_2) &= \left(\sum_{t=1}^M \frac{N_t a_t}{\sum_{t=1}^M N_t a_t} \left(\frac{\mu_2^t}{\mu_2} \right)^{r_2} \right)^{\frac{1}{r_2}}, \\ \kappa(X) &= \frac{\sum_{t=1}^M N_t a_t \sum_{t=1}^M N_t a_t (\mu_1^t)^{r_1} (\mu_2^t)^{r_2}}{\sum_{t=1}^M N_t a_t (\mu_1^t)^{r_1} \sum_{t=1}^M N_t a_t (\mu_2^t)^{r_2}}. \end{aligned}$$

Proof. First, $\delta(X)$ is the proportion of μ_q that is necessary to achieve the same level of welfare if all attributes were distributed equally across types, see Kobus et al. (2020). Formally, let $w_0 = \sum_{t=1}^M N_t U^t(\delta \mu_1, \delta \mu_2)$ denote the welfare level associated with X . Second, let ρ_1 be the

¹ We note this restriction can be easily relaxed.

proportion of μ_1 that is necessary to attain w_0 , if (i) the first attribute was equally distributed across types, and (ii) the distribution of the second attribute across types remained as is. Formally, $w_0 = \sum_{t=1}^M N_t U^t(\rho_1 \mu_1, \mu_2^t)$. Third, let γ_1 be the proportion of μ_1 that is necessary to attain w_0 , if (i) the first attribute was equally distributed across types, and (ii) the second attribute was equally distributed across types. Formally, $w_0 = \sum_{t=1}^M N_t U^t(\gamma_1 \mu_1, \rho_2 \mu_2)$.

It follows that

$$w_0 = \sum_{t=1}^M N_t a_t(\delta \mu_1)^{r_1} (\delta \mu_2)^{r_2} = \sum_{t=1}^M N_t a_t(\gamma_1 \mu_1)^{r_1} (\rho_2 \mu_2)^{r_2}.$$

After modification, we get $\delta^{r_1+r_2} = (\gamma_1)^{r_1} (\rho_2)^{r_2}$, and we obtain

$$\ln(\delta) = \frac{r_1}{r_1 + r_2} \ln(\gamma_1) + \frac{r_2}{r_1 + r_2} \ln(\rho_2) + \frac{1}{r_1 + r_2} \ln(\rho_2/\gamma_2)^{r_2}, \quad (\text{C.3})$$

which is the desired decomposition with $\kappa := (\rho_2/\gamma_2)^{r_2}$.

We now need to derive functional forms of γ_1 , γ_2 and κ .

Note that $w_0 = \sum_{t=1}^M N_t a_t(\gamma_1 \mu_1)^{r_1} (\rho_2 \mu_2)^{r_2} = \sum_{h=1}^M N_t a_t(\mu_1^t)^{r_1} (\rho_2 \mu_2)^{r_2}$. Solving for γ_1 yields:

$$\gamma_1 = \left(\sum_{t=1}^M \frac{N_t a_t}{\sum_{t=1}^M N_t a_t} \left(\frac{\mu_1^t}{\mu_1} \right)^{r_1} \right)^{\frac{1}{r_1}}.$$

Proceeding in analogy, for γ_2 we get:

$$\gamma_2 = \left(\sum_{t=1}^M \frac{N_t a_t}{\sum_{t=1}^M N_t a_t} \left(\frac{\mu_2^t}{\mu_2} \right)^{r_2} \right)^{\frac{1}{r_2}}.$$

Furthermore, we use $w_0 = \sum_{t=1}^M N_t a_t(\mu_1^t)^{r_1} (\rho_2 \mu_2)^{r_2} = \sum_{t=1}^M N_t a_t(\mu_1^t)^{r_1} (\mu_2^t)^{r_2}$ to obtain

$$\rho_2 = \left(\frac{\sum_{t=1}^M N_t a_t(\mu_1^t)^{r_1} (\mu_2^t)^{r_2}}{\sum_{t=1}^M N_t a_t(\mu_1^t)^{r_1} (\mu_2)^{r_2}} \right)^{\frac{1}{r_2}}.$$

Finally, substituting the expressions for γ_2 and ρ_2 into $\kappa := (\rho_2/\gamma_2)^{r_2}$ we get:

$$\kappa = \frac{\sum_{t=1}^M N_t a_t \sum_{t=1}^M N_t a_t(\mu_1^t)^{r_1} (\mu_2^t)^{r_2}}{\sum_{t=1}^M N_t a_t(\mu_1^t)^{r_1} \sum_{t=1}^M N_t a_t(\mu_2^t)^{r_2}}.$$

□

Linear Approximation. Collecting terms and reversing the log-linearization of $\delta(X)$, we obtain the attribute decomposition of $I(X)$ displayed in Equation (C.1):

$$I(X) = 1 - (\gamma_1)^{\frac{r_1}{r_1+r_2}} (\gamma_2)^{\frac{r_2}{r_1+r_2}} (\kappa)^{\frac{1}{r_1+r_2}}. \quad (\text{C.4})$$

Applying a linear approximation around the point of perfect equality (i.e., $\gamma_1 = \gamma_2 = \kappa = 1$), we get the linear decomposition displayed in Equation (3.4):

$$\begin{aligned} I(X) &= \frac{r_1}{r_1+r_2}(1 - \gamma_1) + \frac{r_2}{r_1+r_2}(1 - \gamma_2) + \frac{1}{r_1+r_2}(1 - \kappa) + R, \\ &= \frac{r_1}{r_1+r_2}I_1 + \frac{r_2}{r_1+r_2}I_2 + \frac{1}{r_1+r_2}\kappa_I + R. \end{aligned} \quad (\text{C.5})$$

C.3 Additional Figures and Tables

Table C.1: Descriptive Statistics

	Income	Wealth	Family Background				Age	N
			Educ.	Occ.	Race	Region		
<i>Panel (A): Intergenerational Sample</i>								
	55,745	279,508	2.26	2.32	0.87	0.26	50	1,366
<i>Panel (B): Re-weighted Intergenerational Sample</i>								
	49,150	205,851	2.20	2.28	0.75	0.34	46	1,366
<i>Panel (C): Individual Sample</i>								
1983	34,763	160,272	1.75	1.87	0.84	0.32	41	5,368
1988	42,258	171,767	1.87	1.94	0.82	0.31	40	5,357
1993	40,562	166,160	1.95	2.00	0.81	0.31	41	5,070
1998	44,563	183,894	2.04	2.09	0.79	0.37	42	4,213
2000	46,136	200,777	2.06	2.12	0.78	0.37	43	4,106
2002	46,414	199,987	2.05	2.13	0.77	0.38	43	4,238
2004	48,897	228,384	2.05	2.14	0.77	0.36	43	5,197
2006	49,406	254,948	2.07	2.15	0.76	0.36	43	5,250
2008	48,349	214,078	2.08	2.16	0.76	0.36	44	5,079
2010	45,490	189,976	2.09	2.18	0.73	0.35	44	5,039
2012	46,377	167,962	2.10	2.19	0.72	0.36	44	5,047
2014	46,373	178,185	2.12	2.20	0.71	0.35	44	5,013
2016	46,837	188,240	2.12	2.21	0.70	0.35	44	4,957

Data: PSID.

Note: This table displays summary statistics for the *intergenerational sample* (Panel [A]), the re-weighted *intergenerational sample* (Panel [B]) and the *individual sample* (Panel [C]). Income is defined as annual disposable household income, wealth as household net worth. Both income and wealth are scaled by the modified OECD equivalence scale and expressed in constant 2015 USD. We furthermore drop observations with negative income/wealth and set zero amounts to 1 USD. The family background variables Educ. (Occ.) show the average education (occupation) level of the parent with the highest education (occupation) status measured on a 3-point scale. Race displays the share of whites; region the share of respondents who grew up in the US Census region South. Age refers to the average age in the sample. The last column shows the number of observations.

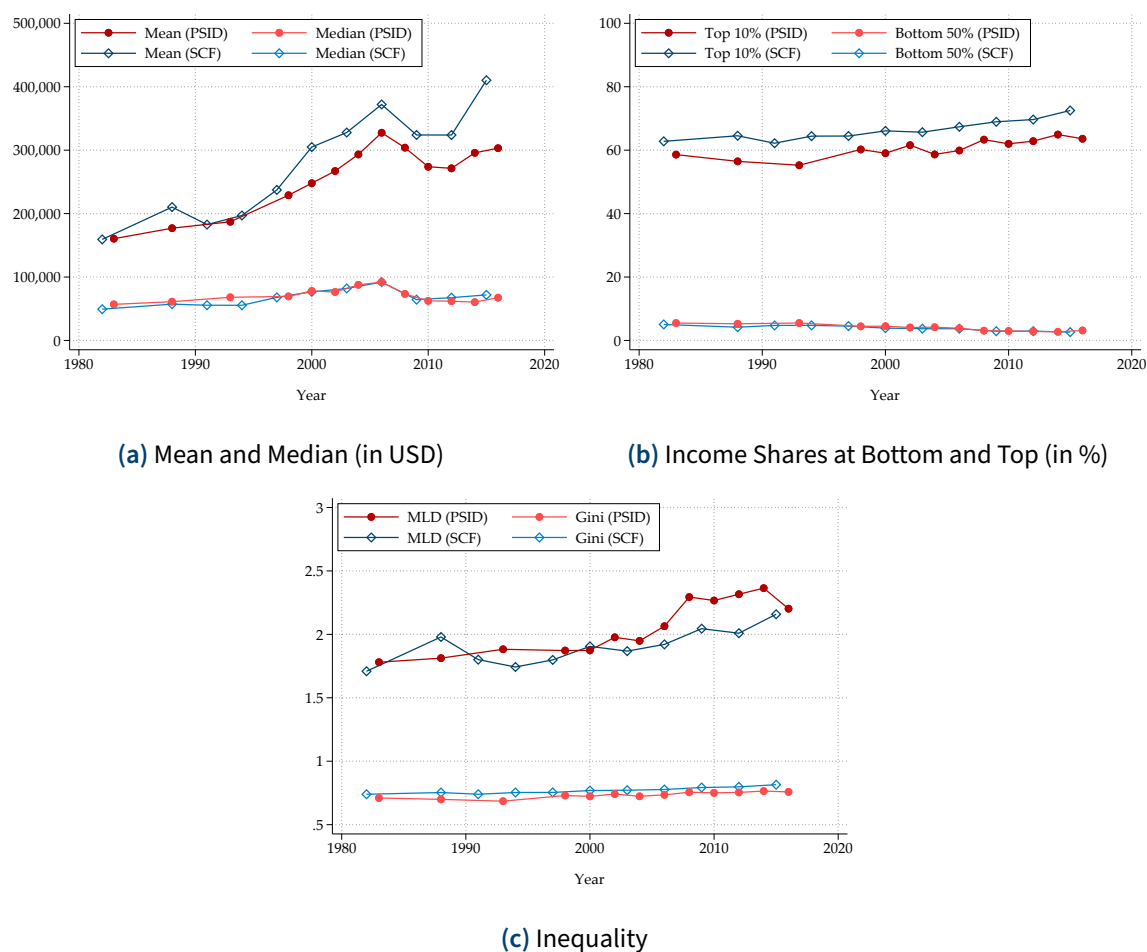
Table C.2: Observation Loss due to Sample Restrictions

Time Period	Sample Restriction	Observations	Share of Original Sample
<i>Panel (A): Intergenerational Sample</i>			
2010-2016		8,824	100%
2010-2016	(i)	8,061	91.4%
2010-2016	(ii)	7,187	89.3%
2010-2016	(iii)	1,366	15.5%
<i>Panel (B): Individual Sample</i>			
1983-2016		80,918	100%
1983-2016	(i)	69,267	81.4%
1983-2016	(ii)	63,934	79.0%
1983		6,257	100%
1983	(i)	5,732	91.6%
1983	(ii)	5,368	85.8%
1993		6,012	100%
1993	(i)	5,384	89.6%
1993	(ii)	5,070	84.3%
2004		6,392	100%
2004	(i)	5,610	87.8%
2004	(ii)	5,197	81.3%
2014		6,896	100%
2014	(i)	5,569	80.8%
2014	(ii)	5,013	72.7%

Data: PSID.

Note: This table shows the loss in observations when defining our analysis samples. For the *intergenerational sample*, we require (i) non-missing and non-negative information on income and wealth, and (ii) non-missing information on parental education, parental occupation, race, and region of upbringing. We further require that (iii) information on the income of both parents is non-missing. For the *individual sample*, we require (i) and (ii) only.


Figure C.1: Wealth in PSID and SCF, 1983-2016



Data: PSID, SCF+ (Kuhn et al., 2020).

Note: This figure compares wealth distributions between the PSID and the Survey of Consumer Finances (SCF). In both data sources, wealth is defined as equivalized household net worth (see Section 3.3); we drop negative values, replace zero values with 1 USD, and winsorize from above at the 99.9 percentile. Samples are restricted to household heads. All figures are expressed in constant 2015 USD.


Figure C.2: Parental Income by Family Background Type



Data: PSID.

Note: This figure shows the mean, 5th-percentile, and 95th-percentile of income distributions within different types of family background characteristics in the *intergenerational sample*. Blue circles and whiskers refer to self-reported family background characteristics. Red triangles and whiskers refer to parental income ranks. Note that bins of family background characteristics tend to be of unequal size, whereas bins of parental income ranks tend to be of equal size (in absence of ties). This feature and resulting differences in the weighting of types explain that averages based on parental income ranks tend to be slightly higher than corresponding averages based on self-reported family background characteristics although the underlying income distributions are the same.

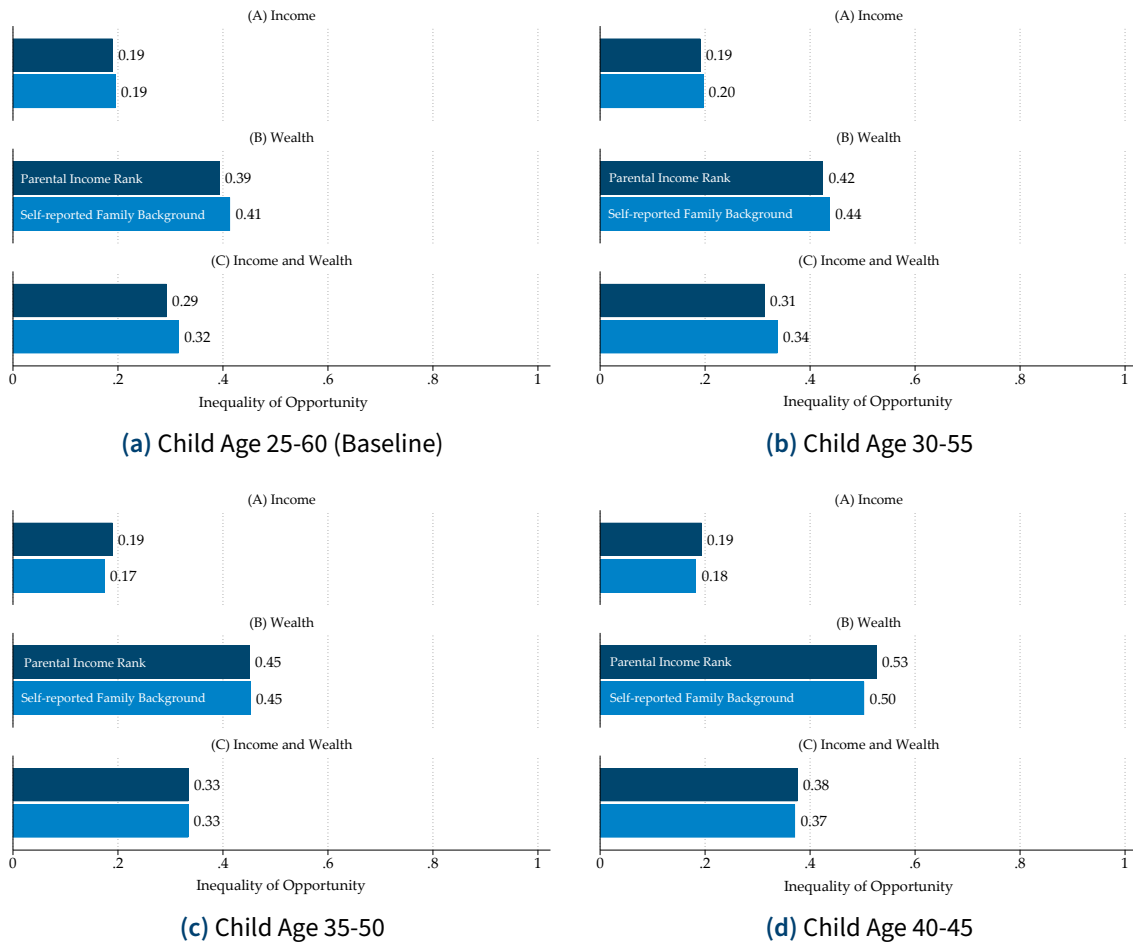
Figure C.3: Inequality of Opportunity in the US
Re-weighted Sample



Data: PSID.

Note: This figure shows the sensitivity of inequality of opportunity in the US when accounting for selective sample attrition in the *intergenerational sample*. In particular, we re-weight the *intergenerational sample* to match the *individual sample* in the observation period 2010-2016 concerning age, parental education, parental occupation, race, and region of upbringing. All estimates are computed based on Equation (3.3) with dimension weights $r_{Income} = r_{Wealth} = -0.2$ and use parental income rank as a proxy for family background.


**Figure C.4: Equality of Opportunity in the US
Varying Age Restrictions**




Data: PSID.

Note: This figure shows the sensitivity of inequality of opportunity in the US to different sample restrictions regarding the age of children. Panel (A) replicates our baseline estimates from Figure 3.2. In Panels (B)-(D) we sequentially narrow the age restriction to 30-55, 35-50, and 40-45. All estimates are computed based on Equation (3.3) with dimension weights $r_{Income} = r_{Wealth} = -0.2$.

Figure C.5: Equality of Opportunity in the US
Extended Family Background Characteristics



(a) Adding Parental Income Ranks



(b) Adding Parental Wealth Ranks

Data: PSID.

Note: This figure shows the sensitivity of inequality of opportunity in the US to extended sets of eligible family background characteristics. The blue bars replicate our baseline estimates from Figure 3.2 based on self-reported family background characteristics (parental education, parental occupation, race, region of upbringing). The red bars show estimates when adding 100 parental income ranks (Panel A) or 100 parental wealth ranks (Panel B) to self-reported family background characteristics and selecting types via a regression tree algorithm. All estimates are computed based on Equation (3.3) with dimension weights $r_{Income} = r_{Wealth} = -0.2$.

Figure C.6: Inequality of Opportunity in the US, 1983-2016
Sensitivity to Data Choices



Data: PSID.

Note: This figure shows the sensitivity of inequality of opportunity in the US for the *individual sample* over the period 1983-2016. In Panel (A), we take a 5-year moving average of income and wealth. In Panel (B), we let a regression tree determine the underlying type partition. In Panel (C), we keep zero income and wealth without adjustment (solid line) or drop individuals with zero income or wealth (dashed line). Panel (D) displays our estimates for the sub-components of labor income and wealth net of active savings in the period of interest. Estimates are computed based on Equation (3.3) with dimension weights $r_{Income} = r_{Wealth} = -0.2$.

**Table C.3: Attribute Decomposition
Alternative Parameterization**

r_{Income}	r_{Wealth}	Contribution of		
		Income	Wealth	Association
-0.1	-0.1	42%	54%	2%
-0.2	-0.2	43%	51%	6%
-0.3	-0.3	43%	49%	12%
-0.4	-0.4	44%	46%	17%

Data: PSID.

Note: This table displays the relative contribution of I_{Income} , I_{Wealth} , and κ_I to the increase in multidimensional inequality of opportunity over the time period 1983-2016. The decomposition is based on the attribute decomposition derived in Appendix C.2. Due to the linear approximation error, the percentages do not exactly add up to 100.

D Appendix to Chapter 4


D.1 Data Appendix

D.1.1 Additional Data Description SOEP

In this section, I first provide additional descriptive statistics about career preferences. Second, I describe additional variables from the SOEP used as covariates in my analysis. Third, I describe how the analysis sample is defined.

Descriptive Statistics on Career Preferences. As discussed in the main body of the paper, Figure D.1 shows the average assessment of each career aspect, ordered from high to low. Table D.1 reports the point estimates of the unconditional and conditional correlations that are displayed in Figure 4.2. Finally, D.2 displays the original question including all answer categories in German, along with the English translation.

Figure D.1: Ordered Career Preferences



Notes: This figure shows the mean desirability of each of the twelve career aspects elicited in the SOEP. The y-axis shows the four-point Likert scale on which respondents rate the career preferences.

Source: German Socio-Economic Panel (SOEP)

Table D.1: Correlation Table of Career Preferences

	1	2	3	4	5	6	7	8	9	10	11	12
1 Job Security	1	-	-	-	-	-	-	-	-	-	-	-
	1	-	-	-	-	-	-	-	-	-	-	-
2 High Income	0.27	1	-	-	-	-	-	-	-	-	-	-
	0.08	1	-	-	-	-	-	-	-	-	-	-
3 Chances of Promotion	0.29	0.41	1	-	-	-	-	-	-	-	-	-
	0.03	0.24	1	-	-	-	-	-	-	-	-	-
4 Respect and Recognition	0.23	0.25	0.38	1	-	-	-	-	-	-	-	-
	-0.08	0.00	0.10	1	-	-	-	-	-	-	-	-
5 Leisure	0.09	0.27	0.12	0.15	1	-	-	-	-	-	-	-
	-0.13	0.14	-0.12	-0.12	1	-	-	-	-	-	-	-
6 Interesting Job	0.07	0.01	0.12	0.16	0.09	1	-	-	-	-	-	-
	-0.18	-0.20	-0.16	-0.14	-0.09	1	-	-	-	-	-	-
7 Working Independently	0.10	0.05	0.16	0.20	0.06	0.34	1	-	-	-	-	-
	-0.16	-0.17	-0.11	-0.10	-0.14	0.17	1	-	-	-	-	-
8 Contact to Others	0.15	0.01	0.16	0.23	0.05	0.21	0.24	1	-	-	-	-
	-0.17	-0.29	-0.18	-0.13	-0.22	-0.05	-0.01	1	-	-	-	-
9 Important for Society	0.16	0.02	0.19	0.34	0.03	0.16	0.21	0.43	1	-	-	-
	-0.17	-0.30	-0.18	0.02	-0.28	-0.13	-0.08	0.16	1	-	-	-
10 Health/Safety Conditions	0.27	0.12	0.18	0.20	0.16	0.19	0.13	0.18	0.24	1	-	-
	0.01	-0.14	-0.15	-0.16	-0.07	-0.07	-0.16	-0.16	-0.10	1	-	-
11 Family Time	0.19	0.13	0.12	0.14	0.35	0.16	0.10	0.18	0.18	0.36	1	-
	-0.08	-0.11	-0.22	-0.23	0.18	-0.09	-0.18	-0.14	-0.18	0.12	1	-
12 Helping Others	0.15	-0.03	0.11	0.23	0.03	0.17	0.18	0.47	0.55	0.20	0.22	1
	-0.16	-0.33	-0.25	-0.12	-0.25	-0.10	-0.10	0.25	0.34	-0.12	-0.09	1

Notes: This table shows pairwise Pearson correlation coefficients between all twelve career preferences. The first row of each preference displays the raw correlation coefficients, the second row displays correlations that are adjusted by individual fixed effects computed by averaging over all preferences, as displayed in Figure 4.2.

Source: German Socio-Economic Panel (SOEP)

Figure D.2: Career Question in the SOEP Questionnaire

52. Für die Arbeit und die Wahl des Berufs können einem unterschiedliche Dinge wichtig sein.

Wie wichtig ist für Ihre Berufswahl ...	Sehr wichtig	Wichtig	Weniger wichtig	Ganz unwichtig
– eine sichere Berufsstellung?.....	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
– ein hohes Einkommen?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
– gute Aufstiegsmöglichkeiten?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
– ein Beruf, der anerkannt und geachtet wird?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
– ein Beruf, der einem viel Freizeit lässt?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
– eine interessante Tätigkeit?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
– eine Tätigkeit, bei der man selbständig arbeiten kann?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
– viel Kontakt zu anderen Menschen?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
– ein Beruf, der für die Gesellschaft wichtig ist?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
– sichere und gesunde Arbeitsbedingungen?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
– ein Beruf, der einem genügend Zeit für familiäre Verpflichtungen lässt?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
– ein Beruf, bei dem man anderen Menschen helfen kann?	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Notes: This figure shows a screenshot of the question on career preferences in the SOEP questionnaire. The official English translation in the SOEP documentation, available [here](#), reads as follows:

Different things may be important to people in choosing a career. Please state how important each of the following is to you—very important, important, not so important, completely unimportant. How important for your career is....

- a secure job?
- a high income?
- good chances of promotion?
- a respected, recognized career?
- a job that leaves a lot of free time?
- an interesting job, career?
- a job that allows you to work independently?
- a job that provides a lot of contact with other people?
- a job that's important for society?
- a job with good health and safety conditions?
- a job that leaves time for family commitments?
- a job where you can help others?

Values. In addition to career preferences, I use a related question on personal values, which reads: *Different things are important to different people. How important are the following things to you (Very important, important, less important or quite unimportant)?* The question with eight answer categories is asked in the main questionnaire, and therefore also observed for most parents in my sample. This allows me to document intergenerational persistence in values (Figure 4.7) and to test channels by which career preferences are linked to socioeconomic status. I impute missing values using chained equations if up to three categories per respondent are missing.

Beliefs. In studying career preferences, it is important to consider the role of beliefs. The reason is the following: If, for example, children from low SES families believe at age 17 to do comparatively worse in terms of ability or opportunities, they might expect to earn below average later in the labor market or not to be able to choose certain careers. In trying to avoid cognitive dissonance, these individuals may update their preferences to better align them to their beliefs. If this is the case, then what I measure as preferences may actually partially reflect beliefs. For most career aspects studied in this paper, however, it does not seem plausible that individuals hold strong corresponding beliefs about the likelihood of achieving them. For example, job security can also be achieved in many low-paying jobs, and it is not clear why some individuals should face stronger barriers in finding a job that is interesting, that offers more time for leisure or family, or a job where one can help others. For the domains where this is theoretically more plausible, I do not find supporting evidence in the data: for example, children from low income families are actually more likely to desire earning a high income. To nevertheless address this issue empirically, I draw on a question from the youth questionnaire, where respondents are asked about their subjective beliefs of how likely they are to achieve their desired career. While this probability indeed increases in parental income rank, I show in Table 4.1 that including it into the analysis makes little difference.

Trust, Risk Aversion, and the Big Five Personality Traits. To demonstrate that career preferences measure something which is distinct from more fundamental traits and preferences, I control for trust, risk aversion, and the Big Five personality traits in some specifications. All these variables are elicited in the main questionnaire of the SOEP at ages 16-25.¹ Table

¹ The age restriction ensures that these preferences are measured as early in life as possible. The reason I cannot measure trust, risk aversion, and the Big Five personality traits at age 17 for all respondents is that they are not asked for in every wave of the main questionnaire. More recently, these variables are additionally elicited in the

Table D.2: Big Five Scales in the SOEP

Dimension	Description	Direction
Openness	original, someone who comes up with new ideas	+
	someone who values artistic, aesthetic experiences	+
	has an active imagination	+
Conscientiousness	does a thorough job	+
	tends to be lazy	-
	does things effectively and efficiently	+
Extraversion	is communicative, talkative	+
	is outgoing, sociable	+
	is reserved	-
Agreeableness	is sometimes somewhat rude to others	-
	has a forgiving nature	+
	is considerate and kind to others	+
Neuroticism	worries a lot	-
	gets nervous easily	+
	is relaxed, handles stress well	-

Notes: This table shows how the 15 personality questions in the main questionnaire of the SOEP are mapped into the five dimensions of the big five personality traits. Starting in 2009, openness was extended by an additional item (“is eager for knowledge”), which I do not use to ensure comparability over time.

D.2 provides an overview of the personality questions and their mapping into the Big Five dimensions. If up to three input questions for the Big Five dimensions were not answered by a respondent, I impute missing values using chained equations.

Sample Definition. I restrict my sample to all respondents in private households aged 15 to 55 that I can link to at least one of their parents in the data. I further require that the respondent answered the career preference question in the youth questionnaire, leaving me with 8,185 individuals. Whenever I look at the labor market outcomes of these children, such as income or occupation, I additionally require that children are at least 28 years old, focusing on the five-year interval from age 28 to 32. This ensures that (most) children have already left education and entered the labor market, and results in a sample of 825 individuals. In Table 4.1, I further drop all respondents with missing values in any of the covariates, resulting in a sample of

youth questionnaire. However, since these items were not included in the youth questionnaire in the early 2000s where most of my variation is coming from, I cannot use this information without drastically reducing sample size.

787 individuals. The cutoff value of 28 is chosen to balance the tradeoff between sample size and lifecycle bias.² Unless indicated otherwise, all analyses use the SOEP's individual level sampling weights.


D.1.2 Additional Data Description British Cohort Study

In this section, I provide additional descriptive statistics about career preferences in the British Cohort Study (BCS) and describe the definition of the analysis sample.

Descriptive Statistics on Career Preferences. Figure D.3 shows the average assessment of each career aspect, ordered from high to low. Noticeably, respondents rated most of the aspects that were asked for only in the BCS but not in the SOEP to be of mostly minor importance in choosing a career. Instead, the most desired career aspects are an interesting job, an understanding boss and long term security, similar to the most highly rated preferences in the SOEP (compare Figure D.1). Figure D.4 displays the factor loadings of each career preference on the first two components of a principal component analysis (PCA). Here, helping others stands out as sharing relatively little variance with the other preferences along both dimensions. As in the SOEP, however, a Kaiser-Meyer-Olkin measure of sampling adequacy of 0.69 suggests that the different career preferences are not easily reduced to a few common components. Finally, Figure D.5 provides a screenshot of the original questionnaire.


Sample Definition. As in the SOEP, my main sample restriction is that I only retain individuals with non-missing career preferences in my sample. This reduces the sample from 11,620 to 5,618 respondents. One reason for the drop in sample size is a teacher's strike in 1986 that resulted in many subjects not receiving their questionnaires. I further require that respondents reported own income at least once, leaving me with 4,784 individuals. Among those, 419 have missing information on parental income. Dropping these respondents results in a final sample of 4,365 individuals. In Table 4.3, dropping respondents with missing values in any of the covariates further reduces the sample to 3,165 individuals.

² Since the youth questionnaire was introduced only in 2000, the oldest respondents are 38 years old in the last survey wave of the SOEP.

Figure D.3: Ordered Career Preferences in the BCS

Notes: This figure shows the mean desirability of each of the 16 career aspects elicited in the BCS. The y-axis shows the three-point Likert scale on which respondents rate the career preferences.

Source: British Cohort Study (BCS)

Figure D.4: Principal Component Analysis Career Preferences in the BCS

Notes: This figure shows a scatter plot of the factor loadings of the first and second component of a principal component analysis (PCA) of all 16 career preferences elicited in the BCS.

Source: British Cohort Study (BCS)

Figure D.5: Career Question in the BCS Questionnaire

WHATS IN A JOB?

5^D

INSTRUCTIONS

This section consists of a list of questions concerning things which people of your age think to be important in deciding what sort of career they want in the future. We are asking you to indicate for each whether it matters very much to you, matters somewhat or doesn't matter.

We have labelled an example below to show you exactly how to do this:

QUESTION	Matters very much (a)	Matters somewhat (b)	Doesn't matter (c)
How much will it matter to me to work with my hands?	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>

Answer (c) means that you think it will matter somewhat for your job or career that you work with your hands.


QUESTIONS

How much does it matter to you:

1. To be able to help other people (C5D1)
2. To have high earnings/wages? (C5D2)
3. To have an understanding boss? (C5D3)
4. To work outside in the open? (C5D4)
5. To work for myself? (C5D5)
6. To have an interesting job with variety. (C5D6)
7. Not to have to work too hard? (C5D7)
8. To get promotion so I can get ahead? (C5D8)
9. To work with figures? (C5D9)
10. To get trained for a trade or profession? (C5D10)
11. To have a quiet life? (C5D11)
12. To have long term security? (C5D12)
13. To get a job with a real challenge? (C5D13)
14. To have a chance to travel? (C5D14)
15. To make or build things? (C5D15)
16. To have a job with regular hours? (C5D16)

Please now turn to page 4 of the Student Score Form. On that page, in section 5^D, you will find a set of lozenges headed (a), (b) and (c). Record in these lozenges your answers to each of the questions listed here about what might be important for a job or career. You should record your answers to 1-16 by filling in questions on the score form either lozenge (a), (b) or (c), in a similar way to the example above. Remember not to put your answers on this Test Booklet but in the Student Score Form. Fill in only one lozenge in answer to each question.

KEEP THIS PAGE OPEN, LOOK AT THE FIRST QUESTION No. 1 ABOVE AND THEN FILL IN YOUR ANSWERS ON PAGE 4 OF THE STUDENT SCORE FORM. THEN PROCEED TO QUESTION 2 . . . AND SO ON.




To Work
For Yourself?

Notes: This figure shows a screenshot of the question on career preferences in the BCS questionnaire.

D.2 Additional Figures and Tables

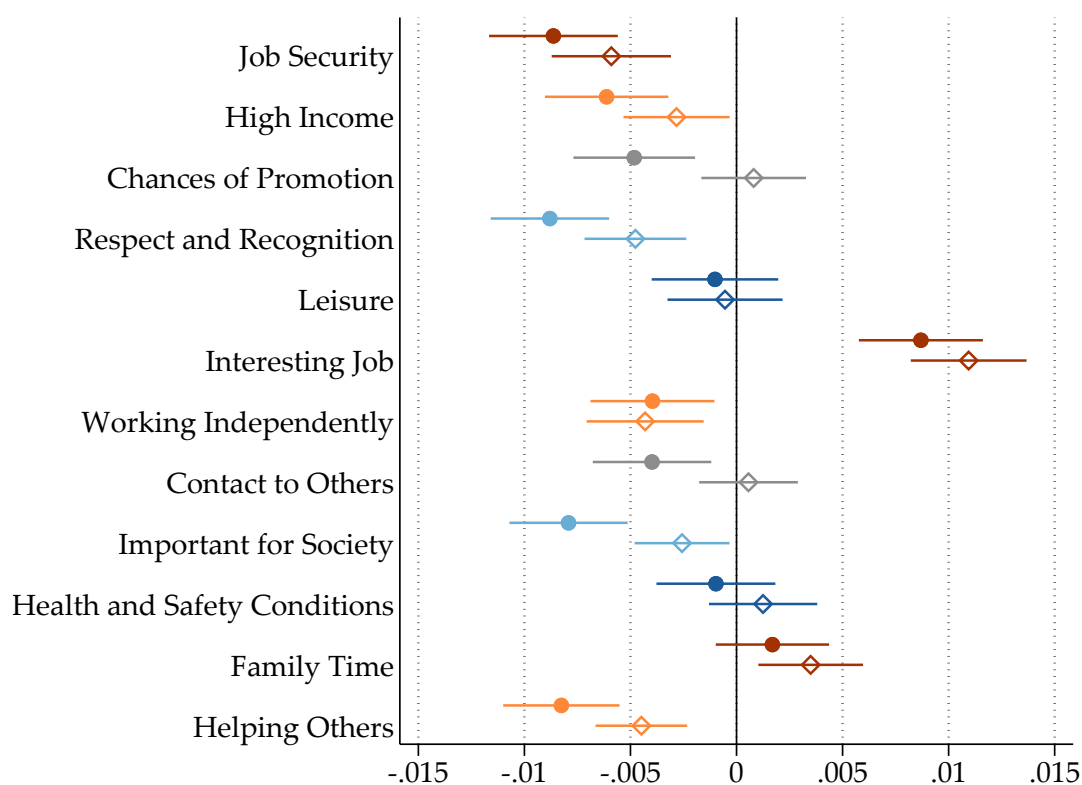
Figure D.6: Variation in Career Preferences by Parental Years of Education



Notes: This figure shows point estimates and the corresponding 95% confidence intervals (robust standard errors) of separate regressions of one of the twelve career preferences on parental years of education. The first estimate (circle) in each category is obtained by regressing the respective preference on years of education of the more educated parent, the second estimate (diamond) by additionally controlling for all other eleven career preferences. The sample consists of 8,162 parent-child pairs.

Source: German Socio-Economic Panel (SOEP)

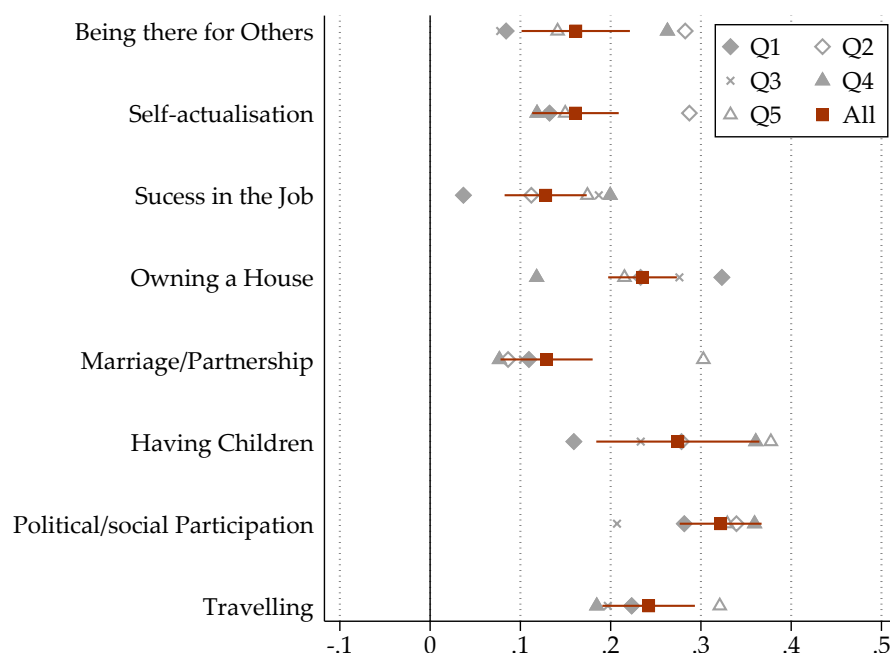
**Figure D.7: Variation in Career Preferences by Parental Occupation
(Standard International Occupational Prestige Scale)**



Notes: This figure shows point estimates and the corresponding 95% confidence intervals (robust standard errors) of separate regressions of one of the twelve career preferences on the parental Standard International Occupational Prestige Scale (SIOPS) value. The first estimate (circle) in each category is obtained by regressing the respective preference on the mean parental SIOPS value, the second estimate (diamond) by additionally controlling for all other eleven career preferences. The sample consists of 8,054 parent-child pairs.

Source: German Socio-Economic Panel (SOEP)

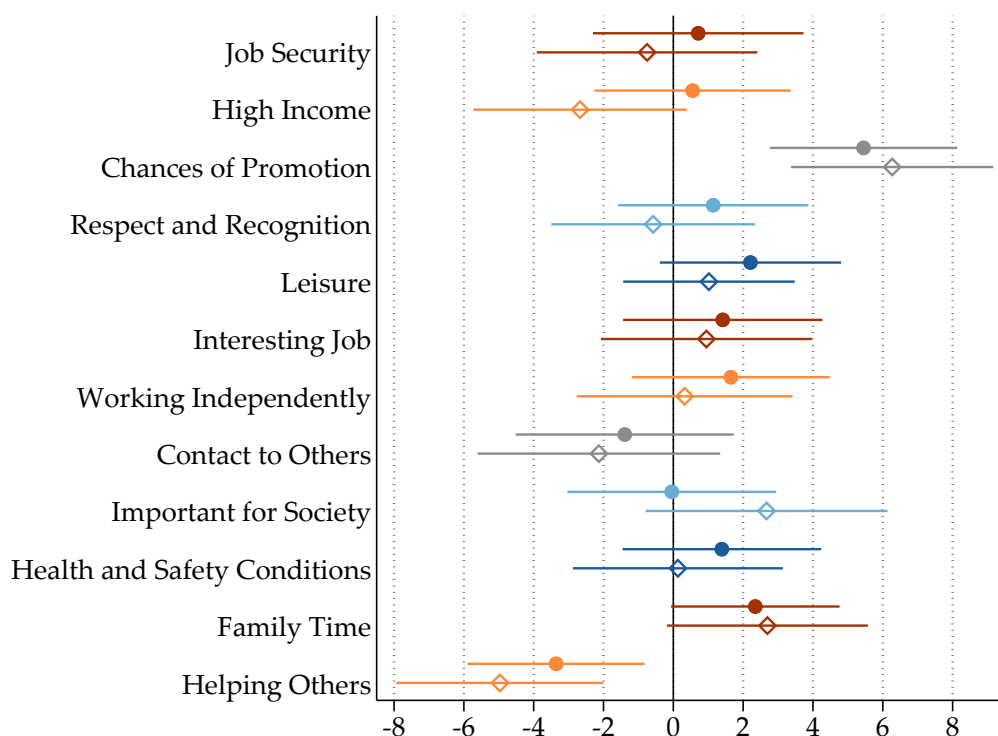
**Figure D.8: Intergenerational Persistence in Preferences and Values –
By Quintile of Parental Income**



Notes: This figure shows point estimates of separate regressions of eight different values on the respective average measure of their parents, based on 4,479 parent-child pairs. The red square and the corresponding 95% confidence intervals (robust standard errors) denote the estimate among all children, whereas the other estimates denote coefficient estimates within each quintile of parental gross household income. The values of the children are measured when children are between 16-25 years old, the values of the parents when children are between 15-19 years old.

Source: German Socio-Economic Panel (SOEP)

Figure D.9: Career Preferences Predict Future Earnings



Notes: This figure shows point estimates and the corresponding 95% confidence intervals (robust standard errors) of separate regressions of the own earnings rank at ages 28 to 32 on one of the twelve career preferences. The first estimate (circle) in each category is obtained by regressing the earnings rank on the respective preference, the second estimate (diamond) by additionally controlling for all other career preferences. The sample consists of 825 parent-child pairs.

Source: German Socio-Economic Panel (SOEP)

Table D.3: Annual Work Hours and Past Career Preferences

	Annual Work Hours					
	(1)	(2)	(3)	(4)	(5)	(6)
Job Security	-9.42 (58.90)	-9.09 (58.80)	-21.63 (57.32)	-8.25 (57.37)	-4.74 (55.37)	-21.24 (53.52)
High Income	-27.27 (48.25)	-20.92 (48.19)	-16.52 (47.42)	3.21 (46.55)	1.84 (45.66)	10.54 (45.86)
Chances of Promotion	100.84** (48.90)	95.50* (49.29)	90.91* (48.69)	99.39** (47.27)	96.71** (46.90)	76.09* (45.86)
Respect and Recognition	-17.37 (42.10)	-12.90 (42.17)	-17.53 (41.31)	-15.71 (42.24)	-6.58 (41.97)	1.80 (41.95)
Leisure	12.48 (41.29)	7.90 (42.46)	7.57 (42.46)	15.54 (43.75)	30.43 (42.11)	38.68 (41.23)
Interesting Job	11.82 (46.96)	9.04 (47.38)	14.07 (46.72)	15.15 (43.66)	28.40 (42.09)	25.36 (42.01)
Working Independently	64.14 (41.23)	64.45 (41.38)	56.15 (40.59)	39.71 (38.95)	44.00 (39.07)	38.35 (39.42)
Contact to Others	-0.61 (52.16)	0.21 (52.14)	9.25 (51.82)	-1.21 (48.11)	-26.79 (47.13)	-22.40 (48.56)
Important for Society	22.57 (59.11)	23.77 (59.04)	17.15 (58.29)	9.72 (54.91)	1.57 (53.56)	-1.69 (50.22)
Health and Safety Conditions	44.09 (55.90)	44.93 (55.82)	47.12 (54.92)	25.22 (51.47)	20.70 (48.91)	8.87 (46.83)
Family Time	-7.45 (51.48)	-12.02 (50.86)	-8.01 (50.49)	-13.44 (50.44)	-22.43 (46.90)	-17.51 (44.81)
Helping Others	-18.63 (51.37)	-16.67 (50.93)	-22.52 (50.15)	-32.83 (49.84)	-27.16 (49.86)	-23.57 (50.98)
Parental Income Rank		1.40 (1.56)	3.06* (1.74)	2.57 (2.13)	2.70 (1.94)	2.62 (1.83)
Parental Years of Education			-37.75** (16.33)	-37.39** (17.25)	-45.17*** (16.96)	-33.42* (17.37)
Probability Desired Career				28.39 (20.41)	28.01 (19.13)	26.77 (20.06)
Gender, State of Birth	-	-	-	✓	✓	✓
Grades, Tracking Recommendation	-	-	-	-	✓	✓
Trust/Risk Preferences, Big Five	-	-	-	-	-	✓
Observations	787	787	787	787	787	787
R^2_{adj}	0.020	0.021	0.030	0.076	0.105	0.126

Notes: This table shows estimates of six separate regressions of annual hours worked for children aged 28 to 32 on past career preferences reported at age 17. Parental income rank refers to gross household income, parental education to years of education of the more educated parent. Column (4) additionally includes gender and dummies for the state of birth, Column (5) dummies for the recommended school track after primary school, the grade average, and interactions between the grade average and the track recommendation. In Column (6), I further control for trust and risk preferences, and measures of the Big Five personality traits. Robust standard errors in parentheses, * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Source: German Socio-Economic Panel (SOEP)

Table D.4: Years of Education and Past Career Preferences

	Years of Education					
	(1)	(2)	(3)	(4)	(5)	(6)
Job Security	0.03 (0.17)	0.02 (0.15)	0.12 (0.15)	0.02 (0.14)	0.09 (0.12)	0.08 (0.12)
High Income	-0.12 (0.18)	0.04 (0.17)	0.04 (0.16)	0.07 (0.15)	0.14 (0.13)	0.17 (0.12)
Chances of Promotion	0.08 (0.18)	-0.02 (0.17)	0.02 (0.17)	0.07 (0.16)	0.05 (0.14)	0.05 (0.14)
Respect and Recognition	-0.23 (0.17)	-0.15 (0.16)	-0.11 (0.15)	-0.03 (0.15)	0.07 (0.12)	0.07 (0.12)
Leisure	0.10 (0.15)	-0.03 (0.14)	-0.01 (0.14)	-0.07 (0.13)	-0.02 (0.12)	-0.03 (0.12)
Interesting Job	0.44*** (0.17)	0.38** (0.17)	0.33** (0.16)	0.26* (0.15)	0.23** (0.11)	0.20* (0.11)
Working Independently	-0.01 (0.17)	-0.08 (0.16)	-0.02 (0.15)	-0.07 (0.15)	-0.06 (0.13)	-0.08 (0.13)
Contact to Others	0.22 (0.17)	0.27* (0.16)	0.19 (0.16)	0.24 (0.18)	0.21 (0.14)	0.24* (0.14)
Important for Society	0.03 (0.22)	0.06 (0.20)	0.12 (0.19)	0.14 (0.19)	-0.01 (0.16)	-0.01 (0.15)
Health and Safety Conditions	-0.25 (0.17)	-0.16 (0.18)	-0.17 (0.18)	-0.17 (0.16)	-0.21 (0.14)	-0.16 (0.14)
Family Time	0.38** (0.16)	0.27* (0.15)	0.22 (0.14)	0.22* (0.13)	0.20* (0.11)	0.19* (0.11)
Helping Others	-0.55** (0.21)	-0.49** (0.20)	-0.42** (0.20)	-0.54*** (0.19)	-0.44*** (0.16)	-0.46*** (0.16)
Parental Income Rank		0.03*** (0.01)	0.02*** (0.01)	0.02*** (0.01)	0.01 (0.01)	0.01 (0.01)
Parental Years of Education			0.33*** (0.06)	0.31*** (0.06)	0.17*** (0.05)	0.15*** (0.05)
Probability Desired Career				-0.00 (0.07)	-0.02 (0.06)	-0.01 (0.06)
Gender, State of Birth	-	-	-	✓	✓	✓
Grades, Tracking Recommendation	-	-	-	-	✓	✓
Trust/Risk Preferences, Big Five	-	-	-	-	-	✓
Observations	646	646	646	646	646	646
R^2_{adj}	0.062	0.163	0.230	0.311	0.483	0.487

Notes: This table shows estimates of six separate regressions of years of education for children aged 28 to 32 on past career preferences reported at age 17. Parental income rank refers to gross household income, parental education to years of education of the more educated parent. Column (4) additionally includes gender and dummies for the state of birth, Column (5) dummies for the recommended school track after primary school, the grade average, and interactions between the grade average and the track recommendation. In Column (6), I further control for trust and risk preferences, and measures of the Big Five personality traits. Robust standard errors in parentheses, * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Source: German Socio-Economic Panel (SOEP)

Table D.5: Career Preferences Mediating Intergenerational Income Mobility

	Percent Explained			
	Ind-Ind	Ind-Gross	Gross-Gross	Net-Net
Job Security	0.5	0.1	-0.1	-0.4
High Income	2.8	2.4	3.6	5.6
Chances of Promotion	4.1	6.0	3.7	3.1
Respect and Recognition	-0.6	-0.1	-2.0	-3.4
Leisure	1.0	0.8	-1.4	-1.2
Interesting Job	0.5	0.6	0.1	-0.3
Working Independently	-0.4	0.2	-0.3	0.1
Contact to Others	-0.6	0.1	0.0	0.4
Important for Society	-3.6	-3.0	-2.1	-3.2
Health and Safety Conditions	-0.2	0.0	0.0	-0.1
Family Time	3.6	4.3	3.8	4.9
Helping Others	5.2	3.2	1.8	2.1
Total	12.2	14.7	7.2	7.7

Notes: This table reports estimates of a descriptive mediation analysis, decomposing the association between child and parent income rank into a direct effect and twelve indirect effects via career preferences, separately for four intergenerational rank-rank correlations. “Ind” refers to individual labor earnings, “Gross” to gross household income, and “Net” to net household income. All numbers are in percent. For example, a value of 2 means that two percent of income persistence is mediated via the respective preference. Negative signs imply that c.p. the preference increases (rather than decreases) income persistence. The last row shows the combined indirect effect of all twelve career preferences jointly.

Source: German Socio-Economic Panel (SOEP)

Table D.6: Career Preferences and Intergenerational Income Mobility – Robustness to Different Weighting Schemes

	Individual Labor Earnings				Net HH Income		Gross HH Income	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<i>Panel (A): in Ranks, Re-weighted</i>								
Ind. Labor Earnings Father	0.236*** (0.082)	0.189*** (0.063)						
Gross HH Income			0.282*** (0.065)	0.237*** (0.057)	0.397*** (0.071)	0.362*** (0.061)		
Net HH Income							0.385*** (0.074)	0.338*** (0.067)
Preferences	-	✓	-	✓	-	✓	-	✓
Observations	700	700	825	825	825	825	825	825
<i>Panel (B): in Ranks, Unweighted</i>								
Ind. Labor Earnings Father	0.150*** (0.033)	0.137*** (0.034)						
Gross HHk Income			0.199*** (0.030)	0.185*** (0.031)	0.319*** (0.031)	0.313*** (0.032)		
Net HH Income							0.337*** (0.031)	0.329*** (0.031)
Preferences	-	✓	-	✓	-	✓	-	✓
Observations	765	765	904	904	904	904	904	904

Notes: This table shows estimates of separate regressions of child on parental income rank. Three different income concepts are used: gross individual labor earnings, net household income and gross household income. In Panel A, I reweight the sample on the basis of the SOEP survey weights with respect to gender (2 categories), migration background (2 categories) and parental education (3 categories). In Panel B, no survey weights are used at all. The higher number of observations results from individuals receiving a weight of zero under the standard sampling frame of the SOEP. In every second column, I additionally control for all twelve career preferences. Robust standard errors in parentheses, * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Source: German Socio-Economic Panel (SOEP)

Table D.7: Career Preferences and Intergenerational Income Mobility – Controlling for Childhood Characteristics

	Individual Labor Earnings				Gross HH Income		Net HH Income	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<i>Panel (A): in Ranks</i>								
Ind. Labor Earnings Father	0.0760 (0.066)	0.0526 (0.061)						
Gross HH Income			0.200*** (0.057)	0.172*** (0.053)	0.238*** (0.056)	0.215*** (0.055)		
Net HH Income							0.269*** (0.059)	0.233*** (0.057)
Full Set of Table 4.1 Controls	✓	✓	✓	✓	✓	✓	✓	✓
Preferences	-	✓	-	✓	-	✓	-	✓
Observations	671	671	787	787	787	787	787	787
<i>Panel (B): in Logs</i>								
Ind. Labor Earnings Father	0.0725 (0.091)	0.0489 (0.083)						
Gross HH Income			0.142 (0.090)	0.102 (0.088)	0.278*** (0.087)	0.239*** (0.088)		
Net HH Income							0.357*** (0.081)	0.309*** (0.078)
Full Set of Table 4.1 Controls	✓	✓	✓	✓	✓	✓	✓	✓
Preferences	-	✓	-	✓	-	✓	-	✓
Observations	624	624	748	748	781	781	787	787

Notes: This table shows estimates of separate regressions of child on parental income. Three different income concepts are used: gross individual labor earnings, gross household income and net household income. For individual labor earnings of the parents, I focus only on earnings of the father, as mothers display large variation at the extensive margin of labor supply. In Panel A, both child and parent incomes are measured in 100 percentile ranks and the estimates represent rank-rank slopes. In Panel B, incomes are measured in logarithmic form and the estimates represent the intergenerational elasticity (IGE). In all columns, I control for the full set of controls in Column (6) of Table 4.1, i.e. parental gross household income, parental years of education of the more educated parent, dummies for gender, dummies for the state of birth and the recommended school track after primary school, the grade average, interactions between the grade average and the track recommendation, trust and risk preferences, and measures of the Big Five personality traits. In every second column, I additionally control for all 12 career preferences. Robust standard errors in parentheses, * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Source: German Socio-Economic Panel (SOEP)

Table D.8: Career Preferences Mediating Intergenerational Income Mobility in the British Cohort Study

	Percent Explained	
	IGE	Rank-Rank
Helping Others	2.6	2.5
High Income	-0.3	-0.2
Understanding Boss	0.2	0.1
Working Outside	0.0	-0.0
Working Independently	0.6	0.9
Interesting Job	0.9	0.0
Not Work Hard	0.1	0.3
Chances of Promotion	3.3	4.2
Work with Figures	0.5	-0.1
Get Trained	-0.0	-0.0
Quiet Life	0.3	0.5
Long Term Security	1.7	2.6
Real Challenge	2.4	2.7
Travel	-0.8	-0.9
Make/build Things	0.1	0.1
Regular Hours	6.7	7.0
Total	18.2	19.6

Notes: This table reports estimates of a descriptive mediation analysis in the BCS, decomposing the association between child and parent incomes into a direct effect and 16 indirect effects via career preferences. All numbers are in percent. For example, a value of 2 means that two percent of income persistence is mediated via the respective preference. Negative signs imply that c.p. the preference increases (rather than decreases) income persistence. The last row shows the combined indirect effect of all 16 career preferences jointly.

Source: British Cohort Study (BCS)

**Table D.9: Career Preferences and Intergenerational Income Mobility
in the British Cohort Study Controlling for Childhood Characteristics**

	Log Gross Earnings		Gross Earnings Rank	
	(1)	(2)	(3)	(4)
Log Gross Income Parents	0.203*** (0.035)	0.186*** (0.034)		
Gross Income Rank Parents			0.105*** (0.018)	0.0969*** (0.018)
Full Set of Table 4.3 Controls	✓	✓	✓	✓
Preferences	-	✓	-	✓
Observations	2934	2934	3165	3165

Notes: This table shows estimates of separate regressions of child on parental income. Income of children is measured as gross weakly individual labor earnings between ages 30-46, parental income as gross weakly household income when children are 10-16 years old. In the first two columns, both child and parent incomes are measured in logs, whereas in the last two columns incomes are measured in 100 percentile ranks. In all columns, I control for the full set of controls in Column (6) of Table 4.3, i.e. parental gross household income rank, parental education (indicating if father and/or mother have an A-level degree), age 10 test scores for language, reading, math and matrices, the number of attended career talks by age 16, and a set of further beliefs on what helps in advancing careers. In Columns (2) and (4), I additionally control for all 16 career preferences. Robust standard errors in parentheses, * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Source: British Cohort Study (BCS)

Bibliography

Abeler, Johannes, Armin Falk, and Fabian Kosse, “Malleability of Preferences for Honesty,” *CEPR Discussion Paper*, 2021, 16164.

Abramitzky, Ran, Leah Boustan, Elisa Jacome, and Santiago Perez, “Intergenerational Mobility of Immigrants in the United States over Two Centuries,” *American Economic Review*, 2021, 111 (2), 580–608.

Abul Naga, Ramses H. and Pierre-Yves Geoffard, “Decomposition of Bivariate Inequality Indices by Attributes,” *Economics Letters*, 2006, 90 (3), 362–367.

Acciari, Paolo, Alberto Polo, and Giovanni L. Violante, “And Yet It Moves: Intergenerational Mobility in Italy,” *American Economic Journal: Applied Economics*, 2022, 14 (3), 118–163.

Adda, Jérôme, Christian Dustmann, and Katrien Stevens, “The Career Costs of Children,” *Journal of Political Economy*, 2017, 125 (2), 293–337.

Alan, Sule, Nazli Baydar, Teodora Boneva, Thomas F. Crossley, and Seda Ertac, “Transmission of Risk Preferences from Mothers to Daughters,” *Journal of Economic Behavior & Organization*, 2017, 134, 60–77.

Alesina, Alberto, Stefanie Stantcheva, and Edoardo Teso, “Intergenerational Mobility and Preferences for Redistribution,” *American Economic Review*, 2018, 108 (2), 521–54.

Algan, Yann and Pierre Cahuc, “Inherited Trust and Growth,” *American Economic Review*, 2010, 100 (5), 2060–2092.

Almeida, Heitor, Murillo Campello, and Michael S. Weisbach, “The Cash Flow Sensitivity of Cash,” *Journal of Finance*, 2004, 59 (4), 1777–1804.

Almås, Ingvild, Alexander W. Cappelen, and Bertil Tungodden, “Cutthroat Capitalism versus Cuddly Socialism: Are Americans More Meritocratic and Efficiency-Seeking than Scandinavians?,” *Journal of Political Economy*, 2020, 128, 1753–1788.

Altmejd, Adam, “Inheritance of Fields of Study,” *IFAU Working Paper*, 2023, 11.

- Arcidiacono, Peter, Esteban M. Aucejo, and V. Joseph Hotz**, “University Differences in the Graduation of Minorities in STEM Fields: Evidence from California,” *American Economic Review*, 2016, 106 (3), 525–562.
- Aristei, David and Cristiano. Perugini**, “Inequality Aversion in Post-communist Countries in the Years of the Crisis,” *Post-Communist Economies*, 2016, 28 (4), 436–448.
- Armstrong, Timothy B., Michal Kolesár, and Mikkel Plagborg-Møller**, “Robust Empirical Bayes Confidence Intervals,” *Econometrica*, 2022, 90 (6), 2567–2602.
- Arneson, Richard J.**, “Equality and Equal Opportunity for Welfare,” *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, 1989, 56 (1), 77–93.
- , “Four Conceptions of Equal Opportunity,” *The Economic Journal*, 2018, 128, F152–F173.
- Atkinson, Anthony B.**, “On the Measurement of Inequality,” *Annual Review of Economics*, 1970, 2 (3), 244–263.
- Attanasio, Orazio, Áureo de Paula, and Alessandro Toppeta**, “The Persistence of Socio-emotional Skills: Life Cycle and Intergenerational Evidence,” *NBER Working Paper*, 2020, 27823.
- Auerbach, Alan J.**, “Who Bears the Corporate Tax? A Review of What We Know,” *Tax Policy and the Economy*, 2006, 20, 1–40.
- **and Yuriy Gorodnichenko**, “Output Spillovers from Fiscal Policy,” *American Economic Review*, 2013, 103 (3), 141–46.
- Babyak, Michael A.**, “What You See May Not Be What You Get: A Brief, Nontechnical Introduction to Overfitting in Regression-Type Models,” *Psychosomatic Medicine*, 2004, 66 (3), 411–421.
- Bachmann, Rüdiger and Peter Zorn**, “What Drives Aggregate Investment? Evidence from German Survey Data,” *Journal of Economic Dynamics and Control*, 2020, 115, 103873.
- , **Steffen Elstner, and Atanas Hristov**, “Surprise, Surprise – Measuring Firm-level Investment Innovations,” *Journal of Economic Dynamics and Control*, 2017, 83, 107–148.
- Barattieri, Alessandro, Susanto Basu, and Peter Gottschalk**, “Some Evidence on the Importance of Sticky Wages,” *American Economic Journal: Macroeconomics*, 2014, 6 (1), 70–101.

- Barsky, Robert B., F. Thomas Juster, Miles S. Kimball, and Matthew D. Shapiro**, “Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Study,” *The Quarterly Journal of Economics*, 1997, 112 (2), 537–579.
- Bayer, Patrick and Kerwin Kofi Charles**, “Divergent Paths: A New Perspective on Earnings Differences Between Black and White Men Since 1940,” *The Quarterly Journal of Economics*, 2018, 133, 1459–1501.
- Becker, Sascha O., Lukas Mergele, and Ludger Woessmann**, “The Separation and Reunification of Germany: Rethinking a Natural Experiment Interpretation of the Enduring Effects of Communism,” *Journal of Economic Perspectives*, 2020, 34 (2), 143–171.
- Belloni, Alexandre and Victor Chernozhukov**, “Least Squares after Model Selection in High-Dimensional Sparse Models,” *Bernoulli*, 2013, 19 (2), 521–547.
- Berman, Yonatan and Branko Milanovic**, “Homoploutia: Top Labor and Capital Incomes in the United States, 1950–2020,” *Review of Income and Wealth*, 2023.
- Bertrand, Marianne, Magne Mogstad, and Jack Mountjoy**, “Improving Educational Pathways to Social Mobility: Evidence from Norway’s Reform 94,” *Journal of Labor Economics*, 2021, 39 (4), 965–1010.
- Best, Kathinka L., Ulrike Sanwald, Susanne Ihlen, and Angela Ittel**, “Gender and STEM in Germany: Policies Enhancing Women’s Participation in Academia,” *International Journal of Gender, Science and Technology*, 2013, 5 (3), 292–304.
- Betthäuser, Bastian**, “Fostering Equality of Opportunity? Compulsory Schooling Reform and Social Mobility in Germany,” *European Sociological Review*, 2017, 33 (5), 633–644.
- Biewen, Martin and Madalina Tapalaga**, “Life-Cycle Educational Choices in a System with Early Tracking and ‘Second Chance’ Options,” *Economics of Education Review*, 2017, 56, 80–94.
- Black, Sandra E. and Paul J. Devereux**, “Recent Developments in Intergenerational Mobility,” in David Card and O Ashenfelter, eds., *Handbook of Labor Economics*, Vol. 4B, Amsterdam: Elsevier, 2011, chapter 16, pp. 1487–1541.
- Blesse, Sebastian, Philipp Doerrenberg, and Anna Rauch**, “Higher Taxes on Less Elastic Goods? Evidence from German Municipalities,” *Regional Science and Urban Economics*, 2019, 75, 165–186.

- Boar, Corina and Danial Lashkari**, “Occupational Choice and the Intergenerational Mobility of Welfare,” *NBER Working Paper*, 2021, 29381.
- Bolt, Uta, Eric French, Jamie Hentall Maccuish, and Cormac O’Dea**, “The Intergenerational Elasticity of Earnings: Exploring the Mechanisms,” *IFS Working Paper*, 2021, 21/07.
- Boneva, Teodora and Christopher Rauh**, “Socio-Economic Gaps in University Enrollment: The Role of Perceived Pecuniary and Non-Pecuniary Returns,” *CESifo Working Paper*, 2017, 6756.
- Bonin, Holger, Thomas Dohmen, Armin Falk, David Huffman, and Uwe Sunde**, “Cross-sectional Earnings Risk and Occupational Sorting: The Role of Risk Attitudes,” *Labour Economics*, 2007, 14 (6), 926–937.
- Borusyak, Kirill, Xavier Jaravel, and Jann Spiess**, “Revisiting Event Study Designs: Robust and Efficient Estimation,” *The Review of Economic Studies*, forthcoming.
- Boserup, Simon H., Wojciech Kopczuk, and Claus T. Kreiner**, “The Role of Bequests in Shaping Wealth Inequality: Evidence from Danish Wealth Records,” *American Economic Review*, 2016, 106 (5), 656–61.
- Bowles, Samuel**, “Endogenous Preferences: The Cultural Consequences of Markets and Other Economic Institutions,” *Journal of Economic Literature*, 1998, 36 (1), 75–111.
- Bratberg, Espen, Jonathan Davis, Bhashkar Mazumder, Martin Nybom, Daniel Schnittlein, and Kjell Vaage**, “A Comparison of Intergenerational Mobility Curves in Germany, Norway, Sweden, and the US,” *Scandinavian Journal of Economics*, 2017, 119 (1), 72–101.
- Brunello, Giorgio and Lorenzo Rocco**, “The Effect of Immigration on the School Performance of Natives: Cross Country Evidence using PISA Test Scores,” *Economics of Education Review*, 2013, 32, 234–246.
- Brunori, Paolo, Paul Hufe, and Daniel Mahler**, “The Roots of Inequality: Estimating Inequality of Opportunity from Regression Trees and Forests,” *The Scandinavian Journal of Economics*, 2023, 125 (4), 900–932.
- Burbano, Vanessa, Nicolas Padilla, and Stephan Meier**, “Gender Differences in Preferences for Meaning at Work,” *IZA Discussion Paper*, 2020, 13053.
- Büchler, Theresa**, “Schulstruktur und Bildungsungleichheit: Die Bedeutung von bundesland-spezifischen Unterschieden beim Übergang in die Sekundarstufe I für den Bildungserfolg,” *KZfSS Kölner Zeitschrift für Soziologie und Sozialpsychologie*, 2016, 68 (1), 53–87.

- Bütikofer, Aline and Giovanni Peri**, “How Cognitive Ability and Personality Traits Affect Geographic Mobility,” *Journal of Labor Economics*, 2021, 39 (2), 559–595.
- Cantoni, Davide, Felix Hagemeister, and Mark Westcott**, “Persistence and Activation of Right-Wing Political Ideology,” *CRC Discussion Paper*, 2019, 143.
- Cappelen, Alexander W., Astri Drange Hole, Erik Ø. Sørensen, and Bertil Tungodden**, “The Pluralism of Fairness Ideals: An Experimental Approach,” *American Economic Review*, 2007, 97 (3), 818–827.
- Charles, Kerwin K. and Erik Hurst**, “The Correlation of Wealth Across Generations,” *Journal of Political Economy*, 2003, 111 (6), 1155–1182.
- Chen, Zhao, Xian Jiang, Zhikuo Liu, Juan Carlos Suárez Serrato, and Daniel Xu**, “Tax Policy and Lumpy Investment Behavior: Evidence from China’s VAT Reform,” *The Review of Economic Studies*, 2022, 90 (2), 634–674.
- Chetty, Raj and Nathaniel Hendren**, “The Impacts of Neighborhoods on Intergenerational Mobility I: Childhood Exposure Effects,” *The Quarterly Journal of Economics*, 2018, 133 (3), 1107–1162.
- , —, **Patrick Kline, and Emmanuel Saez**, “Where is the Land of Opportunity? The Geography of Intergenerational Mobility in the United States,” *The Quarterly Journal of Economics*, 2014, 129 (4), 1553–1623.
- , —, —, —, **and Nicholas Turner**, “Is the United States still a Land of Opportunity? Recent Trends in Intergenerational Mobility,” *American Economic Review*, 2014, 104 (5), 141–47.
- Chuard, Patrick and Veronica Grassi**, “Switzer-Land of Opportunity: Intergenerational Income Mobility in the Land of Vocational Education,” *University of St. Gallen Discussion Paper*, 2020, 2020-11.
- Cloyne, James, Ezgi Kurt, and Paolo Surico**, “Who Gains from Corporate Tax Cuts?,” *NBER Working Paper*, 2023, 31278.
- Cohen, Gerald A.**, “On the Currency of Egalitarian Justice,” *Ethics*, 1989, 99 (4), 906–944.
- Corak, Miles**, “The Canadian Geography of Intergenerational Income Mobility,” *The Economic Journal*, 2020, 130 (631), 2134–2174.
- **and Patrizio Piraino**, “The Intergenerational Transmission of Employers,” *Journal of Labor Economics*, 2011, 29 (1), 37–68.

- Council of Economic Advisers**, *Corporate Tax Reform and Wages: Theory and Evidence*, Washington, DC: Executive Office of the President of the United States, 2017.
- Cunha, Flavio, James J. Heckman, Lance Lochner, and Dimitriy V. Masterov**, “Interpreting the Evidence on Life Cycle Skill Formation,” in Eric Hanushek and Finis Welch, eds., *Handbook of the Economics of Education*, Vol. 1, Elsevier, 2006, chapter 12, pp. 697–812.
- Curtis, E. Mark, Daniel G. Garrett, Eric C. Ohn, Kevin A. Roberts, and Juan Carlos Suárez Serrato**, “Capital Investment and Labor Demand,” *NBER Working Paper*, 2021, 29485.
- Dahrendorf, Ralf**, *Bildung ist Bürgerrecht. Plädoyer für eine aktive Bildungspolitik.*, Hamburg: Nannen-Verlag, 1966.
- Damm, Anna P. and Christian Dustmann**, “Does Growing Up in a High Crime Neighborhood Affect Youth Criminal Behavior?,” *American Economic Review*, 2014, 104 (6), 1806–1832.
- Davis, Jonathan and Bhashkar Mazumder**, “The Decline in Intergenerational Mobility after 1980,” *Federal Reserve Bank of Chicago Working Paper*, 2017.
- de Chaisemartin, Clément and Xavier D’Haultfoeulle**, “Two-Way Fixed Effects and Differences-in-Differences with Heterogeneous Treatment Effects: A Survey,” *The Econometrics Journal*, 06 2022, 26 (3), C1–C30.
- De Philippis, Marta**, “STEM Graduates and Secondary School Curriculum: Does Early Exposure to Science Matter?,” *Journal of Human Resources*, 2021, p. 1219.
- Demirel, Ufuk D.**, “The Short-term Effects of Tax Changes: The Role of State Dependence,” *Journal of Monetary Economics*, 2021, 117, 918–934.
- Derenoncourt, Ellora and Claire Montialoux**, “Minimum Wages and Racial Inequality,” *The Quarterly Journal of Economics*, 2021, 136, 169–228.
- Dettling, Lisa J, Sarena Goodman, and Jonathan Smith**, “Every Little Bit Counts: The Impact of High-Speed Internet on the Transition to College,” *Review of Economics and Statistics*, 2018, 100 (2), 260–273.
- Deutscher, Nathan and Bhashkar Mazumder**, “Intergenerational Mobility across Australia and the Stability of Regional Estimates,” *Labour Economics*, 2020, 66, 101861.
- Dobbins, Laura and Martin Jacob**, “Do Corporate Tax Cuts Increase Investments?,” *Accounting and Business Research*, 2016, 46 (7), 731–759.

- Dodin, Majed, Sebastian Findeisen, Lukas Henkel, Dominik Sachs, and Paul Schüle**, “Social Mobility in Germany,” *Journal of Public Economics*, 2024, 232, 105074.
- Doepke, Matthias and Fabrizio Zilibotti**, “Parenting With Style: Altruism and Paternalism in Intergenerational Preference Transmission,” *Econometrica*, 2017, 85 (5), 1331–1371.
- Dohmen, Thomas, Armin Falk, David Huffman, and Uwe Sunde**, “Homo Reciprocans: Survey Evidence on Behavioural Outcomes,” *The Economic Journal*, 2009, 119 (536), 592–612.
- , —, —, and —, “The Intergenerational Transmission of Risk and Trust Attitudes,” *The Review of Economic Studies*, 2012, 79 (2), 645–677.
- Domar, Evsey D. and Richard A. Musgrave**, “Proportional Income Taxation and Risk-Taking,” *The Quarterly Journal of Economics*, 1944, 58 (3), 388–422.
- Dube, Arindrajit, Daniele Girardi, Òscar Jordà, and Alan M. Taylor**, “A Local Projections Approach to Difference-in-Differences Event Studies,” *NBER Working Paper*, 2023, 31184.
- Dustmann, Christian**, “Parental Background, Secondary School Track Choice, and Wages,” *Oxford Economic Papers*, 2004, 56 (2), 209–230.
- , **Patrick Puhani, and Uta Schönberg**, “The Long-Term Effects of Early Track Choice,” *The Economic Journal*, 2017, 127, 1348–1380.
- Dworkin, Ronald**, “What is Equality? Part 2: Equality of Resources,” *Philosophy & Public Affairs*, 1981, 10 (4), 283–345.
- Eberharter, Veronika V.**, “The Intergenerational Dynamics of Social Inequality – Empirical Evidence from Europe and the United States,” *SOEPpapers*, 2013, 588-2013.
- Eisenhauer, Philipp and Friedhelm Pfeiffer**, “Assessing Intergenerational Earnings Persistence among German Workers,” *Journal of Labour Market Research*, 2008, 2/3, 119–137.
- Elinder, Mikael, Oscar Erixson, and Daniel Waldenström**, “Inheritance and Wealth Inequality: Evidence from Population Registers,” *Journal of Public Economics*, 2018, 165, 17–30.
- Erosa, Andrés, Luisa Fuster, Gueorgui Kambourov, and Richard Rogerson**, “Hours, Occupations, and Gender Differences in Labor Market Outcomes,” *American Economic Journal: Macroeconomics*, 2022, 14 (3), 543–590.

- FDZ der Statistischen Ämter des Bundes und der Länder**, *Mikrozensus*, DOI: 10.21242/12211.1997.00.00.1.1.0 to 10.21242/12211.2018.00.00.1.1.3, own calculations, 1997–2018.
- Finkelstein, Amy, Matthew Gentzkow, and Heidi Williams**, “Sources of Geographic Variation in Health Care: Evidence from Patient Migration,” *The Quarterly Journal of Economics*, 2016, 131 (4), 1681–1726.
- Fong, Christina**, “Social Preferences, Self-interest, and the Demand for Redistribution,” *Journal of Public Economics*, 2001, 82 (2), 225–246.
- Foremny, Dirk and Nadine Riedel**, “Business Taxes and the Electoral Cycle,” *Journal of Public Economics*, 2014, 115, 48–61.
- Fouarge, Didier, Ben Kriechel, and Thomas Dohmen**, “Occupational Sorting of School Graduates: The Role of Economic Preferences,” *Journal of Economic Behavior & Organization*, 2014, 106, 335–351.
- Frick, Joachim R. and Markus M. Grabka**, “Missing Income Data in the German SOEP: Incidence, Imputation and its Impact on the Income Distribution,” *SOEP Survey Papers*, 2014, 225.
- Fuchs-Schündeln, Nicola, Dirk Krueger, and Mathias Sommer**, “Inequality Trends for Germany in the Last Two Decades: A Tale of Two Countries,” *Review of Economic Dynamics*, 2010, 13 (1), 103–132.
- Fuest, Clemens, Andreas Peichl, and Sebastian Siegloch**, “Do Higher Corporate Taxes Reduce Wages? Micro-Evidence from Germany,” *American Economic Review*, 2018, 108 (2), 393–418.
- Furlong, Andy and Andy Biggart**, “Framing ‘Choices’: A Longitudinal Study of Occupational Aspirations Among 13- to 16-Year-Olds,” *Journal of Education and Work*, 1999, 12 (1), 21–35.
- Furno, Francesco**, “The Macroeconomic Effects of Corporate Tax Reforms,” *Available at: <https://ssrn.com/abstract=4023666>*, 2022.
- Gallagher, Ryan, Robert Kaestner, and Joseph Persky**, “The Geography of Family Differences and Intergenerational Mobility,” *Journal of Economic Geography*, 2018, 19 (3), 589–618.
- Garrett, Daniel G., Eric Ohn, and Juan Carlos Suárez Serrato**, “Tax Policy and Local Labor Market Behavior,” *American Economic Review: Insights*, 2020, 2 (1), 83–100.

- Ghassibe, Mishel and Francesco Zanetti**, “State Dependence of Fiscal Multipliers: The Source of Fluctuations Matters,” *Journal of Monetary Economics*, 2022, 132, 1–23.
- Giroud, Xavier and Joshua Rauh**, “State Taxation and the Reallocation of Business Activity: Evidence from Establishment-Level Data,” *Journal of Political Economy*, 2019, 127 (3), 1262–1316.
- Gneezy, Uri, John A. List, Jeffrey A. Livingston, Xiangdong Qin, Sally Sadoff, and Yang Xu**, “Measuring Success in Education: The Role of Effort on the Test Itself,” *American Economic Review: Insights*, 2019, 1 (3), 291–308.
- Goebel, Jan, Markus M. Grabka, Stefan Liebig, Martin Kroh, David Richter, Carsten Schröder, and Jürgen Schupp**, “The German Socio-Economic Panel (SOEP),” *Jahrbücher für Nationalökonomie und Statistik*, 2019, 239 (2), 345–360.
- Golsteyn, Bart H., Hans Grönqvist, and Lena Lindahl**, “Adolescent Time Preferences Predict Lifetime Outcomes,” *The Economic Journal*, 2014, 124 (580), F739–F761.
- Gouveia, Miguel and Robert P. Strauss**, “Effective federal individual income tax functions: an exploratory empirical analysis,” *National Tax Journal*, 1994, 47 (2), 317–338.
- Gregg, Paul, Lindsey Macmillan, and Claudia Vittori**, “Moving Towards Estimating Sons’ Lifetime Intergenerational Economic Mobility in the UK,” *Oxford Bulletin of Economics and Statistics*, 2017, 79 (1), 79–100.
- Guceri, Irem and Maciej Albinowski**, “Investment Responses to Tax Policy under Uncertainty,” *Journal of Financial Economics*, 2021, 141 (3), 1147–1170.
- Guryan, Jonathan, Erik Hurst, and Melissa Kearney**, “Parental Education and Parental Time with Children,” *Journal of Economic Perspectives*, 2008, 22 (3), 23–46.
- Gysbers, Norman C.**, “Comprehensive School Guidance Programs in the United States: A Career Profile,” *International Journal for Educational and Vocational Guidance*, 2005, 5 (2), 203–215.
- Gärtner, Karla**, “Differentielle Sterblichkeit. Ergebnisse des Lebenserwartungssurvey des BiB,” *Zeitschrift für Bevölkerungswissenschaft*, 2002, 27 (2), 185–211.
- Hadjar, Andreas and Rolf Becker**, “Bildungsexpansion – Erwartete und Unerwartete Folgen,” in “Die Bildungsexpansion,” VS Verlag für Sozialwissenschaften., 2006, pp. 11–24.

Haider, Steven and Gary Solon, “Life-Cycle Variation in the Association Between Current and Lifetime Earnings,” *American Economic Review*, 2006, 96 (4), 1308–1320.

Hall, Robert E. and Dale W. Jorgenson, “Tax Policy and Investment Behavior,” *American Economic Review*, 1967, 57 (3), 391–414.

Hanushek, Eric, Lavinia Kinne, Philipp Lergetporer, and Ludger Woessmann, “Patience, Risk-Taking, and Human Capital Investment Across Countries,” *The Economic Journal*, 2022, 132 (646), 2290–2307.

Harju, Jarkko, Aliisa Koivisto, and Tuomas Matikka, “The Effects of Corporate Taxes on Small Firms,” *Journal of Public Economics*, 2022, 212, 104704.

Hartley, Robert Paul, Carlos Lamarche, and James P. Ziliak, “Welfare Reform and the Intergenerational Transmission of Dependence,” *Journal of Political Economy*, 2022, 130 (3), 523–565.

Hassett, Kevin A. and R. Glenn Hubbard, “Tax Policy and Business Investment,” in Alan J. Auerbach and Martin Feldstein, eds., *Handbook of Public Economics*, Vol. 3, Elsevier, 2002, pp. 1293–1343.

Hausner, Karl Heinz, Doris Söhnlein, Brigitte Weber, and Enzo Weber, “Qualifikation und Arbeitsmarkt: Bessere Chancen mit mehr Bildung,” *IAB Kurzbericht*, 2015, 11/2015.

Hayo, Bernd and Sascha Mierzwa, “State-Dependent Effects of Tax Changes in Germany and the United Kingdom,” *MAGKS Papers on Economics*, 2021, 25.

Heckman, James J. and Rodrigo Pinto, “Econometric Mediation Analyses: Identifying the Sources of Treatment Effects from Experimentally Estimated Production Technologies with Unmeasured and Mismeasured Inputs,” *Econometric Reviews*, 2015, 34 (1-2), 6–31.

—, **Lance J. Lochner, and Petra E. Todd**, “Earnings Functions, Rates of Return and Treatment Effects: The Mincer Equation and Beyond,” in Eric Hanushek and Finis Welch, eds., *Handbook of the Economics of Education*, Vol. 1, Elsevier, 2006, chapter 7, pp. 307–458.

Helbig, Marcel and Rita Nikolai, *Die Unvergleichbaren: Der Wandel der Schulsysteme in den deutschen Bundesländern seit 1949*, Verlag Julius Klinkhardt, 2015.

Helsø, Anne-Line, “Intergenerational Income Mobility in Denmark and the United States,” *Scandinavian Journal of Economics*, 2021, 123 (2), 508–531.

- Hendren, Nathaniel and Ben Sprung-Keyser**, “A Unified Welfare Analysis of Government Policies,” *The Quarterly Journal of Economics*, 2020, 135 (3), 1209–1318.
- Hilger, Nathaniel G.**, “The Great Escape: Intergenerational Mobility in the United States Since 1940,” *NBER Working Paper Series*, 2015, 21217.
- Hill, Andrew J.**, “State Affirmative Action Bans and STEM Degree Completions,” *Economics of Education Review*, 2017, 57, 31–40.
- House, Christopher L. and Matthew D. Shapiro**, “Temporary Investment Tax Incentives: Theory with Evidence from Bonus Depreciation,” *American Economic Review*, 2008, 98 (3), 737–768.
- Hufe, Paul, Martyna Kobus, Andreas Peichl, and Paul Schüle**, “Multidimensional Equality of Opportunity in the United States,” *CESifo Working Paper*, 2022, 9630.
- , **Ravi Kanbur, and Andreas Peichl**, “Measuring Unfair Inequality: Reconciling Equality of Opportunity and Freedom from Poverty,” *The Review of Economic Studies*, 2022, 89 (6), 3345–3380.
- Isphording, Ingo, Andreas Lichter, Max Löffler, Thu-Van Nguyen, Felix Pöge, and Sebastian Siegloch**, “Profit Taxation, R&D Spending, and Innovation,” *CEPR Discussion Paper*, 2021, 16702.
- Ivanov, Ivan, Luke Pettit, and Toni M. Whited**, “Taxes Depress Corporate Borrowing: Evidence from Private Firms,” Available at: <https://papers.ssrn.com/abstract=3694869>, 2022.
- IVS, Ifo Investment Survey Industry 1964-2019**, LMU-ifo Economics & Business Data Center, Munich, doi: 10.7805/ebdc-ivs-ind-2019, 2019.
- Jácome, Elisa, Ilyana Kuziemko, and Suresh Naidu**, “Mobility for All: Representative Intergenerational Mobility Estimates over the 20th Century,” *NBER Working Paper*, 2021, 29289.
- Jensen, Peter and Astrid Würtz Rasmussen**, “The Effect of Immigrant Concentration in Schools on Native and Immigrant Children’s Reading and Math Skills,” *Economics of Education Review*, 2011, 30 (6), 1503–1515.
- Jones, Paul M., Eric Olson, and Mark E. Wohar**, “Asymmetric Tax Multipliers,” *Journal of Macroeconomics*, 2015, 43, 38–48.

- Jähnen, Stefanie and Marcel Helbig**, “Der Einfluss schulrechtlicher Reformen auf Bildungungleichheiten zwischen den deutschen Bundesländern,” *KZfSS Kölner Zeitschrift für Soziologie und Sozialpsychologie*, 2015, 67 (3), 539–571.
- Keane, Michael P. and Kenneth I. Wolpin**, “The Role of Labor and Marriage Markets, Preference Heterogeneity, and the Welfare System in the Life Cycle Decisions of Black, Hispanic, and White Women,” *International Economic Review*, 2010, 51 (3), 851–892.
- Kindermann, Fabian, Lukas Mayr, and Dominik Sachs**, “Inheritance Taxation and Wealth Effects on the Labor Supply of Heirs,” *Journal of Public Economics*, 2020, 191, 104127.
- Klein, Markus, Katherin Barg, and Michael Kühhirt**, “Inequality of Educational Opportunity in East and West Germany: Convergence or Continued Differences?,” *Sociological Science*, 2019, 6, 1–26.
- Kline, Patrick and Enrico Moretti**, “People, Places, and Public Policy: Some Simple Welfare Economics of Local Economic Development Programs,” *Annual Review of Economics*, 2014, 6 (1), 629–662.
- Kobus, Martyna, Marek Kapera, and Vito Peragine**, “Measuring Multidimensional Inequality of Opportunity,” *ECINEQ Working Paper Series*, 2020, 528.
- Kolm, Serge C.**, “The Optimal Production of Social Justice,” in J. Margolis and H. Guitton, eds., *Public Economics: An Analysis of Public Production and Consumption and their Relations to the Private Sectors*, Macmillan, 1969, pp. 145–200.
- Kopczuk, Wojciech**, “Incentive Effects of Inheritances and Optimal Estate Taxation,” *American Economic Review*, 2013, 103 (3), 472–477.
- Kosse, Fabian and Friedhelm Pfeiffer**, “Impatience Among Preschool Children and Their Mothers,” *Economics Letters*, 2012, 115 (3), 493–495.
- and —, “Quasi-Hyperbolic Time Preferences and Their Intergenerational Transmission,” *Applied Economics Letters*, 2013, 20 (10), 983–986.
- , **Thomas Deckers, Pia Pinger, Hannah Schildberg-Hörisch, and Armin Falk**, “The Formation of Prosociality: Causal Evidence on the Role of Social Environment,” *Journal of Political Economy*, 2020, 128 (2).
- Kourtellos, Andros, Christa Marr, and Chih Ming Tan**, “Robust Determinants of Intergenerational Mobility in the Land of Opportunity,” *European Economic Review*, 2016, 81, 132–147.

- Kramarz, Francis and Oskar N. Skans**, “When Strong Ties are Strong: Networks and Youth Labour Market Entry,” *The Review of Economic Studies*, 2014, 81 (3), 1164–1200.
- Kuhn, Moritz and José-Víctor Ríos-Rull**, “2013 Update on the U.S. Earnings, Income, and Wealth Distributional Facts: A View from Macroeconomics,” *Quarterly Review*, 2016, 37, 1–73.
- , **Moritz Schularick, and Ulrike I. Steins**, “Income and Wealth Inequality in America, 1949–2016,” *Journal of Political Economy*, 2020, 128 (9), 3469–3519.
- Landersø, Rasmus and James J. Heckman**, “The Scandinavian Fantasy: The Sources of Intergenerational Mobility in Denmark and the US,” *Scandinavian Journal of Economics*, 2017, 119 (1), 178–230.
- Langenmayr, Dominika and Martin Simmler**, “Firm Mobility and Jurisdictions’ Tax Rate Choices: Evidence from Immobile Firm Entry,” *Journal of Public Economics*, 2021, 204, 104530.
- **and Rebecca Lester**, “Taxation and Corporate Risk-Taking,” *Accounting Review*, 2017, 93 (3), 237–266.
- Lee, Chul-In and Gary Solon**, “Trends in Intergenerational Income Mobility,” *Review of Economics and Statistics*, 2009, 91 (4), 766–772.
- Lerche, Adrian**, “Investment Tax Credits and the Response of Firms,” *IZA Discussion Paper No. 15668*, 2022.
- Link, Sebastian, Andreas Peichl, Christopher Roth, and Johannes Wohlfart**, “Information Frictions among Firms and Households,” *Journal of Monetary Economics*, 2023, 135, 99–115.
- , **Manuel Menkhoff, Andreas Peichl, and Paul Schüle**, “Downward Revision of Investment Decisions after Corporate Tax Hikes,” *American Economic Journal: Policy*, forthcoming.
- Ljungqvist, Alexander and Michael Smolyansky**, “To Cut or Not to Cut? On the Impact of Corporate Taxes on Employment and Income,” *NBER Working Paper*, 2018, 20753.
- , **Liandong Zhang, and Luo Zuo**, “Sharing Risk with the Government: How Taxes Affect Corporate Risk Taking,” *Journal of Accounting Research*, 2017, 55 (3), 669–707.
- Lochner, Lance**, “Nonproduction Benefits of Education: Crime, Health, and Good Citizenship,” *Handbook of the Economics of Education*, 2011, 4, 183–282.

- Maffini, Giorgia, Jing Xing, and Michael P. Devereux**, “The Impact of Investment Incentives: Evidence from UK Corporation Tax Returns,” *American Economic Journal: Economic Policy*, 2019, 11 (3), 361–389.
- Manoli, Day and Nicholas Turner**, “Cash-on-Hand and College Enrollment: Evidence from Population Tax Data and the Earned Income Tax Credit,” *American Economic Journal: Economic Policy*, 2018, 10 (2), 242–271.
- Mas, Alexandre and Amanda Pallais**, “Valuing Alternative Work Arrangements,” *American Economic Review*, 2017, 107 (12), 3722–3759.
- Mazumder, Bhashkar**, “Intergenerational Mobility in the United States: What We Have Learned from the PSID,” *The Annals of the American Academy of Political and Social Science*, 2018, 680 (1), 213–234.
- Mertens, Karel and Morten O. Ravn**, “The Dynamic Effects of Personal and Corporate Income Tax Changes in the United States,” *American Economic Review*, 2013, 103 (4), 1212–1247.
- Meyer, Bruce D., Wallace K.C. Mok, and James X. Sullivan**, “The Under-Reporting of Transfers in Household Surveys: Its Nature and Consequences,” *mimeo*, 2015.
- Mogstad, Magne, Joseph P. Romano, Azeem Shaikh, and Daniel Wilhelm**, “Inference for Ranks with Applications to Mobility across Neighborhoods and Academic Achievements across Countries,” *The Review of Economic Studies*, 2024, 91, 476–518.
- Morris, Carl N.**, “Parametric Empirical Bayes Inference: Theory and Applications,” *Journal of the American Statistical Association*, 1983, 78 (381), 47–55.
- Müller, Maximilian W.**, “Intergenerational Transmission of Education: Internalized Aspirations versus Parent Pressure,” *mimeo*, 2021.
- Neumark, David and Helen Simpson**, “Place-Based Policies,” in “Handbook of Regional and Urban Economics,” Vol. 5, Elsevier, 2015, pp. 1197–1287.
- Nicodème, Gaëtan**, “Computing Effective Corporate Tax Rates: Comparisons and Results,” *MPRA Working Paper*, 2001, 3808.
- Niehues, Judith and Andreas Peichl**, “Upper Bounds of Inequality of Opportunity: Theory and Evidence for Germany and the US,” *Social Choice and Welfare*, 2014, 43 (1), 73–99.
- Nyblom, Martin and Jan Stuhler**, “Heterogeneous Income Profiles and Lifecycle Bias in Intergenerational Mobility Estimation,” *Journal of Human Resources*, 2016, 51 (1), 239–268.

- Ohrn, Eric**, “The Effect of Corporate Taxation on Investment and Financial Policy: Evidence from the DPAD,” *American Economic Journal: Economic Policy*, 2018, 10 (2), 272–301.
- Oreopoulos, Philip and Kjell G. Salvanes**, “Priceless: The Nonpecuniary Benefits of Schooling,” *Journal of Economic Perspectives*, 2011, 25 (1), 159–84.
- Pfeffer, Fabian T. and Alexandra Killewald**, “Generations of Advantage. Multigenerational Correlations in Family Wealth,” *Social Forces*, 2018, 96 (4), 1411–1442.
- , **Robert F. Schoeni, Arthur Kennickell, and Patricia Andreski**, “Measuring Wealth and Wealth Inequality: Comparing Two US Surveys,” *Journal of Economic and Social Measurement*, 2016, 41 (2), 103–120.
- Picht, Georg**, *Die deutsche Bildungskatastrophe: Analyse und Dokumentation*, Walter Verlag, 1964.
- Piketty, Thomas and Gabriel Zucman**, “Wealth and Inheritance in the Long Run,” in Anthony B. Atkinson and François Bourguignon, eds., *Handbook of Income Distribution*, Vol. 2 of *Handbook of Income Distribution*, Amsterdam: Elsevier, 2015, chapter 15, pp. 1303–1368.
- Pistolesi, Nicolas**, “Inequality of Opportunity in the Land of Opportunities, 1968–2001,” *Journal of Economic Inequality*, 2009, 7 (4), 411.
- Pokropek, Artur, Gary N. Marks, and Francesca Borgonovi**, “How Much do Students’ Scores in PISA Reflect General Intelligence and How Much do They Reflect Specific Abilities?,” *Journal of Educational Psychology*, 2022, 114 (5), 1121.
- Ramey, Valerie A. and Sarah Zubairy**, “Government Spending Multipliers in Good Times and in Bad: Evidence from US Historical Data,” *Journal of Political Economy*, 2018, 126 (2), 850–901.
- Rauh, Christopher**, “Voting, Education, and the Great Gatsby Curve,” *Journal of Public Economics*, 2017, 146, 1–14.
- Rawls, John**, *A Theory of Justice*, Cambridge: The Belknap Press of Harvard University Press, 1971.
- Riedel, Nadine and Martin Simmler**, “Large and Influential: Firm Size and Governments’ Corporate Tax Rate Choice,” *Canadian Journal of Economics/Revue Canadienne d’Économie*, 2021, 54 (2), 812–839.

- Rindermann, Heiner**, “The G-factor of International Cognitive Ability Comparisons: The Homogeneity of Results in PISA, TIMSS, PIRLS and IQ-tests Across Nations,” *European Journal of Personality*, 2007, 21 (5), 667–706.
- Riphahn, Regina and Guido Heineck**, “Intergenerational Transmission of Educational Attainment in Germany – The Last Five Decades,” *Journal of Economics and Statistics (Jahrbücher für Nationalökonomie und Statistik)*, 2009, 229 (1), 36–60.
- **and Parvati Trübswetter**, “The Intergenerational Transmission of Educational Attainment in East and West Germany,” *Applied Economics*, 2013, 45 (22), 3183–3196.
- Robinson, Sarah and Alisa Tazhitdinova**, “What Drives Tax Policy? Political, Institutional and Economic Determinants of State Tax Policy,” *NBER Working Paper*, 2022, 31268.
- Roemer, John E.**, *Equality of Opportunity*, Cambridge: Harvard University Press, 1998.
- **and Alain Trannoy**, “Equality of Opportunity: Theory and Measurement,” *Journal of Economic Literature*, 2016, 54 (4), 1288–1332.
- Romer, Christina D. and David H. Romer**, “A New Measure of Monetary Shocks: Derivation and Implications,” *American Economic Review*, 2004, 94 (4), 1055–1084.
- **and —**, “The Macroeconomic Effects of Tax Changes: Estimates Based on a New Measure of Fiscal Shocks,” *American Economic Review*, 2010, 100 (3), 763–801.
- Rothbaum, Jonathan**, “Sorting and Geographic Variation in Intergenerational Mobility,” *Unpublished Manuscript*, 2016.
- Rothstein, Jesse**, “Inequality of Educational Opportunity? Schools as Mediators of the Intergenerational Transmission of Income,” *Journal of Labor Economics*, 2019, 37 (1), 85–123.
- Royston, Patrick and Ian R. White**, “Multiple Imputation by Chained Equations (MICE): Implementation in Stata,” *Journal of Statistical Software*, 2011, 45, 1–20.
- San, Shmuel**, “Who Works Where and Why? Parental Networks and the Labor Market,” *Available at SSRN: <https://ssrn.com/abstract=3726993>*, 2020.
- Sanchis-Guarner, Rosa, José Montalbán, and Felix Weinhardt**, “Home Broadband and Human Capital Formation,” *CESifo Working Paper*, 2021, 8846.
- Sauer, Stefan and Klaus Wohlrabe**, “ifo Handbuch der Konjunkturumfragen,” *ifo Beiträge zur Wirtschaftsforschung*, 2020, 88.

Schmillen, Achim and Heiko Stüber, “Lebensverdienste nach Qualifikation - Bildung Lohnt sich ein Leben Lang,” *IAB Kurzbericht*, 2014, 1/2014.

Schnitzlein, Daniel, “How Important is the Family? Evidence from Sibling Correlations in Permanent Earnings in the USA, Germany, and Denmark,” *Journal of Population Economics*, 2014, 27, 69–86.

—, “A New Look at Intergenerational Mobility in Germany Compared to the US,” *Review of Income and Wealth*, 2016, 62 (4), 650–667.

Schüle, Paul, “Career Preferences and Socio-Economic Background,” *ifo Working Paper*, 2023, 395.

Sen, Amartya K., *On Economic Inequality*, Oxford: Clarendon Press, 1973.

Seymour, Elaine and Nancy M. Hewitt, *Talking About Leaving*, Vol. 34, Westview Press, Boulder, CO, 1997.

Shorrocks, Anthony F., “Decomposition Procedures for Distributional Analysis: A Unified Framework Based on the Shapley Value,” *Journal of Economic Inequality*, 2013, 11 (1), 99.

Solis, Alex, “Credit Access and College Enrollment,” *Journal of Political Economy*, 2017, 125 (2), 562–622.

Solon, Gary, “Intergenerational Income Mobility in the United States,” *American Economic Review*, 1992, pp. 393–408.

Statistisches Bundesamt, *Mikrozensus 2018. Qualitätsbericht 2018*.

—, *Konsumausgaben von Familien für Kinder: Berechnungen auf der Grundlage der Einkommens- und Verbrauchsstichprobe 2018* 2021.

Strobl, Carolin, Anne-Laure Boulesteix, Thomas Kneib, Thomas Augustin, and Achim Zeileis, “Conditional Variable Importance for Random Forests,” *BMC Bioinformatics*, 2008, 9 (1), 307.

Sun, Liyang and Sarah Abraham, “Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects,” *Journal of Econometrics*, 2021, 225 (2), 175–199.

Sunde, Uwe, Thomas Dohmen, Benjamin Enke, Armin Falk, David Huffman, and Gerrit Meyerheim, “Patience and Comparative Development,” *The Review of Economic Studies*, 2022, 89 (5), 2806–2840.

- Suárez Serrato, Juan Carlos and Owen Zidar**, “Who Benefits from State Corporate Tax Cuts? A Local Labor Markets Approach with Heterogeneous Firms,” *American Economic Review*, 2016, 106 (9), 2582–2624.
- Valet, Peter, Carsten Sauer, and Jochem Tolsma**, “Preferences for Work Arrangements: A Discrete Choice Experiment,” *PLOS ONE*, 2021, 16 (7), e0254483.
- Ward, Zachary**, “Intergenerational Mobility in American History: Accounting for Race and Measurement Error,” *NBER Working Paper*, 2021, 29256.
- White, Ian R., Patrick Royston, and Angela M. Wood**, “Multiple Imputation Using Chained Equations: Issues and Guidance for Practice,” *Statistics in Medicine*, 2011, 30 (4), 377–399.
- Winberry, Thomas**, “Lumpy Investment, Business Cycles, and Stimulus Policy,” *American Economic Review*, 2021, 111 (1), 364–396.
- Wiswall, Matthew and Basit Zafar**, “Preference for the Workplace, Investment in Human Capital, and Gender,” *The Quarterly Journal of Economics*, 2018, 133 (1), 457–507.
- Wolff, Edward N.**, “Inheritances and Wealth Inequality, 1989-1998,” *American Economic Review*, 2002, 92 (2), 260–264.
- , “Household Wealth Trends in the United States, 1962 to 2016: Has Middle Class Wealth Recovered?,” *NBER Working Paper*, 2017, 24085.
- Xu, Qiping and Eric Zwick**, “Tax Policy and Abnormal Investment Behavior,” *NBER Working Paper*, 2022, 27363.
- Young, H. Peyton**, “Progressive Taxation and Equal Sacrifice,” *American Economic Review*, 1990, 80 (1), 253–266.
- Zumbuehl, Maria, Thomas Dohmen, and Gerard Pfann**, “Parental Involvement and the Intergenerational Transmission of Economic Preferences, Attitudes and Personality Traits,” *The Economic Journal*, 2020, 131 (638), 2642–2670.
- Zwick, Eric and James Mahon**, “Tax Policy and Heterogeneous Investment Behavior,” *American Economic Review*, 2017, 107 (1), 217–248.