Essays in Empirical Macroeconomics

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

vorgelegt von

Lukas Henkel

im Frühjahrssemester 2019

Abteilungssprecher:	Prof. Dr. Jochen Streb
Referent:	Prof. Klaus Adam, Ph.D.
Korreferent:	Prof. Dr. Sebastian Findeisen

Tag der Verteidigung: 17. Mai 2019

Eidesstattliche Erklärung

Hiermit erkläre ich, die vorliegende Dissertation selbstständig angefertigt und mich keiner anderen als der in ihr angegebenen Hilfsmittel bedient zu haben. Insbesondere sind sämtliche Zitate aus anderen Quellen als solche gekennzeichnet und mit Quellenangaben versehen.

Mannheim, 17. April 2019

Curriculum Vitae, Lukas Henkel

Since 09/2014	PhD Studies in Economics, Center for Doctoral Studies in Economics, University of Mannheim.
09/2015 – 05/2016	Visiting Graduate Student, University of California, Berkeley.
10/2012 – 08/2014	M.Sc. in Quantitative Economics, University of Tübingen.
09/2009 – 10/2012	B.Sc. in International Economics, University of Tübingen.
08/2011 – 01/2012	Visiting Student, University of Gothenburg, Sweden.
09/2000 – 06/2009	Abitur, Gymnasium Rutesheim

Acknowledgements

I would like to thank Klaus Adam for his dedication as a mentor and supervisor, for valuable feedback, for giving me the freedom to pursue my research interests, supporting me throughout the work on this dissertation and for encouraging me to venture into the realms of central banking.

My coauthor Sebastian Findeisen has also been an excellent supervisor. I thank him for his inspiration, valuable feedback, encouragement to present, for providing guidance and motivation and for valuable discussions about our joint work.

I am also thankful to my other coauthors, Dominik Sachs and Majed Dodin, for the constructive discussions on the conceptual framework and results presented in the second Chapter of this Dissertation. The work my dissertation is based on has furthermore benefited from numerous comments by the Mannheim Macro group, by participants of the Mannheim Macro Lunch and the Public Finance Seminar, Klaus Adam's Macro Colloquium and participants of various conferences.

I want to thank all the friends I have made while writing this dissertation, not only for a great number of interesting discussions about economics and life in general, but also for making my time spent in Mannheim an enjoyable one. I wanted to mention all of you by name, but the list is too long.

I would also like to thank Marion Lehnert, Sylvia Rosenkranz and Susanne Putz for their patient assistance in many administrative matters. I gratefully acknowledge the hospitality of the StaLa BW, and thank Kristin Nowak, Stefan Seitz and Isabella Brucker for always answering my requests very clearly and quickly.

My family has supported me not only throughout this dissertation, but my whole life. I am grateful that I can always rely on you. Lastly, I want to thank Berenike for supporting me throughout this dissertation and sharing life with me. You put a smile on my face every single day.

Contents

List of Figures

List of Tables

General Introduction

This Dissertation is composed of two Chapters which can be read independently. Each Chapter consists of a self-contained empirical study of an economic issue from the field of macroeconomics. While the questions addressed in the two Chapters stem from different fields of research - the first Chapter addresses a question from the field of monetary economics, whereas the second Chapter contains an empirical study on intergenerational mobility - they share a similar methodological approach.

The first Chapter of my Dissertation is concerned with the transmission of monetary policy decisions to real economic activity. I show that US manufacturing industries react vastly different to monetary policy shocks. I use the heterogeneity in output responses to monetary policy shocks across industries to test whether industries with a higher degree of price stickiness react more strongly to monetary policy shocks, as predicted by New Keynesian models of the monetary transmission mechanism.

The second Chapter of my Dissertation provides new evidence on the extent and the geography of intergenerational mobility in Germany. This chapter is coauthored with Majed Dodin, Sebastian Findeisen and Dominik Sachs. We show that intergenerational mobility, measured by the association of the educational attainment of children in secondary school and their parents income, differs substantially across regions within Germany.

The Appendix collects additional material for each Chapter, e.g. additional data descriptions and additional Tables and Figures. The Dissertation closes with a bibliography collecting the references that I used in writing my Dissertation.

While the two Chapters address rather different questions, they share a similar methodological approach. The empirical approach taken in both Chapters follows the same general steps. Both Chapters document heterogeneity in the relationship between two variables across different economic units. The first Chapter documents heterogeneity in the relationship between output and monetary policy shocks across industries, whereas the second

Chapter documents heterogeneity in the relationship between parent income and the educational attainment of children across local labor markets. Both Chapters use this heterogeneity across units to investigate the economic mechanisms linking the two variables. This common methodological approach unifies the two Chapters.

In the remainder of this Introduction, I present an overview over each Chapter.

Chapter 1: Sectoral Output Effects of Monetary Policy: Do Sticky Prices Matter?

This Chapter studies the role of sticky prices for the monetary transmission mechanism. New Keynesian models of the monetary transmission mechanism greatly emphasize the role of nominal rigidities, in particular price stickiness, for the monetary transmission mechanism (Galí (2015)). If sticky prices play an important role in the monetary transmission mechanism, the output of industries with a higher degree of price stickiness should react more strongly to monetary policy shocks than the output of industries with a lower degree of price stickiness. I investigate this testable implication of New Keynesian models using disaggrated US industrial production data. I estimate the dynamic output responses of 205 US manufacturing industries to identified monetary policy shocks in a panel VAR framework. I find substantial heterogeneity in the output responses of industries to monetary policy shocks. I show that an industry's output response to monetary policy shocks is systematically related to an industry's degree of price stickiness as measured by the average frequency of price adjustment. The size of the differential reaction between sticky price and flexible price industries is economically large and statistically significant. This result is robust to controlling for further industry-level variables, the identification strategy of monetary policy shocks and the estimation method. The results suggest that sticky prices play an important role in the transmission of monetary policy to real economic activity, consistent with New Keynesian macroeconomic models.

Chapter 2: The Geography of Intergenerational Mobility in Germany

This Chapter is co-authored with Majed Dodin, Sebastian Findeisen and Dominik Sachs. We present new empirical evidence on the level and the geography of intergenerational mobility in Germany using administrative microcensus data from the years 2009 to 2015. This data contains detailed information on the educational attainment of 268.000 children aged 16 to 22 and their parents socioeconomic status. In our main approach, we measure intergenerational mobility by the association between the educational attainment of a child and her parents income. Our main child outcome variable is a binary variable, indicating whether a child has obtained an A-Level degree ("Abitur") (or is in the process of obtaining this degree) or not. We focus on parental income to parsimoniously summarize the socioeconomic status of parents.

The relationship between a child's probability of obtaining an A-Level degree and the percentile rank of parents in the national income distribution is well approximated by a linear function at the national level. This allows us to summarize intergenerational mobility by the coefficients of a linear regression of the A-Level dummy of children on the parents percentile rank in the national income distribution within our sample. Based on the linear fit, we find that a 10 percentile point increase in parent income rank is associated with a 4.8 percentage point increase in the probability of obtaining an A-Level degree for children at the national level.

We calculate different measures of absolute and relative intergenerational mobility to compare intergenerational mobility across local labor markets within Germany. We document substantial variation in these measures across local labor markets. For example, the probability that a child from a family in the bottom quintile of the national income distribution obtains an A-Level degree is 27% in Dresden, but only 16% in Leipzig.

We show that regional differences in intergenerational mobility are not explained by differences in family characteristics across local labor markets within our sample. Lastly, we take a first step at describing which local labor market characteristics correlate with intergenerational mobility.

Chapter 1

Sectoral Output Effects of Monetary Policy: Do Sticky Prices Matter?

1.1 Introduction

Do sticky prices matter for the monetary transmission mechanism? Although there is growing consensus that prices are fixed in the short run at the micro level, the macroeconomic implications of this micro-level price stickiness are heavily debated. On the one hand, New Keynesian sticky price models postulate that firms face costs of price adjustment, which causes prices to be sticky in response to real and nominal shocks and that it is this feature that gives rise to monetary non-neutrality. On the other hand, the observed price rigidity does not necessarily imply that nominal shocks have real effects. For example Caplin and Spulber (1987) and Golosov and Lucas (2007) present theoretical models in which prices are sticky and money is neutral or the response to nominal shocks is only very limited.

The goal of this chapter is to empirically test whether sticky prices matter for the monetary transmission mechanism. If price stickiness plays an important role in the monetary transmission mechanism, the output of industries with a higher degree of price stickiness should react more strongly to monetary policy shocks than the output of industries with a lower degree of price stickiness. This Chapter investigates this testable implication of New Keynesian models using disaggrated US industrial production data.

The main finding of this Chapter is that there is a statistically significant association between an industry's output response to monetary policy shocks and the industry's degree of price stickiness, a finding that is new to the literature. The drop in output is estimated to be larger for sticky price industries (compared to flexible price industries) in response to a contractionary monetary policy shock. The size of this association is also economically significant. The cumulative drop in total industrial production is estimated at 0.38% two years after a one standard deviation contractionary monetary policy shock. For the most sticky price industries in the sample[1] the cumulative drop in output is estimated to be twice as large as the drop in total industrial production at the same horizon in response to the same monetary policy shock. At the same time, the cumulative drop in output is estimated at only half the size of the drop in total industrial production for the most flexible price industries in the sample[2] at the same horizon in response to the same monetary policy shock. Indeed, sticky price industries experience a larger change in output in response to

[1]Defined as industries at the 10th percentile of the in-sample distribution of the frequency of price adjustment.

[2]Defined as industries at the 90th percentile of the in-sample distribution of the frequency of price adjustment.

monetary policy shocks, consistent with a New Keynesian price stickiness channel in the monetary transmission mechanism.

This conclusion is reached in several steps. First, I estimate the output responses of 205 manufacturing industries to monetary policy shocks. Monetary policy shocks are identified using the financial market based identification of Barakchian and Crowe (2013). Industry-level output responses to the identified monetary policy shocks are estimated in a Panel VAR framework. There is substantial heterogeneity in the output responses of industries to monetary policy shocks. The cumulative drop in total industrial production is 0.38% two years after an unexpected one standard deviation increase in the policy measure. For some industries the drop in output is as large as 2%, whereas other industries increase output by around 0.9% in response to the same shock. In a next step, this heterogeneity in output responses to monetary policy shocks is used to investigate the transmission channels of monetary policy to real economic activity.

The goal of this chapter is to asses the role of price stickiness in the monetary transmission mechanism. Following the empirical literature on price rigidities (e.g. Bils and Klenow (2004) and Gorodnichenko and Weber (2016)), price stickiness is measured via the monthly frequency of price adjustment at the industry level. The frequency of price adjustment is calculated for the manufacturing industries in the sample using monthly PPI micro data over the period from 2005 to 2011.[3] The frequency of price adjustment differs greatly between industries. The most sticky price industries adjust only 4% of prices per month on average over the sample period. On the other hand, the most flexible price industries adjust 87% of prices per month. The median frequency of price adjustment of the industries in the sample is 19%.

In order to asses the association between an industry's output response to monetary policy shocks and the industries frequency of price adjustment, the cross-section of industry responses (at different horizons after the shock has happened) is regressed on the industry-level (log of the) frequency of price adjustment. More flexible prices are associated with a less strong drop in output in reaction to contractionary monetary policy shocks. A 10% increase in the frequency of price adjustment is associated with a 0.035 percentage point reduction in the cumulative output drop in response to the policy shock (2 years after the shock). Compared to the cumulative drop in total industrial production index, which is 0.38%, the size of this association is economically (and statistically) significant. A 10%

[3] I am grateful to Michael Weber for providing the industry-level frequency of price adjustment to me.

increase in the frequency of price adjustment is associated with a reduction in the output drop in response to a contractionary policy shock that is as large as 10% of the average drop in output. Hence an increase in the frequency of price adjustment is associated with a less strong reaction to monetary policy shocks, consistent with a New Keynesian price stickiness channel of monetary transmission.

In a next step several additional industry characteristics are added to the regression of the industry output responses to monetary policy shocks on the log of the frequency of price adjustment. Industries do not only differ along their frequency of price adjustment but also along other dimensions. Controlling for additional industry characteristics helps to disentangle the effect of price stickiness from other potentially confounding factors. The additional industry characteristics considered here include e.g. measures of external financial dependence or industry cyclicalty. When controlling for other industry characteristics, the association between the frequency of price adjustment and the strength of the reaction to monetary policy shocks becomes even larger, providing further support for a price stickiness channel in the monetary transmission mechanism.

The remainder of this introduction presents an overview over the related literature and a roadmap of the Chapter. The analysis in this Chapter is related to different strands of the literature on the monetary transmission mechanism. First, it is realted to a number of papers such as Peersman and Smets (2005), Ganley and Salmon (1997), Hayo and Uhlenbrock (2000) and Dedola and Lippi (2005) that examine the industry effects of monetary policy shocks. All these papers find considerable cross-industry heterogeneity in the output reaction to monetary policy shocks (identified from SVARs). While Ganley and Salmon (1997) and Hayo and Uhlenbrock (2000) focus on the UK and Germany, respectively, Dedola and Lippi (2005) and Peersman and Smets (2005) analyze cross-country differences in the monetary policy transmission mechanism as well. These papers find that the durability of industry output is a significant determinant of cross-industry heterogeneity in the output responses to monetary policy shocks, with durable goods producing industries reacting more strongly. In the US, Carlino and DeFina (1998) find substantial heterogeneity across regions in the response to monetary policy shocks. None of these papers has considered heterogeneity in price stickiness as explanatory variable in their analysis of cross-industry heterogeneity in responses to monetary policy shocks.

Gorodnichenko and Weber (2016) show that after monetary policy announcements the conditional volatility of stock market returns rises more for firms with stickier prices than for firms with more flexible prices. Their finding indicates that menu costs are an important factor causing nominal price rigidities, as presumed by New Keynesian models.

Furthermore, the findings of this Chapter speak to the literature on multi-sector New Keynesian models of the monetary transmission mechanism. A common finding in the literature on multi-sector New Keynesian models is that heterogeneity in the degree of price stickiness across sectors enhances monetary non-neutrality, e.g. Bouakez et al. (2013) or Pasten et al. (2018b). In these types of models, industries with a higher degree of price stickiness (ceteris paribus) react more strongly to monetary policy shocks (see e.g. Bouakez et al. (2013) and Ghassibe (2018)). This finding is not limited to models that use nominal price rigidities modeled as in Calvo (1983), but can also be extended to multi-sector menu cost models. For example in the calibrated multi-sector menu cost model of Nakamura and Steinsson (2010), the output of sectors with a lower frequency of price adjustment reacts more strongly to nominal demand shocks. The analysis at hand speaks to this literature by testing this testable implication of multi-sector models of the monetary transmission mechanism.

It should be noted that testing this prediction of multi-sector New Keynesian models speaks to New Keynesian models of the monetary transmission mechanism in general. New Keynesian models greatly emphasize the role of nominal rigidities, in particular price stickiness, for the monetary transmission mechanism and as the source of monetary non-neutrality policy (Galí (2015)). Hence the test whether price stickiness has a role in explaining differential industry reactions to monetary policy shocks is an important test of the New Keynesian paradigm.

The rest of the Chapter is organized as follows. Section 1.2 presents the data used in the analysis. Section 1.3 presents the response of (total) industrial production to monetary policy shocks. Section 1.4 presents the output responses of 205 manufacturing industries to monetary policy shocks. Section 1.5 presents the relationship between an industries output response and the frequency of price adjustment. In Section 1.6 presents this relation when controlling for additional industry characteristics. Section 1.7 presents various robustness checks. Section 1.8 summarizes the results and concludes.

1.2 Data

This Section presents the data used in the main part of the analysis. First, the sectoral industrial production data is described. Next, I describe how price stickiness is measured at the industry level. Lastly, the monetary policy shock series used in the analysis is presented.

1.2.1 Sectoral Industrial Production Data

Sectoral industrial production data is obtained from the Board of Governors of the Federal Reserve System. At the most disaggregated level, the Board of Governors publishes 213 monthly industrial production index series, measuring the monthly real output of the industries covered by these series. These series are the basis used to construct the commonly used (aggregate) monthly industrial production index for the US. The 213 series cover the whole manufacturing sector plus part of the mining and utilities sectors. Industries covered by the industrial production index are classified in the 2007 version of the North American Industrial Classification System (NAICS)[4]. All series are non-overlapping in terms of NAICS 6-digit industries, i.e. every NAICS 6-digit industry is contained in at most one industrial production series. The level of disaggregation used here is the finest level of disaggregation at which monthly industrial production data is published. The analysis here is confined to industries covered by the industrial production index due to a lack of (disaggregated) monthly output data for other sectors of the economy (like services). All series are seasonally adjusted.

The series used in the industrial production index cover the whole manufacturing sector (NAICS groups 31 - 33), plus those industries that have traditionally been considered to be manufacturing, namely NAICS 1133 (logging) and 5111 (newspaper publishing)), mining and oil extraction (NAICS groups 211-213) and electrical power generation and gas distribution (NAICS 2211 and 2212).

[4]NAICS is the North American Industry Classification System. The finest level of disaggregation available in this classification scheme is the 6-digit classification. NAICS is a hierarchical classification scheme, with the digit-length of the NAICS code indicating the level of disaggregation. For example, several NAICS 6-digit industries can share the same 5-digit industry code.

The Board of Governors does not publish a separate output series for every single NAICS 6-digit industry (which is the finest level of disaggregation in the NAICS classification system). Instead, some NAICS 6-digit industries are grouped together with other (similar) NAICS 6-digit industries that share the same 4- or 5-digit NAICS code.

This leads to the fact that 213 separate industrial production series are available from the Board of Governors[5]. I exclude eight series from the following analysis due to missing data on the frequency of price adjustment for these series. Hence in the following analysis, 205 monthly industrial production series are used.

1.2.2 Data on Price Stickiness

Price stickiness is measured by the (monthly) frequency of price adjustment (FPA) at the level of the industrial production series. This follows the standard approach in the empirical literature (e.g. Bils and Klenow (2004), Nakamura and Steinsson (2008) and Gorodnichenko and Weber (2016)) to measure the degree of price stickiness via the (monthly) frequency of price adjustment. A high frequency of price adjustment implies low observed price stickiness, and vice versa.

The frequency of price adjustment is calculated at the industry level from confidential microdata underlying the US producer price index (PPI) from the Bureau of Labor Statistics (BLS)[6]. The PPI measures selling prices from the perspective of producers, in contrast to the CPI which measures prices from the perspective of consumers. For the analysis here, using the data underlying the PPI is desirable, as prices and output are measured at the same unit (i.e. the producer of the good).

The PPI tracks monthly prices of all goods-producing industries, such as manufacturing and mining. Every month, the BLS surveys the prices of around 100,000 individual items to construct the PPI. The PPI seeks to measure the entire marketed output of US producers (Goldberg and Hellerstein (2009)). The BLS uses a multi-stage sampling procedure to select the items included in the PPI. The sampling procedure is summarized here, based on information given in the BLS Handbook of Methods, Chapter 14.[7] Similar summaries of the sampling procedure for the PPI used by the BLS can be found in e.g. Gorodnichenko

[5]There are 472 unique 6-digit NAICS industries in the manufacturing sector in the 2007 NAICS classification.

[6]I am grateful to Michael Weber for providing the data on the frequency of price adjustment at the industry level to me.

[7]Available under https://www.bls.gov/opub/hom/pdf/ppi-20111028.pdf.

and Weber (2016), Nakamura and Steinsson (2008), Pasten et al. (2018a) and Goldberg and Hellerstein (2009).

In an initial step, the BLS selects the producers, so-called price-forming units, to be included in the PPI. Selection of price-forming units is stratified by industry and based on the following procedure. First, for a given industry, the BLS compiles a list of all establishments (in that industry), based on the information given in the Unemployment Insurance System[8]. In the next step, establishments are clustered into so-called price-forming units. Price-forming units are establishments belonging to the same company, within the same industry. This ensures that prices are collected at the level relevant for price setting, as several establishments owned by a single company may be operated as a cluster and constitute a profit-maximizing center. Finally, a sample of price-forming units is selected to be included in the PPI, with the probability of selection being proportional to its employment size[9].

After a price-forming unit has been selected (and agreed to participate) in the PPI survey, the BLS selects which items produced by the price-forming unit are included in the PPI. Selection of individual items is based on a probability sampling technique called dissaggregation. In the dissaggregation procedure, BLS field economists first combine individual items of a price-forming unit into categories, and assign sampling probabilities to each category proportional to the value of shipments within the reporting unit. Next, the selected categories are broken into additional detail in subsequent stages, until unique items are identified. If the same item is sold at more than one price, then the all price-determining characteristics — for example size and unit of shipments, freight type, type of buyer or color of the item — also must be selected on the basis of probability. This method for identifying the exact transaction terms and price-determining characteristics ensures that the same type of transaction is priced over time.

In line with this procedure, the PPI defines prices as "net revenue accruing to a specified producing establishment from a specified kind of buyer for a specified product shipped under specified transaction terms on a specified day of the month". Taxes and fees collected on behalf of the (federal, state, or local) government are not included in the price. Sales and temporary reductions are reflected in collected prices in so far as they reduce the revenue generated by a specific item received by the producer.

[8]Most employers are legally required to participate in the Unemployment Insurance System.

[9]Possibly within several strata defined by the BLS for a given industry

The BLS collect prices from around 25,000 establishments for approximately 100,000 individual items on a monthly basis. Prices are collected by means of a survey that is e-mailed or faxed to participating establishments. An establishment will remain in the sample for seven years, until a new sample is selected to account for changes in industry structure and changing product market conditions within the industry.

The prices of individual items (defined as above) reported in the PPI database are used to calculate the frequency of price adjustment (FPA) at the industry-level. The FPA at the industry level is calculated using the same method as Pasten et al. (2018a) and Gorodnichenko and Weber (2016). This method is described here.

First, the FPA is calculated for every single item in the data. The FPA at the item level is calculated as the ratio of the number of price changes to the number of sample months. To illustrate this, consider the following (hypothetical) example. Suppose an item is observed in the data for 5 months. The observed price path of this item is \$10 for two months and then \$15 for another three months. Here, one price change occurs during five months, hence the frequency of price adjustment is 1/5 for this item. Price changes due to sales are uncommon in PPI data, but are excluded in the calculation of the FPA, following Pasten et al. (2018a) and Gorodnichenko and Weber (2016). To calculate the FPA at the industry-level, the item-based frequencies are aggregated to the industry level, giving equal weight to all products produced by this industry. The FPA of an industry is hence given as the average FPA of all items produced by this industry.

Industry-level frequencies are calculated at the NAICS 6-digit, 5-digit, 4-digit and 3-digit level. The industry-level frequencies of price adjustment are matched to the corresponding industrial production series.[10]

The PPI sample used to calculate frequency of price adjustment ranges from the years 2005 to 2011. The average monthly frequency of price adjustment across all industrial production series is 23.37%, implying an average price duration, $-1/ln(1-FPA)$, of 3.7 months. Substantial heterogeneity is present in the frequency across sectors, ranging from

[10]If an industrial production series consists of multiple NAICS 6-digit industries, but does not cover the whole (corresponding) NAICS 5-digit industry, the series is assigned the mean of the FPA of the included NAICS 6-digit industries (giving equal weight to all industries). Consider the following example. There are 4 different NAICS 6-digit industries included in a specific NAICS 5-digit industry. The Board of Governors reports an output series for the first NAICS 6-digit industry and a series reporting the combined output of the remaining three NAICS 6-digit industries. The first series, consisting of a single industry, is assigned the frequency of price adjustment of the corresponding NAICS 6-digit industry. The other series is assigned the average of the reported FPA of the three NAICS 6-digit industries included in the series. This method is used to calculate the FPA for 33 industrial production series.

as low as 4.01% (for Semiconductor Machinery Manufacturing, NAICS 333295) to as high as 87.5 % (for Crude Petroleum and Natural Gas Extraction, NAICS 211111). Detailed summary statistics for the frequency of price adjustment can be found in Table A.1 in the Appendix. Figure A.1 in the appendix shows a histogram of the distribution of the frequency of price adjustment (left Panel) and the distribution of the log of the frequency of price adjustment (right Panel). Around 50% of industries in the sample have a frequency of price adjustment between 15% and 25% per month (corresponding to an average duration between 6 and 3.5 months).

1.2.3 Monetary Policy Shocks

Identification of unanticipated, presumably exogenous shocks to monetary policy is a widely discussed topic in the macroeconomic literature. This paper does not propose a new identification scheme for monetary policy shocks, but uses an existing measure of monetary policy shocks from the literature, namely the one proposed by Barakchian and Crowe (2013). Monetary policy shocks in Barakchian and Crowe (2013) are identified as the (financial-market based) 'surprise' component of monetary policy actions, estimated using movements in Fed Funds futures contract prices on the day of monetary policy announcements following meetings of the Federal Open Market Committee (FOMC).

Here, I describe the identification of monetary policy shocks in Barakchian and Crowe (2013) in detail, following their exposition closely. Barakchian and Crowe (2013) identify monetary policy shocks using high-frequency data on Federal Funds futures contracts, financial derivatives whose payoff is calculated based on the effective federal funds rate. Federal Funds futures contracts have been traded since October 1988 (see e.g. Söderström (2001)). The price of a futures contract for month $m+h$ (i.e. at a horizon h from the current month m) is a bet on the monthly average effective Fed Funds rate in month $m+h$, here denoted by $\bar{r}^e_{m+h}$. As Barakchian and Crowe (2013) point out, the average effective Funds Rate might differ from the average target Fed Funds rate ($\bar{r}_{m+h}$, the policy rate set by the Fed) due to implementation errors on part of the Fed:

$$\bar{r}^e_{m+h} = \bar{r}_{m+h} + \epsilon_{m+h} \tag{1.1}$$

where ϵ_{m+h} is the average targeting error for month $m+h$. The futures rate on day d in month m with horizon h is then given by

$$f_d^h = E_d(\bar{r}^e_{m+h}) + \rho_d^h \tag{1.2}$$

where ρ_d^h is a possible risk premium. Under the assumption of an unchanged risk premium and no change in the expected average targeting error during subsequent calendar months ($h \geq 1$), the change in the expected target rate following a policy announcement on a day d of month m is given by

$$\Delta E_d \bar{r}_{m+h} = f_d^h - f_{d-1}^h \tag{1.3}$$

The change in the remainder of the current month (with length M days) is given by

$$\Delta E_d \bar{r}_m = \frac{M}{M-d}(f_d^0 - f_{d-1}^0) \tag{1.4}$$

so the change in the expected target rate is proportional to the (scaled) jump in the futures rate around the policy announcement.

Barakchian and Crowe (2013) calculate the change in the futures rate by comparing the end of day price on the day following the (last) day of an FOMC meeting with that on the meeting day for meetings occurring before February 1994. After February 1994, the change in the futures rate is calculated by comparing the end of day price on the meeting day with the end of day price on the day before the meeting. The analysis is confined to days with FOMC meetings, inter-meeting changes in the target rate are not considered.

The change in the futures rate is calculated for 6 different maturities, starting with the future contract maturing in the current month ($h = 0$), up to the future contract maturing five months after the meeting ($h = 5$). The monetary policy shock measure is then defined as the first principal component of the jump in the futures rate of all 6 maturities. This approach has several advantages over just considering a single maturity. First, this approach minimizes the effect of noise in a specific maturity. Second, as policy decisions are persistent over time, a policy change in the current period will also affect the futures rate several periods ahead. Hence taking into account longer maturities might reveal information of the persistence of the shock. This is important as persistent shocks should have a greater impact on economic activity. It should be noted that financial market based identification schemes of monetary policy shocks, like the one used here, assume that financial market participants beliefs about the Fed's information set prior to the announcement of monetary policy actions are correct, i.e. that unexpected changes[11] in the federal funds rate are indeed due to monetary policy shocks, and not due to superior information of the Fed. To assess this assumption, Barakchian and Crowe (2013) regress their monetary policy shock measure on the difference between the Fed's Greenbook forecasts and high-quality private sector (Blue Chip) forecasts for the 17 variables used in Romer and Romer (2004), where this difference in forecasts is used as a proxy for the Fed's internal information. They find little evidence of superior information of the Fed compared to financial market participants[12]. This suggests that the shock measure used here should be relatively uncorrelated with the Fed's exclusive information, and superior information on the side of the Fed should therefore not be a significant problem. In Section 1.7, I consider a different identification scheme of monetary policy shocks that explicitly controls for the Fed's information set in the identification of monetary policy shocks and find very similar results as in the baseline analysis using the shock measure of Barakchian and Crowe (2013).

The shock series is available from December 1988 onwards at monthly frequency. By construction, the policy shock has mean zero and a standard deviation of one. A detailed overview over the identification of monetary policy shocks can be found in Ramey (2016). The data used here can be downloaded from the website of Valerie Ramey, available at `http://econweb.ucsd.edu/~vramey/research.html#data`. A graph depicting the time series of the shock measure can be found in the Appendix in Figure A.2.

1.3 Aggregate Effects of Monetary Policy Shocks

Before turning to the reaction of different industries to monetary policy shocks, I estimate the reaction of (aggregate) industrial production (and other aggregate variables) to monetary policy shocks in this Section. To estimate the dynamic effects of the identified policy shocks on aggregate variables, I include the cumulated identified shock measure in a VAR. This approach is similar to Romer and Romer (2004) and common in the empirical literature on monetary policy transmission (Ramey (2016)). The specification of the VAR follows Coibion (2012) and includes the same set of variables used in Coibion (2012). The variables included are Industrial Production (in logs), the unemployment rate, the CPI (in logs), a

[11]Unexpected from the viewpoint of financial market participants

[12]The joint hypothesis of zero coefficients on all 17 variables cannot be rejected at the 10% level.

commodity price index[13] (in logs) (all seasonally adjusted) and the cumulated shock series. The VAR is estimated at monthly frequency from December 1988 to December 2007.[14] The VAR includes 12 lags and a constant.

Following the literature on monetary policy shocks in structural VARs, the recursive identification scheme of Christiano et al. (1999) is employed. Monetary policy is assumed to respond to, but not affect, the other variables contemporaneously. The specification used here employs the measure of policy shocks of Barakchian and Crowe (2013), whereas Christiano et al. (1999) use the actual funds rate. Since standard VARs enter the federal funds rate in levels, the shock series is cumulated to produce a comparable series. The same estimation procedure (with a different measure of policy shocks) is used in e.g. Romer and Romer (2004) or Coibion (2012). Similar specifications are commonly used in the literature (see Ramey (2016) for a detailed review). It should be noted that the recursiveness assumption is not necessary for identification of monetary policy shocks in the analysis at hand, as an identified shock series is used. In order to be comparable with the previous literature, in the baseline analysis the recursiveness assumption is used. The recursiveness assumption is relaxed in Section 1.7.

[13]Taken from Valerie Rameys website described above.

[14]I end the sample in December 2007 to abstract from issues related to the Zero Lower Bound.

Figure 1.1: Response of Industrial Production to a Contractionary Monetary Policy Shock

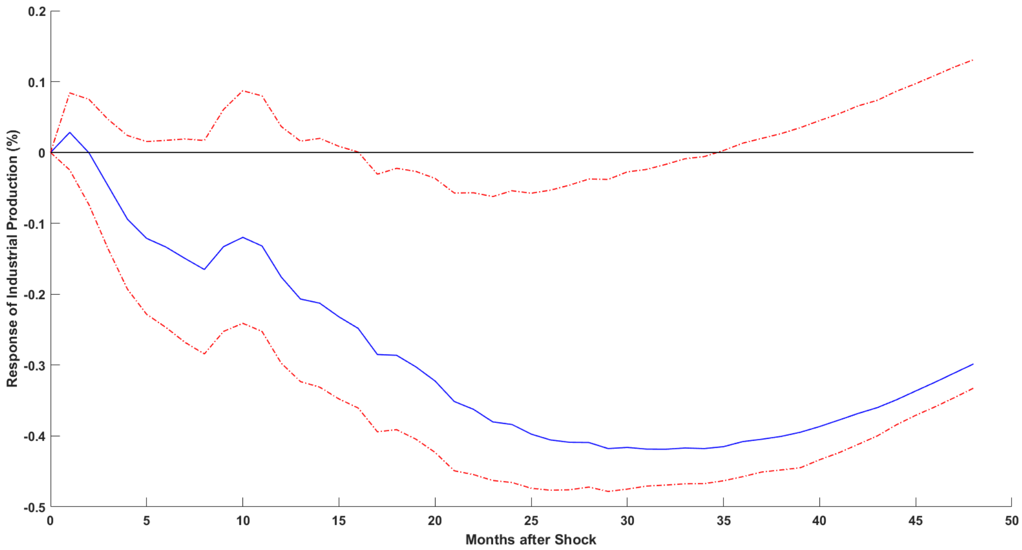

Structural VAR (Monthly data, 5 endogenous variables, 12 lags). Variables ordered as industrial production (in logs), unemployment rate, consumer price index (in logs), commodity price index (in logs) (all seasonally adjusted) and the cumulated shock measure of Barakchian and Crowe (2013). The graph shows response of industrial production to a one standard deviation increase in the policy measure. Structural shocks obtained via Cholesky decomposition. Bootstrapped 90% Confidence Intervals based on 5000 Bootstrap Replications are shown in red.

Figure 1.1 shows the (cumulative) response of the aggregate industrial production index to a one standard deviation increase in the policy measure (in percent). After a one standard deviation increase in the policy measure, industrial production drops by around 0.38% 2 years after the shock has occurred. Qualitatively, the sign and the speed of the reaction are in line with other estimates found in the literature, e.g. Coibion (2012), Romer and Romer (2004) and Ramey (2016). The timing of the response is very similar to Barakchian and Crowe (2013), who use a different specification for their VAR. The size of the response is slightly larger than the estimated response in Barakchian and Crowe (2013) (whose specification implies a drop in output of around 0.3% 2 years after a one standard deviation contractionary shock). The size of the reaction is also comparable to other studies that use financial market based measures of monetary policy shocks. For example, Gertler and Karadi (2015) report a drop in industrial production of around 0.4% two years after a one standard deviation increase in their policy instrument. Also the timing of the response in Gertler and Karadi (2015) is very similar to the results reported here.

As the main focus of this of this study is the output response of different industries, the reactions of the other variables are relegated to Figure A.3 in the Appendix.

1.4 The Industry Effects of Monetary Policy Shocks

This Section describes the output responses of the 205 industries to the identified monetary policy shocks. First, the econometric specification is discussed. Then, the results of the estimation are presented.

1.4.1 Econometric Specification

To estimate the effect of monetary policy shocks on the output of a particular industry i, the (log) output (seasonally adjusted) of industry i is added to the VAR described in the Section 1.3 as additional (sixth) variable.

Several additional identifying restrictions are imposed on the VAR, compared to the VAR in Section 1.3. The purpose of these restrictions is to ensure that the sequence of monetary policy shocks is the same for all industries i.

Let $Y_{t,i} = [IP_t, UNEMP_t, CPI_t, PCOM_t, CSHOCK_t, OUT_{t,i}]'$ denote the variables included in the VAR for industry i. The first five variables are the same as described in the previous Section and $OUT_{i,t}$ denotes the log of the industrial production index of industry i.

Denote the reduced-form VAR for industry i as[15]

$$Y_{t,i} = \sum_{p=1}^{P} \Phi_{p,i} Y_{t-p,i} + e_{t,i} \tag{1.5}$$

where $\Phi_{p,i}$ are the reduced-form VAR coefficients (with $P = 12$ in the specification used here) and $e_{t,i}$ denoting the mean zero reduced-form VAR residuals with variance-covariance matrix $E(e_{t,i}e'_{t,i}) = \Omega_i$.

In order to ensure that the sequence of (aggregate) shocks is the same for each industry i, several restrictions are imposed.

Assume that the structural form of the model for every industry i is given by

$$A_i(L)Y_{t,i} = v_{t,i} \tag{1.6}$$

where $A_i(L) = A_{0,i} - A_{1,i}L - \ldots - A_{P,i}L^P$ is a (invertible) lag polynomial of order P and

[15]The constant c is omitted here, i.e. the data is demeaned.

L denotes the lag operator. The mutually uncorrelated structural innovations are denoted by $v_{t,i}$ with diagonal variance-covariance matrix Σ_i.

The goal of the following restrictions is to identify the reaction to the fifth element of the vector of structural shocks $v_{t,i}$, which is the innovation to the policy measure. Other structural shocks ordered before the policy variable are left unspecified (but common to all industries).

The relation between the structural parameters and the reduced-form coefficients is hence given by:

$$\Phi_{p,i} = A_{0,i}^{-1} A_{p,i} \tag{1.7}$$

$$\Omega_i = A_{0,i}^{-1} \Sigma_i (A_{0,i}^{-1})^T \tag{1.8}$$

In line with the recursiveness assumption made in the previous section, the matrix $A_{0,i}$ is assumed to be lower triangular with an additional zero restriction:

$$A_{0,i} = \begin{bmatrix} * & 0 & 0 & 0 & 0 & 0 \\ * & * & 0 & 0 & 0 & 0 \\ * & * & * & 0 & 0 & 0 \\ * & * & * & * & 0 & 0 \\ * & * & * & * & * & 0 \\ * & * & * & * & 0 & * \end{bmatrix} \tag{1.9}$$

where $*$ denotes an unrestricted coefficient. In economic terms, these restrictions imply the same recursiveness assumption that was invoked in the preceding section: Monetary policy shocks (ordered fifth) have no contemporaneous impact on all other variables in the system, including the output of industry i. The additional zero in the fifth column in the last row of matrix $A_{0,i}$ ensures that policy shocks have no contemporaneous impact on the output of industry i.

Additionally, the following restrictions are imposed on the structural VAR parameters for $p \neq 0$:

$$A_{p,i} = \begin{bmatrix} * & * & * & * & * & 0 \\ * & * & * & * & * & 0 \\ * & * & * & * & * & 0 \\ * & * & * & * & * & 0 \\ * & * & * & * & * & 0 \\ * & * & * & * & * & * \end{bmatrix} \tag{1.10}$$

Under these restrictions, the reduced-form parameter matrices $\Phi_{p,i}$ have zeros in the same place as the structural VAR parameter matrices $A_{p,i}$, the system has a block recursive structure. First, an aggregate block, containing the five aggregate variables, whose dynamics are the same for every industry i, and the same as the VAR described in Section 1.3. Second, an industry-specific block, whose coefficients are different for every industry i, which contains the output of industry i as only variable. It should be noted that these restrictions imply that sector-specific movements in industry i's output are constrained to affect the variables in the common subsystem (the aggregate variables) in proportion to the sector's share of total industrial production[16].

Imposing the structural restrictions described here ensures that the sequence of monetary policy shocks (and of the other, unspecified, common shocks) is the same for all industries i, and equal to the sequence of policy shocks in the estimation without industry i. Similar restrictions are imposed in Davis and Haltiwanger (2001), who refer to this set-up as "near-VAR". In fact, the whole system can be described as a restricted panel VAR, with a common block of macroeconomic variables and an industry-specific block (which here contains a single variable, industry i's output). One appealing feature of this estimation is that it allows sectoral responses to monetary policy shocks to vary freely, while keeping the same sequence of policy shocks and dynamics of aggregate variables for every industry i.

[16] i.e. the feedback of industry i's output, $OUT_{t,i}$, on total industrial production IP_t and the other aggregate variables is not explicitly estimated, but the coefficients on (lagged) values of $OUT_{t,i}$ for the aggregate variables are constrained to be zero. However, movements in the output of industry i still have an effect on the aggregate system, due to the fact that total industrial production IP_t is a (weighted) sum of the output of all industries i.

1.4.2 Results

The panel VAR described in Section 1.4.1 is estimated at monthly frequency using data from the period of December 1988 to December 2007. The VAR includes 12 lags and a constant, similar to the specification of the aggregate VAR described in Section 1.3.

Figure 1.2 shows the distribution of the dynamic output responses of the 205 industries to a one standard deviation contractionary monetary policy shock. The red line shows the median of the (cross-sectional) output responses (in percent) of the 205 industries at each horizon after the shock. The blue lines show the 25th and 75th percentile of the distribution of responses at each horizon after the shock, respectively. The green lines show the fifth and 95th percentile of the distribution of responses at each horizon.

Figure 1.2: Industry Responses to a Contractionary Monetary Policy Shock

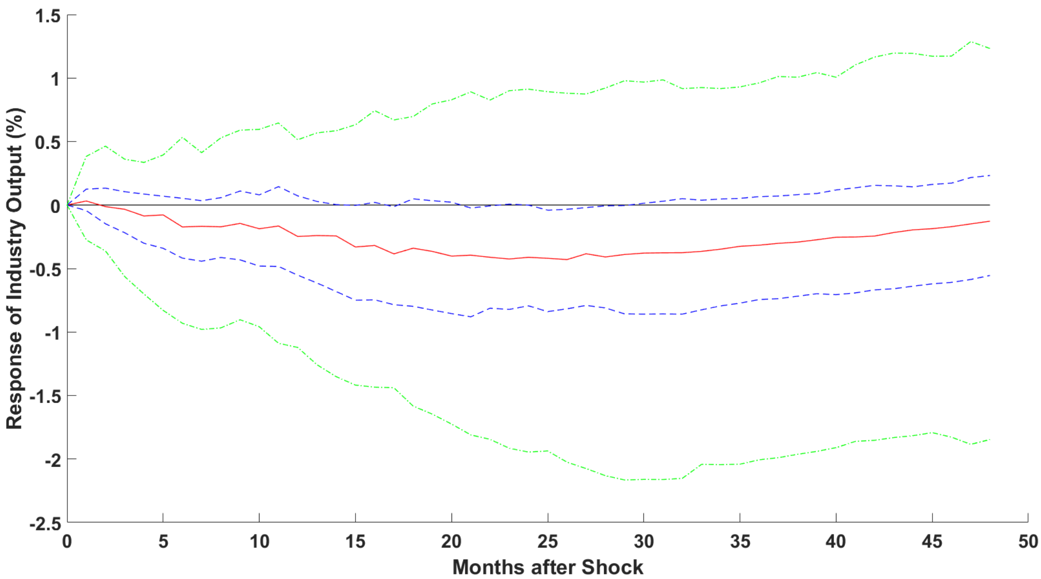

This Figure shows the distribution of output responses of the 205 different industries to a one standard deviation contractionary policy shock identified as in Barakchian and Crowe (2013). Structural shocks are identified via Cholesky decomposition. Policy shocks and aggregate dynamics are common across industries. The red line shows the median response of all industries at each horizon. The blue lines show the 25th and 75th percentile of the distribution of the industries output responses at each horizon. The green lines show the fifth and 95th percentile of the distribution of the industries output responses at each horizon.

The median response of industries (red) tracks the response of aggregate industrial production well: 2 years after the shock has happened, the median industries output has fallen by around 0.4%. However, there is substantial heterogeneity present in the industries output responses. The interquartile range of responses is about 0.9% 2 years after the shock[17].

For some industries the fall in output is as large as 2% two years after the shock, i.e. around 5 times as large as the drop in total output. On the other hand, around 25% of industries even experience an increase in output in response to a contractionary policy shock. The fact that a substantial share of industries increase their output in response to contractionary monetary policy shocks deserves mentioning. In the multi-sector New Keynesian models of Ghassibe (2018) and Bouakez et al. (2013), sectoral output responses are all negative in response to contractionary monetary shocks. In an extension of his model, Ghassibe (2018) shows that positive output reactions to a contractionary policy shock are only possible under an elasticity of substitution between sectors that is greater than one. The fact that a non-negligible share of sectors increases output after a contractionary policy shock is hence not consistent with the most basic versions of multi-sector New Keynesian models. In the multi-sector menu cost model of Nakamura and Steinsson (2010), several industries experience a drop in output in response a positive nominal demand shock[18], which is consistent with the empirical reactions presented here.

This result shows that there is indeed substantial heterogeneity in the output responses of industries to monetary policy shocks. A natural question to ask is which industry characteristics determine these differential reactions. The following Section investigates the role of heterogeneity in price stickiness across industries in determining these differential reactions.

[17] The difference between the 25 percentile and the 75 percentile of responses, which corresponds to the distance between the blue lines in Figure 1.2.

[18] which is the equivalent of monetary policy shocks in their model

1.5 The Role of Price Stickiness in Explaining Industry Output Reactions to Monetary Policy Shocks

The building blocks of New Keynesian Macroeconomics imply that the observed heterogeneity in the output responses to monetary policy shocks across industries should be systematically related to the degree of price stickiness of these industries. Industries with a higher degree of price stickiness should systematically react more strongly to monetary policy shocks (ceteris paribus). This prediction is made by various multi-sector New Keynesian models, for example Bouakez et al. (2013) and Ghassibe (2018), but can also arise from multi-sector menu cost models like for example Nakamura and Steinsson (2010). Testing this prediction does not only speak to multi-sector New Keynesian models, but more generally to New Keynesian models of the monetary transmission mechanism. New Keynesian models greatly emphasize the role of nominal rigidities, in particular price stickiness, for the monetary transmission mechanism and the real effects of monetary policy (Galí (2015)). Hence the test whether price stickiness has a role in explaining differential industry reactions is an important test of the New Keynesian paradigm. This Section asses whether this proposed "price stickiness channel" is supported by the data.

To asses the correlation between price stickiness and the strength of the response to monetary policy shocks, I run the following regression:

$$IRF_i^h = \alpha^h + \beta^h log(FPA_i) + e_i^h \tag{1.11}$$

where IRF_i^h is the output response of industry i to an unexpected one standard deviation increase in the policy measure h months after the shock (measured in percent) and FPA_i is the monthly frequency of price adjustment of industry i. The frequency of price adjustment enters in logs rather than in levels to estimate the (semi) elasticity of the frequency of price adjustment.[19]

Note that a higher frequency of price adjustment means that prices are *less* sticky. The results of this regression for different horizons (18, 24 and 30 months) after the monetary policy shock can be found in Table 1.1. I focus on these horizons to be comparable with the

[19] All results are robust to using the frequency of price adjustment in levels rather than logs as independent variable in the regression. The choice of using FPA in logs rather than levels is also motivated by the fact that, as can bee seen in Figure A.1 in the Appendix, the distribution of the frequency of price adjustment is very skewed to the right, whereas the distribution of the logged frequency of price adjustment is more symmetrical.

previous literature on sectoral differences in the monetary policy transmission mechanism, e.g. Dedola and Lippi (2005), and as the peak response of aggregate industrial production is reached between these horizons. Table 1.1 also reports the reaction of the (total) industrial production index h months after the shock.

It should be noted that the coefficient β^h here should not be interpreted as a causal effect, as price stickiness is not randomly assigned. β^h could only be interpreted as a causal effect if one assumes that there are no other factors that affect both, an industries reaction to monetary policy shocks, and this industries degree of price stickiness. In Section 1.6 I investigate how the estimated coefficient on the frequency of price adjustment changes once additional industry characteristics are added to the regression.

Table 1.1: Price Stickiness and Monetary Policy Responses

	(1) $h = 18$ Months	(2) $h = 24$ Months	(3) $h = 30$ Months
Log(FPA)	0.223	0.341**	0.339**
	(0.143)	(0.147)	(0.136)
Observations	205	205	205
R-squared	0.015	0.031	0.034
$I\bar{R}F^h$	-0.286	-0.383	-0.416

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

This Table reports the coefficient on log(FPA) estimated in Equation 1.11 at different horizons h. The regression also includes a constant (not reported). $I\bar{R}F^h$ denotes the response of the (total) industrial production index h months after an unexpected one standard deviation increase in the policy measure. Robust Standard Errors are reported in Parenthesis.

The scatter plot underlying the regression for $h = 24$ can be found in the Appendix in Figure A.4. For the horizons $h = 24$ and $h = 30$ months after the monetary policy shock, there is a statistically significant association between the strength of the response to policy shocks and the degree of price stickiness.

Qualitatively, the sign of the coefficient on the FPA is consistent with the predictions made by New Keynesian models. In response to the monetary policy shock, there is a drop in (total) industrial production (e.g. by 0.38% for $h = 24$ months). The positive sign on the regression coefficient means that, on average, industries with a higher frequency of price adjustment, i.e. more flexible prices, experience a *less* negative drop in output. This reaction is qualitatively in line with the prediction of (Multi-Sector) New Keynesian models: Industries with more flexible prices should react less strongly to monetary policy surprises. The positive sign on the coefficient shows that this is indeed the case: The magnitude of the reaction to policy shocks is lower for flexible price industries.

Quantitatively, the size of the coefficient is also economically significant. Consider the horizon of $h = 24$ months after the policy shock. A 10% increase in the frequency of price adjustment is associated with a 0.034 percentage point increase in the cumulative output response to the shock. Compared to the drop in the (total) industrial production index at the same horizon, which is 0.38, the size of this effect is economically meaningful: a 10% increase in the frequency of price adjustment leads to a less negative output response to the policy shock that is (approximately) as large as 10% of the total drop in output. The implied differential reaction between the most sticky price industries and the most flexible price industries is large. The tenth percentile of the in-sample frequency of price adjustment is given by 11.52 (meaning that on average 11.52% of prices are changed per month), implying a log(FPA) of 2.44. The 90th percentile of the in-sample frequency of price adjustment is given by 41.68, implying a log(FPA) of 3.73. The regression coefficient for the horizon of $h = 24$ months implies that the drop in output is $0.341 \times (3.73 - 2.44) = 0.44$ percentage points larger for industries at the tenth percentile of the in-sample frequency of price adjustment compared to industries at the 90th percentile of the in-sample frequency of price adjustment. Compared to the drop in total industrial production, which is 0.38 percentage points at the same horizon, this implied differential reaction is sizable.

The relative size of the effect is similar for the horizon of $h = 30$ months. For the horizon of $h = 18$ months, the sign and relative size of the coefficient are comparable to the other two horizons considered, but the coefficient is (marginally) insignificant (the p-value based on robust standard errors is $p = 0.12$).

Figure 1.3 shows the coefficient on log(FPA) estimated from Equation 1.11 for all horizons from $h = 0$ up to $h = 48$. Additionally, Figure 1.3 shows the inverted impulse-response function of total industrial production (i.e. the IRF is multiplied with (-1)). This Figure shows that for the first year after the shock there is no differential reaction between sticky and flexible price industries. However, when total industrial production starts to fall (around 1 year after the shock), the drop is stronger for sticky price industries than it is for flexible price industries. In fact, the coefficient on log(FPA) moves very much in parallel with the IRF of total industrial production (starting from $h = 12$ months after the shock), differences between sticky and flexible price industries coincide with the drop in aggregate output and do not seem to be driven by differences in the speed of the reaction.

Figure 1.3: Regression Coefficient on log(FPA) for different Horizons h

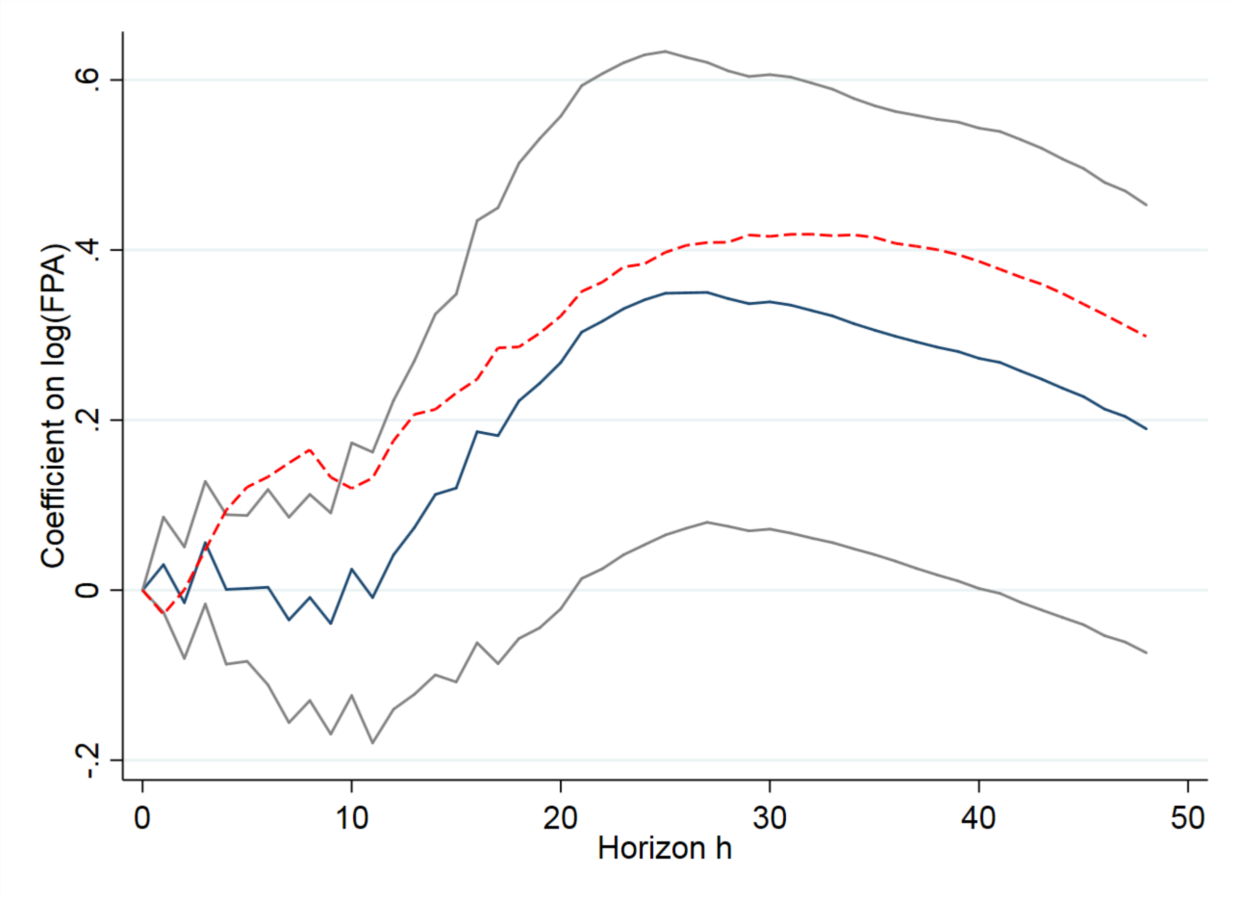

This Figure shows the time series of the estimated coefficient on log(FPA) from Equation 1.11 for all horizons from $h = 0$ up to $h = 48$ months after a contractionary policy shock. The blue line shows the estimated coefficient on log(FPA) at each horizon h. The gray lines are 95% confidence bands of the coefficient based on robust standard errors. The red, dashed line shows the inverted reaction of total Industrial Production in response to the policy shock in percent (i.e. the IRF is multiplied with (-1)).

1.6 Adding Further Industry Characteristics

The analysis so far has focused on differences in the frequency of price adjustment as source of heterogeneity in output responses between industries. However, the frequency of price adjustment is not the only (potential) factor determining an industries reaction to monetary policy shocks. For example, more cyclical industries exhibit a higher frequency of price adjustment (Klenow and Malin (2010)). At the same time, cyclical industries might exhibit a larger drop in output following a contractionary monetary policy shock, as they react more strongly to swings in economic activity. Not controlling for the cyclicality of an industry could hence lead to omitted variable bias in the estimated relationship between the frequency of price adjustment and the reaction to monetary policy shocks[20].

The goal of this Section is to assess to what extent the result established in the previous Section is affected by other (so far omitted) industry-level characteristics. To asses this further industry-level characteristics are added as additional control variables to Regression 1.11. The goal of this approach is to disentangle the role of the frequency of price adjustment from other potentially confounding factors in the monetary transmission mechanism. Regression 1.11 is hence changed to:

$$IRF_i^h = \alpha^h + \beta^h log(FPA_i) + \gamma^h X_i + e_i^h \tag{1.12}$$

where X_i denotes the additional industry-level control variables that will be added in this Section. Additional industry-level control variables X_i are added to the regression in order to test the robustness of the main result established in the Section 1.5. It should be noted that the goal of this analysis is not to explicitly test for the importance of other industry characteristics for the transmission of monetary policy, like e.g. the analysis of Peersman and Smets (2005) or Dedola and Lippi (2005). Hence the sign and size of the coefficients on the other industry-level control variables considered in this Section are not discussed in detail.

Before presenting the results of this exercise, I present an overview over the additional control variables included in Equation 1.12. The inclusion of the control variables is motivated either by the fact that an explicit link between the control variable and the frequency of price adjustment has been suggested, or by the fact that the variable has been found to be a significant determinant of industry-level responses to monetary policy shocks in the previous literature[21].

The full list of additional control variables is reported in Table 1.2. Whenever possible, variables are calculated as industry averages over the same time period that is used in Section 1.3 and 1.4. Summary statistics for the variables can be found in the Appendix in Table A.2. Additional details on the calculation of the variables can be found in Section A.1 in the Appendix.

[20]In this specific example, failing to control for industry cyclicality should attenuate the coefficient on FPA towards zero.

[21]Even if an additional control variable should be orthogonal to the frequency of price adjustment, but influences an industries reaction to monetary policy shocks, the inclusion of this control variable will help to obtain more precise estimates of the effect of the frequency of price adjustment.

Table 1.2: Additional Control Variables and Data Sources

Variable	Source & Time Period	Frequency
Inventory over Sales	NBER-CES: 1988 - 2007	Yearly
Labor Cost over Sales	NBER-CES: 1988 - 2007	Yearly
Capital Intensity	NBER-CES: 1988 - 2007	Yearly
Average Firm Size (Number of Employees)	Economic Census 2007	Yearly
Interest Rate Burden	Compustat: 1988 - 2007	Yearly
Leverage Ratio	Compustat: 1988 - 2007	Yearly
Short-Term Debt Ratio	Compustat: 1988 - 2007	Yearly
Standard Dev. of Output Growth	Ind. Prod. Data: 1988 - 2007	Monthly
Cyclicality	Ind. Prod. Data: 1988 - 2007	Monthly
Durable Goods Dummy	BLS	Fixed

This Table presents an overview of the additional industry-level control variables considered in this Section.

The data sources used to calculate the additional control variables are the NBER-CES manufacturing database (denoted by NBER-CES in Table 1.2), the Compustat North America Fundamentals Annual database (denoted by Compustat in Table 1.2), the industrial production data form the Fed Board of Governors (described in Section 1.2, denoted by Ind. Prod. Data in Table 1.2) and the Durable Goods Producer definition of the Bureau of Labor Statistics (denoted by BLS in Table 1.2). In the following, the variables presented in Table 1.2 are explained in detail.

Inventory over Sales & Labor Cost over Sales: Industry-level differences in the reaction to monetary policy shocks might be driven by industry-level differences in the dependence on external funding (following Bernanke et al. (1999)). At the same time, Balleer et al. (2017) document a link between financial constraints and the frequency of price adjustment at the firm level (using German survey data). To control for this potential link of external financial dependence and the frequency of price adjustment, the external financial dependence of an industry is added as additional control variable. Following Raddatz (2006), two measures of external financial dependence are calculated: The industry-level ratio of inventories over sales and the industry-level ratio of labor costs over sales. Industries with higher ratios can finance less of ongoing costs through revenues and hence might dependent more on external financing. Both measures are calculated from the NBER-CES manufacturing database.

Capital Intensity: More capital-intensive industries are expected to be more sensitive to changes in the user cost of capital, which itself will depend on changes in interest rates (Peersman and Smets (2005) and Bouakez et al. (2013)). To control for this cost of capital channel, an industry's capital intensity is added as additional control variable. Following Peersman and Smets (2005), capital intensity is calculated as the ratio of capital expenditures over sales, using data from the NBER-CES manufacturing database over the time period 1988 to 2007.

Average Firm Size: Gertler and Gilchrist (1994) argue that small firms are more strongly affected by financial frictions than large firms and hence small firms are affected more strongly by monetary policy shocks than large firms. At the same time, large firms exhibit a higher frequency of price adjustment (Goldberg and Hellerstein (2009)). To control for industry-level differences in firm size, average firm size is added as additional control variable[22]. Average firm size is measured as the average number of employees per firm at the industry-level, calculated from the Economic Census 2007.

Interest-Rate Burden: Following Dedola and Lippi (2005), industries with higher interest rate expenses should be more exposed to changes in the interest rate. These industries should experience larger changes in costs following changes in interest rates. To control for this interest rate expense channel, the interest burden is added as additional control variable. The interest burden is calculated as the ratio of interest expenses over sales, using Compustat data from 1988 to 2007.

Leverage Ratio: D'Acunto et al. (2018) show that flexible price firms, on average, have a higher leverage ratio than sticky price firms. At the same time, Ottonello and Winberry (2018) show that at the firm level low leverage is associated with stronger (investment) responses to monetary policy shocks. In order to control for industry-level differences in leverage, I add industry-level leverage as additional control variable. Leverage is calculated as the ratio of total debt over total assets using Compustat data from 1988 to 2007.

Short-Term Debt Ratio: Industries with a larger share of short-term debt should be more exposed to changes in interest rates than industries with longer debt maturities, as they need to refinance debt more often. Hence industries with a larger share of short-term debt should experience a relatively larger change in the user cost of capital following changes in the interest rate. Following Dedola and Lippi (2005), the ratio of short-term debt over total assets is added as additional control variables. This variable is calculated using Compustat data from 1988 to 2007.

Standard Deviation of Output Growth: The size and frequency of idiosyncratic shocks

[22]This variable enters regression 1.12 in logs, rather than in levels.

might differ along industries as well. Industries facing more volatile idiosyncratic shocks might adjust prices more often, stay closer to their optimal price level and hence react less severely to monetary shocks (see e.g. Nakamura and Steinsson (2010)). Following Gorodnichenko and Weber (2016) the standard deviation of (monthly) output growth is added as additional control variable to control for the variance of idiosyncratic shocks. The standard deviation of output growth is calculated using the industry-level production data described in Section 1.2, over the time period December 1988 to December 2007.

Cyclicality: A further factor that might confound the effect of the frequency of price adjustment is the cyclicality of an industry (as noted in the beginning of this Section). When a contractionary policy shock causes economic activity to drop, more cyclical industries might experience a larger drop in output. At the same time, cyclical industries change prices more often (Klenow and Malin (2010)). To control for the cyclicality of an industry, the coefficient on total output growth estimated from a regression of (demeaned) monthly industry-level output growth on (demeaned) monthly total output growth is added as additional control variable. This coefficient is calculated using the monthly industrial production data described in Section 1.2 over the time period from 1988 to 2007.

Durable Goods Producers: Lastly, following Dedola and Lippi (2005), a dummy variable for industries producing durable goods is added, as durable goods producers might face more cyclical and more interest-rate sensitive demand.

I control for the additional variables described here in two separate ways. First, I report the results of a regression that includes only one of the control variables in addition to the frequency of price adjustment (and no other control variables), separately for every control variable. Furthermore, I estimate Equation 1.12 including all additional control variables jointly.

Table 1.3 and 1.4 shows the result for each of the regressions for the timing of $h = 24$ months after the shock.

Table 1.3: Price Stickiness and Monetary Policy Responses - Additional Controls

	(1)	(2)	(3)	(4)	(5)	(6)
Log(FPA)	0.545***	0.534**	0.471***	0.377**	0.348**	0.336**
	(0.181)	(0.228)	(0.172)	(0.164)	(0.149)	(0.146)
Inventory/Sales	2.374***					
	(0.795)					
Labor Costs/Sales		1.070				
		(1.464)				
Capital Intensity			-5.554			
			(4.534)			
Firm Size				-0.0595		
				(0.0767)		
Interest Expense Ratio					-1.489	
					(1.853)	
Leverage						0.227
						(0.598)
Observations	187	187	187	202	204	204
R-squared	0.060	0.044	0.047	0.033	0.034	0.032

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

This Table reports the estimation results when estimating Equation 1.12 at a horizon of $h = 24$ months, including the additional control variables described in the text one at a time.

Table 1.4: Price Stickiness and Monetary Policy Responses - Additional Controls Cont.

	(1)	(2)	(3)	(4)	(5)
Log(FPA)	0.337**	0.319**	0.356**	0.257*	0.612***
	(0.148)	(0.147)	(0.150)	(0.154)	(0.235)
Inventory/Sales					2.880***
					(0.841)
Labor Costs/Sales					2.410
					(1.777)
Capital Intensity					-6.844
					(4.306)
Firm Size					0.0434
					(0.0804)
Interest Expense Ratio					-3.058
					(2.307)
Leverage					1.283*
					(0.748)
Short-Term Debt Ratio	-1.591				-2.096
	(1.690)				(2.013)
Cyclicality		-0.161**			-0.269***
		(0.0677)			(0.0813)
Std(Output Growth)			-0.0215		0.0416
			(0.0321)		(0.0499)
Durability				-0.245*	-0.243
				(0.138)	(0.175)
Observations	204	205	205	205	186
R-squared	0.034	0.055	0.033	0.045	0.149

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

This Table reports the estimation results when estimating Equation 1.12 at a horizon of $h = 24$ months, including the additional control variables described in the text one at a time. In the last column all control variables described in the text are included jointly.

Table 1.3 and 1.4 show that the main finding established in Section 1.5 is not only robust to the inclusion of further industry characteristics, but the estimated coefficient on the frequency of price adjustment becomes even larger in most cases. When controlling for all other control variables jointly (Column 5 in Table 1.4), the coefficient on the frequency of price adjustment becomes nearly twice as large compared to the case when no other control variables are used (0.612 when including all other control variables vs 0.341 when no other control variables are used).

Table 1.5 shows that this is also the case for the horizons of $h = 18$ months and $h = 30$ months after the shock. Table 1.5 reports the results of Equation 1.12 when all other control variables are included jointly for the horizons of $h = 18$ months and $h = 30$ months. Similar to the case of $h = 24$ months after the shock, the coefficient on the frequency of price adjustment is now substantially larger than before (0.496 when including all other control variables vs 0.223 when no other control variables are included for the horizon $h = 18$ months after the shock, and 0.548 when including all other control variables vs 0.339 when no other control variables are included for the horizon $h = 30$ months after the shock).

Table 1.5: Price Stickiness and Monetary Policy Responses - Additional Controls at Different Horizons

	(1) $h = 18$ Months	(2) $h = 30$ Months
Log(FPA)	0.496**	0.548**
	(0.242)	(0.226)
Inventory/Sales	2.557***	2.340***
	(0.806)	(0.881)
Labor Costs/Sales	2.532	2.242
	(1.584)	(1.828)
Capital Int.	-5.785	-7.000
	(3.980)	(4.792)
Firm Size	0.0597	0.0606
	(0.0770)	(0.0788)
Interest Expense Ratio	-2.975*	-3.515
	(1.754)	(2.476)
Leverage	0.627	1.469**
	(0.711)	(0.741)
Short-Term Debt Ratio	-1.245	-2.419
	(1.824)	(2.008)
Cyclicality	-0.199***	-0.283***
	(0.0751)	(0.0796)
Std(Output Growth)	0.0144	0.0679
	(0.0478)	(0.0452)
Durability	-0.227	-0.167
	(0.161)	(0.175)
Observations	186	186
R-squared	0.112	0.152

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

This Table reports the estimation results when estimating Equation 1.12 at horizons $h = 18$ months (column (1)) and $h = 30$ months (column (2)), respectively. All control variables described in the text are included jointly.

The results presented in this Section provides further support for the hypothesis that sticky price industries react more strongly to monetary policy shocks than flexible price industries. The results reported in Tables 1.3, 1.4 and 1.5 show that the differential output reactions in response to monetary policy shocks between sticky- and flexible-price industries established in Section 1.5 is not spuriously caused by omitted variable bias induced by a wide range of other industry characteristics controlled for in this Section. This finding is consistent with the prediction made by (multi-sector) New Keynesian models and suggests that sticky prices indeed play an important role in the monetary transmission mechanism. The results reported here alleviate concerns that the (cross-industry) correlation between the output response to monetary policy shocks and the frequency of price adjustment is spuriously caused by other industry characteristics. In fact, the findings reported in Tables 1.3, 1.4 and 1.5 rather suggest the opposite: when not controlling for the additional industry-level characteristics considered in this Section (as in Equation 1.11), the correlation between the output response and the frequency of price adjustment is attenuated towards zero.

Qualitatively, the result is unchanged compared to Section 1.5 - Industries with more sticky prices react more strongly to monetary policy shocks. Quantitatively, the results reported here suggest that the differential reaction between sticky and flexible price industries is even stronger as suggested by the results reported in Section 1.5.

1.7 Robustness

This Section further investigates the robustness of the results in two respects. First, I consider the robustness of the results with respect to the minimum delay restriction (i.e. the restriction that monetary shocks have no effect on industrial production on impact) imposed in the baseline estimation scheme. Furthermore, I consider the robustness of the results with respect to the identification of monetary policy shocks. To asses the robustness of the results in this dimension, I repeat the analysis using the identification scheme for monetary policy shocks of Miranda-Agrippino (2016).

As one robustness check, I relax the minimum delay restriction imposed in the baseline estimation scheme described in Sections 1.3 and 1.4. In this robustness check, I allow for an immediate effect of monetary policy shocks on the other variables in the system (most importantly on aggregate and sectoral industrial production). This is achieved by ordering the monetary policy shock as first variable in the VAR, allowing for an immediate reaction of all other variables[23]. To allow for an immediate reaction of (sectoral) output of industry i, matrix $A_{0,i}$ is changed to:

$$A_{0,i} = \begin{bmatrix} * & 0 & 0 & 0 & 0 & 0 \\ * & * & 0 & 0 & 0 & 0 \\ * & * & * & 0 & 0 & 0 \\ * & * & * & * & 0 & 0 \\ * & * & * & * & * & 0 \\ * & * & * & * & * & * \end{bmatrix} \tag{1.13}$$

which allows for an immediate reaction of industry i's output to a policy shock (and for an immediate reaction of (aggregate) industrial production in response to a monetary policy shock). In the following, I will refer to this robustness check as 'Alternative Ordering'.

As further robustness check, I repeat the analysis using the monetary policy shock series of Miranda-Agrippino (2016). The identification scheme of Miranda-Agrippino (2016) is a hybrid approach of identifying monetary policy shocks, combining market-based measures of monetary policy shocks (like Barakchian and Crowe (2013) used above, or e.g. Gürkaynak et al. (2005)) and the narrative approach of Romer and Romer (2004). One concern with

[23] In terms of the notation used before, $Y_{t,i}$ is now given by $Y_{t,i} = [CSHOCK_t, IP_t, UNEMP_t, CPI_t, PCOM_t, OUT_{t,i}]^T$

the financial market based identification approach of monetary shocks used in the baseline analysis is that the central bank might have superior information compared to financial market participants (as noted in Section 1.2). If this is the case, monetary policy decisions that are surprising to financial market participants might not be pure monetary shocks, but also convey new information to agents in the economy (see e.g. Jarocinski and Karadi (2018)). To control for this information effect, Miranda-Agrippino (2016) constructs a financial market based measure of monetary policy shocks that explicitly controls for the information set of the central bank. This approach combines the high-frequency identification of monetary policy shocks of Gertler and Karadi (2015) with the narrative identification of Romer and Romer (2004). Monetary policy shocks in Miranda-Agrippino (2016) are constructed as the residual of a regression of the high-frequency monetary policy shocks of Gertler and Karadi (2015) on the Fed Greenbook Forecast variables used by Romer and Romer (2004)[24]. In the robustness check using this identification of monetary policy shocks, the estimation of (aggregate and industry-specific) responses to monetary shocks is carried out as described in Section 1.3 and Section 1.4, replacing the shock series of Barakchian and Crowe (2013) with the shock series of Miranda-Agrippino (2016). In the following, I will refer to this robustness check as 'Alternative Shock Identification'.

[24]The monetary policy shock series constructed this way can be downloaded from the webpage of Silvia Miranda-Aggripino, avaiable at http://silviamirandaagrippino.com/code-data.

Figure 1.4: Response of Industrial Production to a Contractionary Monetary Policy Shock - Robustness Checks

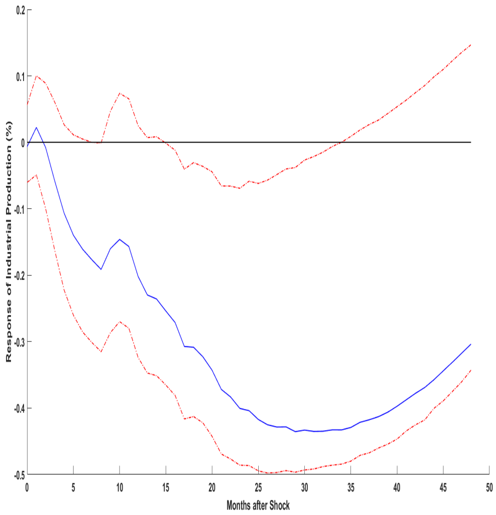

(a) Alternative Ordering

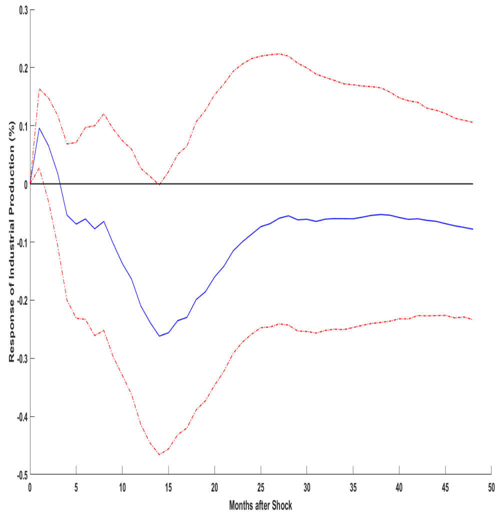

(b) Alternative Shock Identification

These Figures show the response of (aggregate) industrial production to a one standard deviation contractionary policy shock. The left Panel shows the results when the shock measure of Barakchian and Crowe (2013), but no minimum delay restriction is used. The right Panel shows the estimation results when the shock measure of Miranda-Agrippino (2016) is used. In both Panels, the blue line shows the point estimate of the IRF. The red, dashed lines show the 90% Confidence Interval based on 5000 Bootstrap replications.

Figure 1.4 depicts the estimated response of (aggregate) industrial production to a one standard deviation contractionary policy shock estimated in both robustness checks. The estimated responses of the other aggregate variables in the two robustness checks can be found in the Appendix in Figure A.5 (alternative ordering) and Figure A.6 (alternative shock identification), respectively. Relaxing the minimum delay restriction has barely any impact on the estimated response of industrial production, compared to the baseline estimation scheme. Regarding the estimation using the shock measure of Miranda-Agrippino (2016), two things are noteworthy. First, the peak response of industrial production to a contractionary monetary policy shock is reached faster (after approximately 15 months), compared to when the shock measure of Barakchian and Crowe (2013) is used. Second, the magnitude of the peak drop in (aggregate) industrial production in response to a contractionary shock is (comparatively) smaller. The estimated peak in the drop in industrial production in response to a one standard deviation contractionary shock is around 0.25% in this robustness check, smaller than the drop of 0.38% found when using the monetary shock series of Barakchian and Crowe (2013).

Figure 1.5 shows the distribution of sectoral (output) responses to a one standard deviation contractionary monetary policy shock estimated in both robustness checks described here. The left Panel in Figure 1.5 shows the distribution of estimated industry output responses when using monetary policy shocks identified as in Barakchian and Crowe (2013) and not imposing a minimum delay restriction. The right Panel in Figure 1.5 shows the distribution of estimated industry output responses when using monetary policy shocks identified as in Miranda-Agrippino (2016) and the baseline estimation scheme described in Section 1.4 is used.

Figure 1.5: Industry Responses to a Contractionary Monetary Policy Shock - Robustness Checks

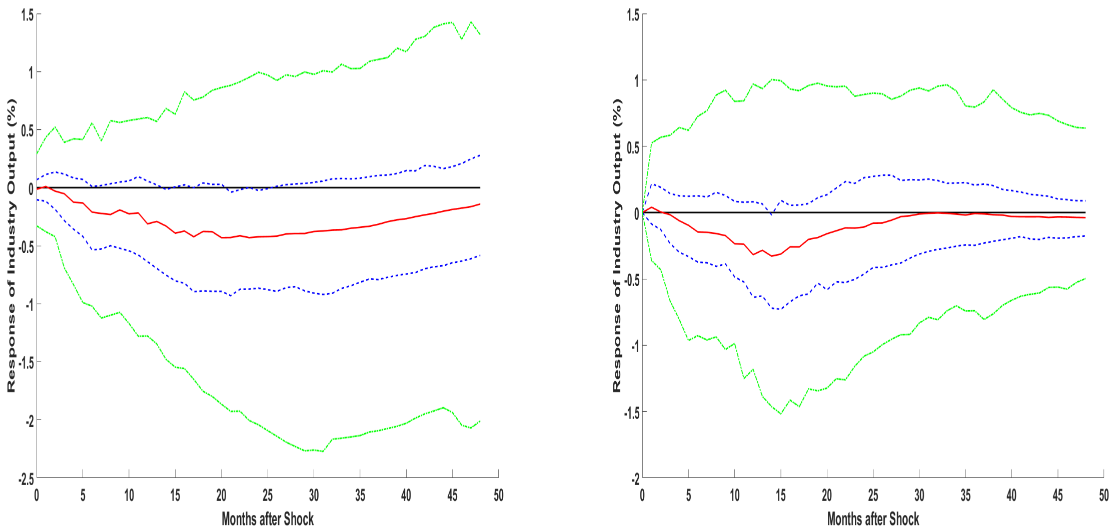

(a) Alternative Ordering **(b)** Alternative Shock Identification

These Figures show the output response of the 205 different industries to a one standard deviation contractionary policy shock. The left Panel shows the results when the shock measure of Barakchian and Crowe (2013), but no minimum delay restriction is used. The right Panel shows the estimation results when the shock measure of Miranda-Agrippino (2016) is used. In both Panels, the red line shows the median response of all industries at each horizon. The blue lines show the 25th and 75th percentile of the distribution of the industries output responses at each horizon. The green lines show the fifth and 95th percentile of the distribution of the industries output responses at each horizon.

Figure 1.5 shows that in both cases there is substantial heterogeneity present in the output response across industries. As before, roughly 25% of industries experience an increase in output in response to a contractionary policy shock in both robustness checks. When identifying monetary policy shocks as in Miranda-Agrippino (2016), industry output responses are smaller in magnitude and the peak response is reached more quickly compared to the results when using the shock measure of Barakchian and Crowe (2013), mirroring to the difference in the estimated reaction of aggregate industrial production across both estimation schemes described above. Otherwise, industry responses are very similar across the different estimation schemes discussed here. The cross-sectional correlation of industry responses between the baseline estimation in Section 1.4 and the two estimation schemes presented here is $\rho = 0.98$ (alternative ordering) and $\rho = 0.59$ (different shock identification) at a horizon of $h = 24$ months, respectively.

The regression results obtained in the two robustness checks when including no other control variables (i.e. as described in Section 1.5) can be found in Table 1.6. Consistent with the results reported in Section 1.5, I report the estimates using the cross-section of industry responses $h = 18$, 24 and 30 months after a contractionary policy shock. Columns labeled (1), (2) and (3) in Table 1.6 report the results of the alternative ordering robustness check. Columns labeled (4), (5) and (6) in Table 1.6 report the results of the alternative shock identification robustness check.

Table 1.6: Price Stickiness and Monetary Policy Responses - Robustness Checks

	(1) $h = 18$	(2) $h = 24$	(3) $h = 30$	(4) $h = 18$	(5) $h = 24$	(6) $h = 30$
Log(FPA)	0.255*	0.364**	0.357**	0.573***	0.361***	0.0620
	(0.152)	(0.158)	(0.147)	(0.169)	(0.106)	(0.0780)
Observations	205	205	205	205	205	205
R-squared	0.017	0.030	0.033	0.131	0.066	0.002
$I\bar{R}F^h$	-0.308	-0.404	- 0.433	-0.199	-0.087	- 0.061

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

This Table reports the coefficient on log(FPA) in regression 1.11 at different horizons h after the shock. The constant term is not reported. Columns labeled (1), (2) and (3) report the results of the alternative ordering robustness check. Columns labeled (4), (5) and (6) report the results of the alternative shock identification robustness check. $I\bar{R}F^h$ denotes the response of the (total) industrial production index h months after an unexpected one standard deviation increase in the policy measure, estimated in the same robustness check. Robust standard errors are reported in Parenthesis.

The results obtained in the alternative ordering robustness check (reported in the Columns labeled (1), (2) and (3) of Table 1.6) are nearly identical to the baseline results reported in Table 1.1. The results obtained when identifying monetary policy shocks as in Miranda-Agrippino (2016) (reported in the Columns labeled (4), (5) and (6) in Table 1.6) suggest a stronger association between the frequency of price adjustment and the output response to monetary policy shocks, compared to the baseline results. In this robustness check, the estimated coefficient on the frequency of price adjustment is larger, while the drop in industrial production is smaller.[25]

Table 1.7: Price Stickiness and Monetary Policy Responses - Additional Controls in Robustness Checks

	(1) $h = 18$	(2) $h = 24$	(3) $h = 30$	(4) $h = 18$	(5) $h = 24$	(6) $h = 30$
Log(FPA)	0.547**	0.653***	0.590**	0.717**	0.530***	0.208*
	(0.252)	(0.250)	(0.241)	(0.306)	(0.180)	(0.120)
Observations	186	186	186	186	186	186
R-squared	0.115	0.149	0.155	0.203	0.131	0.100
Add. Controls	Y	Y	Y	Y	Y	Y

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

This Table reports the coefficient on log(FPA) in regression 1.12 at different horizons h after the shock. The constant term is not reported. Columns labeled (1), (2) and (3) report the results of the alternative ordering robustness check. Columns labeled (4), (5) and (6) report the results of the alternative shock identification robustness check. In all results the full set of industry-level control variables described in Section 1.6 is included. Robust standard errors are reported in Parenthesis.

Table 1.7 reports the results obtained in the respective robustness checks when con-

[25]When comparing the results obtained in the different shock identification robustness check with the baseline results, it should also be kept in mind that when using the identification scheme of Miranda-Agrippino (2016), the peak in the drop of industrial production is reached faster. For the sake of comparing the size of the coefficients obtained in this (alternative shock identification) robustness check to the baseline results, it is also helpful to consider the standardized regression coefficient obtained from regression 1.11 (i.e. standardizing all variables to have mean zero and a standard deviation of one before estimating the regression). The standardized regression coefficients obtained in the baseline results are $\beta^{h=18} = .123$ and $\beta^{h=24} = .175$ for $h = 18$ and $h = 24$ months, respectively. The interpretation of e.g. $\beta^{h=24} = .175$ is that a one standard deviation increase in the log of the frequency of price adjustment is associated with a 0.175 standard deviation increase in the output reaction to a contractionary monetary policy shock at a horizon of 24 months after the shock (in the baseline results reported in Section 1.5). In the alternative shock identification robustness check, the standardized regression coefficients are estimated at $\beta^{h=18} = .361$ and $\beta^{h=24} = .257$ at the two horizons, respectively. Using this standardized measure, the (relative) size of the association between the frequency of price adjustment and the output reaction to monetary policy shocks is approximately twice as large in the alternative shock identification robustness check, compared to the baseline result.

trolling for all additional industry-level variables described in Section 1.6.[26] These results confirm the finding established in Section 1.6: When controlling for additional industry characteristics, the correlation between the frequency of price adjustment and an industry's output response to monetary policy shocks becomes larger compared to the case where no additional industry-level variables are considered.

Both robustness checks confirm a significant association between the frequency of price adjustment and the output response to monetary policy shocks at the industry level. The results obtained in the robustness checks suggest that the association between the frequency of price adjustment and the output response to monetary policy shocks at the industry level is rather under- than overestimated in the main specification. This alleviates concerns that the established association is spuriously caused by the estimation method, the identification of monetary policy shocks or by omitted industry characteristics. The results suggest that sticky prices indeed play a central role the monetary transmission mechanism, consistent with the predictions of (multi-sector) New Keynesian models.

1.8 Conclusion

Do sticky prices matter for the monetary transmission mechanism? This Chapter provides new evidence on this question by studying the role of price stickiness for the monetary transmission mechanism using disaggregated industry-level data from 205 US manufacturing industries. The output reactions of different industries to a (common) contractionary monetary policy shock differ substantially. Two years after a one standard contractionary monetary policy shock, (total) industrial production is estimated to drop by approximately 0.38%. Some industries experience a drop in output as large as 2%, while other industries even increase output by 0.9% in reaction to the same shock, at the same horizon.

I show that an industry's output response to monetary policy shocks is systematically related to the industry's degree of price stickiness. Price stickiness is measured via the industry-level frequency of price adjustment, calculated from PPI microdata. Industries with a higher frequency of price adjustment (i.e. less sticky prices) experience a smaller drop in output than industries with a lower frequency frequency of price adjustment (i.e. more sticky prices) in reaction to the same contractionary monetary policy shock. The

[26] Table A.3 (alternative ordering) and Table A.4 (alternative shock identification) in the Appendix report the coefficients on all variables included in the estimation.

association between an industry's frequency of price adjustment and the output reaction to monetary policy shocks is statistically significant and the size of the differential reaction is economically relevant. A 10% increase in the frequency of price adjustment is associated with a reduction in the output drop in response to a contractionary monetary policy shock that is approximately as large as 10% of the drop in total industrial production. This result is robust to the inclusion of various industry-level control variables, intended to capture alternative transmission channels of monetary policy.

Qualitatively, the results established in this paper are consistent with predictions of multi-sector New Keynesian models. Quantitatively, the results provide empirical support for the New Keynesian view that sticky prices indeed play an important quantitative role in the transmission of monetary policy to real economic activity. Sticky prices indeed matter for the monetary transmission mechanism. While the association between an industries degree of price stickiness and the reaction to monetary policy shocks documented in this Chapter provides no direct evidence on the degree of aggregate monetary non-neutrality, the results established in this paper can be used to discipline multi-sector New Keynesian models to provide new evidence on this classical question.

Chapter 2

The Geography of Intergenerational Mobility in Germany

2.1 Introduction

Intergenerational mobility is a key indicator for both fairness and economic efficiency in a society. When intergenerational mobility is low, a child's chances of success in life are largely pre-determined by their parents, which goes against a basic notion of fairness in most societies. Also, low mobility can lead to misallocation of resources, as talented children from disadvantaged families are excluded from opportunities that favor those born in more fortunate circumstances rather than those with the greatest potential. Despite the overwhelming importance of intergenerational mobility, and recent progress in the empirical literature, empirical evidence on the extent and drivers of intergenerational mobility is still rather thin. In this Chapter we contribute to the literature on intergenerational mobility by providing a comprehensive description of intergenerational mobility in Germany. We are the first to describe differences in intergenerational mobility in Germany at a detailed regional level.

The institutional setting in Germany makes it a particularly interesting case to study intergenerational mobility. Historically, Germany has had smaller levels of income and wage inequality than the US, the UK, or many southern European countries. While Germany has seen a notable uptick in inequality in the last 30 years, this increase has been much less pronounced compared to other countries (Fuchs-Schündeln et al. (2010)). Moreover, Germany is often portrayed as a country with particularly inclusive labor market institutions, as indicated by the fact that Germany has Europe's lowest youth unemployment rate. These considerations suggest that in Germany inequality is rather low and that equality of opportunity and intergenerational mobility should be high. But the public perception in Germany is a different one[1]. The public debate seems to have converged on the view that equality of opportunity is "too low" in Germany, and that the German schooling system is failing the disadvantaged and is exacerbating differences in family background instead of "leveling the playing field"[2]. The practice of tracking students at the relatively early age of 10 or 12 is often cited as a major culprit of this, and many pundits believe that it is a major impediment to intergenerational mobility in Germany, although academic studies have, so far, not detected negative causal long-term effects of tracking (Dustmann et al. (2017)).

[1] See, for example, https://www.welt.de/politik/deutschland/article113082544/In-Deutschland-bleibt-Chancengleichheit-ein-Traum.html.

[2] See, for example, http://www.tagesspiegel.de/politik/chancenungleichheit-welche-chance-hat-ein-kind-in-deutschland/9556636.html.

Is Germany a country of high intergenerational mobility and equality of opportunity? In this Chapter, we present a new set of facts on this question. We present new empirical evidence on the level and the geography of intergenerational mobility in Germany using administrative microcensus data from the years 2009 to 2015. This data contains detailed information on the educational attainment of 268.000 children aged 16 to 22 and their parents socioeconomic status. In our main approach, we measure intergenerational mobility by the association between the educational attainment of a child and her parents income.

We focus on the educational attainment of children in our measure of intergenerational mobility as the particular institutional setting of the German education system makes educational attainment in secondary school a strong measure of social status itself and is a powerful predictor of economic and social outcomes in later life. Furthermore, the education system is generally seen as one of the most important factors in shaping intergenerational mobility. Focusing on the educational attainment of children allows us to measure intergenerational mobility for very recent cohorts (the youngest children in our sample are born in 1999), providing a very timely measure of intergenerational mobility. Our main child outcome variable is hence a binary variable, indicating whether a child has obtained an A-Level degree (”Abitur”) (or is in the process of obtaining this degree) or not.

We focus on income as main parental characteristic to parsimoniously summarize the economic status and available resources of parents in our main analysis, but consider other parental characteristics, in particular parental education, in our analysis as well. Following a series of recent papers (e.g. Chetty et al. (2014) and Acciari et al. (2019)), we rank parents based on their incomes relative to other parents in our sample and use the percentile rank of parents in the national income distribution within our sample to measure of parental income.

We find that the relationship between a child's probability of obtaining an A-Level degree and the parents income rank is well described by a linear function at the national level. This allows us to summarize intergenerational mobility at the national level by the slope coefficient of a linear regression of the A-Level dummy of children on the parents percentile rank. We refer to this slope coefficient (multiplied by 100) as parental income gradient. Based on the linear fit, we find that a 10 percentile point increase in parent income rank is associated with a 4.8 percentage point increase in the probability of obtaining an A-Level degree for children.

To provide a comprehensive description of intergenerational mobility in Germany, we analyze differences in intergenerational mobility between different groups at the national level. We find no pronounced differences in the parental income gradient between girls and boys. We also do not find pronounced differences in parental income gradient between children with different parental education levels. However, we find pronounced differences in upward mobility between these groups. We measure upward mobility by the probability of obtaining an A-Level degree for a child with parents from the bottom quintile of the national income distribution. We refer to this measure as (absolute) upward mobility, or $Q1$ measure. Girls have a nine percentage point higher $Q1$ measure than boys and children from households where at least one parent has an A-Level degree have a thirty percentage point higher $Q1$ measure compared to children from households where neither parent has an A-Level degree. When comparing migrants and natives, we find that the parental income gradient is lower for migrants compared to natives, and that migrants have higher $Q1$ measure than natives. Lastly, we find that the parental income gradient is higher and upward mobility is lower in East Germany compared to West Germany.

Next, we characterize variation in intergenerational mobility across local labor markets (LLM) in Germany. Local labor markets are aggregations of German counties, defined on the basis of the distribution of local economic activity. We assign children to local labor markets based on their place of residence at the time they are observed in our sample. We calculate different measures of absolute and relative intergenerational mobility, intended to capture different normative concepts, for every local labor market in our sample to compare mobility across local labor markets. First, we estimate the parental income gradient, defined as above, separately for every LLM in our sample. This lets us compare (absolute) differences in the outcomes between children from top versus bottom income families within a LLM and serves as a first measure of relative intergenerational mobility. Second, we calculate the relative likelihood of obtaining an A-Level degree between children from families in the top quintile of the national income distribution and children from families in the bottom quintile of the national income distribution for every local labor market as an additional measure of relative mobility. We refer to this measure as $Q5/Q1$ ratio. The $Q5/Q1$ ratio lets us compare relative differences in outcomes between children from rich and poor families, taking potential differences in the baseline probability of obtaining an A-Level degree across local labor markets into account. Furthermore, we calculate the probability of obtaining an A-Level degree for children with parents from the bottom quintile of the national income distribution, the $Q1$ measure, separately for every LLM. This lets us compare the (absolute) outcomes of children from poor families across LLMs and serves as a measure of absolute mobility.

We find substantial variation in our measures of absolute and relative mobility across LLMs. Relative mobility, summarized by the parental income gradient, is lowest for children in East Germany, especially in Saxony and Thuringia. The parental income gradient varies roughly between 0.1 (meaning that a 10 percentile point increase in the parental income rank is associated with a 1 percentage point increase in the probability of obtaining an A-Level degree) and 0.7 (meaning that a 10 percentile point increase in the parental income rank is associated with a 7 percentage point increase in the probability of obtaining an A-Level degree) across LLMs in Germany.

We also find pronounced differences in relative mobility as measured by the $Q5/Q1$ ratio. The $Q5/Q1$ ratio varies roughly between 2, meaning that in some local labor markets children from the top quintile of the national income distribution are twice as likely to obtain an A-Level degree than children from the bottom quintile of the national income distribution, and 7, meaning that in other local labor markets they are seven times as likely. The $Q5/Q1$ ratio is especially high in Bavaria and Thuringia, and low in North Rhine-Westphalia.

We find substantial variations in absolute mobility as well. The $Q1$ measure[3] varies between 5% and 40% across LLMs, revealing substantial differences in absolute mobility across LLMs. Absolute mobility is especially low in Bavaria, and high in North Rhine-Westphalia.

Absolute and relative mobility are only mildly correlated across LLMs. While on average LLMs with high absolute mobility (measured by the $Q1$ measure) are also LLMs of high relative mobility (measured by the parental income gradient), the correlation between both measures across local labor markets is rather weak. There exists a substantial number of LLMs described by low relative mobility (measured by a high parental income gradient), but high absolute mobility (as measured by the $Q1$ measure), and vice versa. This highlights the importance of the normative concept (and hence mobility measure) one uses to compare mobility across local labor markets.

Our results are robust across different sample specifications, alternative measures of

[3]Which is the probability of obtaining an A-Level degree for children with parents from the bottom quintile of the national income distribution.

parental income and alternative measures of intergenerational mobility. We find that our parental income-based measures of intergenerational mobility generally exhibit a high correlation with measures of intergenerational mobility that are based on parental education levels.

We document substantial and robust geographical variation in different measures of intergenerational mobility across LLMs in Germany. The next question we address is which factors drive the observed differences in intergenerational mobility. We investigate to what extent observed differences in intergenerational mobility measures across LLMs can be explained by differences in other (non-income) household characteristics across LLMs within our sample, for example differences in the education level of parents across LLMs. We show that explicitly controlling for additional household characteristics in the estimation of LLM specific mobility measures has barely any effect on the order of LLM-specific estimates of intergenerational mobility we obtain. The correlation between the LLM-specific estimates of mobility we obtain in our baseline approach and the LLM-specific mobility estimates we obtain when controlling for additional household characteristics in the calculation of the mobility measures is substantial. This suggests that differences in intergenerational mobility across LLMs are not simply driven by differences in household characteristics across LLMs and that neighborhood effects play a role in shaping intergenerational mobility. Our approach is complementary to the approach of Chetty and Hendren (2018), who use children moving across geographical units to identify neighborhood effects.

Lastly, we take a first step at describing which local labor market characteristics correlate with intergenerational mobility. We do not claim that these correlations should be interpreted as a causal relationship, but they can help to guide future research in the search for causal determinants of intergenerational mobility. We find that intergenerational mobility differs between urban and rural local labor markets in our sample. Furthermore, our results suggest that labor market variables (within an LLM) affect intergenerational mobility. The results also suggest that (broadly defined) social capital and social cohesion might affect intergenerational mobility.

The remainder of the Chapter is structured as follows. Section 2.2 presents a detailed overview over the related literature. Section 2.3 presents an overview over the institutional background of the German education system. Section 2.4 presents the data used in the analysis and defines the variables used in the analysis. Section 2.5 presents and discusses different measures of intergenerational mobility. Section 2.6 presents the results of our analysis at the national level. Section 2.7 compares intergenerational mobility between different groups. Section 2.8 presents the results for regional differences in intergenerational mobility across local labor markets. Section 2.9 investigates to what extent differences in household characteristics across LLMs can explain geographical variation in our measures of intergenerational mobility. Section 2.10 presents local labor market characteristics correlated with our measures of intergenerational mobility. Section 2.11 concludes.

2.2 Related Literature

At the most general level, intergenerational mobility describes to which extent a child's social or economic opportunities depend in her parent's characteristics, like social status, education level or income. As opportunities are hard to measure, empirical studies of intergenerational mobility resort to the study of how a child's outcomes are associated with parental characteristics (Chetty et al. (2014)). A popular measure of intergenerational mobility is the degree of intergenerational income mobility, the association between a child's income and her parents income. A further widely used measure of intergenerational mobility is intergenerational educational mobility, the association between the educational attainment of parents and the educational attainment of their children (Black and Devereux (2011)). However, the study of intergenerational mobility is not limited to studying the association between the same outcome for both, parents and children. As described below, there are also studies investigating e.g. the association between parental income and the educational attainment of children.

A comparatively rich literature exists on the measurement of intergenerational mobility in the United States. Much of this literature is interested in measuring the degree of intergenerational income mobility, but other measures of intergenerational mobility are studied as well.

Two early seminal paper in this literature are Zimmerman (1992) and Solon (1992). Both papers present estimates of the degree of intergenerational income mobility in the United States. In both papers, intergenerational income mobility is measured via the intergenerational elasticity of earnings (hereafter IGE). Both papers find estimates of the IGE coefficient of around 0.4, meaning a one percent increase in parental income is associated with a 0.4 percent increase in the income of the child.

Mazumder (2005) suggests that previous estimates of the IGE in the United States have been biased downwards due to various measurement issues. Correcting for these measurement issues Mazumder (2005) estimates the IGE in the United States to be around 0.6, substantially higher than previous estimates suggested.

Dahl and DeLeire (2008) estimate measures of intergenerational mobility using administrative earnings data from the US Social Security administration records. They show that estimates of the IGE are quite sensitive to the exact specification and sample used and can range from 0.26 to 0.63 using US data. They propose a new measure of intergenerational income mobility, the association between fathers' and children's relative positions (ranks) in their respective earnings distributions. This measure is summarized by the regression of the (relative) rank of the child on the (relative) rank of the parents in their respective income distributions. This rank-rank slope is much more robust to specification and sample choices than the IGE. Dahl and DeLeire (2008) estimate a rank-rank slope of roughly 0.3 for males, meaning that a 10 percentile point increase in a father's relative position is associated with roughly a 3 percentile point increase in his son's relative earnings position.

Using administrative tax return data, Chetty et al. (2014) confirm the findings of Dahl and DeLeire (2008). The reported estimates of the IGE obtained from tax return data in Chetty et al. (2014) range from 0.26 to 0.7, covering the whole range of estimates reported in the previous literature. The rank-rank slope is much more robust to the exact specification and is estimated to be 0.34, rather similar to Dahl and DeLeire (2008). Chetty et al. (2014) are the first to document substantial geographical variation in the rank-rank slope across regions in the US. Different regions in the US are described by different degrees of intergenerational mobility, with the rank-rank slope ranging from 0.07 to 0.5 across regions. Similar geographical variation is found in a measure of absolute mobility, the expect income rank of a child from the 25 percentile of the parental income distribution.

Furthermore, Chetty et al. (2014) find a strong association between educational outcomes of children (measured e.g. by the probability to attend college at the age of 19) and the parental income rank. They refer to the slope coefficient on the parental income rank in a regression of a dummy variable indicating whether a child attends college at the age of 19 on the parental income rank as college attendance gradient. At the national level, this college attendance gradient is estimated to be 0.675, meaning that a 10 percentile point increase in the parental income position is associated with an increase in the probability of college attendance of 6.75 percentage points. This college attendance gradient exhibits substantial geographical variation as well, varying between 0.3 and 0.9 across US commuting zones. The correlation between the college attendance gradient and the rank-rank slope for income is 0.68 (0.72 population weighted) between commuting zones, indicating a tight link between income mobility of children and their educational outcomes.

Chetty and Hendren (2018) study the causal effect of the place of residence (which they refer to as neighborhoods) on the intergenerational mobility of children by comparing the later life outcomes of children moving to different neighborhoods at different ages. They find that the regional variation in intergenerational mobility documented in Chetty et al. (2014) is largely due to the causal effect of places rather than differences in the characteristics of residents across regions. Furthermore, they find that places matter for intergenerational mobility largely due to differences in childhood environment, rather than the differences in labor market conditions in later life. Lastly, they find that every year spent in a place as a child matters roughly equally.

Chetty et al. (2017) study the role of colleges for intergenerational mobility in greater detail. Access to college is highly associated with parental income. However, children from low- and high-income families have similar earnings outcomes conditional on the quality of the college they attend, indicating the importance of educational attainment in shaping intergenerational mobility.

Compared to the US, research on the degree of intergenerational mobility in Germany is more scarce. Schnitzlein (2016) compares intergenerational income mobility between Germany and the US, based on comparable survey data sets from both countries. He estimates intergenerational income mobility via the IGE as described above. The IGE coefficient is estimated to be lower in Germany compared to the US (meaning that intergenerational income mobility is higher in Germany). However, Schnitzlein (2016) notes that even a reasonable degree of variation in sampling rules leads to similar estimates in both countries, preventing a clear conclusion which country exhibits a higher degree of intergenerational mobility.

Bratberg et al. (2017) compare intergenerational income mobility between Germany, Norway, Sweden and the US. Bratberg et al. (2017) measure intergenerational income mobility using rank-rank slopes. The rank-rank slope for Germany is estimated at 0.245, comparable to rank-rank slopes in Norway and Sweden and lower than in the US. However, results for

Germany are based on only around 1000 children born between 1957 and 1976. No difference in rank-rank slopes is found between North- and South Germany.

More evidence exists on the degree of intergenerational educational mobility in Germany.

Riphahn and Heineck (2009) investigate intergenerational educational mobility in Germany over 50 years, for the birth cohorts from 1929 to 1978. Intergenerational educational mobility is measured via the extent to which children's secondary school attainment depends on parental characteristics, in particular parental education. The authors find no evidence of a significant change in the role of the parental educational attainment for child educational outcomes over the last decades.

Riphahn and Trübswetter (2013) extent the study of Riphahn and Heineck (2009) to more recent cohorts (with the youngest cohort being born in 1987) and compare intergenerational educational mobility between East and West Germany. They find that intergenerational educational mobility is lower in East Germany than it is in West Germany.

Klein et al. (2019) compare educational mobility between East- and West Germany for birth cohorts between 1929 to 1992. They measure educational mobility as the probability of a child of obtaining an A-Level degree (Abitur) conditional of parental education level. They find no pronounced differences in educational persistence between East- and West Germany for Birth cohorts since 1987. For cohorts born since 1987, they find that the probability to obtain an A-Level degree is approximately 40% higher if at least one parent has an A-Level degree herself (compared to children where neither parent has an A-Level degree) in both East- and West Germany.

2.3 Institutional Background

We focus on the educational attainment of children conditional on parental characteristics (in particular parental income) as our measure of intergenerational mobility. The particular institutional setting of the German education system makes educational attainment in secondary education a well-suited measure to study intergenerational mobility. This allows us to measure intergenerational mobility for very recent cohorts (the youngest children in our sample are born in 1999), providing a (comparatively) timely measure of intergenerational mobility. Furthermore, children typically still live at home while completing secondary education, which allows us to use (micro)census data to measure intergenerational mobility, where we can link children and parents living in the same household. This section explains the institutional setting of the German education system and provides evidence on the relevance of our measure.

2.3.1 The German Education System

This section provides an overview over the secondary schooling system in Germany. The overview presented here follows Dustmann et al. (2017) and describes the general structure of the secondary schooling system in Germany.

Before describing the general structure, it should be noted that the responsibility for the education system falls under the jurisdiction of the German states, not under the Federal Governments jurisdiction. The general structure of the education system is the same in all German states, but there are differences in the exact institutional details between German states.

For the first 4 years of school (age 6 to 10)[4] all children attend primary school. After primary school, students are allocated to one of three different tracks, namely Hauptschule, Realschule and Gymnasium. The three tracks lead to different degrees, described below. Following Dustmann et al. (2017), we refer to these tracks as "low", "medium" and "high" track.

There is no strict rule to determine which track children can attend after elementary school, but elementary school teachers do make recommendations. In most states parents have the final decision which track their child should attend following primary school. In a few states, track choice is limited by the marks of the child in the final year of primary school.

Education in the low and medium track lasts five, respectively six years (until year 9/10), whereas the high track lasts eight or nine years (until year 12/13, dependent on the birth cohort and the federal state[5]). The low and medium track prepare students for vocational training in blue-collar (e.g. crafts) and white-collar (e.g. medical assistant or office clerk) occupations. The high track is comparatively more academically oriented and prepares

[4]In the states of Berlin and Brandenburg primary school includes the first 6 years of school (age 6 to 12).

[5]An overview over the different regulation within different German states can be found on the website of the Standing Conference of State Education secretaries (Kultusministerkonferenz) available at https://www.kmk.org/themen/allgemeinbildende-schulen/bildungswege-und-abschluesse/sekundarstufe-ii-gymnasiale-oberstufe-und-abitur.html

students for tertiary education at academic institutions. Successful completion of the high track results in the award of an A-Level degree (Abitur).

Furthermore, there exists a type of school that does not follow the tracking system ("integrierte Gesamtschule"). At this type of school, all students attend the same classes together until year 10, but the level of classes is possibly differentiated on a subject-by-subject basis.[6] Important for our analysis, this type of school is only non-tracking until grade 10. At these schools, if students stay in school after year 10, they can only attend classes at the level of the high track ("Gymnasiale Oberstufe"). Students attending such a school after year 10 do so with the goal of graduating with an A-Level degree, same as students at a regular high track school.

The German schooling system allows students to switch tracks. After completing the low or medium track, students have the possibility to upgrade their degree and attend the next higher track. A student graduating from the low track can stay in school one more year and complete the medium track. Similarly, a student completing the medium track can switch to a regular high track school and obtain an A-Level degree. Furthermore, after year 10, students have the possibility to switch to a specialized high track. This specialized high track leads to an A-Level degree as well, but possibly limits the field of study ("Fachgebundene Hochschulreife"). A comprehensive analysis of track switching in the German schooling system can be found in Biewen and Tapalaga (2017).

For the purpose of our analysis, it is important that a student who is enrolled in a regular school[7] after year 10 (enrolled in school years 11 to 13) is studying with the goal of obtaining an A-Level degree.

Due to the fact that educational policy in Germany falls in the jurisdiction of the German states, there are state-level differences in the education system. At the stage of years 11 to 13 in the high track, the state-level regulations differ e.g. in the type of classes students need to attend (e.g. some states require students to take 4 weekly hours of math, whereas students can choose between 3 or 5 hours weekly hours of math in other states). However, the Standing Conference of State Education secretaries ("Kultusministerkonferenz") has the stated goal to ensure a high degree of comparability of educational qualifications across German States. There are no legal differences between the A-Level degrees issued from dif-

[6]So a child can e.g. attend English classes at the level of the low track and at the same time math classes at the level of the high track.

[7]Here meaning not enrolled in a special education school, "Förderschule", or enrolled in a vocational school

ferent German States, all states legally recognize A-Level degrees awarded by other German states. Furthermore, differences between different school tracks are more pronounced than differences between the same track across different states. Dustmann et al. (2017) report that 80% of the variation in school-level mean test scores in ninth grade can be explained by track choice alone, suggesting that schools of the same track are fairly homogeneous, even across states.

However, there exist state-level differences in the institutional setting. To our knowledge, there exists no comprehensive academic study investigating whether the return to an A-Level degree depends on the state where the degree was obtained. Furthermore, we are also not aware of any study comparing standardized test scores of children enrolled in grades 11 to 13 across German states (comparable studies are limited to earlier grades, e.g. grade 9, as noted above). We present a first step to compare A-Level degrees across states in Section B.1 in the Appendix. In Section B.1 in the Appendix, we calculate state-specific A-Level wage premia for full-time workers based on the current place of residence of the worker. We find that there is a substantial A-Level wage premium in every state, and that differences in the A-Level wage premium across states are rather small, compared to the size of A-Level wage premium.

2.3.2 Benefits of an A-Level degree

The particular institutional setting in the German education system described above makes the graduation from the high-level track a meaningful outcome to measure intergenerational mobility of children. First, graduation from the high track grants access to the (free) national university system in Germany, offering students who graduate from this track more educational and occupational choices than students graduating from the other tracks.

Also, individuals with an A-Level degree have substantially better economic outcomes than persons without an A-Level degree on average. There is a substantial wage premium associated with an A-Level degree. Using data on full-time workers aged 30 to 45 in the waves 2009 to 2014 in the German Mikrozensus (described below), we find an A-Level wage premium of around 44% for monthly net income. Details on the calculation can be found in the Appendix in Section B.1. In Section B.1, we also present evidence that state-level differences in the A-Level wage premium are rather small, especially compared to the size of the A-Level wage premium. Furthermore, Schmillen and Stüber (2014) report a substantial

A-Level premium for total (gross) lifetime earnings. The total lifetime earnings of individuals with an A-Level degree (but no university degree) are estimated to be 44% higher than the total (gross) lifetime earnings of individuals without an A-Level degree (and no vocational training). This figure becomes even higher if individuals additionally have obtained a college degree[8], which usually requires an A-Level degree first.[9]

An A-Level degree is not only associated with higher income levels, but also with a lower risk of being unemployed. In the year 2013, the unemployment rate of persons with an A-Level degree was about only half as high as the unemployment rate of persons with a medium track degree (5.6% vs 10%), and much lower than the unemployment rate of people with a low track degree (17.6%), see Hausner et al. (2015).

An A-Level degree is also a beneficial factor for obtaining vocational training in certain (white-collar) occupations, like bank clerk (Klein et al. (2019)). Furthermore, an A-Level degree associated with a higher life expectancy (Gärtner (2002)). As a side note, we also note that anecdotal evidence suggests that an A-Level degree is still highly regarded within the German public. The fact that the German politician Martin Schulz, former front-runner of the German Social Democratic Party, never obtained an A-Level degree sparked a substantial debate in the German media[10], indicating that among certain groups in Germany, an A-Level degree is viewed as an important mark of distinction.

All of this indicates that an A-Level degree is indeed a meaningful measure of economic success. Finally, one should not forget that many people hold the firm belief that education has substantial worth in itself, regardless of the economic benefits associated with education.

2.4 Data

2.4.1 Mikrozensus

The primary dataset used in our analysis is the German microcensus (Mikrozensus, hereafter MZ), an annual representative survey of the German population, administered by the statistical office of Germany (Destatis). The survey is comparable to, but more detailed

[8]University of "Fachhochschulabschluss"

[9]Schmillen and Stüber (2014) further report that total (gross) lifetime earnings of individuals with an A-Level degree (but no university degree) are estimated to be 17.8% higher than lifetime earnings of individuals with vocational training, but no A-Level degree. Lifetime Earnings of University graduates are estimated to be 75% higher than lifetime earnings of individuals with vocational training, but no A-Level degree.

[10]see e.g. http://www.spiegel.de/politik/deutschland/martin-schulz-debatte-um-abitur-ungehoerig-kommentar-a-1124658.html

than, the American Community Survey. The MZ is the largest survey program of official statistics in Germany. In West Germany the first MZ was administered in 1957, in East Germany in 1991. Results of the MZ are widely used in official statistics in Germany and in various EU-wide surveys. In our analysis, we pool the waves 2009 to 2015 of the MZ.

Every year approximately 1% of the population living in Germany is randomly selected to participate in the survey. In recent years, around 800,000 individuals participate in the survey every year. The sampling units are residential properties in Germany, rather than individuals. Sampling is based on the registry of all addresses in Germany. The sampling scheme is designed to ensure that the sample is representative of the German population. If a dwelling is selected to be included in the survey, all persons living in this particular location are legally required to participate in the survey. Once selected, participation is mandatory and refusal to participate can result in fines or even incarceration. An overview over the MZ can be found on the website of the German federal statistical office[11].

The questions in the MZ cover a wide range of topics, including family status, citizenship, living conditions, labor market status, income and educational attainment.

Once a dwelling is selected to be included in the MZ, it will remain in the survey for 4 subsequent years, before it is rotated out. Hence the MZ is a revolving panel, with roughly one quarter of respondents changing from year to year. However, the MZ does not contain pre-defined panel identifiers, preventing us from tracking the same individuals over time.

2.4.2 Linking Children and Parents

If an address is selected to be included in the MZ, information on all household members, including underage children, must be provided. A set of questions in the survey is designed to describe the (family) relationship between all household members. These questions allow us to link parents and children as long as they are living in the same household.

Specifically, the MZ asks whether the father or the mother of the respondent is living in the same household. We identify children living with their parents as individuals who state that they live in the same household as at least one of their parents.[12] Furthermore, the data also allows to link the responses of children to the responses provided by their parents.

[11] Under this link:https://www.destatis.de/DE/Methoden/Qualitaet/Qualitaetsberichte/Bevoelkerung/mikrozensus-2017.html.

[12] Parents in the sense of the MZ are not only biological parents, but also step-, adoptive and foster parents.

The format of the data imposes two constraints on our analysis. First, we are able to match parents and children as long as they live in the same household. However, we are not able to link children who moved out of their parental household to their parents. Second, we observe parents and children at the same point in time. Hence we are not able to use lagged parent or household characteristics as explanatory variables in our analysis.

Table 2.1 denotes the share of children living with their parents by age of the child, calculated from our data. Practically all children younger than 15 still live with their parents and are not reported in the Table. We can clearly see that the share of children living with their parents (i.e. children we can link to their parents) is decreasing with child age. While 83% of 19 year olds are living with their parents, only 36% of 24 year olds are living with their parents.

Table 2.1: Share of Children Living With Their Parents by Child Age

Child Age	Share living with parents
15	0.987
16	0.978
17	0.962
18	0.921
19	0.831
20	0.720
21	0.615
22	0.525
23	0.439
24	0.363

This Table reports the fraction of children living with their parents among all children in the waves 2009 to 2015 of the German microcensus by age of the children.

2.4.3 Educational Attainment of Children

The primary measure for educational attainment of children in our analysis is a binary variable that is equal to one if a child has already obtained an A-Level degree or if a child is on track to obtain an A-Level degree, and zero otherwise. Specifically, our child outcome variable is equal to one under two conditions. Our outcome variable is equal to one either if

(i) the child has obtained anA-Level degree[13] or (ii) if the child is currently enrolled in the so-called "Gymnasiale Oberstufe" (school years 11 to 13), which leads to an A-Level degree at the successful completion of school. If not stated otherwise, we refer to this outcome variable as A-Level degree.

We choose this definition for several reasons. First of all, defining our child outcome variable like this allows us to include younger children in our sample. This reduces concerns of bias induced by children moving out of their parents household, which is more severe for older children (see Table 2.1). Furthermore, using this broader definition reduces measurement error in our outcome measure for older children (aged around 18 to 19). A-Level degrees are usually awarded towards the end of the school year (roughly in the second quarter of the calendar year). Answers to the MZ are collected throughout the whole calendar year. Hence we have a non-negligible amount of children in our sample that are surveyed before their graduation, but eventually graduate with an A-Level degree in the same 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 variable for around 40% of the graduating cohort in the survey year.[14] On the other hand the share of children failing the final examination in a given year is very low (around 3 percent on average Germany-wide in 2014)[15]. The way our outcome variable is defined hence helps us to reduce measurement issues associated with the nature of our data. We further discuss this point when discussing our primary sample definition in the following paragraph.

2.4.4 Sample Definition and Summary Statistics

In our primary sample, we restrict our analysis to children aged (including) 16 to 22, living in the same household as at least one parent.[16]

[13]Defined here as Abitur and Fachhochschulreife.

[14]Consider the following example. A child is currently enrolled in the last year of school of the high track. The child graduates from the high track in July, obtaining an A-Level degree. The household of the child is surveyed in the MZ in February, when the child has not graduated yet. Hence, if we only take into account completed A-Level degrees, we would miss-measure the A-Level degree for around 40% of the children graduating from the high track in a given year. This calculation is based on the fact that data collection for the MZ is distributed evenly over the year.

[15]An overview over the share of children failing the final examination can be found on the website of the KMK under this link:https://www.kmk.org/dokumentation-statistik/statistik/schulstatistik/abiturnoten.html

[16]We exclude children living in shared accommodation facilities ("Gemeinschaftsunterkünfte") and children living in households in which at least one household members primary occupation is self-employed farmer. For both types of households, it is not possible to measure household income (for these households, household income is always recorded as zero). For this reason, we exclude households with a reported total

The age cut-offs of 16 and 22 are chosen to balance the following trade-off. For older children, our child outcome variable is measured more precisely, i.e. we don't need to rely on enrollment in the high track, but are more likely to observe the completed degree. However, the share of children living with their parents is decreasing in the age of the child, see Table 2.1, and moving out of the parental household is not random. For younger children, moving out of the parental household is a less severe concern, as a larger share of children is still living with their parents. On the other hand, educational outcomes are measured less precisely for younger children.

To balance this trade-off, we focus on children aged 16 to 22 in our primary sample. The lower cut-off of 16 is chosen as children in school year 11 (the first year of "Gymnasiale Oberstufe") are typically between 16 and 17 years old. The upper cut-off of 22 is chosen as, at this age, still more than 50% of children live in their parental household. Furthermore, the choice of the cut-offs ensures that we observe a sufficient amount of children in every local labor market in our regional analysis (described in Section 2.8).

To alleviate concerns related to the sample composition, for our main results we report robustness checks where we only use children aged 16 to 19, 17 to 20 and 20 to 22. It should be noted that, for the purpose of regional comparisons, the possible bias induced through the sample composition is only a concern if this bias differs across geographical units. The regional variation we find is very similar across samples, indicating that possible bias induced through sample composition is not a severe concern for our results.

We present summary statistics for the children in our main sample in Table 2.2.

net household income of zero. These recorded zeros do not indicate an actual income of zero, but rather that no income measurement is available in our data.

Table 2.2: Descriptive Statistics - Children 16-22

	Children 16-22	
$N = 268523$	Mean	SD
Household Size	3.84	1.16
Share of Children with A-Level	0.403	.49
No. of Children in Household	2.01	0,98
Share of Single Parent Households	0.26	0,44
Total Household Income	3210.68	2343.08
Child Age	18.70	1.94
Share of Females	0.46	0.49
Share Parent A-Level	0.33	0.47
Share Parent College Degree	0.21	0,37
Migrantional Background	0.25	0.43
Fathers Age	50.02	6.03
Mothers Age	47.07	5.26

This Table reports the summary statistics for our main sample, children aged 16 to 22 living with their parents in the waves 2009 to 2015 of the German MZ. The share of children with A-Level is calculated according to the definition described in Section 2.4.3.

Here, we elaborate on the summary statistics presented in Table 2.2. The share of children with an A-Level degree is the fraction of children in our sample who have an A-Level degree according to the definition of A-Level degree described in Section 2.4.3. Slightly more than 40% of children in our primary sample have an A-Level degree (or attend the last 2/3 years of school leading to an A-Level degree). As a plausibility check, we compare this number to a comparable statistic calculated using publicly available data from the German Statistical Office. The A-Level share in our sample is slightly lower than the share of persons aged 20 to 24 with a (completed) A-Level degree[17] in the years 2009 to 2015 calculated using publicly available data. The share of persons aged 20 to 24 with completed A-Level degree is 44.85%[18] when pooling the years 2009 to 2015, slightly higher than the share of children with an A-Level degree in our sample. This difference suggests that we (slightly) underestimate the share of children with an A-Level degree. This is due to several reasons. First, not every

[17] completed "Abitur" or "Fachhochschulreife"

[18] This number is calculated using data from DESTATIS, using Table 12211-0040 ("Bevölkerung (ab 15 Jahren): Deutschland, Jahre, Geschlecht, Altersgruppen, Allgemeine Schulausbildung")in Genesis Online.

16 year old child is already enrolled in school year 11, but might be enrolled in school year 10, even if she attends the high level school track[19]. Including 16 year olds in the calculation of the A-Level share hence leads to an underestimation of the share, even when using our (broader) definition of A-Level. Second, while children might still obtain an A-Level degree later in life, it would be very uncommon for someone to lose an already obtained A-Level degree, so for a given cohort the share of persons with an A-Level degree should increase in the age of the cohort (at least in the age range we consider, where differential mortality should not be a significant issue). Lastly, for older children (from 20 onwards), the A-Level share is slightly higher for children that already moved out of the parental household in our sample. These reasons explain why the A-Level share in our sample is slightly lower than comparable numbers calculated using older children from publicly available data.

After having explained our measure of educational attainment of children in detail, we now turn to other household characteristics of our sample. Around 26% of children live in single-parent households. Around 82% of children living in single parent households live with their mother.

Total household income reported in Table 2.2 is monthly total net household income, excluding the income of all children present in the household. Income here is expressed in 2015 Euro (i.e. adjusted for inflation). Median household income is 2795.93 (not reported in Table 2.2). We describe the measurement of income in the MZ in detail in Section 2.4.5.

Around 33% of children live in a household where at least one parent has obtained an A-Level degree and 21% of children live in a household where at least one parent has obtained a University degree ("Universität" or "Fachhochschulreife"). If not stated explicitly otherwise, we always measure parental education as the highest degree obtained by either parent[20].

Around 25% of children in our sample have migrational background ("Migrationshintergrund"). We define migrational background using the same definition as German official statistics. A person is defined to have migrational background if this person was (i) either born without German citizenship herself or (ii) has at least one parent born without German citizenship. We use this definition to define migrational background as it is commonly used in official German statistics.

[19]The share of 16 year olds with an A-Level according to our definition is 20%.

[20]If, for example, the Mother of a child has a University degree and the Father has completed vocational training, we measure parental education at the child level as a University degree.

2.4.5 Measurement of Parental Income

This Section explains the measurement of parental income in detail. Our primary measure of parental income is the monthly total net household income, excluding the monthly total net income of all dependent children in the household. We refer to this variable as parental income. We focus on parental income as our main characteristic of parents as it allows us to parsimoniously summarize the economic status and available resources of parents. This income measure includes all sources of income, including e.g. labor income, firm profits and social security transfers received. Income in the MZ is always reported in terms of net income, rather than gross income. This means that income is reported net of tax payments and social security contributions (like unemployment insurance or health insurance premia paid by the individual). From here on, income always refers to the monthly total net income (of the household or an individual), if not stated explicitly otherwise. We exclude the income of dependent children from our household income measure in order to focus on parental resources. Income reported in the MZ is always the income in the month preceding the survey, but includes all regular annual payments[21].

Our data contains a continuous measure of household income, which we use as our main income measure. This income measure is not asked directly in the survey, but imputed by the Statistical Office. Imputation is based on the reported personal income of all household members. The imputation procedure is described in Section B.2 in the Appendix.

We are aware that the imputation of income leads to measurement error in our income variable. However, this measurement error should be independent of the educational outcomes of children, and independent of the geographical location of the household. Hence, while the imputation leads to imprecision in our estimates, it should not induce a bias.

We use parental income as baseline income measure in our analysis. The results of our regional analysis are robust to using total household income instead of total parental income[22]. Total household income and total parental income exhibit a very high correlation ($\rho = 0.91$), as most children living with their parents have zero or relatively little income.

Adjusting for Household Size: Households do not only differ in terms of income, but also in terms of household size. We consider three different approaches to adjust for differences in

[21]For example, if respondents receive a 13th monthly salary, or a regular annual bonus, $\frac{1}{12}$ of the amount of this payment should be added to the reported monthly income measure.

[22]The difference between the two measures is that total household income also includes the income of all children present in the household. When using total household income as income measure, instead of total parental income, we find lower parental income gradients at the national level, as expected.

household size. First, we do not adjust for differences in household size and measure income as parental income not adjusted for household size. This approach has been pursued by e.g. Chetty et al. (2014) and Hilger (2015). We then consider two approaches to adjust for differences in household composition. The first approach is to divide parental income by the number of household members n and hence measure parental income per capita. We refer to this income measure as income per household member. Lastly, we follow the literature on equivalence scales (e.g. Pollak and Wales (1979)) and use parental income divided by the square root of the number of household members $\sqrt{n}$ as third income measure. We refer to this income measure as equivalence scale income.

The first approach tends to measure how educational attainment of children is associated with the economic status of the parents, whereas the second and third that approach rather tends to measure how educational attainment of children is associated with available economic resources. Figure B.1 in the Appendix shows the average household size in each percentile of the income distribution (not adjusted for household size) in our primary sample. Figure B.1 shows that when we do not adjust income for household size, households in the bottom half of the income distribution are on average smaller than households in upper half of the national income distribution.

Household Income Percentiles: Finally, following Chetty et al. (2014) and Dahl and DeLeire (2008), we assign every household its percentile rank in the income distribution relative to all other households within our sample, separately for every year. Household percentile rank is hence based on the rank of the household relative to all other households within our sample within the same year. By assigning the percentile ranks within the same year (and doing this separately for every year), we avoid adjusting reported income for inflation[23]. We calculate percentile ranks separately for all three incomes measures discussed above.

Figures B.2, B.3 and B.4 in the Appendix show the empirical cumulative distribution function of the three income measures based on wave 2014 of the MZ to illustrate our approach.

[23] When we first adjust incomes for realized inflation and then calculate percentile ranks using the whole sample, the results we obtain are basically identical to our baseline approach.

2.5 Measuring Intergenerational Mobility: Absolute vs Relative Mobility

Before presenting the results of our analysis, we here present how we measure intergenerational mobility. At the most general level, intergenerational mobility is described by the joint distribution of parent and child characteristics. However, the measurement of intergenerational mobility requires to summarize this joint distribution by a parsimonious set of statistics (Chetty et al. (2014)). Following Chetty et al. (2014), we report different measures of intergenerational mobility. The different measures we report aim at capturing two different normative concepts: relative and absolute mobility.

Relative Mobility: As the name suggests, relative mobility is concerned with the difference in outcomes between children from low-income families compared to children from high-income families. In the context of our analysis, relative mobility can be summarized by the question:

"How does the probability of obtaining an A-Level degree differ between children from low-income households and children from high-income households?".

We focus on two measures of relative mobility to answer this question. Our first measure of relative mobility is aimed at measuring the (absolute) difference in the probability of obtaining an A-Level degree between children from rich and poor parents[24]. The second measure we report is aimed at measuring the relative likelihood of obtaining an A-Level degree between children from rich and poor parents[25].

To calculate the first measure of relative mobility, we estimate the slope coefficient β of a regression of our binary outcome variable Y_i on the parental income rank R_i:

$$Y_i = \alpha + \beta R_i + \epsilon_i \tag{2.1}$$

If the relationship between parental income rank and the probability of obtaining an A-Level degree is well described by a linear function (we provide evidence that this is indeed the case later on), the slope coefficient of this regression provides a parsimonious measure

[24]Hence a possible answer to the question asked before could e.g. be: the probability to obtain an A-Level degree is 30 percentage points higher for children from rich families, compared to children from poor families.

[25]hence a possible answer to the question asked before could e.g. be: children from rich families are twice as likely to obtain an A-Level degree, compared to children form poor families.

of relative mobility. Then, the probability of obtaining an A-Level degree of a child from a given income percentile r is given by:

$$E(Y_i = 1|R_i = r) = \alpha + \beta r \tag{2.2}$$

and the difference in probabilities between two income percentiles r and q, $r > q$, is given by:

$$E(Y_i = 1|R_i = r) - E(Y_i = 1|R_i = q) = \beta(r - q) \tag{2.3}$$

Hence the slope coefficient summarizes the difference in the probability of obtaining an A-Level degree between households from different income percentiles well. We use estimates of this slope coefficient β as our first measure of relative mobility. In what follows, we refer to $\beta \times 100$ as (parental) income gradient. For better readability, we always multiply β by 100, i.e. we consistently report $\beta \times 100$. $\beta \times 100$ is the difference in the probability to obtain an A-Level degree between children from the very top of the income distribution and the very bottom of the income distribution, based on the linear fit described in Equation 2.1. Hence our first measure of intergenerational mobility, the parental income gradient is given by:

$$\text{Parental Income Gradient} = \beta \times 100 \tag{2.4}$$

We also report a second measure of relative mobility. One might also be interested in the relative likelihood of obtaining an A-Level degree between children from rich and poor families[26].

We measure the relative likelihood of obtaining an A-Level degree between children from rich and poor families as the ratio of the share of children who obtain an A-Level degree from the top quintile of the income distribution and the share of children who obtain an A-Level degree from the bottom quintile of the income distribution:

$$Q5/Q1 = \frac{E(Y_i = 1|R_i \geq 80)}{E(Y_i = 1|R_i < 20)} \tag{2.5}$$

We refer to this measure as $Q5/Q1$ ratio. This measure can be interpreted as the relative likelihood of obtaining an A-Level degree between both groups, e.g. $Q5/Q1 = 2$ would mean

[26] i.e. how many times more likely are children from rich families to obtain an A-Level degree than children from poor families.

that children from the top quintile of the income distribution are twice as likely to obtain an A-Level degree than children from the bottom quintile of the income distribution.

In the calculation of the $Q5/Q1$ ratio, we do not rely on the linear fit as described in Equation 2.1, but calculate the probability to obtain an A-Level degree non-parametrically for each quintile of the income distribution. We have also calculated the $Q5/Q1$ ratio based on the linear fit described by Equation 2.1 and obtain very similar results.

For both measures of relative mobility, the parental income gradient and the $Q5/Q1$ ratio, a higher value implies lower relative mobility. It should also be noted that, while both measures are measures of relative mobility, they aim at capturing different normative concepts. Hence when comparing mobility measures across regions, it might well be the case that a region is described by a (comparatively) low parental income gradient, but at the same time by a high $Q5/Q1$ ratio[27].

Absolute Mobility:

A different question with regards to intergenerational mobility in the context of our analysis is the following:

"What is the probability of obtaining an A-Level degree for a child from a given (poor) household income level?".

The normative implication of this question is obviously different from the one before. The question posed here is concerned with the absolute outcomes of (disadvantaged) children, regardless of the outcomes of children from more advantaged households.

Following Chetty et al. (2014), we present two measure aimed at answering this question. The first one is the expected probability of obtaining an A-Level degree for a child from the very bottom of the parental income distribution implied by the regression in Equation 2.1, i.e. the constant estimated from this regression. We refer to the constant as baseline probability:

$$\text{Baseline Probability} = E(Y_i = 1 | R_i = 0) = \alpha \tag{2.6}$$

The second measure of absolute mobility we report is the probability of obtaining an

[27] This could be the case if the baseline probability of obtaining an A-Level degree, described under Absolute Mobility, is low in this region as well.

A-Level degree of a child from the bottom quintile of the national income distribution:

$$Q1 = E(Y_i = 1|R_i < 20) \tag{2.7}$$

We refer to this measure as $Q1$ measure. We always calculate the $Q1$ measure non-parametrically as the share of children obtaining an A-Level degree in the bottom quintile of the income distribution. If the relationship between the probability of obtaining an A-Level degree and parental income rank is (close to) linear across the whole income (rank) distribution (we present evidence that this is the case later on), both measures will lead to very similar results. We present both measures to ensure the robustness of the results we obtain. For both measures, the baseline probability and the $Q1$ measure, a high value implies high absolute mobility.

We present different measures of both, relative and absolute mobility, as both sets of measures provide useful information to answer the question which regions are described by a high degree of intergenerational mobility. The answer to the question which region exhibit high intergenerational mobility might depend, after all, on the exact question asked.

2.6 National Results

2.6.1 Comparing Different Measures of Household Income

In this Section, we describe the association between parental income rank and the child's probability of obtaining an A-Level degree at the national level. We present the results at the national level for the different income measures discussed in Section 2.4.5.

Figure 2.1 shows how the share of children with an A-Level degree varies with the position in the parental income distribution. The income measure used in Figure 2.1 is parental income (not adjusted for household size). The linear fit (black line) slightly overestimates the A-Level share between (roughly) the 30th and 80th percentile, and underestimates it in the tails of the income distribution. The linear fit based on this income measure has a constant (baseline probability) of $\alpha = .169$ and a slope coefficient of $\beta = .0045$, i.e. a parental income gradient of $\beta \times 100 = 0.45$. The slope of the linear fit implies that a 10 percentage point increase in parental income rank is associated with an increase in the probability of obtaining an A-Level by 4.5 percentage points.

Figure 2.1: Intergenerational Mobility at the National Level: No Adjustment for Household Size

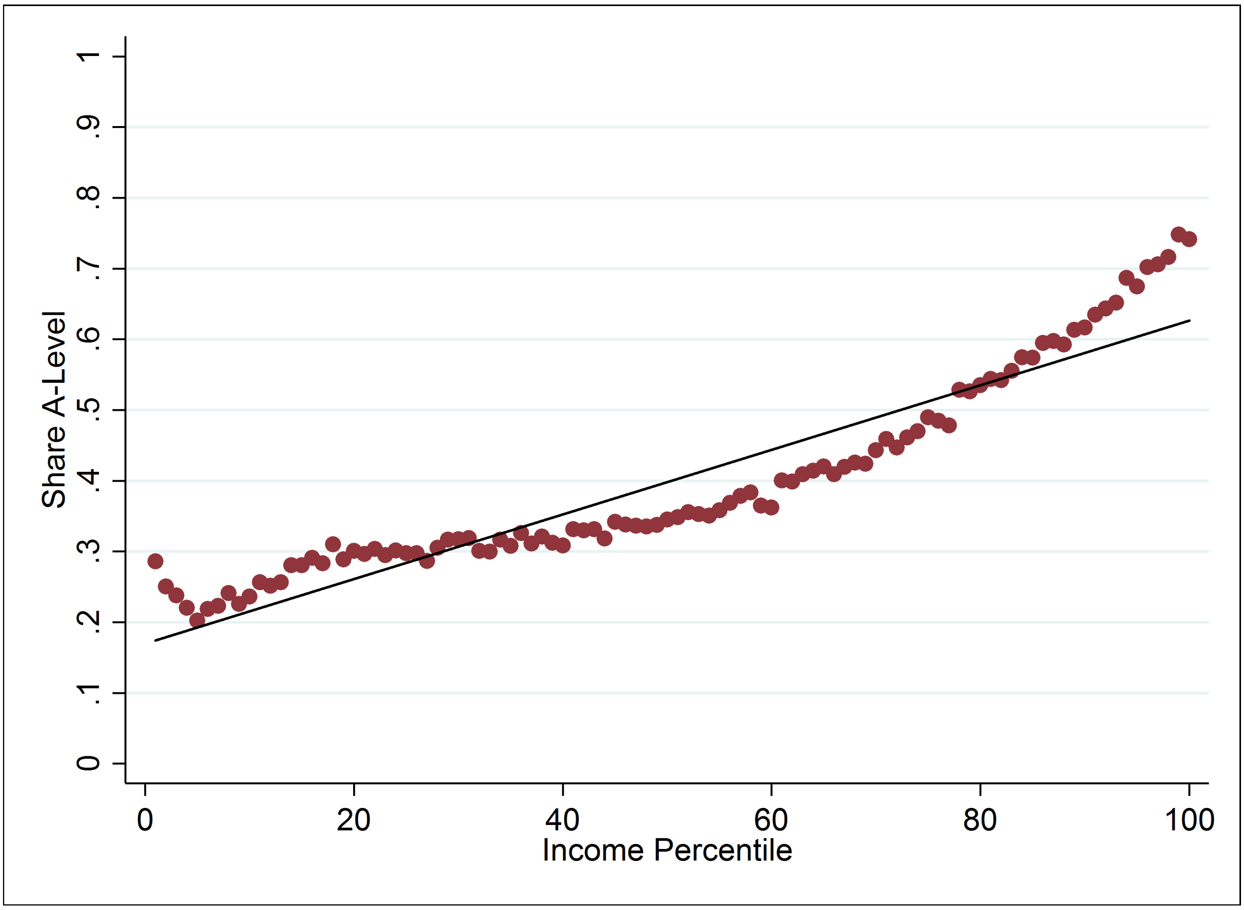

This Figure plots the fraction of children aged 16-22 (y-axis) that are either enrolled in the last two/three years of the A-Level track or already completed the A-Level degree vs. the percentile rank of their parents the national parental income distribution when parental income is not adjusted for household size (x-axis). The Figure is constructed by binning parent rank into one-percentile point bins (so that there are 100 bins) and plotting the fraction of children aged 16-22 currently enrolled in the last two/three years of the A-Level track or having already completed it among all children in this bin vs. the parent rank for each bin. The Figure is based on $N = 268523$ children living with their parents. The OLS regression of the A-Level dummy on the parent income rank in the underlying micro-data (black line) yields a slope coefficient of 0.0045 and a constant of 0.169.

In Figure 2.2, we measure income as parental income divided by the number of household members (i.e. parental income per capita). Interestingly, this results in a relationship that is better approximated by a linear relationship, compared to the case with no adjustment for household size. The linear fit based on this income measure has a constant (baseline probability) of $\alpha = .166$ and a slope coefficient of $\beta = .0048$, i.e. a parental income gradient of $\beta \times 100 = 0.48$. When using this income measure, a 10 percentage point increase in the parental income rank is associated with an increase in the probability of obtaining an A-Level degree of the child of 4.8 percentage points.

Figure B.5 in the Appendix presents the results at the national level when we divide

Figure 2.2: Intergenerational Mobility at the National Level: Income per Household Member

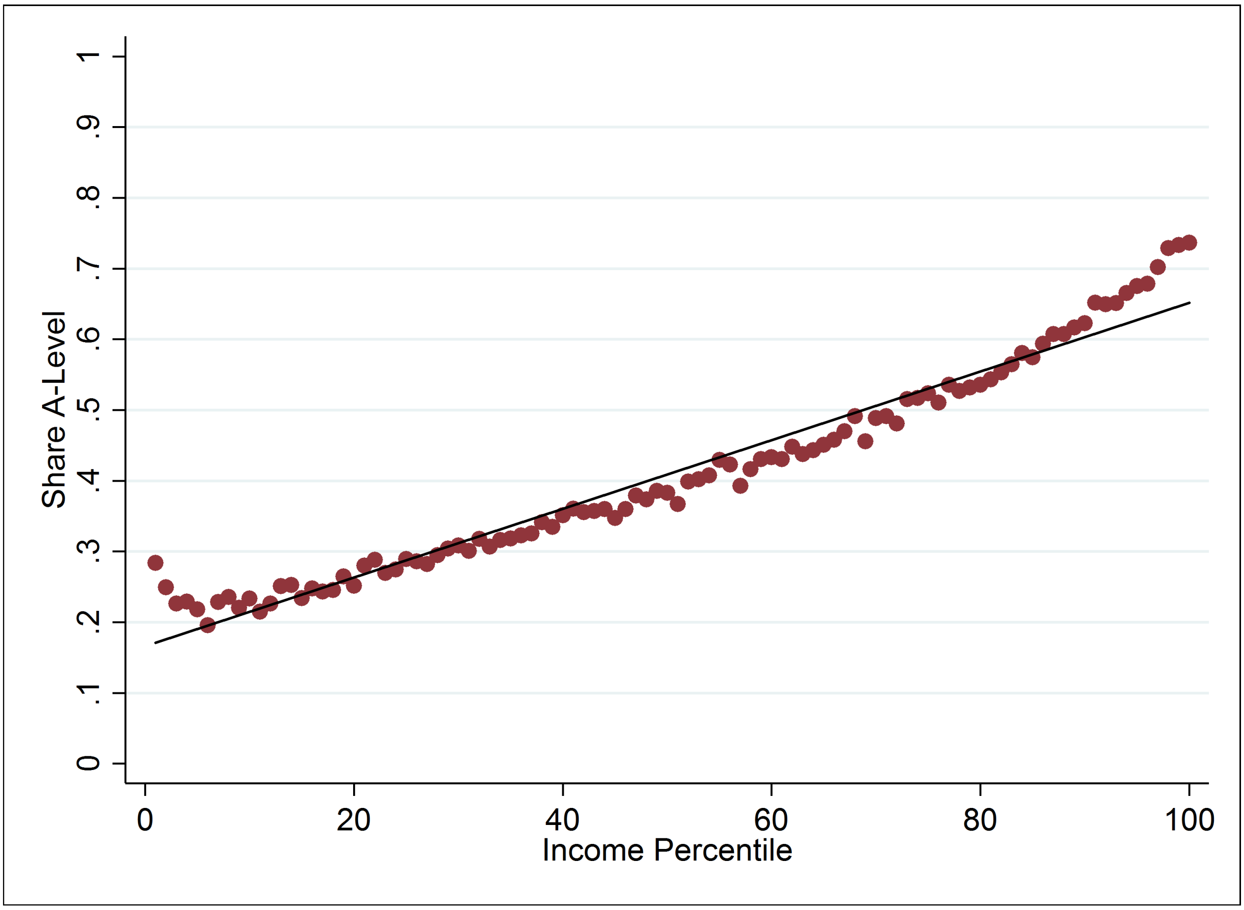

This Figure plots the fraction of children aged 16-22 (y-axis) that are either enrolled in the last two/three years of the A-Level track or already completed the A-Level degree vs. the percentile rank of their parents in the national net income per household member distribution (x-axis). The Figure is constructed by binning parent rank into one-percentile point bins (so that there are 100 bins) and plotting the fraction of children aged 16-22 currently enrolled in the last two/three years of the A-Level track or having already completed it among all children in this bin vs. the parent rank for each bin. The Figure is based on $N = 268523$ children living with their parents. The OLS regression of the A-Level dummy on the parent income rank in the underlying micro-data (black line) yields a slope coefficient of 0.0048 and a constant of 0.166.

parental income by the square root of the number of household members (equivalence scale income). This results in a relationship between parental income rank and the share of children obtaining an A-Level degree that is more linear than the one presented in Figure 2.1, but less linear than the relationship displayed in Figure 2.2. The linear fit based on this income measure has a constant (baseline probability) of $\alpha = .159$ and a slope coefficient of $\beta = .0048$, i.e. a parental income gradient of $\beta \times 100 = 0.48$.

The results presented in Figures 2.1, 2.2 and B.5 show that that linear fit based on the three different income measures results in very similar estimates for the baseline probability (the constant of the regression) and the parental income gradient.

Furthermore, the quintile-based measures of intergenerational mobility are also very similar across specifications. Measuring income as in Figure 2.1 (no adjustment for household size) results in estimates of $Q1 = .257$ and $Q5/Q1 = 2.47$. Measuring income as in Figure 2.2 (income per household member) results in estimates of $Q1 = .237$ and $Q5/Q1 = 2.67$. Measuring income as in Figure B.5 (equivalence scale income) results in estimates of $Q1 = .243$ and $Q5/Q1 = 2.63$.

At the national level, our estimates of measures of intergenerational mobility are hence very similar for the three different income measures. In all Sections following Section 2.6.2, we present all results based on the income measure used in Figure 2.2, parental income per household member. The use of this income measure is motivated mainly by the fact that the relationship between educational attainment of the child and parental income is best described by a linear relationship when using this income measure. This allows us to compare intergenerational mobility across regions parsimoniously based on the fit of Equation 2.1. We obtain very similar results when using the two other income measures.

2.6.2 Comparison to Other Estimates of Intergenerational Mobility

Based on the linear fit in Figure 2.1, we find that a 10 percentage point increase in the parental income rank is associated with an increase in the share of children with an A-Level degree of 4.5 percentage points[28]. Chetty et al. (2014) find that in the United States a 10 percentage point increase in parental income rank is associated with an increase in the share of children who attend college by 6.7 percentage points. This suggests that the association between parental income rank and child educational outcomes is weaker in Germany, compared to the US. However, this comparison is of course rather an indirect one, as both outcomes are different. While college attendance in the US and an A-Level degree in Germany are, of course, not the same outcome, we choose this comparison as both measures focus on the educational attainment of children. Hence this should not be viewed a direct comparison of intergenerational mobility of both countries, as the institutional setting is different. Furthermore, we expect that the parental income gradient for college attendance in Germany should be higher than the parental income gradient for obtaining an

[28] For the comparison here, we use income not adjusted for household size as income measure, to use the same income measure as Chetty et al. (2014) for the comparison with their results.

A-Level degree. The particular institutional setting of our study prevents us to make a clear comparison to other countries. We rather focus on studying how intergenerational mobility differs across different regions in Germany, as the institutional setting is much more similar across regions within Germany.

2.6.3 Robustness of National Results

In Table 2.3, we present the estimates of intergenerational mobility we obtain at the national level when we use different age cut-offs for children to define our sample. When using different age-cutoffs we also re-calculate income percentile ranks, i.e. we define income ranks always as the rank relative to other households within the sample. Here, we always use income per household member as our income measure. In addition to our measures of intergenerational mobility, we also report the A-Level share for the different age groups in our sample, and the sample size. For convenience, we also re-state the results we obtain when using our baseline sample, children aged 16 to 22.

Table 2.3: Intergenerational Mobility Measures using Different Age Cut-Offs

Sample	Sample Size	A-Level Share	$Q1$	α	$\beta \times 100$	$Q5/Q1$
Children aged 16 - 22	268523	.403	.237	.166	.485	2.67
Children aged 16 - 19	171441	.369	.209	.137	.465	2.81
Children aged 17 - 20	163903	.432	.254	.179	.508	2.63
Children aged 20 - 22	97082	.465	.281	.205	.521	2.51

This Table present the estimates of intergenerational mobility we obtain at the national level when we use different age cut-offs for children to define our sample.

When using older children, the (unconditional) A-Level share in our sample becomes higher, as already noted before. Our estimates of the parental income gradient are increasing in the age cut-offs, as well as our estimates of the baseline probability. This suggests that we slightly underestimate both, the parental income gradient and the baseline probability, at the national level in our main sample, compared to estimates we obtain when using older children. However, one should also keep in mind that the estimates based on older children are also based on a more selected sample, as we can link less children to their parents for these age groups (see Table 2.1). We also note that for all age groups, the relationship between parental income rank and the probability of obtaining an A-Level is well approximated by a linear function, as can be seen in Figures B.6, B.7 and B.8 in the Appendix. The main focus of our analysis lies in the comparison of intergenerational mobility across regions. Our regional estimates of intergenerational mobility are robust to the use of different age-cutoffs, as we show in Section 2.8.7.

2.7 Differences in Intergenerational Mobility between Groups

Before turning to regional differences in intergenerational mobility at a granular level, we analyze differences in intergenerational mobility between (broadly defined) groups in Germany at the national level. The goal of this analysis is twofold. First, we want to provide a comprehensive description of intergenerational mobility in Germany. Second, this analysis provides evidence that the relationship between parental income and child educational attainment is indeed well described by a linear approximation, not only at the most aggregate level, but also for different subgroups of the population. As noted before, in all comparisons we use parental income per household member as our income measure. In all comparisons, income percentiles always refer to the percentiles in the national (aggregate) income distributions[29].

2.7.1 The Role of Gender

In Figure 2.3 we separately report estimates of intergenerational mobility for boys and girls. We find a rather large difference in the baseline probability of obtaining an A-Level degree (i.e. the constant of the linear fit). The baseline probability is estimated at $\alpha = 0.215$ for girls, and $\alpha = 0.125$ for boys, respectively. This means that baseline probability of obtaining an A-Level is 9 percentage points higher for girls. This difference is nearly constant (but slightly widening) across the income distribution, as indicated by the fact that parental income gradients are rather similar (but slightly higher for boys): $\beta \times 100 = 0.49$ for girls, compared to $\beta \times 100 = 0.47$ for boys.

A similar picture emerges when we compare quintile-based mobility measures between both groups. We find $Q1 = .284$ for girls and $Q1 = .198$ for boys, again indicating a roughly 9 percentage point difference in the probability of obtaining an A-Level degree between both groups. When calculating our second measure of relative mobility, we find $Q5/Q1 = 2.40$

[29] i.e. percentiles are not calculated based on group-specific income distributions, but the same for all groups.

for girls and $Q5/Q1 = 2.98$ for boys, reflecting the differences in the baseline probability between both groups.

Figure 2.3: Differences by Gender

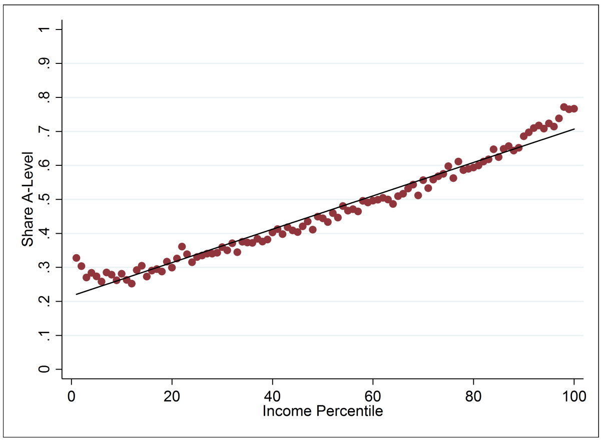

(a) Girls

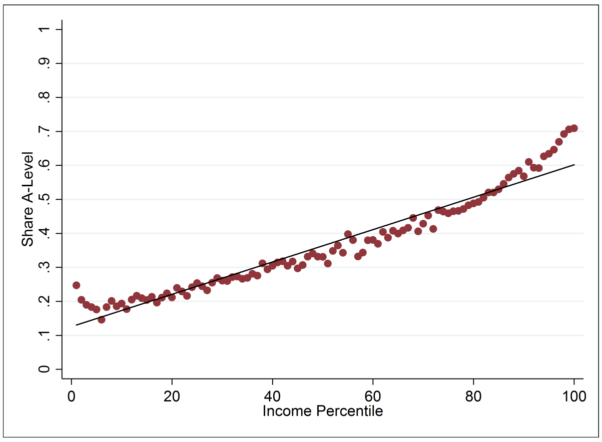

(b) Boys

This Figure plots the fraction of children aged 16-22 (y-axis) that are either enrolled in the last two/three years of the A-Level track or already completed the A-Level degree vs. the percentile rank of their parents in the national net income per household member distribution (x-axis), separate by gender of the child. The Figures are constructed by binning parent rank into one-percentile point bins (so that there are 100 bins) and plotting the fraction of children aged 16-22 currently enrolled in the last two/three years of the A-Level track or having already completed it among all children in this bin vs. the parent rank for each bin. The left Figure (girls) is based on $N = 124121$ observations, the right Figure (boys) is based on $N = 144402$ observations. For girls, the slope and constant of the OLS regression in the underlying micro-data (black line) are 0.0049 and 0.215, for boys the slope and constant are 0.0047 and 0.125.

2.7.2 The Role of Parental Education

In Figure 2.4 we separately report the association between parental income and child educational outcomes for children with different parental education levels. Specifically, we report results separately for children from households where no parent has an A-Level degree (left Panel in Figure 2.4) and for children from households where at least one parent has an A-Level degree (right Panel in Figure 2.4). We find a large difference in the baseline probability of obtaining an A-Level degree between the two groups. For children with parents without A-Level degree, the baseline probability of obtaining an A-Level degree themselves is $\alpha = 0.16$, whereas the baseline probability for children with at least one parent with an A-Level degree is $\alpha = 0.44$, more than twice as high. Remarkably, this difference is basically constant across the whole income distribution, the parental income gradients of the two different groups are nearly identical: $\beta \times 100 = 0.3018$ for children without parental A-Level degree and $\beta \times 100 = 0.3068$ for children with parental A-Level degree.

We obtain very similar results when calculating quintile-based measures of intergenerational mobility. We find $Q1 = .192$ for children where neither parent has an A-Level degree, and $Q1 = .496$ for children where at least one parent has an A-Level degree, again indicating a nearly 30 percentage point gap in the probability of obtaining an A-Level degree between both groups.

Due to the high baseline probability of obtaining an A-Level degree for children where at least one parent has an A-Level degree, our second measure of relative mobility, the $Q5/Q1$ ratio is rather low for these children. For children where at least one parent has an A-Level degree, we estimate $Q5/Q1 = 1.46$. For children where neither parent has an A-Level degree, we estimate $Q5/Q1 = 2.24$, reflecting the flattening of the parental income gradient once parental education is taken into account.

Figure 2.4: Differences by Parental Education

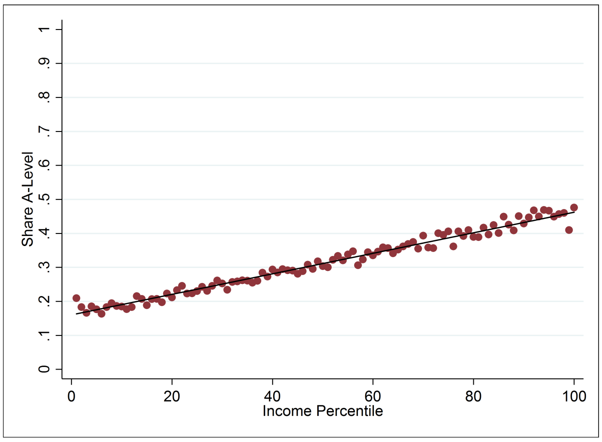

(a) Neither Parent has an A-Level Degree

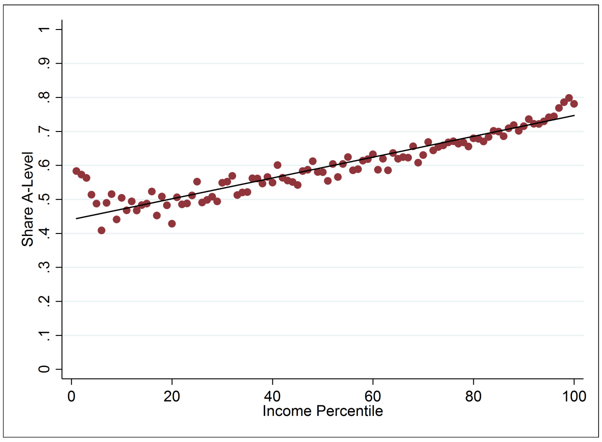

(b) At Least One Parent With A-Level Degree

This Figure plots the fraction of children aged 16-22 (y-axis) that are either enrolled in the last two/three years of the A-Level track or already completed the A-Level degree vs. the percentile rank of their parents in the national net income per household member distribution (x-axis), separate by parental education levels. The Figures are constructed by binning parent rank into one-percentile point bins (so that there are 100 bins) and plotting the fraction of children aged 16-22 currently enrolled in the last two/three years of the A-Level track or having already completed it among all children in this bin vs. the parent rank for each bin. The left Figure (parents without A-Level degree) is based on $N = 177766$ observations, the right Figure (at least one parent has an A-Level degree) is based on $N = 90757$ observations. For children whose parents do not have an A-Level degree, the slope and constant of the OLS regression in the underlying micro-data (black line) are 0.0030 and 0.160, for children where at least one parent has an A-Level degree the slope and constant are 0.0030 and 0.440.

2.7.3 Migrants and Natives

In Figure 2.5, we separately report the association between parental income and child's educational attainment for children with migrational background ("Migrationshintergrund", right Panel in Figure 2.5) and children without migrational background (left Panel in Figure 2.5). Here, we loosely refer to children with migrational background as migrants, and children without migrational background as natives.

The baseline probability of obtaining an A-Level degree is slightly higher for children with migrational background ($\alpha = 0.19$ for migrants and $\alpha = 0.15$ for natives). The reverse is true for the parental income gradient. We estimate $\beta \times 100 = 0.44$ for migrants and $\beta \times 100 = 0.50$ for natives.

Based on the linear fit, the probability of obtaining an A-Level degree is higher for children with migrational background (compared to children without migrational background) in the bottom half of the income distribution, whereas the reverse is true in the upper half of the income distribution.

We obtain similar results when calculating quintile-based measures of intergenerational mobility. We find $Q1 = .248$ for migrants and $Q1 = .229$ for natives, again indicating a slightly higher baseline probability of obtaining an A-Level degree for migrants.

Due to the higher baseline probability and the lower slope for migrants, we estimate a (slightly) lower $Q5/Q1$ measure for migrants, compared to natives. We find $Q5/Q1 = 2.58$ for migrants, and $Q5/Q1 = 2.76$ for natives

It is noteworthy that conditional on parental income, the probability of obtaining an A-Level degree is identical (or even higher) for children with migrational background, compared to children without migrational background.

Figure 2.5: Differences by Migrational Background

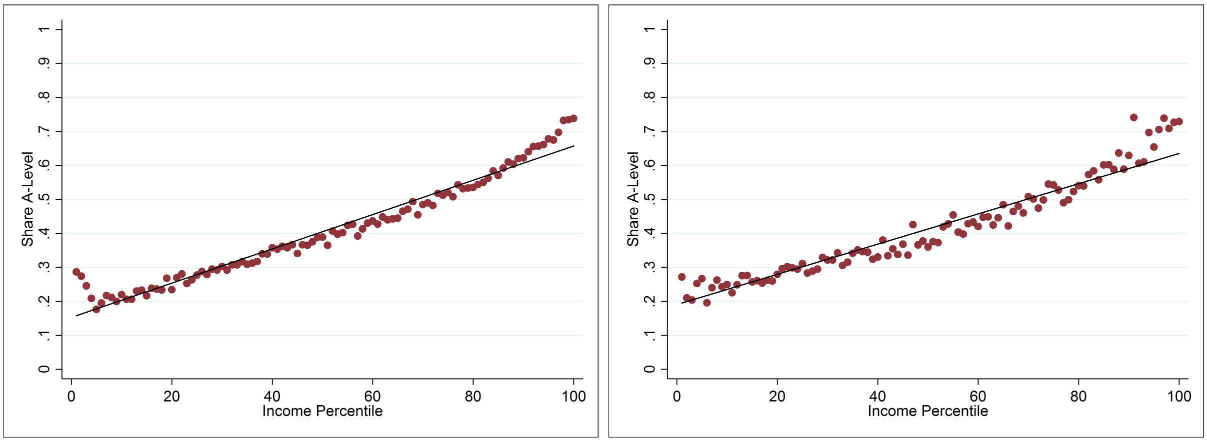

(a) No Migrational Background **(b)** Migrational Background

This Figure plots the fraction of children aged 16-22 (y-axis) that are either enrolled in the last two/three years of the A-Level track or already completed the A-Level degree vs. the percentile rank of their parents in the national net income per household member distribution (x-axis), separate by migrational background. The Figures are constructed by binning parent rank into one-percentile point bins (so that there are 100 bins) and plotting the fraction of children aged 16-22 currently enrolled in the last two/three years of the A-Level track or having already completed it among all children in this bin vs. the parent rank for each bin. The left Figure (no migrational background) is based on $N = 200465$ observations, the right Figure (migrational background) is based on $N = 68058$ observations. For children without migrational background, the slope and constant of the OLS regression in the underlying micro-data (black line) are 0.0050 and 0.152, for children with migrational background the slope and constant are 0.0044 and 0.190.

2.7.4 East Germany and West Germany

Lastly, before turning to a granular regional analysis, we report differences in the association between parental income and child's educational attainment between East- and West Germany[30]. Households are assigned to either part of the country based on the current place of residence.

Results of this comparison are reported in Figure 2.6. The baseline probability is $\alpha = 0.138$ in East Germany and $\alpha = 0.166$ in West Germany. The parental income gradient is $\beta \times 100 = 0.55$ in East Germany and $\beta \times 100 = 0.47$ in West Germany. Remarkably, the parental income gradient is substantially higher in East Germany, compared to West Germany.

Based on the linear fit, the probability of obtaining an A-Level degree is higher for children from West Germany in the bottom tertile of the income distribution, whereas the probability of obtaining an A-Level degree is higher for children from East Germany in the two top tertiles of the national income distribution.

We obtain similar results when calculating quintile-based measures of intergenerational mobility. We find $Q1 = .202$ in East Germany and $Q1 = .240$ in West Germany, again indicating higher levels of absolute mobility in West Germany.

Consistent with the other results, we estimate $Q5/Q1 = 3.22$ in East Germany and $Q5/Q1 = 2.61$ in West Germany, indicating a higher degree of relative mobility in West Germany.

[30]We exclude Berlin from this comparison, as it not possible to distinguish between East- and West Berlin in our data.

Figure 2.6: Differences between East Germany and West Germany

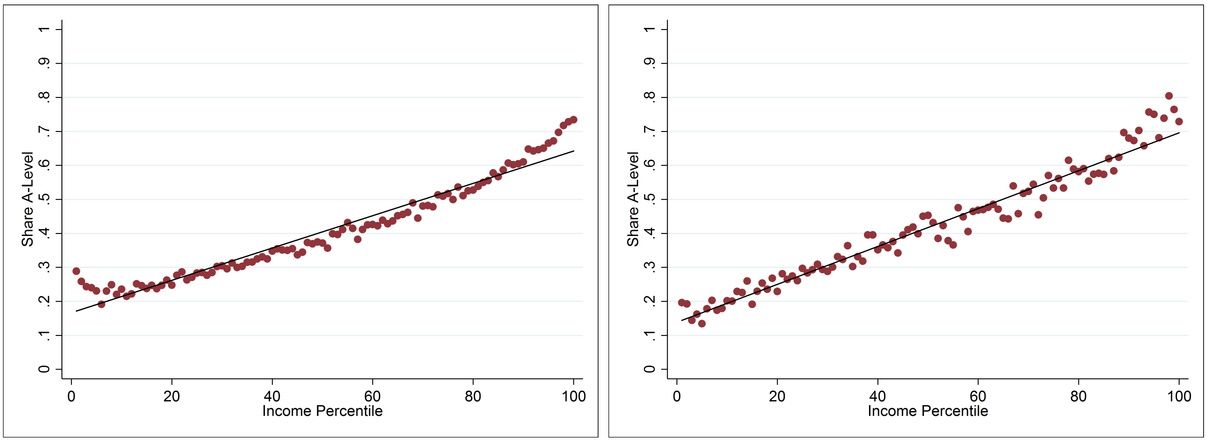

(a) West Germany

(b) East Germany

This Figure plots the fraction of children aged 16-22 (y-axis) that are either enrolled in the last two/three years of the A-Level track or already completed the A-Level degree vs. the percentile rank of their parents in the national net income per household member distribution (x-axis), separate by place of residence. The Figures are constructed by binning parent rank into one-percentile point bins (so that there are 100 bins) and plotting the fraction of children aged 16-22 currently enrolled in the last two/three years of the A-Level track or having already completed it among all children in this bin vs. the parent rank for each bin. The left Figure (West Germany) is based on $N = 230366$ observations, the right Figure (East Germany) is based on $N = 29545$ observations. For children living in West Germany, the slope and constant of the OLS regression in the underlying micro-data (black line) are 0.0047 and 0.166, for children living in East Germany the slope and constant are 0.0055 and 0.138.

2.8 Regional Variation in Intergenerational Mobility

We now turn to a detailed characterization of regional differences in intergenerational mobility across Germany. First, we describe the geographical units of our regional analysis and describe how we assign children to geographical units. Next, we present estimates of absolute and relative mobility for our geographical units.

2.8.1 Geographical Units

We characterize regional differences in intergenerational mobility at the level of local labor markets (LLM). The Federal Institute for Building, Urban Affairs and Spatial Research ("Bundesinstitut für Bau-, Stadt- und Raumforschung") defines 258 local labor market regions in Germany. Local labor markets are a strict aggregation of the 401 German counties ("Kreise"), with each county being contained in exactly one local labor market. Local labor markets are defined on the basis of the distribution of regional economic activity (on the basis of commuting patterns), and strict subsets of German states (with the exceptions of the five local labor markets Bremen, Bremerhaven, Hamburg, Mannheim and Ulm)[31]. The definition is conceptually similar to the definition of commuting zones in the US, used e.g. by Chetty et al. (2014).

We assign households to the local labor market of their current place of main residence reported in our data. That is, we assign households to the local labor market they reside in at the time they are observed in the MZ. We use the current place of residence as our measure of household location, as we are not able to track household movements over time. Conceptually, this is very similar to the approach in Chetty et al. (2014), who use the place of residence of a child at the age of 15 (for most children) as main geographical indicator. The median number of children in our sample (observations) per LLM is 649 (mean: 1040). The lowest number of observations across all LLM's is 162 (in the LLM Sonneberg in Thuringia), the largest number of observations is 9299 (in the LLM Stuttgart in Baden-Württemberg).

In Figure 2.7, we present a heatmap of unconditional A-Level shares for children in our sample, i.e. the fraction of children with an A-Level degree (according to our definition) in our sample calculated separately for every local labor market in Germany. In all heatmaps, we always use dark colors to indicate a low level of intergenerational mobility and light colors to indicate high levels of intergenerational mobility.

Figure 2.7 reveals that there is substantial variation in (unconditional) A-Level shares across local labor markets. In some local markets, the share of children obtaining an A-Level degree is as low as 17%, whereas in other local labor markets more than 50% of children obtain an A-Level degree. While A-Level shares differ between local labor markets within federal states ("Bundesländer"), we observe clustering at the state-level. A-Level shares are rather low in Bavaria and in Saxony-Anhalt and high in Berlin, Brandenburg and North-Rhine Westfalia.

Roughly 50% of the variation in the unconditional A-Level shares across LLMs is due to within-state variation in the unconditional A-Level shares[32]. While these numbers reveal nothing about intergenerational mobility per se, they help to interpret our results for intergenerational mobility reported in the following Sections.

[31] More information can be found on the website of the relevant federal institute, available at https://www.bbsr.bund.de/BBSR/DE/Raumbeobachtung/Raumabgrenzungen/AMR/Arbeitsmarktregionen.html.

[32] This calculation is based on a regression of the A-Level share in each LLM on a set of state-level dummy variables. This analysis uses only LLMs that are strict subsets of states, i.e. the 5 LLMs that cross state boundaries are excluded from this analysis. The R-squared of this regression is $R^2 = 0.507$, meaning that roughly 50% of the variation in unconditional A-Level shares is due to between state variation.

Figure 2.7: A-Level Shares Across LLM's in Germany

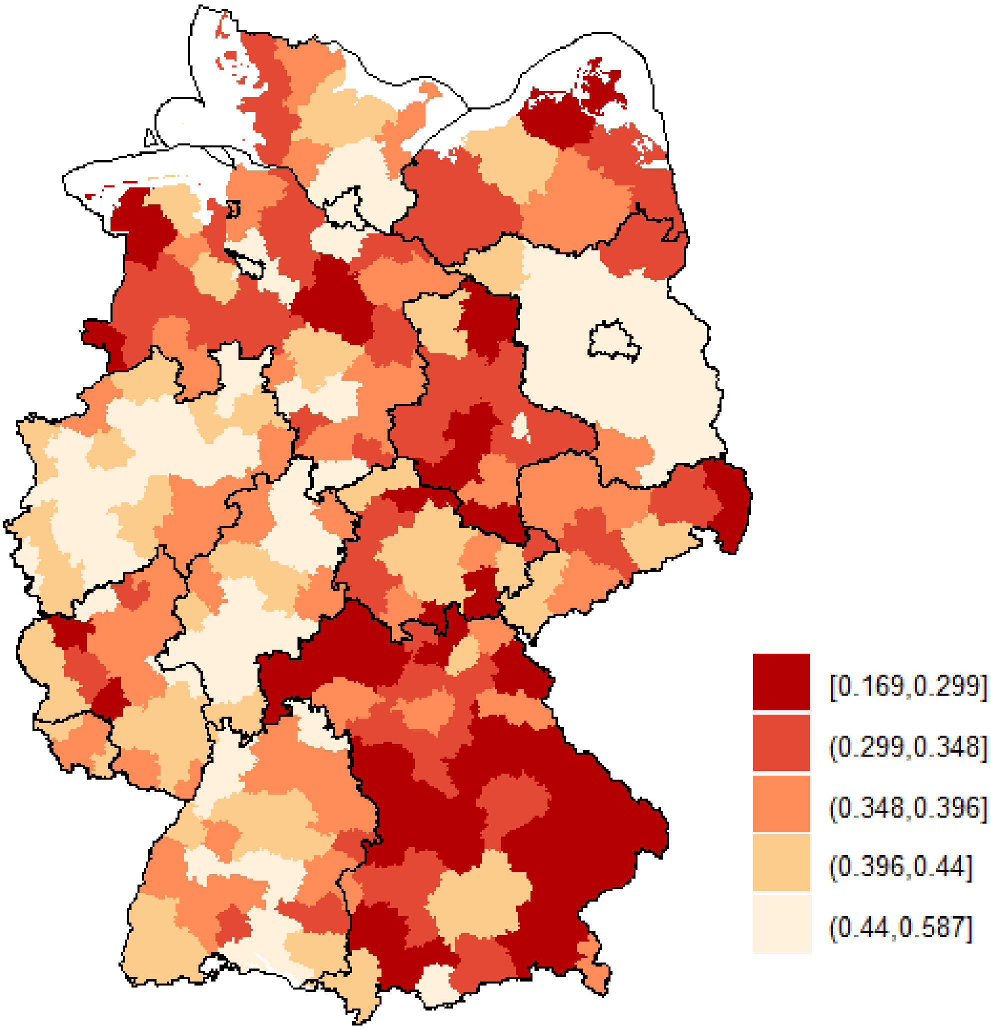

This Figure shows the fraction of children in our sample that either have an A-Level degree or are enrolled in the last two/three years of school leading to an A-Level degree among all children in the sample, calculated separately for every local labor market in Germany.

2.8.2 Relative Intergenerational Mobility

Figure 2.8 presents the regional variation in our first measure of relative mobility, the parental income gradient. The estimates of the parental income gradient at the LLM-level are obtained by estimating Regression 2.1 separately for each LLM l:

$$Y_{i,l} = \alpha_l + \beta_l R_i + \epsilon_{i,l} \tag{2.8}$$

where $Y_{i,l}$ denotes our binary outcome variable of a child i living in LLM l with parental income rank R_i[33].

Figure 2.8 reveals substantial heterogeneity in the parental income gradient across regions: While in some regions, $\beta \times 100$ is estimated to be as low as 0.1, in other regions we estimate $\beta \times 100$ to be as high as 0.7. This means that in some regions the difference in the probability of obtaining an A-Level degree between children from the very top of and the very bottom in the national income distribution is only 10 percentage points, whereas this difference is as high as 70 percentage points in other regions. The magnitude of the regional variation in our estimates of the parental income gradient is roughly comparable to the magnitude of the regional variation in rank-rank slopes and college attendance slopes across commuting zones in the US reported in Chetty et al. (2014).

Our estimates of the parental income gradient are generally higher in East Germany, compared to West Germany, corroborating the finding reported in Figure 2.6. Parental income gradients are especially high in Saxony and in Thuringia. Apart from that, no clear regional pattern at the state-level emerges for the parental income gradient. Even within in the same state, substantial variation in the parental income gradient exists: only 14% of the variation in the parental income gradient across LLMs is due to between state variation[34].

[33]The parental income rank is the rank in the national income distribution. To stress this, we drop the subscript l on the parental income rank.

[34]Calculated as the R^2 of a regression of the parental income at the LLM level on a set of state dummies

Figure 2.8: Parental Income Gradient Across LLM's in Germany

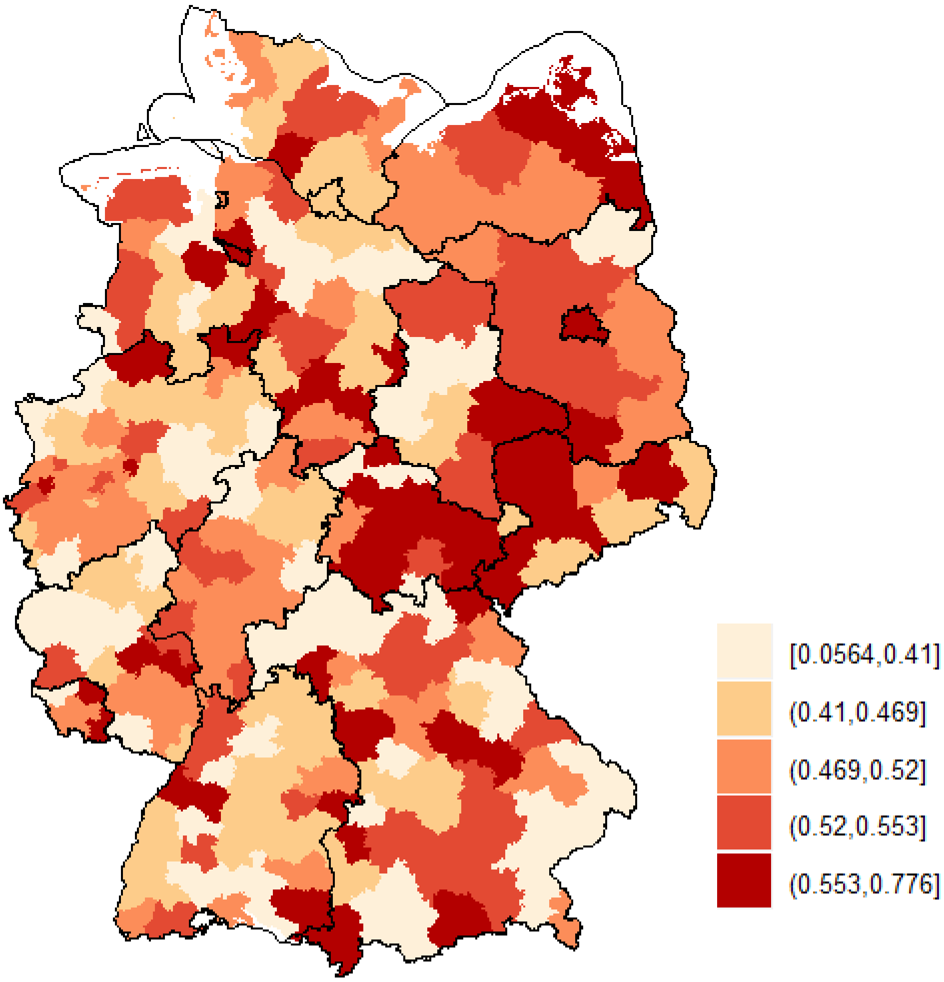

This Figure shows the estimates of the parental income gradient, estimated separately for every local labor market in Germany.

Figure 2.9 presents our second measure of relative Intergenerational Mobility, the $Q5/Q1$ ratio, which is the ratio of the share of children obtaining an A-Level degree from the top quintile of the income distribution and the share of children obtaining an A-Level degree from the bottom quintile of the income distribution. At the national level, the $Q5/Q1$ ratio is 2.67. Again, there is substantial variation across local labor markets. In some local labor markets, the $Q5/Q1$ ratio is as low as 1.3 (meaning that the relative likelihood to obtain an A-Level degree is 1.3 times higher for children from the top quintile of the income distribution compared to children from the bottom quintile of the income distribution), in other local labor markets the $Q5/Q1$ ratio is as high as 7 (meaning that the relative likelihood to obtain an A-Level degree is 7 times higher for children from the top quintile of the income distribution compared to children from the bottom quintile of the income distribution). Around 20% of the observed variation in the $Q5/Q1$ ratio between LLM's is due to between state variation.

Figure 2.9: Q5/Q1 Ratio Across LLM's in Germany

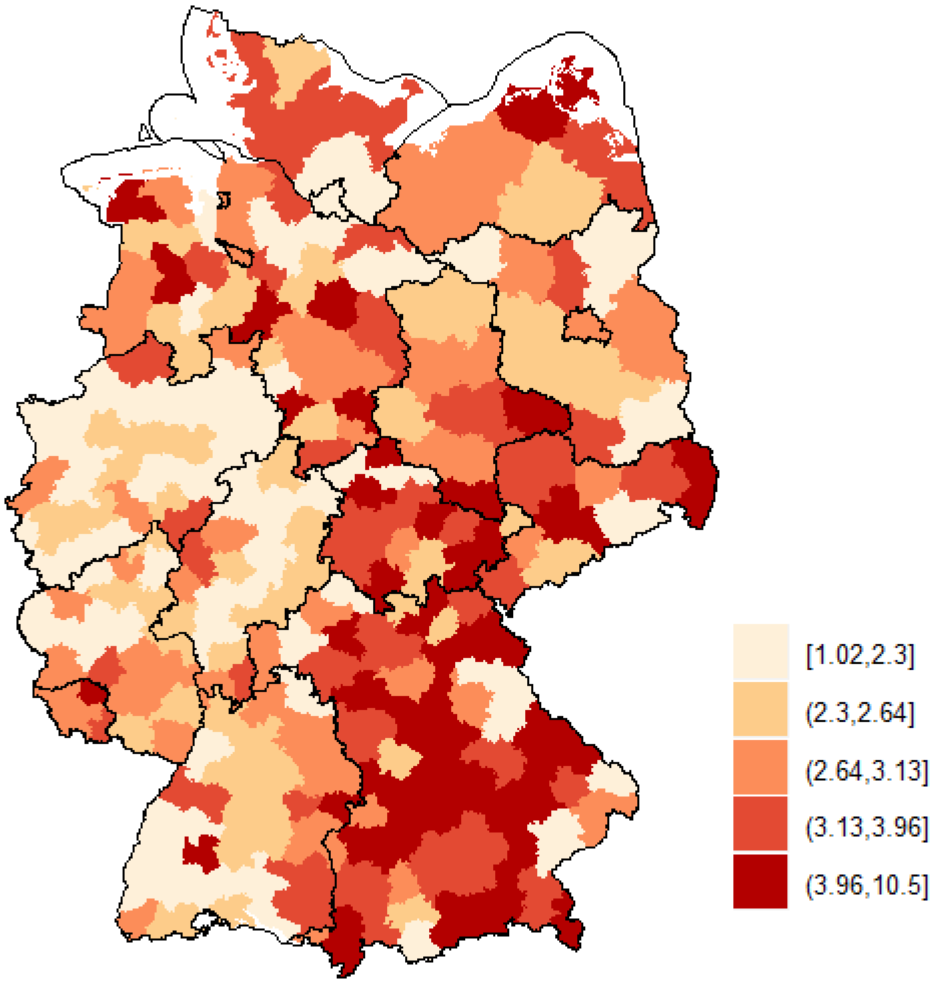

This Figure shows estimates of the Q5/Q1 ratio (i.e. the ratio of the share of children obtaining an A-Level degree among children in the top quintile of the income distribution and the share of children obtaining an A-Level degree among children the bottom quintile of the income distribution), calculated separately for every local labor market in Germany.

The $Q5/Q1$ ratio is generally higher in regions with low (unconditional) A-Level shares in our sample (e.g. Bavaria) and lower in regions with high A-Level shares (e.g. Northrhine-Westfalia). Indeed, there is a negative correlation between the $Q5/Q1$ ratio and the unconditional A-Level share presented in Figure 2.7 of -0.41 across LLMs (the population-weighted correlation is -0.53 and the (Spearman) rank correlation is -0.40). The correlation of the different measures of intergenerational mobility across local labor markets is explored in detail in Section 2.8.5.

2.8.3 Absolute Intergenerational Mobility

This Section reports the results for the estimates of the absolute intergenerational mobility measures. In Figure 2.10, we report the results for the $Q1$ measure, the share of children obtaining an A-Level degree from households in the bottom quintile of the national income distribution. Results for our second measure of absolute mobility, the constants estimated in regression 2.8, are relegated to Figure B.9 in the Appendix, as the correlation between both measures is very high ($\rho = 0.87$).

We find substantial differences in the $Q1$ measure across local labor markets. At the national level, 23% of children from the bottom quintile of the income distribution obtain an A-Level degree (i.e. $Q1 = 0.23$). In some regions, only 5% of children from the bottom quintile of the national income distribution obtain an A-Level degree, a substantially lower share of children than at the national average. In other regions, 40% of children from the bottom quintile of the national income distribution obtain an A-Level degree. Roughly 40% of the variation in the $Q1$ measure across LLM's is due to between state-variation in the $Q1$ measure.

Figure 2.10: $Q1$ Measure Across LLM's in Germany

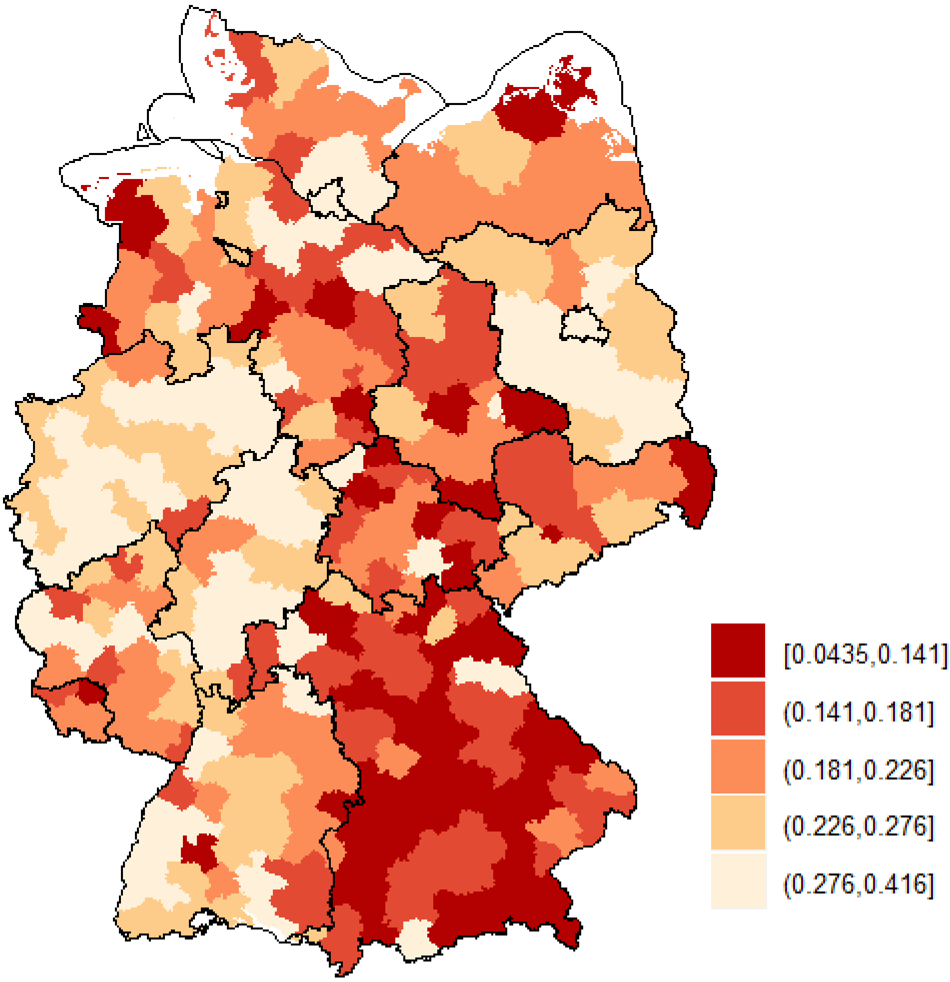

This Figure shows estimates of $Q1$ measure (i.e. the share of children obtaining an A-Level degree among children with parents in the bottom quintile of the national income distribution), calculated separately for every local labor market in Germany.

2.8.4 Intergenerational Mobility in the Largest LLM's

In this Section, we present estimates of absolute and relative mobility for the 25 largest local labor markets[35] in detail in Table 2.4.

Table 2.4 shows that there is considerable variation in our measures of intergenerational mobility across major urban areas in Germany. In the most (absolutely) mobile urban local labor markets, more than 30% of children from the bottom quintile of the income distribution obtain an A-Level degree (Column 5 in Table 2.4, which reports the $Q1$ measure). In the

[35]Measured by total population in 2015. In all local labor markets reported in Table 2.4, we observe more than 1500 children in our sample.

least (absolutely) mobile major urban areas, this share is as low as 9%. Many urban local labor markets in North Rhine Westphalia exhibit high degrees of absolute intergenerational mobility, whereas urban LLM's in Bavaria (and Leipzig in Saxony) exhibit a very low degree of absolute intergenerational mobility.

Table 2.4 also shows that there is substantial variation in the parental income gradient across urban local labor markets (Column 6, which reports the parental income gradient). The parental income gradient varies between .431 (in Dresden, lowest parental income gradient of the 25 local labor markets reported here) and .605 (in Leipzig, highest parental income gradient of the urban LLM's reported here). Coincidentally, Dresden and Leipzig are both located in the federal state of Saxony, illustrating that the parental income gradient exhibits substantial variation even within the same state[36].

Results for the urban LLM's reported in Table 2.4 also illustrate the fact that a high degree of absolute mobility does not necessarily imply a high degree of relative mobility. Out of the 25 urban LLM's reported here, the absolute mobility rank of Dresden (based on the $Q1$ measure) is 17 (out of 25), whereas Dresden has the lowest parental income gradient of the 25 LLM's reported here. Another example is Berlin, which has an absolute mobility rank of 12, but the second highest parental income gradient (out of the 25 LLM's reported here).

The second measure of relative mobility, the $Q5/Q1$ ratio exhibits substantial variation across the largest LLM's as well. In Essen (lowest $Q5/Q1$ ratio out of the 25 LLM's here), children from the top quintile of the income distribution are twice as likely as children from the bottom of the income distribution to have an A-Level degree. In Augsburg, they are 6 times as likely. Differences in the $Q5/Q1$ ratio are mainly driven by differences in the probability to obtain an A-Level degree for children in the bottom quintile of the income distribution for the 25 LLM's reported here. The $Q5/Q1$ ratio goes hand in hand with the $Q1$ measure for the local labor markets reported in Table 2.4, the correlation between both measures is $\rho = -0.918$ for the 25 LLM's reported in the Table.

The results reported in Table 2.4 are limited to large (urban) local labor markets. In the next Section, we discuss the relationship between our different measures of intergenerational mobility across all 258 local labor markets in Germany to shed light on the question whether absolute and relative mobility go hand in hand.

[36] Another example that illustrates this are e.g. Freiburg and Heidelberg.

Table 2.4: Mobility Measures for the 25 Largest Local Labor Markets

(1)	(2)	(3)	(4)	(5)	(6)	(7)
$Q1$ Rank	LLM Name	Population	A-Level Share	$Q1$	$\beta \times 100$	$Q5/Q1$
1	Bonn	915022	0.54	0.35	0.47	2.19
2	Essen	751902	0.50	0.34	0.45	2.08
3	Düsseldorf	1545483	0.53	0.33	0.52	2.24
4	Münster	805871	0.46	0.32	0.45	2.17
5	Frankfurt/Main	2238027	0.51	0.32	0.52	2.30
6	Hamburg	2850118	0.49	0.32	0.45	2.16
7	Duisburg	1164829	0.46	0.31	0.46	2.27
8	Potsdam-Brandenburg	608465	0.59	0.30	0.52	2.52
9	Köln	1809968	0.47	0.30	0.50	2.37
10	Karlsruhe	743596	0.47	0.30	0.53	2.26
11	Freiburg	645818	0.41	0.29	0.45	2.23
12	Berlin	3520031	0.48	0.28	0.56	2.70
13	Dortmund	1161613	0.44	0.28	0.53	2.59
14	Heidelberg	698126	0.50	0.27	0.55	2.62
15	Gelsenkirchen	1151169	0.42	0.27	0.48	2.55
16	Bremen	747366	0.47	0.27	0.56	2.76
17	Dresden	791237	0.40	0.27	0.43	2.15
18	Stuttgart	2482676	0.43	0.26	0.47	2.43
19	Hannover	1144481	0.42	0.23	0.53	2.89
20	Kiel	724185	0.43	0.21	0.52	3.26
21	Saarbrücken	658124	0.38	0.19	0.51	3.09
22	München	2731249	0.40	0.16	0.53	3.59
23	Leipzig	1016485	0.35	0.16	0.61	3.96
24	Nürnberg	1081648	0.30	0.12	0.52	4.32
25	Augsburg	662890	0.32	0.09	0.52	6.01

This Table reports estimates of absolute and relative mobility for the 25 largest LLM's in Germany (measured by total population in 2015). Local labor markets are sorted in descending order by absolute mobility rank (measured using the $Q1$ measure). The first column, labeled $Q1$ Rank, reports the rank of the LLM based on the $Q1$ measure relative to the other LLM's reported in the Table.

2.8.5 The Relationship Between Different Measures of Intergenerational Mobility

In this Section, we discuss the relationship between our measures of relative and absolute mobility across all LLM's in Germany. Are regions of high relative mobility also regions of high absolute mobility? In order to answer this question, we investigate the relationship of our estimates of relative and absolute mobility across local labor markets.

For the purpose of comparing intergenerational mobility across all 258 local labor markets, we additionally report the results for a third measure of relative mobility: the difference in the probability of obtaining an A-Level degree between children from the top quintile of the income distribution and children from the bottom quintile of the income distribution:

$$Q5 - Q1 = E(Y_i = 1|R_i \geq 80) - E(Y_i = 1|R_i < 20) \tag{2.9}$$

We refer to this measure as $Q5-Q1$. Conceptually, this measure is similar to the parental income gradient, but does not rely on the assumption that the relationship between parental income rank and a child's probability to obtain an A-Level degree is well approximated by a linear function in all regions. Hence the $Q5 - Q1$ measure serves as a robustness check for the results we obtain when comparing the parental income gradient across local labor markets.

The correlation between the different income-based measures of intergenerational mobility across LLM's is reported in Table 2.5. In Table 2.5, we report the correlation ρ (unweighted, and weighted by the number of observations in each LLM) and the (spearman) rank correlation r of different intergenerational mobility measures across LLM's (treating each LLM as one observation). In addition to the correlation between our measures of absolute and relative mobility, we also report the correlation of these measures with the unconditional A-Level share in every LLM.

Table 2.5: Correlation between Intergenerational Mobility Measures Across Local Labor Markets

Intergen. Mob. Measure	Correlation	$Q1$	α	$\beta \times 100$	$Q5 - Q1$	$Q5/Q1$	A-Level Share
$Q1$	ρ	1	0.87	-0.26	-0.22	-0.78	0.73
	ρ (weighted)	1	0.91	-0.18	-0.15	-0.82	0.80
	r	1	0.88	-0.24	-0.21	-0.83	0.73
α	ρ	*	1	-0.41	-0.32	-0.76	0.70
	ρ (weighted)	*	1	-0.29	-0.19	-0.78	0.78
	r	*	1	-0.39	-0.31	-0.82	0.71
$\beta \times 100$	ρ	*	*	1	0.86	0.55	0.27
	ρ (weighted)	*	*	1	0.87	0.45	0.26
	r	*	*	1	0.84	0.61	0.23
$Q5 - Q1$	ρ	*	*	*	1	0.55	0.28
	ρ (weighted)	*	*	*	1	0.47	0.27
	r	*	*	*	1	0.65	0.23
$Q5/Q1$	ρ	*	*	*	*	1	-0.41
	ρ (weighted)	*	*	*	*	1	-0.53
	r	*	*	*	*	1	-0.40

This Table reports the correlation between estimates of different measures of intergenerational mobility across LLM's in Germany. ρ denotes the (Pearson) correlation coefficient of two measures across LLM's. ρ (weighted) denotes the (Pearson) correlation coefficient of two measures across LLM's weighted by the number of observations in each LLM. r denotes the (Spearman) rank correlation coefficient of two measures across LLM's.

First, we note that the correlation between quintile-based and regression based measures that are conceptually similar is very high. The correlation between the $Q1$ measure and the baseline probability α estimated from Regression 2.8 is very high across LLM's. Also the correlation between the parental income gradient $\beta \times 100$ and the $Q5 - Q1$ measure is substantial. This reassures us that the relationship between the probability of a child obtaining an A-Level degree and parent income rank is indeed well approximated by a linear function in all LLM's.

The relationship between absolute mobility and relative mobility across LLM's is of moderate size. It is true that, on average, LLM's with a higher degree of upward mobility (high Q_1 or high α) are also LLM's with a higher degree of relative mobility (low parental income gradient, or low $Q5 - Q1$ measure), but this relationship is far from perfect. Across all LLM's, it is not true that absolute and relative mobility go "hand in hand", when measuring relative mobility by the parental income gradient or by the $Q5 - Q1$ measure. The correlations reported here are, for example, lower than the cross-commuting zones correlation of absolute and relative mobility reported in Chetty et al. (2014). Chetty et al.

(2014) report a correlation of -0.61 between their measures of absolute and relative mobility. This means that in Germany the existence of regions described by low absolute mobility and high relative mobility is not uncommon. For example, the local labor market of Burghausen in Bavaria is described by low absolute mobility ($\alpha = 0.11$, $Q1 = 0.15$), but also by high relative mobility, i.e. a low parental income gradient ($\beta \times 100 = .23$, $Q5 - Q1 = .20$). An example of a LLM that exhibits high absolute mobility, but low relative mobility is Berlin. Berlin is described by relatively high absolute mobility ($\alpha = 0.23$, $Q1 = 0.28$), but at the same time by relatively low relative mobility, i.e. a high parental income gradient ($\beta \times 100 = 0.56$, $Q5 - Q1 = 0.47$). These results highlight the importance of the exact normative concept one has in mind when comparing intergenerational mobility across regions. It might well be the case that a local labor market has low relative mobility, but high absolute mobility, or vice versa.

The $Q5/Q1$ ratio exhibits a strong correlation with both, measures of absolute mobility (a negative one) and other measures of relative mobility (a positive one). By construction, the $Q5/Q1$ measure depends on both, so this strong relation is not unexpected.

LLM's with higher unconditional A-Level shares are on average also described by higher measures of absolute mobility - a higher A-Level share "lifts all the boats" (although part of this correlation is, of course, mechanical). Interestingly, LLM's with higher unconditional A-Level shares are also described by higher parental income gradients, but this correlation is rather moderate.

The results reported in Table 2.5 highlight the importance of the exact normative concept one has in mind when comparing intergenerational mobility across LLM's in Germany.

2.8.6 Comparison with Educational Intergenerational Mobility

Our analysis so far has focused on the relationship between parental income and the educational attainment of children. In this Section, we analyze how regional variation in our measures of intergenerational mobility compares to regional variation in measures of intergenerational mobility that are based on the association between parent and child education levels.

For the purpose of this comparison, we define three measures of intergenerational mobility that are based on the association between the educational attainment of children and the educational attainment of their parents.

The first measure we consider is the probability of obtaining an A-Level degree for children where neither parent has obtained an A-Level degree. We refer to this measure as Educational Upward Mobility.

$$\text{Educational Upward Mobility} = E(Y_i = 1|\text{Neither Parent has an A-Level degree}) \quad (2.10)$$

Conceptually, this measure resembles our measures of absolute mobility. We refer to this measure as educational upward mobility. The heatmap presenting the distribution of this measure across LLM's is shown in Figure 2.11. At the national level, our measure of Educational Upward Mobility is given by .283 (i.e. 28.3% of children where neither parent has an A-Level degree obtain an A-Level degree).

Figure 2.11: Educational Upward Mobility Across LLM's

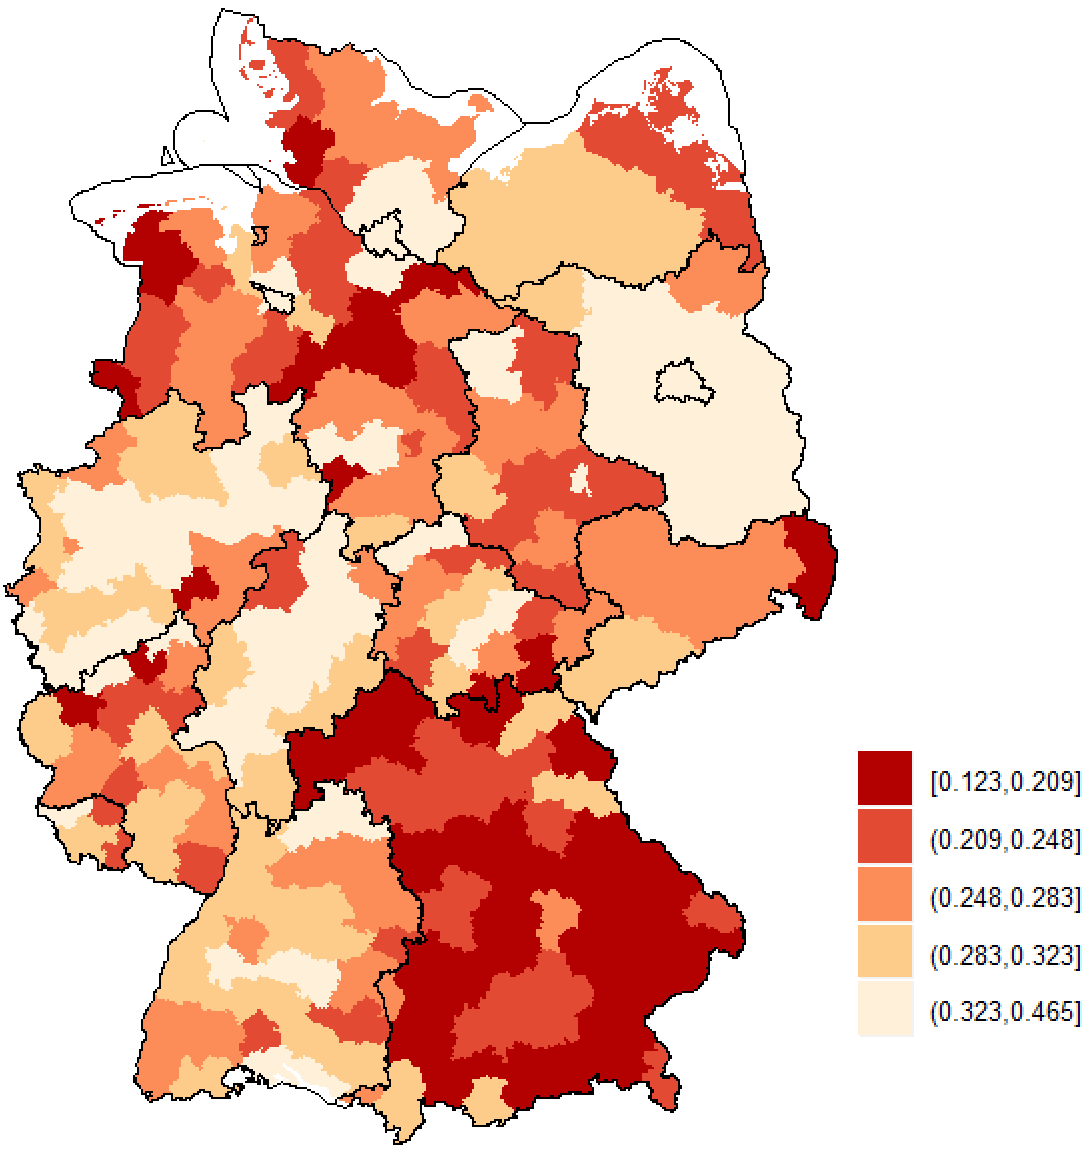

This Figure shows Educational Upward Mobility across LLM's. Educational Upward Mobility is defined as the fraction of children who obtain an A-Level degree among children where neither parent has an A-Level degree.

The second measure we consider is the difference in the probability of obtaining an A-Level degree between children where at least one parent has an A-Level degree and children where neither parent has an A-Level degree. We refer to this measure as Education Gradient.

$$\text{Education Gradient} \\ = E(Y_i = 1|\text{At least one parent has an A-Level}) - E(Y_i = 1|\text{Neither Parent has an A-Level}) \quad (2.11)$$

Conceptually, this measure resembles the parental income gradient (and the $Q5 - Q1$ measure). The heatmap presenting the distribution of this measure across LLM's is shown in Figure 2.12. At the national level, the Education Gradient is given by .355 (i.e. the probability of obtaining an A-Level degree of a child is 35.5 percentage points higher when at least one parent has an A-Level degree, compared to a child where neither parent has an A-Level degree).

Figure 2.12: Education Gradient Across LLM's

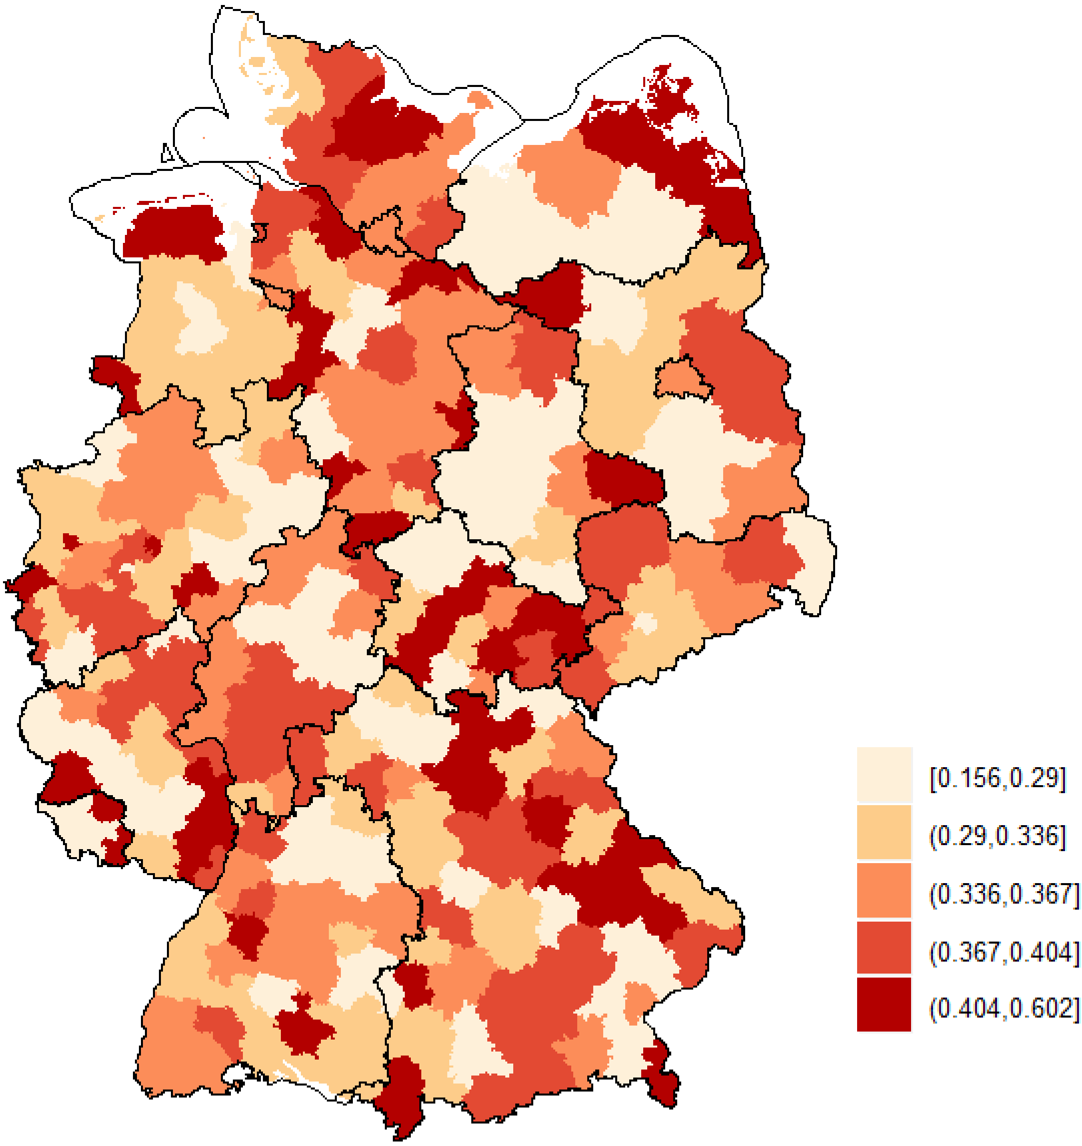

This Figure shows the Education Gradient across LLM's. The Education Gradient is defined as the difference in the probability of obtaining an A-Level degree between children where at least one parent has an A-Level degree and children where neither parent has an A-Level degree.

The third (and last) measure of educational mobility we report is the relative likelihood of obtaining an A-Level degree between children from parents with an A-Level degree and children from parents without an A-Level degree. We refer to this measure as Education Ratio.

$$\text{Education Ratio} = \frac{E(Y_i = 1|\text{At least one parent has an A-Level})}{E(Y_i = 1|\text{Neither Parent has an A-Level degree})} \tag{2.12}$$

Conceptually, this measure resembles our $Q5/Q1$ measure. At the national level, the education ratio is given by 2.25. The heatmap presenting the distribution of this measure across LLM's is shown in Figure 2.13.

Figure 2.13: Education Ratio Across LLM's

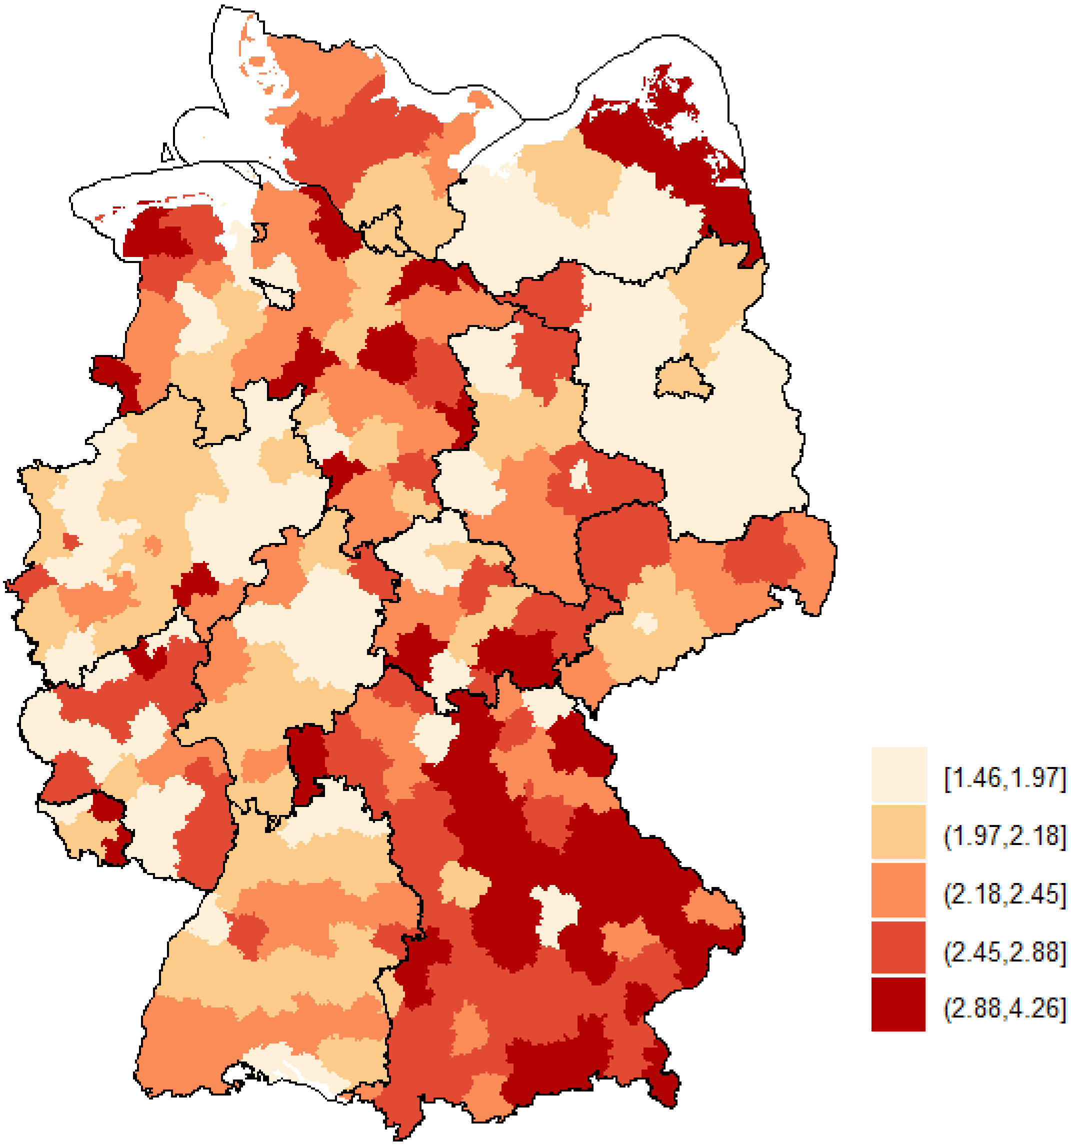

This Figure shows the Education Ratio across LLM's. The Education Ratio is defined as the relative likelihood of obtaining an A-Level degree between children from parents with an A-Level degree and children from parents without an A-Level degree

Generally, the results we obtain for education based measures of intergenerational mobility at the national level are very similar to the results reported in Riphahn and Trübswetter (2013) and Klein et al. (2019) for similar measures. It is also worth noting that the regional variation reported here is larger than the variation in similar measures over time, as reported by Riphahn and Trübswetter (2013) and Klein et al. (2019).

In Table 2.6 we report the correlation between the different income-based and education-based measures of intergenerational mobility across LLM's. In Table 2.6, we report the correlation ρ (unweighted, and weighted by the number of observations in each LLM) and the (Spearman) rank correlation r across LLM's, treating every LLM as one observation. The correlation between measures that are conceptually similar are printed in **bold**.

Table 2.6: Correlation between Income-Based and Education-Based Measures of Intergenerational Mobility Across LLM's

Education-Based Measures	Correlation Measure	$Q1$	α	$\beta \times 100$	$Q5 - Q1$	$Q5/Q1$
Educational Upward Mobility	ρ	**0.71**	**0.76**	0.16	0.17	-0.45
	ρ (weighted)	**0.80**	**0.84**	0.13	0.16	-0.58
	r	**0.73**	**0.77**	0.11	0.12	-0.45
Education Gradient	ρ	-0.06	-0.21	**0.38**	**0.34**	0.21
	ρ (weighted)	-0.05	-0.17	**0.40**	**0.37**	0.19
	r	-0.06	-0.19	**0.40**	**0.33**	0.22
Education Ratio	ρ	-0.53	-0.64	0.10	0.07	**0.47**
	ρ (weighted)	-0.64	-0.72	0.09	0.05	**0.58**
	r	-0.52	-0.63	0.11	0.26	**0.42**

This Table reports the correlation between different parental education-based measures of intergenerational mobility and different parental income-based measures of intergenerational mobility across LLM's in Germany. Correlations printed in **bold** indicate that measures are conceptually similar.

The results reported in Table 2.6 show that regional patterns in intergenerational mobility are generally rather similar, regardless of whether intergenerational mobility measures are based on parental income or on parental education levels. LLM's described by high absolute mobility are - on average - also LLM's desribed by high educational upward mobility. Similarly, LLM's described by high parental income gradient are on average also described by a high education gradient, albeit this correlation is a little weaker compared to the measures of absolute mobility. The same is true for the $Q5/Q1$ ratio and the Education Ratio.

It should also be noted that we also find a rather weak correlation between measures of absolute mobility and relative mobility across LLM's when using education-based measures of intergenerational mobility. The correlation between Educational Upward Mobility and the Education Gradient across LLM's is $\rho = -0.21$ (observation-weighted $\rho = -0.20$, $r = -0.22$). This again highlights the importance of the exact normative concept one has in mind when comparing intergenerational mobility across regions.

2.8.7 Robustness of Regional Patterns

In this Section, we present evidence on the robustness of our results for geographical variation in our measures of intergenerational mobility.

One potential concern with our regional comparisons is that the general price level might differ across regions. As a consequence, real income difference across regions might be masked when not adjusting for price level differences[37], leading to an overestimation of the parental income gradient in regions with low price levels. Weinand and von Auer (2019) calculate regional price level indices across German counties based official German CPI microdata from May 2016. They find that the average price level in the most expensive county (Munich) is 27 percent higher than the price level in the cheapest county and that price level differences are mainly driven by differences in housing costs. To address the issue of regional price level differences, we calculate all income-based measures of intergenerational mobility based on state-specific income distributions, i.e. we assign every household its income rank based on the income distribution in its state of residence (instead of the national income distribution).

Furthermore, we present evidence that regional differences in intergenerational mobility are not sensitive to the exact age cut-offs. We calculate all measures of regional mobility for samples only including children ages 16 to 19, 17 to 20 and 20 to 22, the same samples that are used in Section 2.6.3.

Lastly, we also report all results when using parental income not adjusted for household size in our main sample (instead of using parental income per capita).

The results of all robustness checks are summarized in Table 2.7. In Table 2.7 we present the correlation ρ (unweighted, and weighted by the number of observations in each LLM) and the (Spearman) rank correlation r across LLM's between our baseline results and the results of the robustness checks described here.

[37]For example, a household with a monthly income of 3000 Euro might - in real terms - be richer in Madgeburg than in Munich.

Table 2.7: Robustness of Regional Variation

Robustness Check	Correlation Measure	$\beta \times 100$	$Q1$	$Q5/Q1$	α	$Q5 - Q1$
State-Specific	ρ	0.99	0.97	0.92	0.98	0.94
	ρ (weighted)	0.98	0.98	0.93	0.98	0.94
Income Distribution	r	0.98	0.97	0.96	0.98	0.93
Children Aged 16 to 19	ρ	0.87	0.88	0.74	0.86	0.84
	ρ (weighted)	0.86	0.91	0.83	0.91	0.85
	r	0.85	0.88	0.81	0.85	0.83
Children Aged 17 to 20	ρ	0.90	0.92	0.84	0.94	0.87
	ρ (weighted)	0.89	0.94	0.88	0.96	0.87
	r	0.89	0.92	0.85	0.93	0.84
Children Aged 20 to 22	ρ	0.71	0.82	0.59	0.83	0.65
	ρ (weighted)	0.69	0.85	0.63	0.87	0.63
	r	0.68	0.82	0.72	0.84	0.63
Income not adjusted	ρ	0.88	0.89	0.74	0.93	0.80
	ρ (weighted)	0.88	0.92	0.79	0.95	0.79
for HH Size	r	0.87	0.89	0.82	0.93	0.78

This Table reports the correlation of different measures of intergenerational mobility estimated using different sample specifications with the same measures of intergenerational mobility estimated from our main sample, calculated across LLMs.

The results reported in Table 2.7 show that our results are robust in various dimensions. The geographical variation across LLM's is very similar for all measures of intergenerational mobility for all robustness checks considered in Table 2.7. All this reassures as that differences in intergenerational mobility across LLM's are indeed a robust feature of our data and not simply driven by noise or measurement error.

Next, we turn to the question to what extent differences in household characteristics can explain differences in intergenerational mobility across LLM's.

2.9 Can Household Characteristics Explain Regional Differences in Intergenerational Mobility?

In the previous Section, we have documented substantial and robust geographical variation in different measures of intergenerational mobility across regions in Germany. A natural question to ask is which factors drive the observed differences in intergenerational mobility. This Section investigates to what extent observed differences in intergenerational mobility measures across LLM's can be explained by differences in household characteristics across LLM's.

2.9.1 Household Characteristics and Regional Estimates of Intergenerational Mobility

In the US, Gallagher et al. (2018) show that household characteristics differ greatly across commuting zones used in Chetty et al. (2014). Gallagher et al. (2018) suggest that a substantial share of the geographical variation in the intergenerational mobility measures reported in Chetty et al. (2014) can be explained by differences in household characteristics across commuting zones.

Addressing the question to what extent differences in household characteristics can explain the geographical variation in our mobility measures is highly relevant for the interpretation of our findings. To illustrate the mechanism how regional differences in household characteristics can potentially create spurious geographical variation in intergenerational mobility measures, consider the following scenario.

Suppose the true model of intergenerational mobility is given by Equation 2.13:

$$Y_i = \alpha + \beta R_i + \gamma X_i + \epsilon_i \tag{2.13}$$

where Y_i is our binary child outcome variable, R_i is parental income rank and X_i denotes other household (or child) characteristics that affect the educational attainment of a child, e.g. parental education. Furthermore, suppose that $COV(R_i, X_i) \neq 0$ and $\gamma \neq 0$, i.e. other household characteristics X_i are correlated with parent income and these characteristics have an effect on the educational attainment of a child. Additionally suppose that the distribution of X_i differs across local labor markets l (for example because households sort into different local labor markets l based on household characteristics X_i). Importantly, note that none of the coefficients has a local labor market specific subscript l, i.e. the place of residence has no effect on the outcome Y_i of a child i. If the true model of intergenerational mobility is given by Equation 2.13, but one instead estimates Equation 2.8 (without controlling for additional household characteristics), one can find substantial variation in the local labor market specific coefficient estimates α_l and β_l, due to the fact that omitted variable bias differs across local labor markets[38]. If the true model of intergenerational mobility were given by Equation 2.13, neighborhood effects would play no role in shaping intergenerational mobility and explaining geographical differences in intergenerational mobility would boil down to controlling for differences in family characteristics across local labor markets.

In this Section, we investigate how our local labor market specific estimates of intergenerational mobility change once we control for additional household characteristics. As our data set contains rich information on all household members, it is well suited to study to what extent geographical variation in our measures of intergenerational mobility is driven by differences in household characteristics across LLM's. Before describing the estimation method, Table 2.8 presents an overview over the additional control variables X_i that we consider in this Section.

Table 2.8: Additional Household Characteristics

Control Variable	Encoding
Household Type	Three Categories
Highest Parental Education Level	Six Categories
No. of Children in Househild	Five Categories
Migrational Background	Two Categories
Year Dummies	Seven Categories
Gender of the Child	Two Categories

This Table reports the household characteristics used as additional control variables in this Section. All variables are transformed into dummy variables. Encoding denotes the number of (mutually exclusive) dummy categories for each of the variables.

Here, we describe the variables presented in Table 2.8 in more detail. To control for differences in the household type, we add three dummy variables indicating the type of the household (two parent household, single mother household, single father household).

[38]Keep in mind that we assume that the distribution of X_i differs across local labor markets.

To control for the highest parental education level, we add six dummies indicating the highest parental education level. The six parental education dummies are a University degree dummy, a dummy indicating an A-Level degree + vocational training, a dummy indicating an A-Level degree without vocational training, a dummy indicating vocational training but no A-Level, a dummy indicating a secondary school degree without vocational training and a dummy indicating no degree. Note that this dummy always indicates the highest education level of either parent[39]. To control for differences in the number of siblings, we add five dummies indicating the number of children in the household. These dummies indicate the presence of one, two, three, four or five or more children in the household, respectively. We add a dummy variable indicating migrational background, which is defined as before. Additionally, we add a full set of year dummies, and a dummy variable indicating the gender of the child.

2.9.2 National Results

Before turning to the results for regional variation, we describe how our estimates of intergenerational mobility change at the national level when we control for additional household characteristics.

For the regression-based measures of intergenerational mobility at the national level, we estimate the coefficients of a regression of the A-Level dummy on parental income rank, including the additional control variables describe above, i.e. we estimate the regression described in Equation 2.14:

$$Y_i = \alpha + \beta R_i + \gamma X_i + \epsilon_i \tag{2.14}$$

where X_i denotes the set of additional household-level control variables described above.

Additionally, we calculate the probability of obtaining an A-Level degree for every percentile in the national income distribution when controlling for other household characteristics as specified in Equation 2.15:

$$Y_i = \alpha + \sum_{k=0}^{99} \theta_k I(R_i = k) + \gamma X_i + \epsilon_i \tag{2.15}$$

[39]i.e. if the mother has an university degree and the father has vocational training as highest education level, we assign this household a university degree dummy.

where $I(R_i = k)$ is a dummy variable that is equal to one if child i is in percentile k of the national parental income distribution, and zero otherwise.

Figure 2.14 illustrates how the relationship between parental income rank and child's educational attainment changes at the national level when controlling for additional household characteristics as described above. The red dots in Figure 2.14 show the relationship between parental income rank and the educational attainment of a child when not controlling for other household characteristics, i.e. the results already presented in Figure 2.2. The gray dots in Figure 2.14 show the fraction of children in each income percentile calculated as described in Equation 2.15. Note that in Equation 2.15, the probability to obtain an A-Level degree does not only depend on the parental income percentile, but also on all other household characteristics. In Figure 2.14, other household characteristics are fixed to show a two-parent household, with two children, for a male child, in the year 2009, without migrational background and the highest parental education level being an A-Level degree + vocational training, for all percentiles in the national income distribution. Note that as other household characteristics enter additively, they shift the probability of obtaining an A-Level degree equally for children from all percentile ranks equally.

Figure 2.14: Intergenerational Mobility at the National Level: Controlling for Additional Household Characteristics

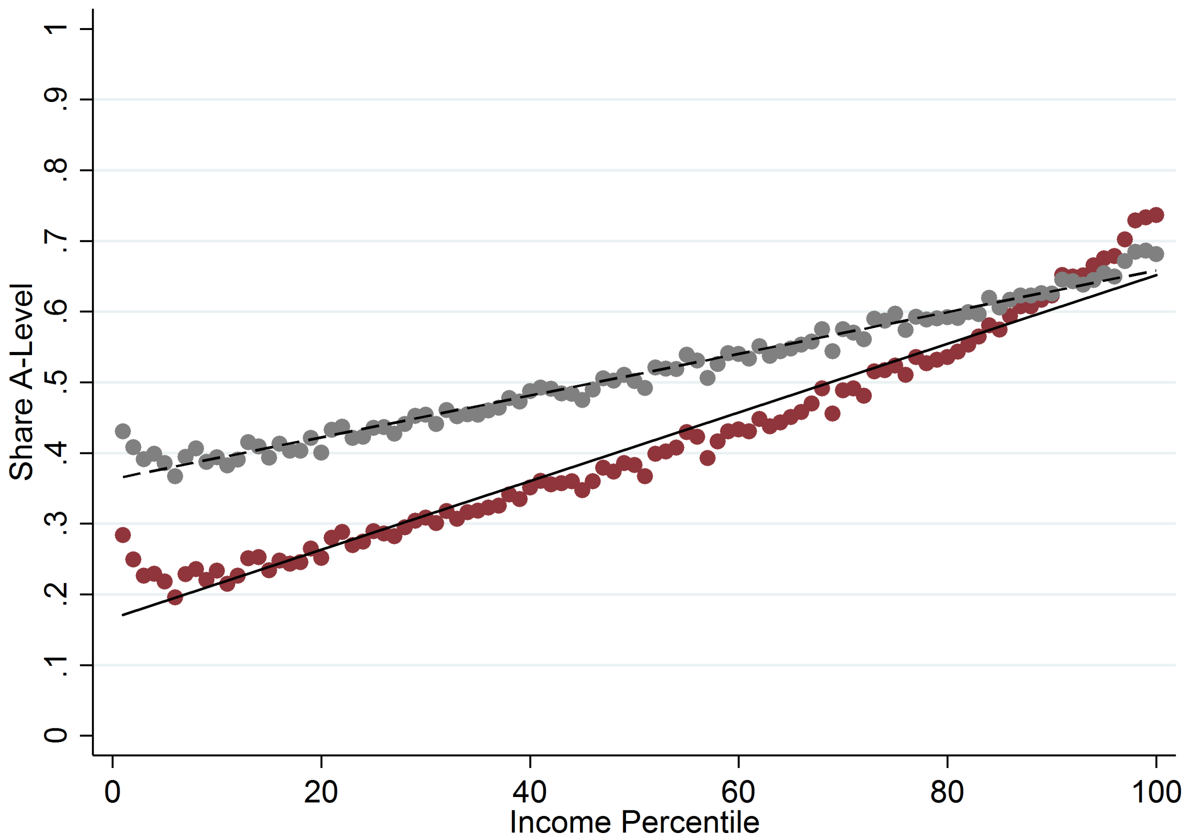

This Figure plots the fraction of children aged 16-22 (y-axis) that are either enrolled in the last two/three years of A-Level track or have already completed it vs. the percentile rank of their parents in the national net income per household member distribution (x-axis). The red dots depict the fraction of children aged 16-22 currently enrolled in the last two/three years of the A-Level track or having already completed it vs. the parent rank in each percentile bin when not controlling for additional household characteristics. The gray dots show the fractional of children ages 16-22 currently enrolled in the last two years of the A-Level track or having already completed it vs. the parent rank in each percentile bin when controlling for additional household characteristics as described in Equation 2.15. The gray dots fix other household characteristics to show a two-parent household, with two children, for a male child, in the year 2009, without migrational background and the highest parental education level being an A-Level degree + vocational training. The Figure is based on $N = 268523$ children living with their parents. The OLS regression of the A-Level dummy on the parent income rank without controlling for other household characteristics (solid black line) yields a slope coefficient of 0.0048 and a constant of 0.166. The OLS regression of the A-Level dummy on the parent income rank when controlling for additional household characteristics (dashed black line) yields a slope coefficient of 0.0028 and a constant of 0.321.

Two things should be noted in Figure 2.14. First, when controlling for additional household characteristics, the relationship between parental income rank and the probability to obtain an A-Level degree becomes even more linear. Second, the parental income gradient shrinks: the parental income gradient $\beta \times 100$ drops from .485 (shown in the solid black line) to .287 (shown in the dashed black line).

2.9.3 Regional Results

Next, we turn to the question how controlling for additional household characteristics changes our estimates of intergenerational mobility at the regional level. We estimate regression-based measures of intergenerational mobility using Equation 2.16:

$$Y_{i,l} = \alpha_l + \beta_l R_i + \gamma X_i + \epsilon_{i,l} \tag{2.16}$$

As before, i is used to denote children, and l denotes LLM's. We constrain the effects of additional household characteristics to be equal in all local labor markets, i.e. γ does not have a local-labor market specific subscript. A similar approach has been employed by Chetty et al. (2018) to study differences in intergenerational mobility between different ethnic groups in the US.

To calculate quintile-based measures of intergenerational mobility, we estimate Equation 2.17:

$$Y_{i,l} = \alpha_l + \sum_{k=1}^{5} \theta_l I(R_i \in Q_k) + \gamma X_i + \epsilon_{i,l} \tag{2.17}$$

where $I(R_i \in Q_k)$ is an indicator function, equal to one if child i is in the $k-th$, $k = 1, 2, 3, 4, 5$, quintile of the national parent income distribution, and zero otherwise.

Figure 2.15 shows the scatterplot of the LLM-Specific parental income gradients estimated when not controlling for additional household characteristics (Equation 2.8) on the X-Axis, and the LLM-specific parental income gradient estimated from Equation 2.16 on the Y-Axis. The red line shows the linear fit.

As can be seen from Figure 2.15, controlling for additional household characteristics effects most of the LLM-specific estimates of the parental income gradient equally. The correlation between the LLM-specific parental income gradients estimated from Equation 2.8 and Equation 2.16 across LLM's is $\rho = 0.92$. However, similar to the results obtained at the national level, controlling for additional household characteristics flattens the parental income gradient. When not controlling for additional household characteristics, the (unweighted) mean of the parental income gradient across LLM's is .47, whereas it shrinks to .30 when controlling for additional household characteristics.

Figure 2.15: Scatterplot of Parental Income Gradients

This Figure shows the scatterplot of LLM-specific parental income gradients estimated when not controlling for additional household characteristics (X-Axis) and LLM-specific Parental Income Gradients estimated when controlling for additional household characteristics as described in Equation 2.16 (Y-Axis)

Table 2.9 summarizes the correlation between the baseline estimates presented in Section 2.8 and the estimates obtained from Equation 2.16 and Equation 2.17 in detail.

Table 2.9: Robustness of Regional Variation - Household Characteristics

Robustness Check	Correlation Measure	$\beta \times 100$	$Q1$	α	$Q5 - Q1$
Controlling for Additional Household Characteristics	ρ	0.92	0.96	0.97	0.93
	ρ (weighted)	0.97	0.98	0.97	0.92
	r	0.91	0.97	0.97	0.91

This Table reports the correlation of different measures of intergenerational mobility estimated when controlling for additional household characteristics with the same measures of intergenerational mobility estimated from our main sample across LLM's.

Table 2.9 shows that controlling for additional household characteristics has barely any effect on the LLM-specific estimates we obtain. The correlation between LLM-specific estimates when not controlling for other household characteristics and LLM-specific estimates when controlling for additional household characteristics is always larger than 0.9, suggesting that differences in intergenerational mobility across LLM's are not simply driven by differences in household characteristics across LLM's.[40] This suggests that neighborhood effects (i.e. LLM characteristics) play a role in shaping intergenerational mobility and that differences in intergenerational mobility measures are simply explained by differences in household characteristics across LLMs. In the next Section, we investigate which LLM characteristics correlate with intergenerational mobility.

[40]The comparison of $Q5/Q1$ measure between both approaches is omitted. All other measures of intergenerational mobility we consider are linear measures, so the comparison between both approaches is independent of the reference group (i.e. the value of other household characteristics). The $Q5/Q1$ ratio is nonlinear in both $Q5$ and $Q1$, hence when comparing it between approaches, one would need to make a comparison for each reference group (i.e. for each possible combination of X_i). This would result in a large amount of comparisons and is hence omitted.

2.10 Spatial Correlates of Intergenerational Mobility

In this Section, we take a first step at describing which local labor market characteristics correlate with our measures of intergenerational mobility. We do not claim that these correlations should be interpreted as a causal relationship, but they can help to guide future research in the search for causal determinants of intergenerational mobility.

To construct local labor market characteristics, we use a large set of regional indicators compiled by the Federal Institute for Building, Urban Affairs and Spatial Research ("Bundesinstitut für Stadt-, Bau- und Raumforschung", in short BBSR). The BBSR maintains a database of regional indicators (called "Inkar"), containing around 600 different regional indicators[41], covering statistics on labor market conditions, housing structure and housing stock, demographics, education, income, debt, environmental indicators, access to medical services, public finances, social security payments, infrastructure and economic activity at the level of local labor markets, using the same definition of local labor markets we employ in our analysis. The data available in the database is collected by the BBSR from various government bodies in Germany (e.g. the German statistical office DESTATIS and the Institute for Employment Research IAB). Most variables in the database are available at annual frequency[42].

We use this database to find local labor market characteristics correlated with our measures of intergenerational mobility. We construct local labor market characteristics as the time averages of variables over the years 2009 to 2015 at the local labor market level (i.e. for every variable, for every local labor market, we calculate the average value of the variable over the years 2009 to 2015 and use this as local labor market characteristic)[43]. Here, we take a first step at describing which local market characteristics are correlated with our measure of intergenerational mobility. To do so, we report the ten local labor market characteristics with the highest (absolute) correlation with our measures of intergenerational mobility[44]. We focus on two measures of intergenerational mobility in this Section: The $Q1$

[41]The database can be reached under the the website https://www.inkar.de/

[42]Not all variables are available at annual frequency. Data on election results is e.g. available at the same frequency as elections take place.

[43]If a variable is available only for a short time period, we construct the time average using all available observations instead.

[44]In reporting the ten local labor market characteristics with the highest correlation, we omit duplications of local labor market characteristics from the analysis, i.e. local labor market characteristics that are very similar. For example, our database contains the population share of 6 to 18 years olds, and the population share of 6 to 20 years olds in every local market. For variables that are very similar, like the two in the

measure and the parental income gradient $\beta \times 100$. The results of this exercise can be found in Table 2.10 for the $Q1$ measure and in Table 2.11 for the parental income gradient.

example here, we only consider one of the variables, namely the one with the higher correlation with the intergenerational mobility measure.

Table 2.10: Local Labor Market Characteristics Correlated with the Absolute Mobility Measure $Q1$

Rank	Local-Labor Market Characteristic	ρ	r
1	Labor Force Participation, Young People	-0.43	-0.43
2	Social Aid Recipients under 18	-0.42	-0.45
3	Share of Employees w/ Vocational Training	-0.40	-0.40
4	Vote-Share "Other Parties"	-0.39	-0.38
5	Long-Term Unemployment Share, Males	0.38	0.38
6	Vocational Training Spots	-0.37	-0.42
7	Naturalizations per capita	0.35	0.33
8	Open-Space Area	-0.35	-0.31
9	Unemployment Rate, Male Foreigners	0.35	.36
10	Long-Term Unemployment Share	0.34	0.34

This Table reports the ten local labor market characteristics available in the database "Inkar" with the highest (absolute) correlation with the $Q1$ measure of intergenerational mobility. Correlations are calculated across local labor markets, treating every local labor market as one observation. The column labeled ρ denotes the Pearson correlation, the column labeled r denotes the Spearman Rank Correlation. Variables are sorted in descending order by absolute correlation ρ with the $Q1$ measure. The fist column, labeled Rank, denotes the rank of the local labor market characteristic based on the absolute correlation with the $Q1$ measure relative to all other local labor market characteristics in the database. Labor Force Participation, Young People denotes the share of 15 to 30 year olds being employed in a regular working position, among all 15 to 30 year olds ("Quote Jüngere Beschäftigte"). Social Aid Recipients under 18 denotes the share of social aid recipients that are younger than 18 among all social aid recipients ("Leistungsempfänger nach SGB XII, Kap. 5-9, unter 18 Jahren"). Share of Employees w/ Vocational Training denotes the share of employess with vocational training (but no college degree) among all employees ("Anteil Beschäftigte mit Berufsabschluss"). Vote Share "Other Parties" denotes the share of votes received by parties not represented in the Federal Parliament, in the election of the federal parliament ("Stimmenanteil Sonstige Parteien"). Long-Term Unemployment Share, Males denotes the share of long-term (more than one year) unemployed males among all unemployed males ("Männliche Langzeitarbeitslose"). Vocational Training Spots denotes the fraction of the number of open and filled vocational training positions over the number of filled vocational training spots and applicants who did not find a vocational training position and are not in education ("Ausbildungsplätze"). Naturalizations per capita denotes the number of naturalizations per capita ("Einbürgerungen je Einwohner"). Open-Space Area denotes the share of the area not occupied by buildings among the total area ("Anteil Freifäche"). Unemployment Rate, Male Foreigners denotes the unemployment rate among male foreigners ("Männliche ausländische Arbeitslose"). Long-Term Unemployment Share denotes the share of unemployed persons with an unemployment duration of more than one year, among all unemployed persons ("Langzeitarbeitslose").

Table 2.11: Local Labor Market Characteristics Correlated with the Parental Income Gradient

Rank	Local-Labor Market Characteristic	ρ	r
1	Vote Share "CDU/CSU"	-0.38	-0.35
2	Average Pension Payments	0.37	0.34
3	Share of Female Employees w/ College Degree	0.36	0.41
4	Share of Dwellings w/ 5 or more Rooms	-0.35	-0.36
5	Share Marginal Employees, 65 and Older	-0.35	-0.36
6	Population Share 6 to 18 Years	-0.35	-0.36
7	Share of Young Unemployed	-0.35	-0.32
8	Daycare Share under 3-Year Olds	0.34	0.32
9	Share of Dwellings in Detached Houses	-0.33	-0.34
10	Share of Employees in Vocational Training	-0.33	-0.34

This Table reports the ten local labor market characteristics available in the database "Inkar" with the highest (absolute) correlation with the parental income gradient. Correlations are calculated across local labor markets, treating every local labor market as one observation. The column labeled ρ denotes the Pearson correlation, the column labeled r denotes the Spearman Rank Correlation. Variables are sorted in descending order by absolute correlation ρ with the parental income gradient. The fist column, labeled Rank, denotes the rank of the local labor market characteristic based on the absolute correlation with the parental income gradient relative to all other local labor market characteristics in the database. Vote Share "CDU/CSU" denotes the share of votes received by the CDU/CSU in the election of the federal parliament ("Stimmenanteil CDU/CSU"). Average Pension Payments denotes the average yearly pension payments received by pensioners over 65 ("Durchschnittlicher Rentenzahlbetrag"). Share of Female Employees w/ College Degree denotes the share of female employees with an academic degree among all female employees ("Quote weibliche Beschäftigte mit akademischem Abschluss"). Share of Dwellings w/ 5 or more Rooms denotes the share of dwelling units with 5 or more rooms among all dwelling units ("Anteil 5- und mehr Raum-Wohnungen"). Share Marginal Employees, 65 and Older denotes the share of employees aged 65 and older, holding a marginal employee position, among all inhabitants aged 65 to 75 ("Quote geringfügige Beschäftigte über 65 Jahre"). Population Share 6 to 18 Years denotes the share of inhabitants aged 6 to 18 years among the total population ("Einwohner von 6 bis unter 18 Jahren"). Share of young Unemployed denotes the share of unemployed persons under 25 years among all unemployed persons ("Anteil jüngere Arbeitslose"). Daycare Share under 3-Year Olds denotes the share of children under 3 years enrolled in full-day daycares among all children under 3 years ("Ganztags-Betreuungsquote Kleinkinder"). Share of Dwellings in Detached Houses denotes the share of dwelling units that are detached houses among all dwelling units ("Anteil Wohnung in Ein- und Zweifamilienhäusern"). Share of Employees in Vocational Training denotes the share of employees currently in vocational training among all employees ("Auszubildende je 1.000 SV Beschäftigte").

While the correlations reported in Table 2.10 and Table 2.11 do not reveal causal channels through which local labor market characteristics affect intergenerational mobility, they are nevertheless interesting.

Here, we discuss the results for the absolute mobility measure $Q1$ reported in Table 2.10 in more detail. The results reported in Table 2.10 suggest that local labor market conditions can potentially explain (part of the) regional differences in absolute mobility. There is a negative correlation between the $Q1$ measure and both, the labor force participation rate of young people, and the share of vocational training spots. While part of this correlation is of course mechanical (a young person can either work or go to regular school, but not both), it hints at the fact that local employment opportunities affect intergenerational mobility. This is also suggested by the positive correlation between the $Q1$ measure and the long-term unemployment share - suggesting that if the local labor market is in a worse shape, children choose to accumulate more human capital. This is also suggested by the comparatively large correlation between the $Q1$ measure and the share of (all) employees with vocational training. However, the correlations do not reveal whether this is supply- or demand-driven. Furthermore, the results suggest that local labor markets with a larger share of a marginalized population display worse outcomes - suggested by the (comparatively) high negative correlation with the share of social aid recipients under 18, and the negative correlation with the vote share of "Other Parties". The results also suggest that absolute intergenerational mobility is lower in rural areas (as indicated by the correlation with the share of open space area in a local labor market). The correlation with naturalizations per capita could suggests that local labor markets with better integration of foreigners (i.e. a less segregated society) also exhibit higher absolute mobility.

Next, we discuss the correlations of local labor market characteristics with the parental income gradient. In interpreting these results, keep in mind that high parental income gradients indicate a low level of relative intergenerational mobility. The results reported in Table 2.11 reveal some interesting correlations of local labor markets characteristics with relative mobility rates (i.e. the parental income gradient). Surprisingly, the variable with the highest (absolute) correlation is the vote share of the conservative parties CDU/CSU in federal elections. This correlation could be explained by a number factors (e.g. lower slopes in rural areas, where the conservative party is traditionally stronger) and we do not attempt to list all potential interpretations of this correlation here. Surprisingly, the average pension payments are correlated with lower rates of relative mobility, as well as the share of female employees with a college degree. We find a negative correlation of the parental income gradient with the share of large dwellings, and the share of dwellings in detached houses, suggesting that rural areas generally have lower parental income gradients (as dwellings are larger on average in rural areas). We again find evidence that local labor market conditions affect intergenerational mobility, as suggested by the correlation with the share of employees in vocational training, and the share of young unemployed. Local labor markets with a larger share of young residents are also described by higher rates of relative mobility. However, one needs to be cautious when interpreting single correlations, as some correlations might reflect e.g. differences between federal states, rather than differences between local labor markets. For example, the (surprisingly) positive correlation between the daycare share of under 3 year olds and the parental income gradients vanishes when one uses only within state variation to calculate the correlation[45]. When controlling for state-level fixed effects, the correlation between both variables drops to $\rho = 0.06$. Furthermore, part of this correlation could also reflect differences between rural and urban areas, or other omitted local labor market characteristics.

Overall, the correlations reported here suggest that intergenerational mobility differs between urban and rural local labor markets. Furthermore, the results suggest that local labor market conditions affect intergenerational mobility. The results also suggest that (broadly defined) social capital and social cohesion might affect intergenerational mobility.

We again point out that the bivariate correlations reported here, although interesting, should be interpreted with caution. The bivariate correlations reported here can serve as a starting point for future research. We suggest various dimensions in which our results can be extended. First of all, we did not calculate the correlation between our measures of intergenerational mobility and measures of (income) inequality at the local level. It would be interesting to investigate whether a "Great Gatsby Curve" exists between German local labor markets. Second, it could be very interesting to conduct a multivariate analysis, using the full information contained in the dataset of local market characteristics. The large number of local labor market characteristics contained in the data leads to the need of parsimoniously selecting the covariates included in a multivariate analysis. This could e.g. be achieved by the regularization methods described in Belloni et al. (2014). Alternatively,

[45]This is achieved by first regressing the variables on a full set of state dummies, and then calculating the correlation between residuals from these regressions.

one can reduce the number of variables into a smaller number of principal components, similar to Acciari et al. (2019).

2.11 Conclusion

We have used official microcensus data to document intergenerational mobility in Germany. We measure intergenerational mobility by the association between a childs probability of obtaining an A-Level degree and her parents income rank in the parental income distribution within our sample. We are the first to present detailed evidence on regional differences in intergenerational mobility across local labor markets within Germany. Intergenerational mobility differs greatly between local labor markets. For example, the probability of obtaining an A-Level degree for a child from the bottom quintile of the national income distribution is 32% in Frankfurt am Main, but only 16% in Munich.

We presented evidence that regional differences in intergenerational mobility can't be simply explained by differences in household characteristics across regions. Even when controlling for a range of household characteristics, regional differences in our measures of intergenerational mobility are nearly unchanged. This suggests that the place of residence indeed shapes the economic outcomes of children. We took a first step at describing which regional characteristics are correlated with our measures of intergenerational mobility. Our results suggest that intergenerational mobility differs between urban and rural areas, that local labor market conditions influence intergenerational mobility and that social capital and social cohesion play a role in shaping intergenerational mobility.

For future research it is a worthwhile endeavor to further investigate which regional characteristics shape intergenerational mobility. The educational attainment of children in school greatly shapes their possibilities later in life and we document that educational attainment in secondary school depends strongly on the place of residence. Understanding through which channels places matter for intergenerational mobility is not only important for normative reasons, but can also help to enhance human capital accumulation and hence economic growth.

Appendix A

Appendix to Chapter 1

A.1 Detail on the Construction of Industry-Level Control Variables

This Section describes the calculation of the industry-level control variables presented in Table 1.2 in Section 1.6. The exact calculation procedure of the variables is explained separately for every data source.

Variables calculated from the NBER-CES manufacturing database and (denoted as NBER-CES in Table 1.2) are calculated in two steps. First, the industry average (of the respective variable) is calculated separately for every year of the data. Industry-level averages of ratios are always calculated by first calculating industry-level totals of variables used in the calculation of the ratio, and then taking the ratio of the industry-level totals of the respective variables. In a second step, the final control variable is calculated as the time average of the yearly industry-averages in the sample.

Variables calculated based on yearly Compustat data are calculated in a similar fashion. First, firm-level observations are aggregated to the industry-level separately for every year of the data. This is done by summing firm-level variables at the industry-level. Ratios are then calculated based on industry-level totals of the variables. The final industry-level control variables are calculated as the time average of the yearly industry-level observations over all sample years. Following Ottonello and Winberry (2018), firms with a leverage ratio larger than 10 are dropped from the sample. Only US-based firms are used to calculate industry-level variables.

A.2 Additional Tables

Table A.1: Detailed Summary Statistics for the Frequency of Price Adjustment

	N	mean	sd	min	max	p5	p10	p25	p50	p75	p90	p95
FPA	205	23.37	14.20	4.01	87.53	9.79	11.52	15.08	19.22	25.89	41.68	52.28

This Table reports the summary statistics for the frequency of price adjustment (FPA) at the industry level. p5 to p95 denote the respective percentiles of the distribution within the sample.

Table A.2: Summary Statistics for Additional Control Variables

	(1) N	(2) mean	(3) sd	(4) p5	(5) p95
Frequency of Price Adjustment	205	23.37	14.20	9.797	52.28
Durable Goods Dummy	205	0.522	0.501	0	1
Capital intensity	187	0.0312	0.0158	0.0114	0.0631
Inventory/Sales	187	0.124	0.0641	0.0443	0.202
Labor Cost over Sales	187	0.163	0.0709	0.0576	0.274
Firm Size	202	116.2	129.8	19.34	362.3
Leverage	200	0.595	0.126	0.386	0.796
Interest Expense Ratio	202	0.0359	0.0387	0.0104	0.122
Short-Term Debt Ratio	202	0.0588	0.0326	0.0236	0.125
Cyclicality	205	0.917	0.948	-0.0770	2.798
Standard Dev. of Output Growth	205	3.272	2.144	1.144	7.866

This Table reports the summary statistics for the additional industry-level control variables.

Table A.3: Price Stickiness and Monetary Policy Responses - Alternative Ordering Robustness Check

	(1) $h = 18$ Months	(2) $h = 24$ Months	(3) $h = 30$ Months
Log(FPA)	0.547**	0.653***	0.590**
	(0.252)	(0.250)	(0.241)
Inventory/Sales	2.817***	3.093***	2.526***
	(0.867)	(0.913)	(0.960)
Labor Costs/Sales	2.699	2.533	2.425
	(1.705)	(1.931)	(1.991)
Capital Int.	-5.000	-6.346	-6.597
	(4.169)	(4.496)	(4.959)
Firm Size	0.0510	0.0276	0.0388
	(0.0833)	(0.0864)	(0.0833)
Interest Expense Ratio	-3.184	-3.468	-4.105
	(1.979)	(2.490)	(2.600)
Leverage	0.773	1.485*	1.712**
	(0.775)	(0.809)	(0.790)
Short-Term Debt Ratio	-1.960	-2.710	-2.805
	(2.074)	(2.269)	(2.257)
Cyclicality	-0.215***	-0.283***	-0.297***
	(0.0821)	(0.0883)	(0.0865)
Std(Output Growth)	0.0129	0.0422	0.0741
	(0.0524)	(0.0541)	(0.0484)
Durability	-0.269	-0.281	-0.194
	(0.175)	(0.189)	(0.187)
Observations	186	186	186
R-squared	0.115	0.149	0.155

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

This Table reports the estimation results when estimating Equation 1.11 at different horizons h in the Alternative Ordering Robustness Check. The constant term is not reported. Robust standard errors are reported in Parenthesis.

Table A.4: Price Stickiness and Monetary Policy Responses - Alternative Shock Identification Robustness Check

	(1) h = 18 Months	(2) h = 24 Months	(3) h = 30 Months
Log(FPA)	0.717**	0.530***	0.208*
	(0.306)	(0.180)	(0.120)
Inventory/Sales	1.685	0.234	-0.175
	(1.126)	(1.085)	(0.720)
Labor Costs/Sales	0.419	1.253	1.884*
	(1.211)	(0.986)	(0.982)
Capital Int.	-4.302	-6.158	-10.57*
	(3.920)	(3.872)	(5.384)
Firm Size	0.0915	0.0733	0.0689
	(0.0951)	(0.0813)	(0.0663)
Interest Expense Ratio	0.595	-1.529	-2.320*
	(2.075)	(1.766)	(1.377)
Leverage	0.579	0.684	0.985**
	(0.642)	(0.550)	(0.461)
Short-Term Debt Ratio	-2.614	-1.566	-2.281
	(1.624)	(1.535)	(1.897)
Cyclicality	-0.101*	-0.0488	-0.0245
	(0.0554)	(0.0549)	(0.0489)
Std(Output Growth)	-0.0309	0.00504	0.0253
	(0.0367)	(0.0317)	(0.0278)
Durability	-0.0765	-0.0935	-0.0849
	(0.111)	(0.113)	(0.109)
Observations	186	186	186
R-squared	0.203	0.131	0.100

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

This Table reports the estimation results when estimating Equation 1.11 at different horizons h in the Alternative Shock Identification Robustness Check. The constant term is not reported. Robust standard errors are reported in Parenthesis.

A.3 Additional Figures

Figure A.1: Distribution of the Frequency of Price Adjustment

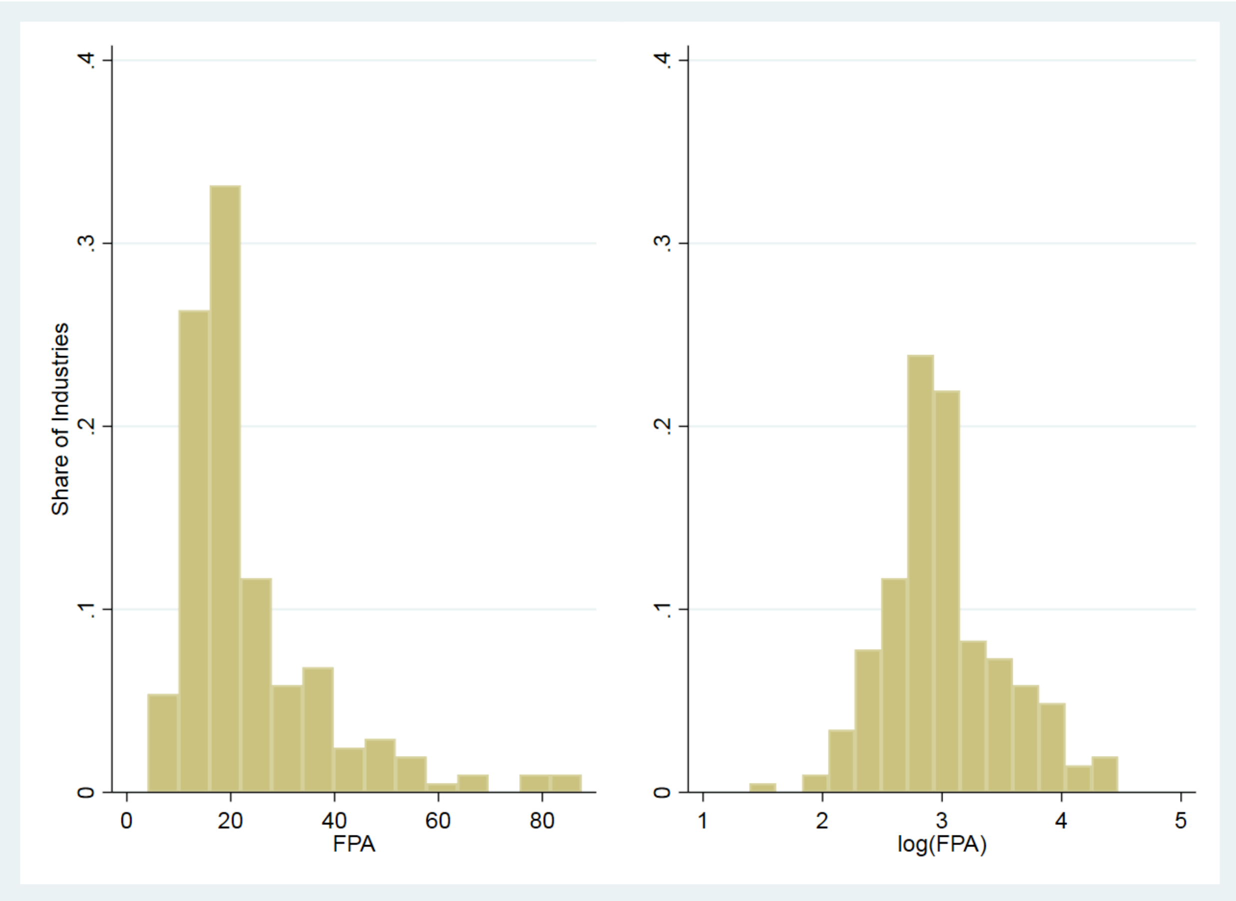

This Figure shows the with-in sample distribution of the frequency of price adjustment (FPA, left Panel) and the log of the frequency of price adjustment (right Panel) of the 205 industries in the sample.

Figure A.2: Monetary Policy Shock Series

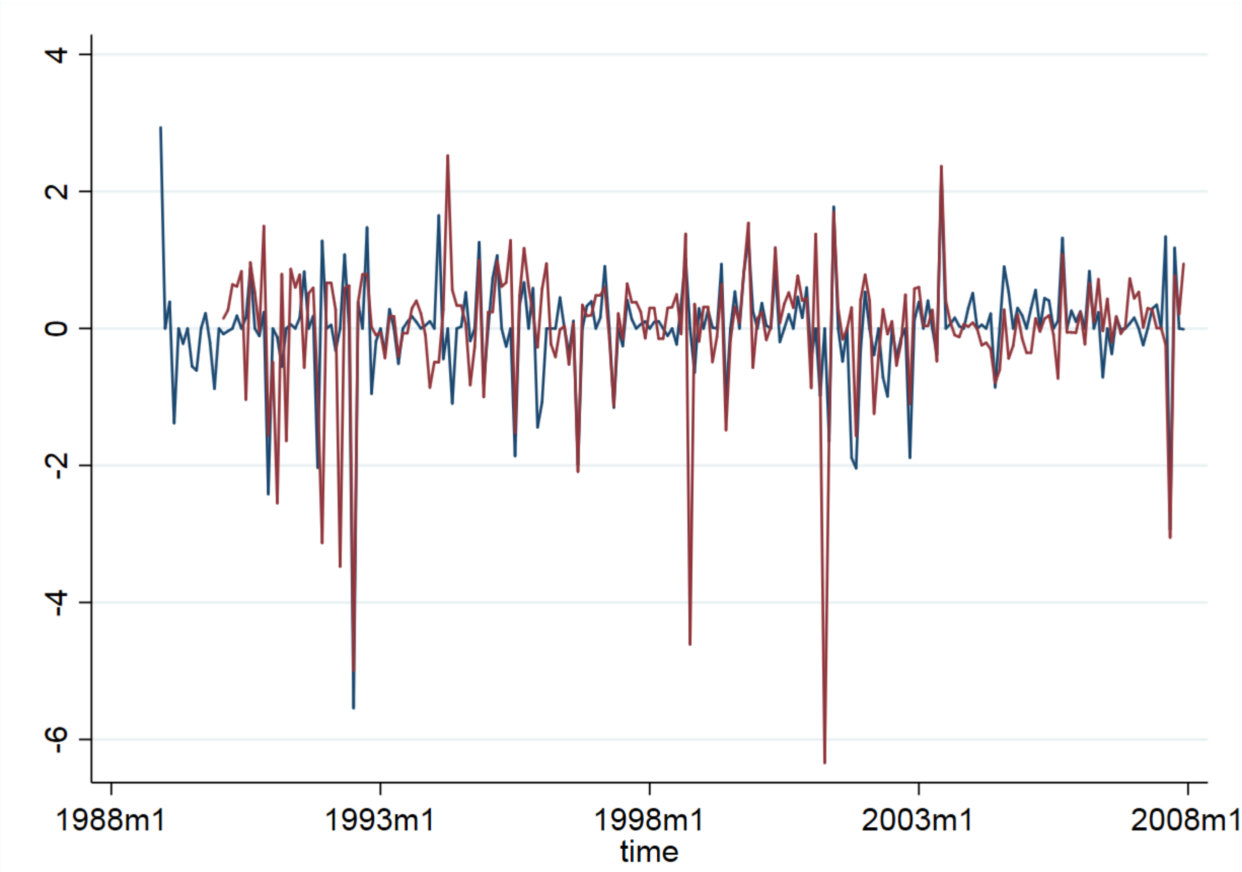

This Figure shows the time series of the monetary policy shock series used in the paper. The blue line is the monetary shock series of Barakchian and Crowe (2013). The red line shows the shock series of Miranda-Agrippino (2016). Both series are standardized with mean zero and a standard deviation of one. The y-axis is expressed in units of standard deviations of the shock.

Figure A.3: Response of other Variables to a Contractionary Monetary Policy Shock

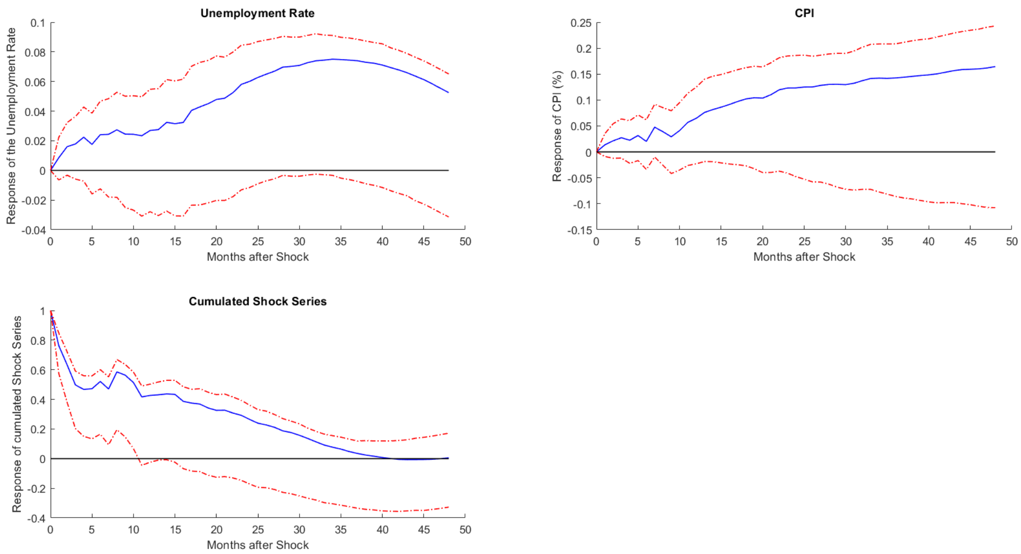

Structural VAR (Monthly data, 5 endogenous variables, 12 lags). Variables ordered as industrial production (in logs), unemployment rate, consumer price index (in logs), commodity price index (in logs) (all seasonally adjusted) and the cumulated shock measure of Barakchian and Crowe (2013). Graphs show the responses of the variables in the system (excluding industrial production) to a one standard deviation increase to the policy measure. Structural shocks obtained via Cholesky decomposition. Bootstrapped 90% Confidence Intervals based on 5000 Bootstrap Replications are shown in the red, dotted lines.

Figure A.4: Scatterplot of Industry Responses 24 Months after a Monetary Policy Shock vs log(FPA)

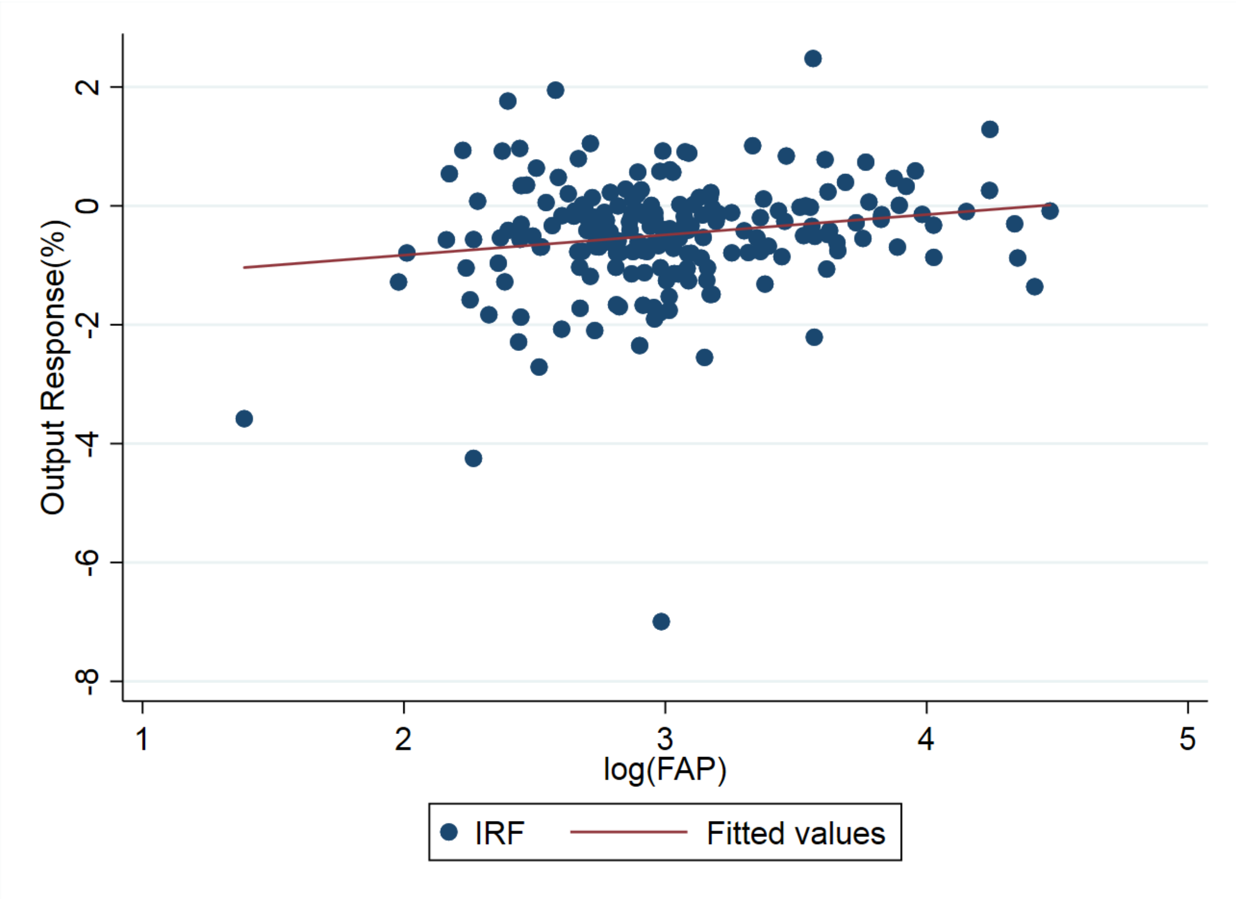

This Figure shows the scatter plot of the cumulative output response of the 205 different industries to a one standard deviation contractionary policy shock at a horizon of 24 Months after the shock plotted against the log of the frequency of price adjustment of the industry. Each blue dot represents a single industry. On the X-axis, the log of the frequency of price adjustment is shown. On the Y-Axis, the cumulative output response, reported in percent, to a one Standard Deviation contractionary Monetary Policy Shock is shown, 24 Months after the shock has happened. The red line shows the linear fit.

Figure A.5: Responses to a Contractionary Monetary Policy Shock - Alternative Ordering

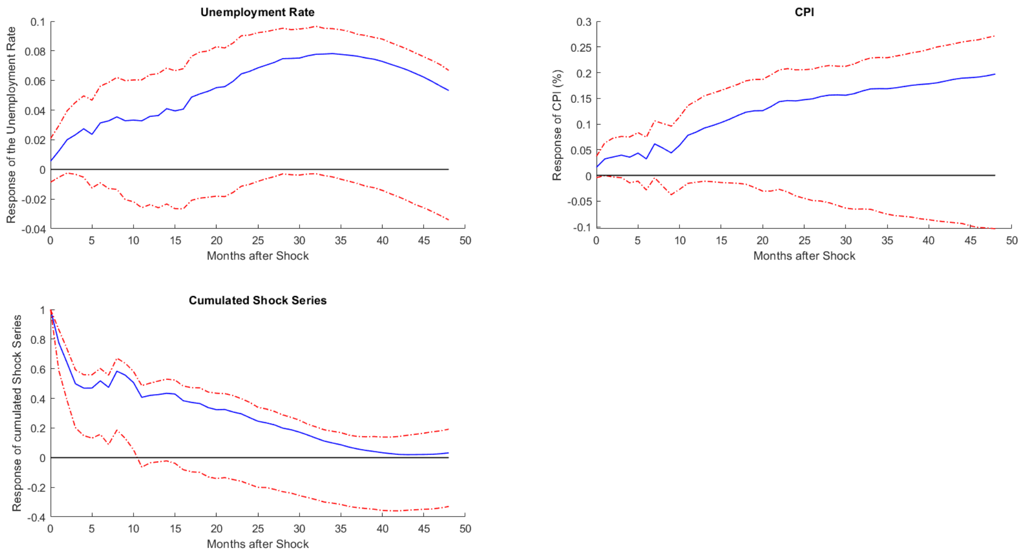

Structural VAR (Monthly data, 5 endogenous variables, 12 lags). Variables ordered as the cumulated shock measure of Barakchian and Crowe (2013), industrial production (in logs), unemployment rate, consumer price index (in logs), and a commodity price index (in logs) (all seasonally adjusted). Graphs show the responses of the variables in the system to a one standard deviation increase to the policy measure. Bootstrapped 90% Confidence Intervals based on 5000 Bootstrap Replications are shown in the red, dotted lines.

Figure A.6: Responses to a Contractionary Monetary Policy Shock - Alternative Shock Identification

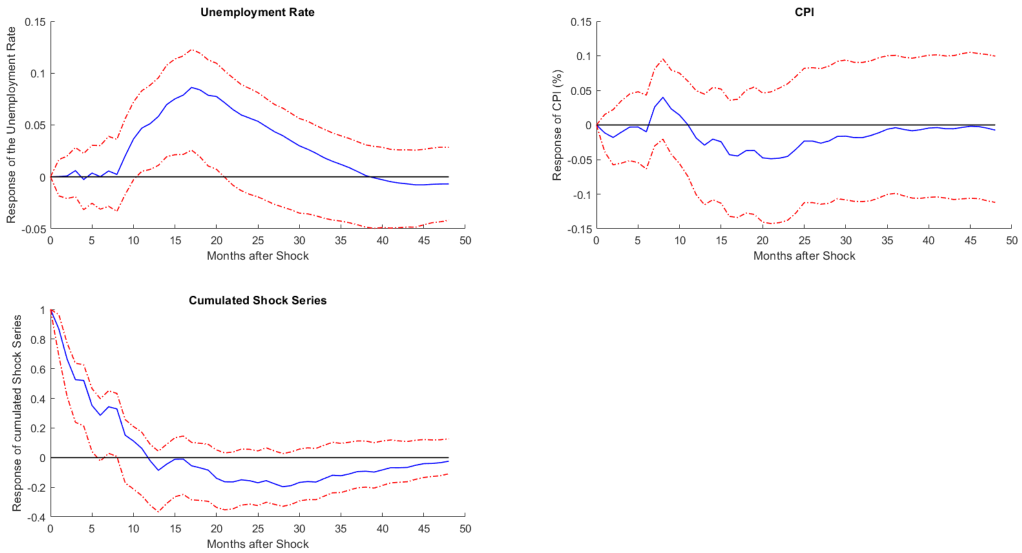

Structural VAR (Monthly data, 5 endogenous variables, 12 lags). Variables ordered as industrial production (in logs), unemployment rate, consumer price index (in logs), commodity price index (in logs) (all seasonally adjusted) and the cumulated shock measure of Miranda-Agrippino (2016). Graphs show the responses of the variables in the system to a one standard deviation increase to the policy measure. Structural shocks obtained via Cholesky decomposition. Bootstrapped 90% Confidence Intervals based on 5000 Bootstrap Replications are shown in the red, dotted lines.

Appendix B

Appendix to Chapter 2

B.1 State-Level Differences in the A-Level Wage Premium

In this Section, we describe how we estimate the A-Level wage premium in our data and how estimates of the A-Level wage premium differ across states ("Bundesländer") in Germany.[1]

In the estimation of A-Level wage premia, we restrict our sample to individuals aged 30 to 45, working full time. All estimations of A-Level wage premia are based on the waves 2009 to 2014 of the MZ. We focus on individuals aged 30 to 45 to estimate the A-Level wage premium for recent cohorts in prime working age.

We measure income as monthly personal net income (of the individual). Income is measured as continuous variable, imputed based on reported binned income (see Section B.2 for details on the imputation procedure). We focus on individuals working full time as labor income should make up the largest share of net income for these individuals.

We express all incomes in 2009 Euros (i.e. we adjust reported income for inflation) and pool observations across all waves. We use log income as the dependent variable in our estimations of A-Level wage premia[2].

[1]We focus on differences in the A-Level wage premium across states rather than differences across LLM's as the school system is regulated at the state-level.

[2]We drop individuals with a reported monthly personal net income of 0. As discussed before, the income of individuals working as farmers is recorded as zero in our data, hence a reported net income of zero rather indicates a farmer, than a real net income of zero. Furthermore, as we focus on individuals working full time, a net income of zero is not plausible.

We estimate the A-Level wage premium using the following specification:

$$log(y_{i,t}) = \alpha + \beta I(A - Level_{i,t}) + \gamma_a I_a + \delta_t I_t + \epsilon_{i,t} \tag{B.1}$$

Here, $log(y_{i,t})$ is the log of the net monthly personal income of individual i observed in wave t, $I(A - Level_{i,t})$ is a binary indicator, indicating whether the person i (observed in wave t) has an A-Level degree or not, I_a is a full set of age dummies (16 age dummies, as we focus on individuals ages 30 to 45) to control for the age profile of income and I_t are wave fixed effects, indicating the wave of the MZ used (5 wave dummies, 2009 is used as base category).

In the full sample, we obtain an estimate of $\beta = .367$, implying an A-Level wage premium(given by $exp(\beta) - 1$) of around 44% for monthly net income.

Next, we estimate the A-Level wage premium separately for each state. In this estimation, we restrict the age profile of income and wave fixed effects to be the same in every state, i.e. we estimate:

$$log(y_{i,t}) = \sum_{k=1}^{16} \alpha_k I(State_i = k) + \sum_{k=1}^{16} \beta_k I(A - Level_{i,t}, State_i = k) + \gamma_a I_a + \delta_t I_t + \epsilon_{i,t} \tag{B.2}$$

where $k = 1, ..., 16$ denotes the subscript for the state of residence, $I(State_i = k)$ is an indicator function equal to one if individual i is living in state k and $I(A - Level_{i,t}, State_i = k)$ is an indicator function equal to one if individual i is living in state k and has an A-Level degree, and zero otherwise. Note that state here refers to the current state of residence (at the point in time when the individual is observed), not the state where the A-Level degree was obtained (this information is not recorded in our data).

The estimates of state-specific A-Level wage premia are reported in Table B.1. We report the estimates of β_k obtained when estimating Equation B.2, and the implied A-Level wage premium $exp(\beta_k) - 1$.

Table B.1 shows that state-level differences in the observed A-Level wage premium exist but are rather small, compared to the size of the national A-Level wage premium. The lowest state-specific A-Level wage premium (Bremen, around 36% A-Level wage premium) is around 8 percentage points lower than the national A-Level wage premium. The highest state-specific A-Level wage premium (Bavaria, round 48%) is around 4 percentage points

Table B.1: State-Specific A-Level Wage Premia

State Name	β_k	A-Level Wage Premium
Baden-Württemberg	.344	.411
Bayern	.397	.488
Berlin	.353	.424
Brandenburg	.358	.430
Bremen	.307	.359
Hamburg	.396	.485
Hessen	.350	.419
Mecklenburg-Vorpommern	.373	.452
Niedersachsen	.332	.394
Nordrhein-Westfalen	.323	.381
Rheinland-Pfalz	.370	.447
Saarland	.311	.365
Sachsen	.379	.461
Sachsen-Anhalt	.353	.423
Schleswig-Holstein	.370	.449
Thüringen	.366	.442

higher than the nation-wide A-Level wage premium. In most states, the observed A-Level wage premium is rather similar.

These results suggest a substantial wage premium associated with an A-Level degree, regardless of the state of residence. Our results provide evidence that differences in the A-Level wage premium are rather small across states, compared to the size of the A-Level wage premium.

To our knowledge, no study exists on differences in the return to an A-Level degree obtained in different German States.

B.2 Details on Income Imputation

This Section describes how total household income is imputed in the MZ, and how our measure of parental income is constructed.

First, every household member reports her personal income (including all sources of income, net of tax and social security contributions) individually. If a household member has no personal income, this is recorded as zero personal income. If personal income is positive, income is reported in binned format. Choosing out of 24 pre-defined bins, individuals state the bin in which her personal income falls. The most common bin sizes[3] are 199 Euro and

[3]bin size means the difference between the upper and the lower bound of a bin

299 Euro. Bin sizes vary between 149 Euro (for the two lowest income bins) and 7999 Euro (for the second highest income bin). The highest income bin (for a monthly income higher than 18000 Euro) is open ended. The exact definition of the income bins can be found in Table B.2.

Table B.2: Income Bin Sizes used in the MZ

Income Bin	Lower Bound (Euro)	Upper Bound (EURO)
1	1	149
2	150	299
3	300	499
4	500	699
5	700	899
6	900	1099
7	1100	1299
8	1300	1499
9	1500	1699
10	1700	1999
11	2000	2299
12	2300	2599
13	2600	2899
14	2900	3199
15	3200	3599
16	3600	3999
17	4000	4499
18	4500	4999
19	5000	5499
20	5500	5999
21	6000	7499
22	7500	9999
23	10000	17999
24	18000	∞

Next, based on the reported (binned) personal income of all household members, total household income is imputed. For every household member with positive personal income, the binned income is transformed into a continuous income variable. The continuous measure of personal income is imputed using the following formula:

$$NE_{i,b} = A_b + Z_i * B_b \tag{B.3}$$

where $NE_{i,b}$ is the imputed personal income of individual i, who reported a personal income within income bin b. A_b and B_b are bin-specific parameters, chosen by the statistical office. Z_i is a random variable, drawn from a discrete uniform distribution from the natural numbers in the interval $[0, 99]$. The bin-specific parameters are reported in Table B.3. The result of this procedure is a continuous measure of personal income.

Total household income is then calculated as the sum of the imputed personal income of all household members. The imputation procedure is carried out by the Statistical Office. The result of this imputation procedure is a continuous measure of household income, which we use in our further analysis.

If at least on member of the household is working as self-employed farmer („Landwirt/-in (selbstständig in der Haupttätigkeit)“), total household income is not reported for this household, and income is imputed as zero. Hence an imputed income of zero indicates a self-employed farmer, rather than a true zero. Furthermore, as income here includes income from all sources, the German social safety net makes it (nearly) impossible to have a net income of zero (after all transfers). Due to this reason, we choose to drop households with total household income recorded zero income from the analysis.

Total household income reported in the MZ is the sum of the personal income of all household members. It includes not only parental income, but also the personal income of all children present in the household. In our analysis, we want to focus on parental income. In order to focus on parental income, we subtract the income of all dependent children present in the household from total household income. For children with positive personal income, we observe the income bin of the child. We assign every child the mid-point of its respective income bin. We sum the income of all children at the household level, call this total children income. We subtract total children income from total household income to obtain total parental income.

Table B.3: Bin-Specific Parameters used for Income Imputation in the MZ

Income Bin	A_b	B_b
1	50	1.0
2	151	1.5
3	301	2.0
4	501	2.0
5	701	2.0
6	901	2.0
7	1101	2.0
8	1301	2.0
9	1501	2.0
10	1702	3.0
11	2002	3.0
12	2302	3.0
13	2602	3.0
14	2902	3.0
15	3203	4.0
16	3603	4.0
17	4004	5.0
18	4504	5.0
19	5004	5.0
20	5504	5.0
21	6014	15.0
22	7515	25.0
23	10040	80.0
24	18002	2.0

B.3 Additional Figures

Figure B.1: Average Household Size per Income Percentile: No Adjustment for Household Size

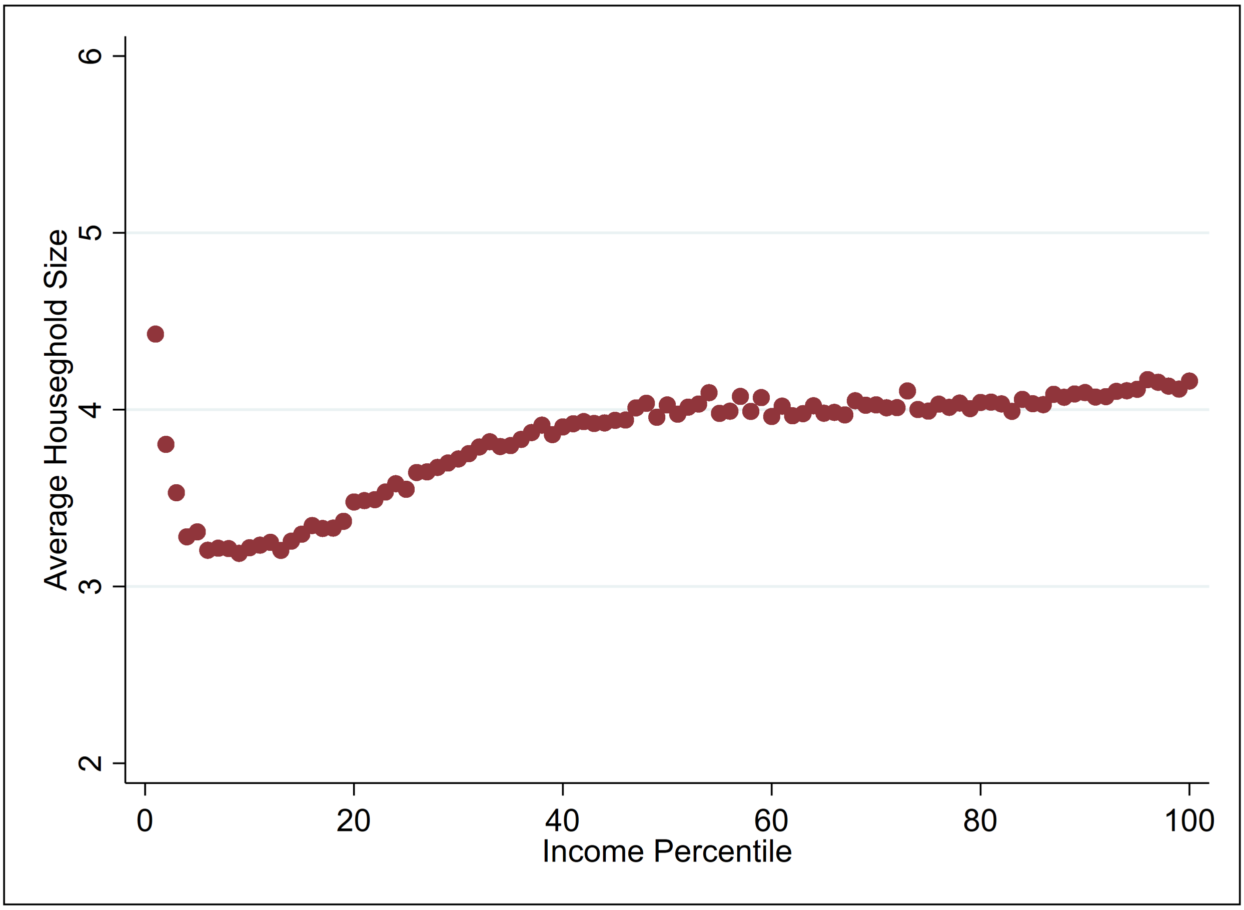

This figure shows the average household size vs. the percentile rank in the national parental income distribution when income is not adjusted for household size. The figure is constructed by binning parent rank into one-percentile point bins (so that there are 100 bins) and plotting the average household size in each bin vs. the parent rank in each bin. The figure is based on $N = 268523$ children living with their parents.

Figure B.2: Household Income 2014 vs Income Percentiles: No Adjustment for Household Size

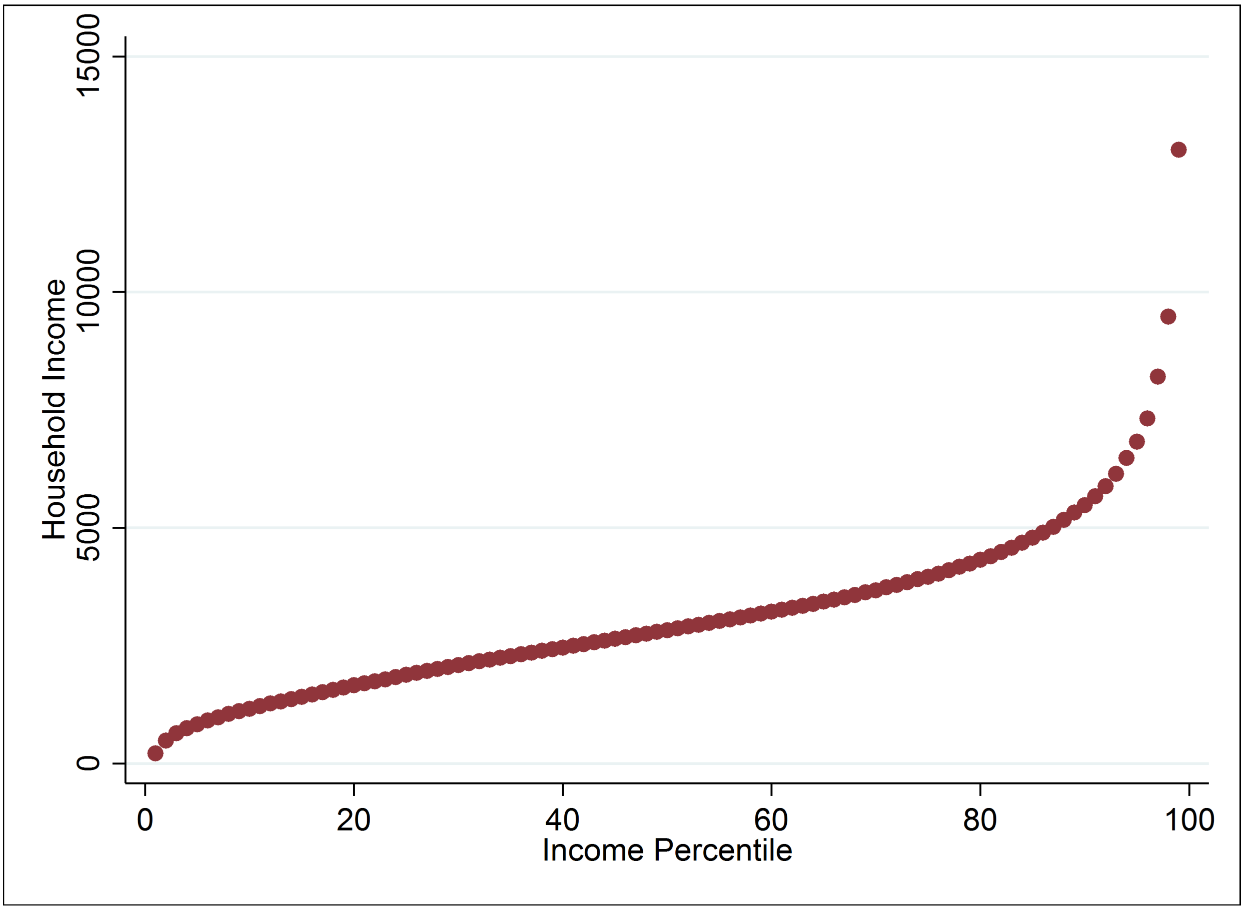

This figure shows the Euro amount of the upper bound of each income percentile bin vs. the percentile rank in the national parental income distribution when income is not adjusted for household size. The figure is constructed by binning parent income rank into one-percentile point bins (so that there are 100 bins) and plotting the upper bound of each percentile bin vs. the parent rank in each bin. This Figure is based on the 2014 wave of the MZ.

Figure B.3: Household Income 2014 vs Income Percentiles: Income per Household Member

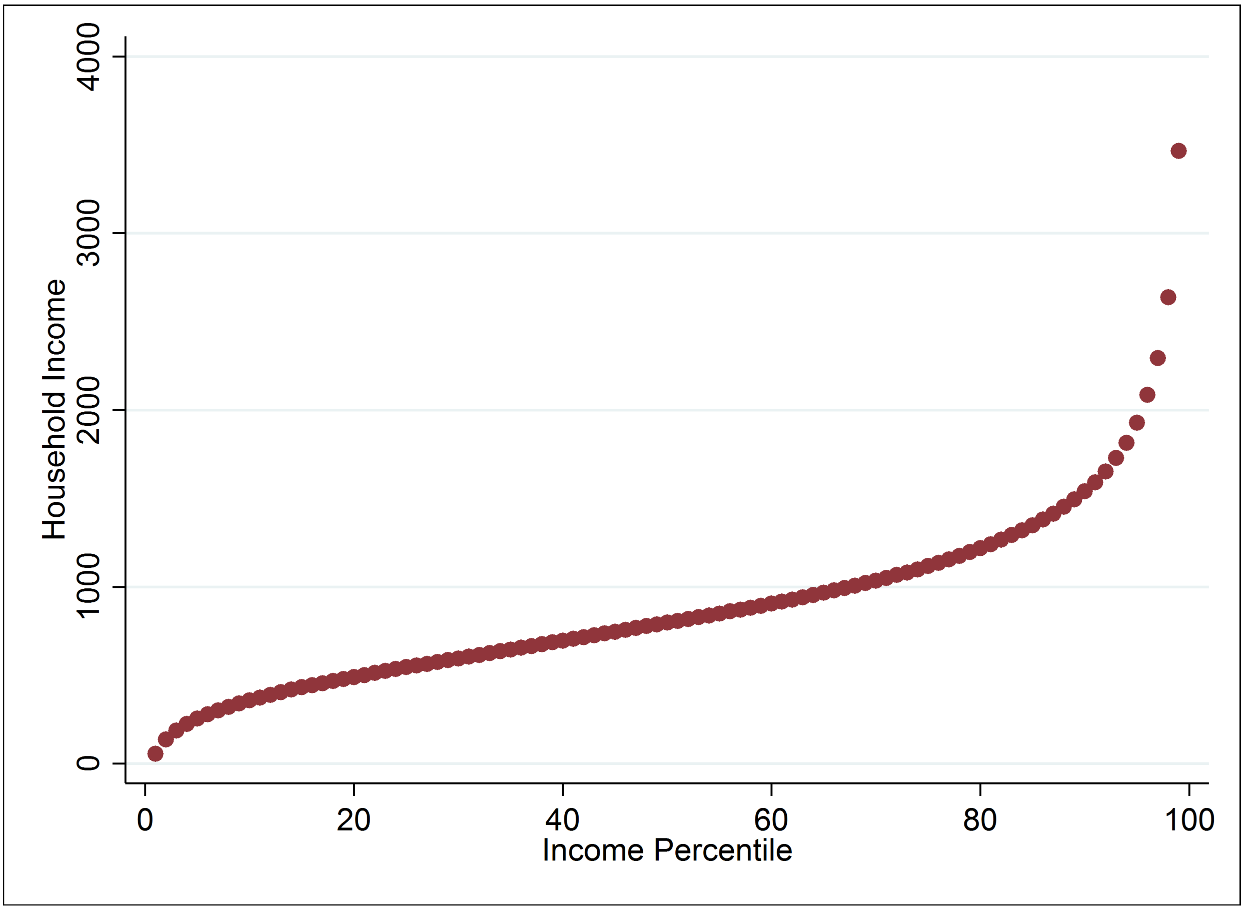

This figure shows the Euro amount of the upper bound of each income percentile bin vs. the percentile rank in the national parental income distribution when income is divided by the number of household members. The figure is constructed by binning parent income rank into one-percentile point bins (so that there are 100 bins) and plotting the upper bound of each percentile bin vs. the parent rank in each bin. This Figure is based on the 2014 wave of the MZ.

Figure B.4: Household Income 2014 vs Income Percentiles: Equivalence Scale Income

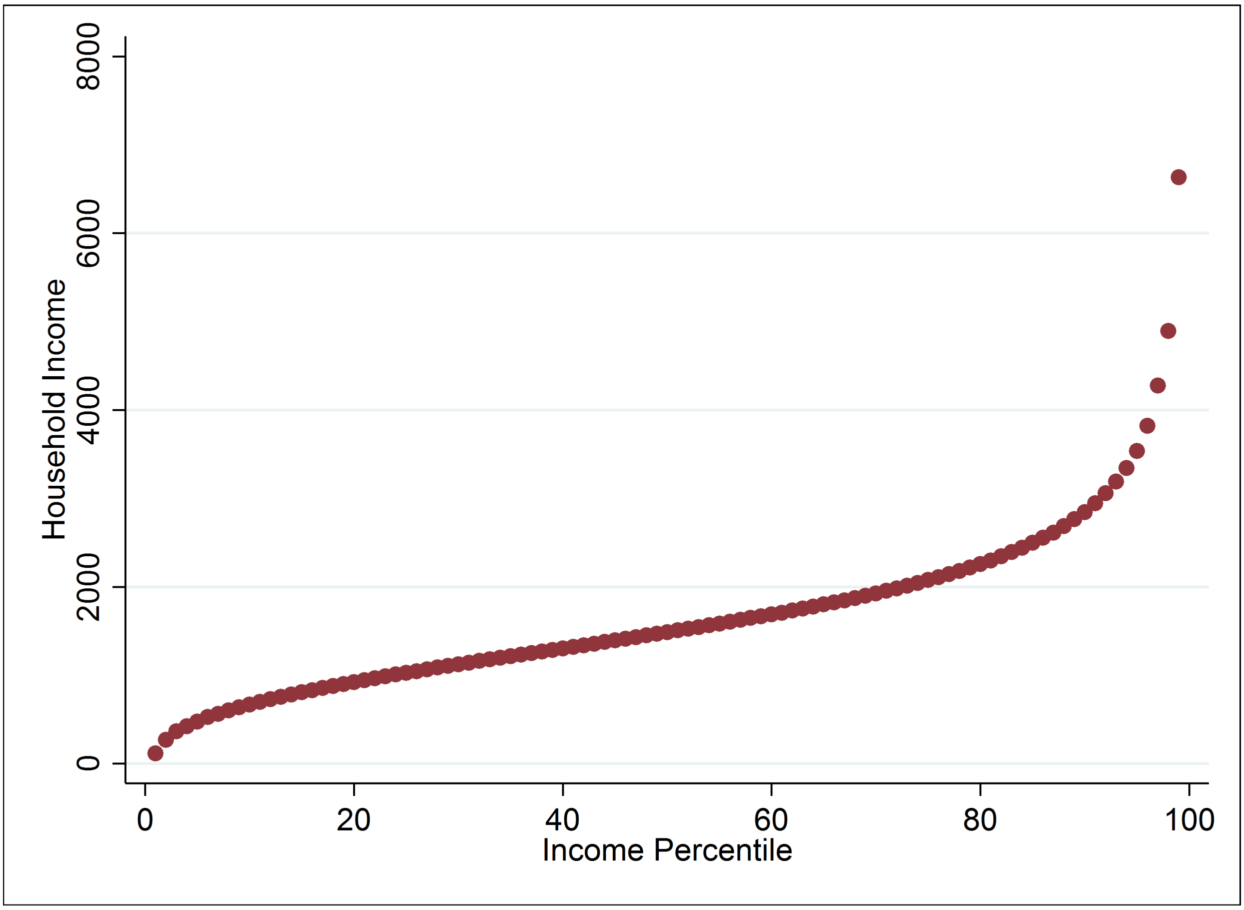

This figure shows the Euro amount of the upper bound of each income percentile bin vs. the percentile rank in the national parental income distribution when income is divided by the square root of the number of household members (equivalence scale approach). The figure is constructed by binning parent income rank into one-percentile point bins (so that there are 100 bins) and plotting the upper bound of each percentile bin vs. the parent rank in each bin. This Figure is based on the 2014 wave of the MZ.

Figure B.5: Intergenerational Mobility at the National Level: Equivalence Scale Income

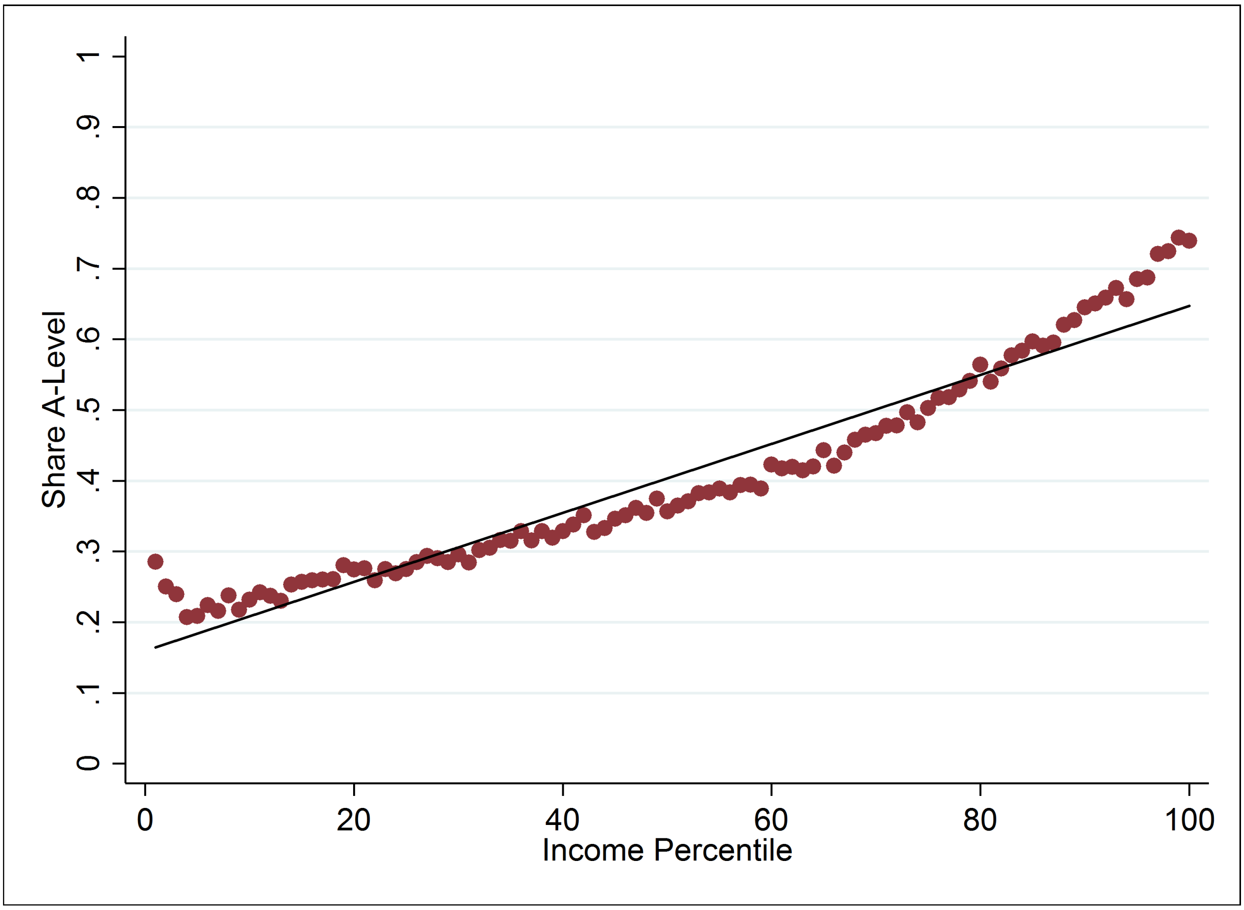

This figure plots the fraction of children aged 16-22 (y-axis) that are either enrolled in the last two/three years of the A-Level track or already completed the A-Level degree vs. the percentile rank in the national parental income distribution when parental income divided by the square root of household size (x-axis). The figure is constructed by binning parent rank into one-percentile point bins (so that there are 100 bins) and plotting the fraction of children aged 16-22 currently enrolled in the last two/three years of the A-Level track or having already completed it among all children in the income bin vs. the parent rank in each bin. The figure is based on $N = 268523$ children living with their parents. The OLS regression of the A-Level dummy on the parent income rank in the underlying micro-data (black line) yields a slope coefficient of 0.0048 and a constant of 0.159.

Figure B.6: Intergenerational Mobility at the National Level: Income per Household Member, Children Aged 16 to 19

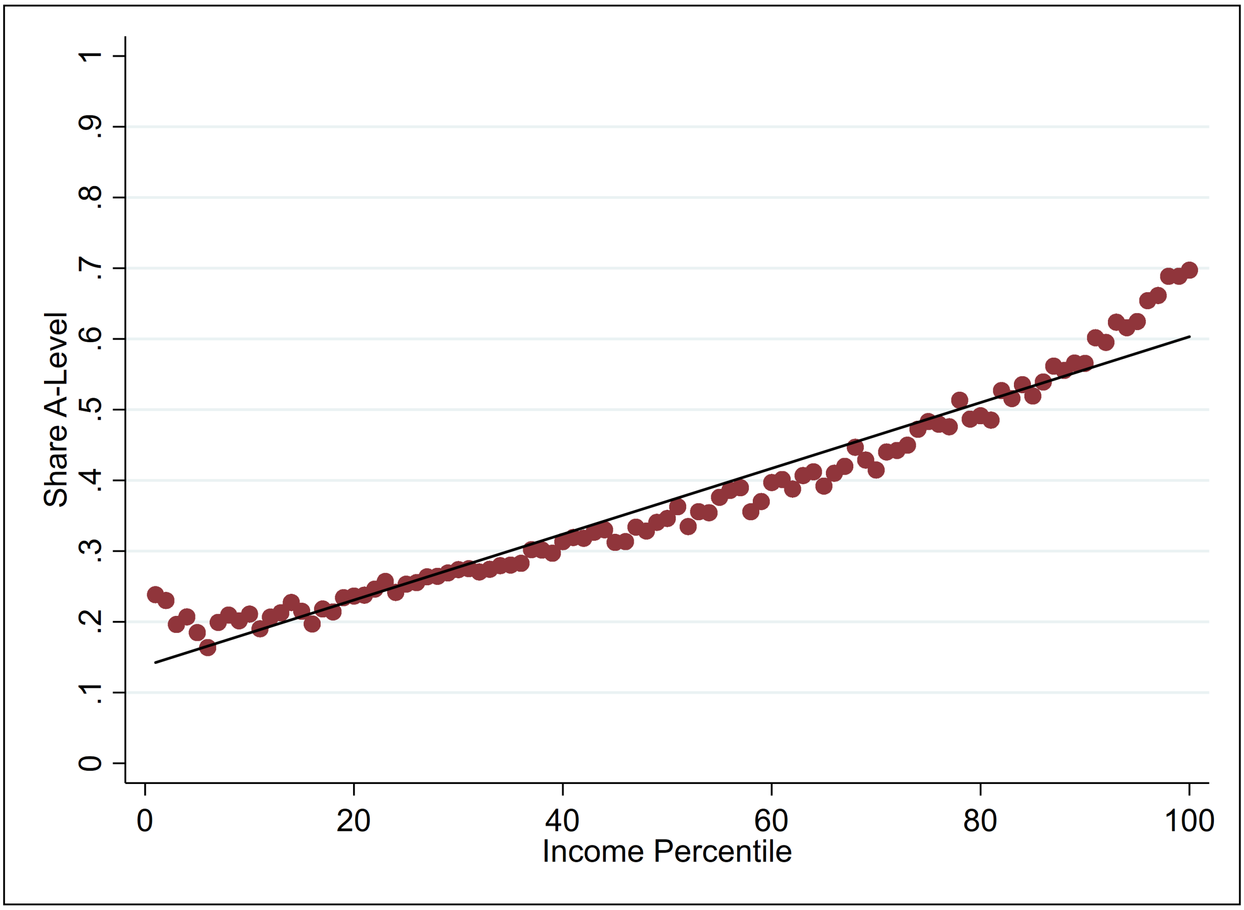

This figure plots the fraction of children aged 16-19 (y-axis) that are either enrolled in the last two/three years of the A-Level track or already completed the A-Level degree vs. the percentile rank in the national parental income distribution when parental income divided by the square root of household size (x-axis). The figure is constructed by binning parent rank into one-percentile point bins (so that there are 100 bins) and plotting the fraction of children aged 16-19 currently enrolled in the last two/three years of the A-Level track or having already completed it among all children in the income bin vs. the parent rank in each bin. The figure is based on $N = 171441$ children living with their parents. The OLS regression of the A-Level dummy on the parent income rank in the underlying micro-data (black line) yields a slope coefficient of 0.0046 and a constant of 0.137.

Figure B.7: Intergenerational Mobility at the National Level: Income per Household Member, Children Aged 17 to 20

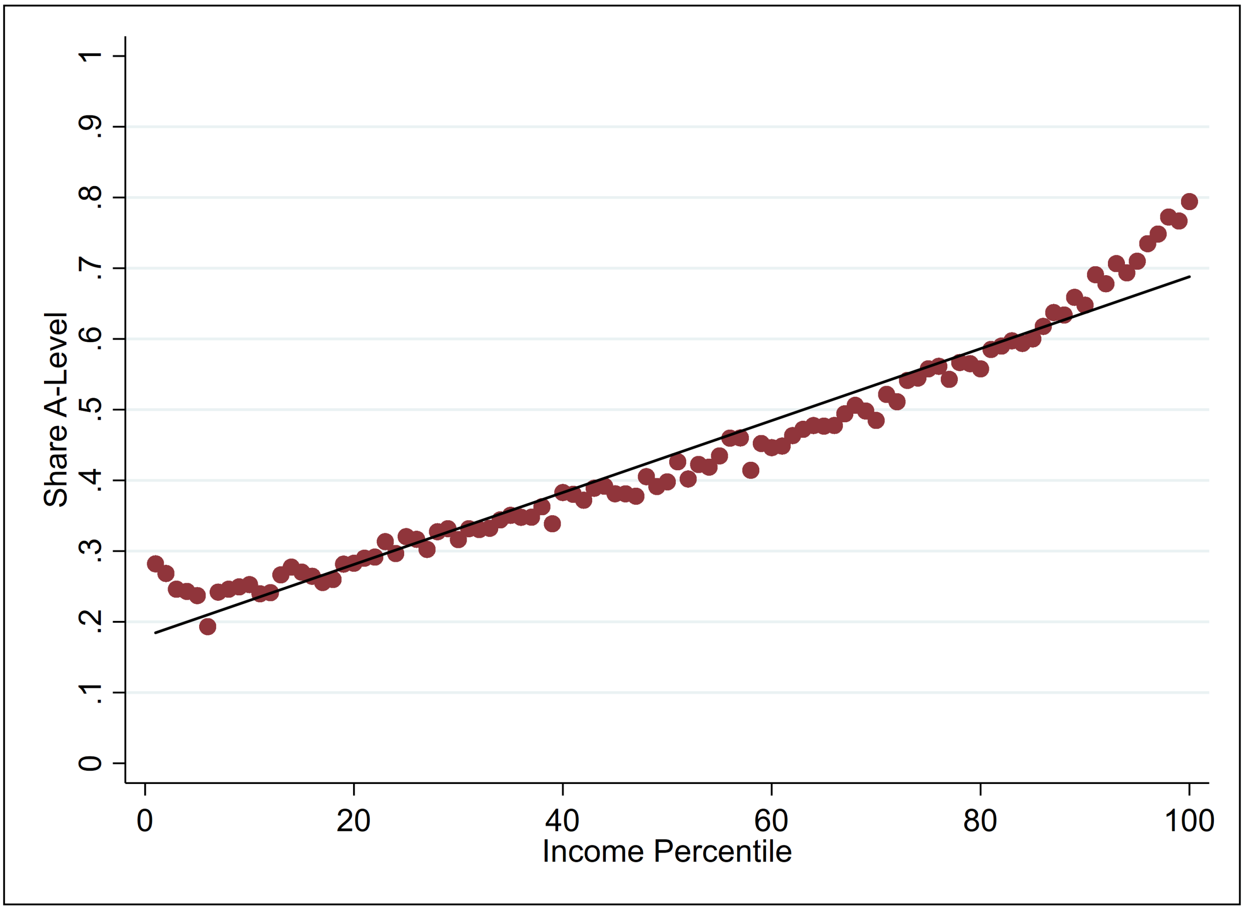

This figure plots the fraction of children aged 17-20 (y-axis) that are either enrolled in the last two/three years of the A-Level track or already completed the A-Level degree vs. the percentile rank in the national parental income distribution when parental income divided by the square root of household size (x-axis). The figure is constructed by binning parent rank into one-percentile point bins (so that there are 100 bins) and plotting the fraction of children aged 17-20 currently enrolled in the last two/three years of the A-Level track or having already completed it among all children in the income bin vs. the parent rank in each bin. The figure is based on $N = 163903$ children living with their parents. The OLS regression of the A-Level dummy on the parent income rank in the underlying micro-data (black line) yields a slope coefficient of 0.0050 and a constant of 0.179.

Figure B.8: Intergenerational Mobility at the National Level: Income per Household Member, Children Aged 20 to 22

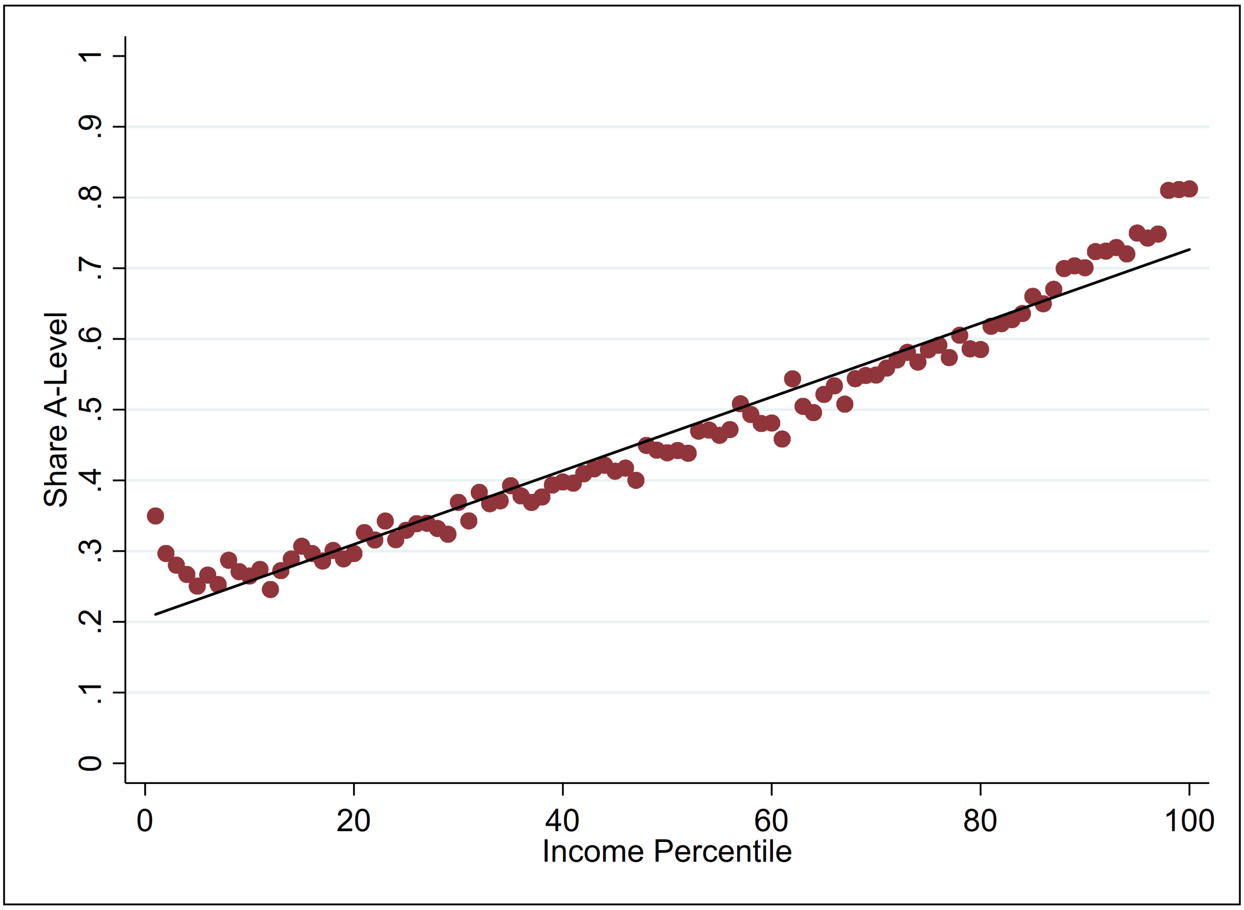

This figure plots the fraction of children aged 20-22 (y-axis) that are either enrolled in the last two/three years of the A-Level track or already completed the A-Level degree vs. the percentile rank in the national parental income distribution when parental income divided by the square root of household size (x-axis). The figure is constructed by binning parent rank into one-percentile point bins (so that there are 100 bins) and plotting the fraction of children aged 20-22 currently enrolled in the last two/three years of the A-Level track or having already completed it among all children in the income bin vs. the parent rank in each bin. The figure is based on $N = 97082$ children living with their parents. The OLS regression of the A-Level dummy on the parent income rank in the underlying micro-data (black line) yields a slope coefficient of 0.0052 and a constant of 0.205.

Figure B.9: Baseline Probability α Across LLM's in Germany

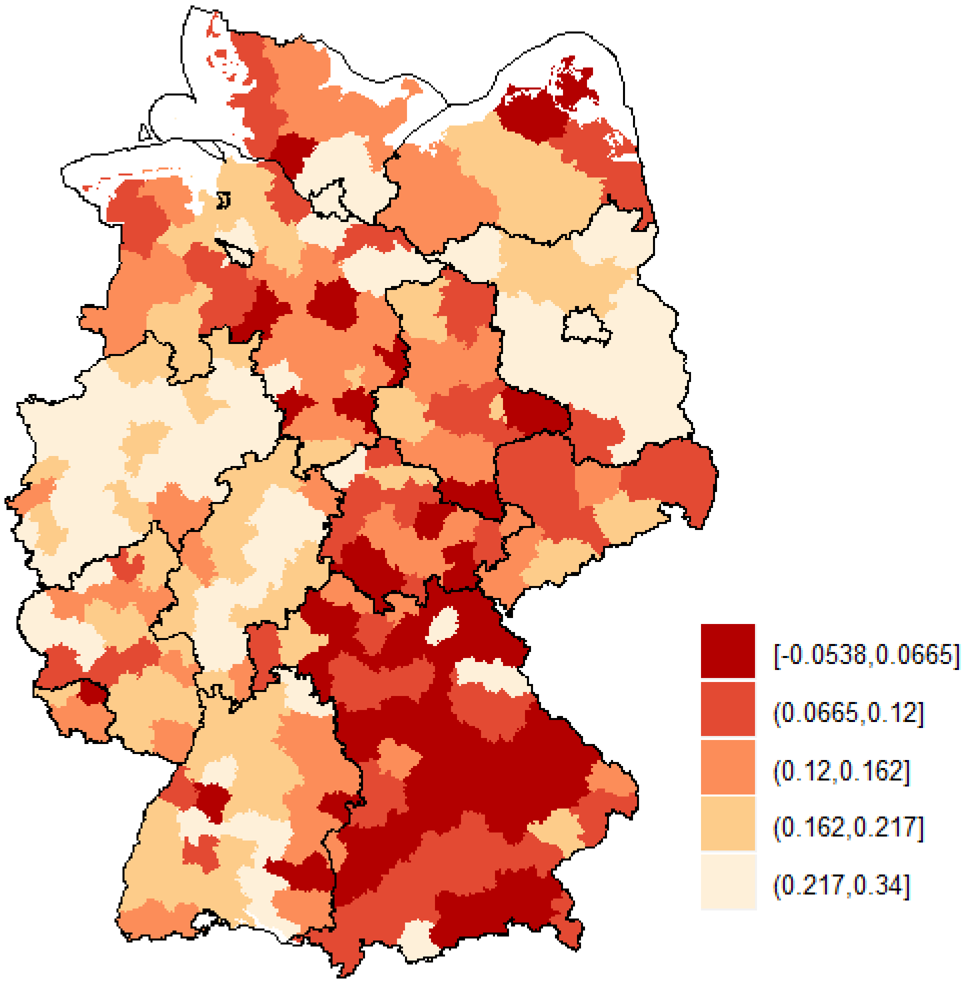

This Figure shows estimates of the baseline probability of obtaining an A-Level degree α (i.e. the constant estimated in Equation 2.8), estimated separately for every local labor market in Germany.

Bibliography

ACCIARI, P., A. POLO, AND G. L. VIOLANTE (2019): “”And Yet, It Moves”: Intergenerational Mobility in Italy,” *NBER Working Paper No. 25732*.

BALLEER, A., N. HRISTOV, AND D. MENNO (2017): “Financial Constraints and Nominal Price Rigidities,” *CESifo Working Paper No. 6309*.

BARAKCHIAN, S. M. AND C. CROWE (2013): “Monetary Policy Matters: Evidence from New Shocks Data,” *Journal of Monetary Economics*, 60, 950 – 966.

BELLONI, A., V. CHERNOZHUKOV, AND C. HANSEN (2014): “High-Dimensional Methods and Inference on Structural and Treatment Effects,” *Journal of Economic Perspectives*, 28, 29 – 50.

BERNANKE, B. S., M. GERTLER, AND S. GILCHRIST (1999): “The Financial Accelerator in a Quantitative Business Cycle Framework,” in *Handbook of Macroeconomics*, ed. by J. B. Taylor and M. Woodford, Elsevier, vol. 1, chap. 21, 1341 – 1393.

BIEWEN, M. AND M. TAPALAGA (2017): “Life-Cycle Educational Choices in a System with Early Tracking and ‘Second Chance’ Options,” *Economics of Education Review*, 56, 80 – 94.

BILS, M. AND P. J. KLENOW (2004): “Some Evidence on the Importance of Sticky Prices,” *Journal of Political Economy*, 112, 947– 985.

BLACK, S. AND P. DEVEREUX (2011): “Recent Developments in Intergenerational Mobility,” in *Handbook of Labor Economics*, ed. by O. Ashenfelter and D. Card, Elsevier, vol. 4, chap. 16, 1487 – 1541.

BOUAKEZ, H., E. CARDIA, AND F. RUGE-MARCIA (2013): “Sectoral Price Rgidity and Aggregate Dynamics,” *European Economic Review*, 65, 1 – 22.

BRATBERG, E., J. DAVIS, B. MAZUMDER, M. NYBOM, D. SCHNITZLEIN, AND K. VAAGE (2017): “A Comparison of Intergenerational Mobility Curvey in Germany, Norway, Sweden, and the US,” *Scandivian Journal of Economics*, 119, 72 – 101.

CALVO, G. (1983): “Staggered Prices in a Utility Maximizing Framework,” *Journal of Monetary Economics*, 12, 383 – 398.

Caplin, A. S. and D. F. Spulber (1987): "Menu Costs and the Neutrality of Money," *The Quarterly Journal of Economics*, 102, 703–726.

Carlino, G. and R. DeFina (1998): "The Differential Regional Effects of Monetary Policy," *Review of Economics and Statistics*, 80, 572–587.

Chetty, R., J. N. Friedman, E. Saez, N. Turner, and D. Yagan (2017): "Mobility Report Cards: The Role of Colleges in Intergenerational Mobility," *NBER Working Paper No. 23618*.

Chetty, R. and N. Hendren (2018): "The Impacts of Neighborhoods on Intergenerational Mobility I: Childhood Exposure Effects," *The Quarterly Journal of Economics*, 133, 1107 – 1162.

Chetty, R., N. Hendren, M. R. Jones, and S. R. Porter (2018): "Race and Economic Opportunity in the United States: An Intergenerational Perspective," *NBER Working Paper No. 24441*.

Chetty, R., N. Hendren, P. Kline, and E. Saez (2014): "Where is the Land of Opportunity? The Geography of Intergenerational Mobility in the United States," *The Quarterly Journal of Economics*, 129, 1553 –1623.

Christiano, L., M. Eichenbaum, and C. Evans (1999): "Monetary Policy Shocks: What Have We Learned and to What End?" in *Handbook of Macroeconomics*, ed. by J. B. Taylor and M. Woodford, Elsevier, vol. 1, chap. 2, 65–148.

Coibion, O. (2012): "Are the Effects of Monetary Policy Shocks Big or Small?" *American Economic Journal: Macroeconomics*, 4, 1–32.

D'Acunto, F., R. Liu, C. Pflueger, and M. Weber (2018): "Flexible Prices and Leverage," *Journal of Financial Economics*, 129, 46–88.

Dahl, M. and T. DeLeire (2008): *The Association Between Children's Earnings and Fathers' Lifetime Earnings: Estimates Using Administrative Data*, Discussion Papers, University of Wisconsin-Madison, Institute for Research on Poverty.

Davis, S. J. and J. Haltiwanger (2001): "Sectoral Job Creation and Destruction Responses to Oil Price Changes," *Journal of Monetary Economics*, 48, 465–512.

Dedola, L. and F. Lippi (2005): "The Monetary Transmission Mechanism: Evidence from the Industries of Five OECD Countries," *European Economic Review*, 49, 1543–1569.

Dustmann, C., P. Puhani, and U. Schönberg (2017): "The Long-Term Effects of Early Track Choice," *The Economic Journal*, 127, 1348 – 1380.

Fuchs-Schündeln, N., D. Krueger, and M. Sommer (2010): "Inequality Trends for Germany in the Last two Decades: A Tale of Two Countries," *Review of Economic Dynamics*, 13, 103–132.

Galí, J. (2015): *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and Its Applications*, Princeton University Press, 2 ed.

Gallagher, R., R. Kaestner, and J. Persky (2018): "The Geography of Family Differences and Intergenerational Mobility," *Journal of Economic Geography.*

Ganley, J. and C. Salmon (1997): "The Industrial Impact of Monetary Policy Shocks: Some Stylised Facts," *Bank of England Working Paper No. 68.*

Gertler, M. and S. Gilchrist (1994): "Monetary Policy, Business Cycles, and the Behavior of Small Manufacturing Firms," *The Quarterly Journal of Economics*, 109, 309 – 340.

Gertler, M. and P. Karadi (2015): "Monetary Policy Surprises, Credit Costs, and Economic Activity," *American Economic Journal: Macroeconomics*, 7, 44–76.

Ghassibe, M. (2018): "Monetary Policy and Production Networks: An Empirical Investigation," *ECB Forum on Central Banking.*

Goldberg, P. K. and R. Hellerstein (2009): "How Rigid are Produer Prices?" *Federal Reserve Bank of New York Staff Report no. 407.*

Golosov, M. and R. E. Lucas (2007): "Menu Costs and Phillips Curves," *Journal of Political Economy*, 115, 171–199.

Gorodnichenko, Y. and M. Weber (2016): "Are Sticky Prices Costly? Evidence from the Stock Market," *American Economic Review*, 106, 165–199.

Gärtner, K. (2002): "Differentielle Sterblichkeit. Ergebnisse des Lebenserwartungssurvey des BiB." *Zeitschrift für Bevölkerungswissenschaft*, 27, 185 – 211.

Gürkaynak, R. S., B. Sack, and E. Swanson (2005): "Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements," *International Journal of Central Banking*, 1, 55 – 93.

Hausner, K. H., D. Söhnlein, B. Weber, and E. Weber (2015): "Bessere Chancen mit mehr Bildung," *IAB Kurzbericht.*

Hayo, B. and B. Uhlenbrock (2000): "Industry Effects of Monetary Policy in Germany," in *Regional Aspects of Monetary Policy in Europe*, ed. by J. von Hagen and C. J. Waller, Springer, 127–158.

Hilger, N. G. (2015): "The Great Escape: Intergenerational Mobility in the United States Since 1940," *NBER Working Paper No. 21217.*

Jarocinski, M. and P. Karadi (2018): "Deconstructing Monetary Policy Surprises: The Role of Information Shocks," *CEPR Disucssion Paper No. 12765.*

Klein, M., K. Barg, and M. Kühhirt (2019): "Inequality of Educational Opportunity in East and West Germany: Convergence or Continued Differences?" *Sociological Science*, 6, 1 – 26.

Klenow, P. and B. Malin (2010): "Microeconomic Evidence on Price-Setting," in *Handbook of Monetary Economics*, ed. by B. M. Friedman and M. Woodford, Elsevier, vol. 3, chap. 6, 231–284.

Mazumder, B. (2005): "Fortunate Sons: New Estimates of Intergenerational Mobility in the United States Using Social Security Earnings Data," *The Review of Economics and Statistics*, 87, 235–255.

Miranda-Agrippino, S. (2016): "Unsurprising Shocks: Information, Premia, and the Monetary Transmission," *Centre for Macroeconomics, Discussion Paper n. 2016-13.*

Nakamura, E. and J. Steinsson (2008): "Five Facts about Prices: A Reevaluation of Menu Cost Models," *The Quarterly Journal of Economics*, 123, 1415–1464.

——— (2010): "Monetary Non-Neutrality in a Multisector Menu Cost Model," *The Quarterly Journal of Economics*, 125, 961 – 1013.

Ottonello, P. and T. Winberry (2018): "Financial Heterogeneity and the Investment Channel of Monetary Policy," *NBER Working Paper No. 24221.*

Pasten, E., R. Schoenle, and M. Weber (2018a): "Price Rigidity and the Origins of Aggregate Fluctuations," *NBER Working Paper No. 23750.*

——— (2018b): "The Propagation of Monetary Policy Shocks in a Heterogeneous Production Economy," *NBER Working Paper No. 25303.*

Peersman, G. and F. Smets (2005): "The Industry Effects of Monetary Policy in the Euro Area," *The Economic Journal*, 115, 319–342.

Pollak, R. A. and T. J. Wales (1979): "Welfare Comparisons and Equivalence Scales," *American Economic Review*, 69, 216–221.

Raddatz, C. (2006): "Liquidty Needs and Vulnerability to Financial Underdeveloptment," *Journal of Financial Economics*, 80, 677–722.

Ramey, V. A. (2016): "Macroeconomic Shocks and Their Propagation," in *Handbook of Macroeconomics*, ed. by J. B. Taylor and H. Uhlig, Elsevier, vol. 2, chap. 2, 71–162.

Riphahn, R. and G. Heineck (2009): "Intergenerational Transmission of Educational Attainment in Germany – The Last Five Decades," *Journal of Economics and Statistics (Jahrbücher für Nationalökonomie und Statistik)*, 229, 36–60.

Riphahn, R. and P. Trübswetter (2013): "The Intergenerational Transmission of Educational Attainment in East and West Germany," *Applied Economics*, 45, 3183–3196.

Romer, C. D. and D. H. Romer (2004): "A New Measure of Monetary Shocks: Derivation and Implications," *American Economic Review*, 94, 1055–1084.

Schmillen, A. and H. Stüber (2014): "Lebensverdienste nach Qualifikation - Bildung lohnt sich ein Leben lang," *IAB Kurzbericht*.

Schnitzlein, D. (2016): "A New Look at Intergenerational Mobility in Germany Compared to the U.S," *The Review of Income and Wealth*, 62, 650–667.

Söderström, U. (2001): "Predicting Monetary Policy with Federal Funds Futures Prices," *Journal of Futures Markets*, 21, 377 – 391.

Solon, G. (1992): "Intergenerational Income Mobility in the United States," *American Economic Review*, 82, 393–408.

Weinand, S. and L. von Auer (2019): "Anatomy of Regional Price Differentials: Evidence from Micro Price Data," *Deutsche Bundesbank Discussion Paper No 04/2019*.

Zimmerman, D. J. (1992): "Regression Toward Mediocrity in Economic Stature," *American Economic Review*, 82, 409–429.