

Discussion Paper Series – CRC TR 224

Discussion Paper No. 662

Project C 01

Local Unemployment, Worker Mobility and Labor Market Outcomes: Evidence from Germany

Johannes Weber¹

March 2025

¹University of Bonn, Email : johannes.weber@uni-bonn.de

Support by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation)
through CRC TR 224 is gratefully acknowledged.

Local Unemployment, Worker Mobility and Labor Market Outcomes: Evidence from Germany*

Johannes Weber[†]

February 27, 2025

Abstract

In most countries, there are large and highly persistent differences in unemployment rates across local labor markets. Such local unemployment rate differences can shape the career outcomes of young who start their careers in different local labor markets. I use high-quality administrative data from Germany to study how workers move between labor markets with different unemployment rates and their resulting lifecycle wage profiles. I find that on average workers who start their careers in lower unemployment regions earn higher wages even when young, experience greater wage growth along the lifecycle and spend less time in unemployment. Even conditional on local price levels and worker fixed effects, I find that between workers from high and workers from low unemployment regions an unexplained wage gap opens up to about 11% until the age of 40. Despite this, I do not find that workers move out of bad labor markets and into good labor markets. Instead, workers spend most of their time in local labor markets with similar relative degrees of unemployment. I find that the differences in wages and unemployment translate into a gap of about 150,000 Euros (adjusted to 2010 level) in real income accumulated until the age of 55.

Keywords: Local Labor Markets, Unemployment, Wages, Worker Mobility

JEL: R10, J31, J61, J64

*I thank Moritz Kuhn and Hans-Martin von Gaudecker for helpful advice and guidance in this project. I also thank participants in lunch seminars at Bonn and Yale as well as participants at various conferences for invaluable feedback and advice. Support by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) through CRC TR 224 (Project C01) is gratefully acknowledged.

[†]University of Bonn, johannes.weber@uni-bonn.de, 53113 Bonn, Germany

1 Introduction

It is well documented that the level of unemployment matters for the development of workers' labor market outcomes. Entering the labor market in a recession at a time of high unemployment can be detrimental to workers' wages even in the long run (e.g. Kahn (2010a), Oreopoulos et al. (2012a), Altonji et al. (2016) and van den Berge (2018)). However, the level of unemployment that young workers face at the beginning of their careers does not only change over time. In most developed economies, there are large and highly persistent differences in unemployment between regional labor markets as well.¹

Policymakers have long recognized this as a source of inequality. Various policy measures aimed at supporting bad local labor markets are put in place by local and national governments in most countries.²

Despite its evident relevance, the effect of entering a career in a bad local labor market has, unlike the effects of starting a career at a bad time, received only scant attention in the literature. This is likely due to data limitations because data that allows researchers to track workers and their workplace over their careers is hard to come by. In this study, I use administrative data from Germany to fill this gap and ask whether it matters for the careers of workers if they enter the labor market in a place of low or high unemployment. There are two components to this question. On the one hand, it is a question about differences in labor market outcomes, i.e. time in unemployment, wages, and ultimately lifetime income. On the other hand, it is also a question about the spatial mobility of workers between local labor markets, since in principle, workers are free to move and look for a job in an area with more favorable conditions. To answer both of these components, I make use of high-quality administrative data from Germany that allow me to track workers over their careers. I study several cohorts of West German male workers born between 1960 and 1980 and their labor market outcomes over the course of their careers, depending on where they entered the labor market.

Germany provides the ideal setting to study this research question. It is a large country with many local labor markets and considerable differences in local unemployment. Behind

¹For documentation of these differences see e.g. Topel (1986) or Kline and Moretti (2013) for the U.S., Bilal (2023) for France, Kuhn et al. (2021) for the UK, Germany, and the U.S. and OECD (2005) for an overview over OECD economies.

²In Germany, too, a set of measures (Gesamtdeutsches Fördersystem) aimed at supporting structurally weak areas relies on local unemployment as a means of determining which area is structurally weak and eligible for support including subsidies for young firms and plants in structurally weak areas.

the overall unemployment rate of 6.7% in 2014, for example, there were local unemployment rates as high as 14.7% in Gelsenkirchen, a city in the industrial Ruhr area, and as low as 2.6% in the Bavarian town of Dillingen a. d. Donau.

The main contribution of this study is to document two novel main facts: Firstly, worker mobility between good and bad labor markets is fairly limited and workers spend most of their careers in areas with similar relative degrees of unemployment. Secondly, workers from better labor markets earn higher wages, experience greater wage growth, and spend less time unemployed. I find that these differences translate into substantial differences in earnings accumulated over workers' careers.

As a first step, I turn to worker mobility and ask whether workers systematically leave bad local labor markets. This turns out not to be the case. Instead, I find that about 40% of workers never leave the commuting zone they entered the labor market in and that workers who move to jobs in other commuting zones usually do not move very far away. Close to 90% of workers never move to a commuting zone that is 200 km or further away from their initial commuting zone. As a result, workers who do move usually move to areas with similar relative degrees of unemployment and spend most of their careers in local labor markets of similar levels of unemployment.

With this in mind, I compare how much time workers spend in unemployment depending on how high unemployment is in their initial local labor market. I find that workers who enter the labor market in worse local labor markets lose their jobs more often and have longer unemployment spells on average. Up to the age of 55 workers who enter the labor market in a commuting zone of the highest unemployment quintile spend about a year more (or about twice as much) time in unemployment than their counterparts from the lowest unemployment quintile.

I then focus on wages. I compare the wages of workers and find that the differences in wage levels and wage growth persist over the lifecycle. The wage gap for full-time employed workers amounts to about 7-10% at 25 and opens up to roughly 20% towards the end of the lifecycle, depending on the cohort. These differences cannot fully be explained by local price levels. In particular, I find that even controlling for worker fixed effects in a regression employing an event study framework leaves an unexplained wage gap that opens up to about 11% over the lifecycle.

Taking both, differences in wages and differences in unemployment together I calculate average accumulated earnings up to the age of 55. I find that the difference between the highest and the lowest quintile amounts to about 250,000 Euros (at 2010-CPI level adjusted

to the regional price level of Bonn) on average. Controlling for local price levels and taking away the part of accumulated earnings explained by observables (education, a worker's occupation, and industry and local urbanization) leaves an average residual gap in these accumulated earnings of about 50,000 Euros (at 2010-CPI level).

I conclude that it does indeed matter considerably, whether workers enter the labor market in a good or in a bad local labor market. If anything, these differences are even more important for younger cohorts.

The remainder of the paper is structured as follows: The following paragraph summarizes related literature. Section 2 describes the data used for the analysis. Section 3 gives an overview of local unemployment rates in Germany and local labor markets. Section 4 documents mobility patterns between the different local labor markets. Section 5 describes differences in labor market outcomes and accumulated earnings. Section 6 concludes.

Related Literature

While labor market outcomes over the full lifecycle in the context of spatial unemployment differences have not been comprehensively documented before, there is a body of literature that this work relates to. There are three main branches of related literature.

Firstly, there are papers concerned with spatial unemployment. Focussing on excessive hiring costs and different local productivity levels Kline and Moretti (2013) discuss circumstances under which place-based hiring subsidies may be beneficial. Bilal (2023) stresses the role of firms and attributes spatial unemployment gaps in France to the agglomeration of productive firms in certain places leading to unproductive worker-firm matches and high job-losing probabilities in bad labor markets. On the empirical side papers often exploit regional employment shocks (e.g. Hornbeck and Moretti (2022)). Perhaps most relevant to my work, Blanchard et al. (1992) document responses to adverse regional employment shocks in the context of U.S. states and find only transitory effects of unemployment shocks on wages. A main difference between these papers and my work is that I can track individual workers over their careers. Blanchard et al. (1992) for instance rely on state aggregate wage data.

More broadly, there is a literature on spatial economic activity that does not necessarily focus on unemployment. Moretti (2011) gives an overview of what is known about gaps in local labor market outcomes and possible explanations for them. Other papers develop models and conditions to explain economic activity differences across space (e.g. Allen

and Arkolakis (2014), Gaubert (2018)). Redding and Rossi-Hansberg (2017) provide an overview of modeling tools and model building blocks for spatial models. Another well-studied area of spatial economics is the one concerned with wage and productivity gaps between rural and urban areas and urban wage premia (e.g. Harris and Todaro (1970), Glaeser and Mare (2001), Gould (2007), Baum-Snow and Pavan (2012), Young (2013) and, notably, for Germany, Dauth et al. (2022)). Local mismatch is another subject more directly related to local unemployment that has received attention. Şahin et al. (2014) and Marinescu and Rathelot (2018) focus on the mismatch of workers and local labor markets and its contribution to overall unemployment for example.

Secondly, there is a literature on worker mobility. Going back to Sjaastad (1962), a branch of spatial economic research focuses on workers' location decisions and worker migration and migration frictions. For example, Roback (1982) presents evidence that local amenities are reflected in local labor market outcomes such as wages. More recently, Amior and Manning (2018) attribute the persistence of local unemployment despite a strong moving response by workers to persistent labor demand shocks. Focussing on moving frictions Bryan and Morten (2019) and Morten and Oliveira (2016) estimate large effects of facilitating worker relocation on productivity in the context of developing countries. A number of papers exploit various employment shocks and study the mobility response (e.g. Levy et al. (2017), Blanchard et al. (1992)).

An influential part of the mobility literature focuses on moving to or away from opportunity. Huttunen et al. (2018) find that displaced workers who decide to move close to their families or rural areas after displacement experience income losses. Nakamura et al. (2016) use a volcano eruption in Iceland to study the inter-generational effect of worker relocation. In terms of inter-generational social mobility a series of influential papers highlights the importance of neighborhoods for children's future careers (e.g. Chetty and Hendren (2018a), Chetty and Hendren (2018b) and Bergman et al. (2024))

The difference between this literature and my work is, that I follow individual workers from all over the country and their location and outcomes over their whole careers as employees. Thirdly, a branch of literature concerns itself with wage structure and the German wage structure in particular. Dustmann et al. (2009) document that wage inequality has been on the rise in Germany in the 1980s and 1990s and Bönke et al. (2015) document a rise of intra-cohort wage inequality. Card et al. (2013) find that the dispersion of task and firm-specific premiums has gone up in West Germany. Relating to regional labor market differences, a number of papers study differences between East and West Germany (e.g. Uhlig (2008),

Uhlig (2006), Spitz-Oener (2007)). My work adds to this literature by documenting the additional spatial component of wage inequality in West Germany.

2 Data

To analyze labor market outcomes over the lifecycles of workers depending on the location a worker works in I need data that is both regional and allows me to track workers and their outcomes and whereabouts over time. The main data source that I choose for this purpose is the Sample of Integrated Labor Market Histories (SIAB) - a 2% sample of workers covered by German social security records. Additionally, the SIAB contains information on short-term unemployment benefits workers receive. Since in Germany only employed (and unemployed) workers are covered by social security, the sample does not contain self-employed workers or workers who are civil servants in government positions. However, a large share of the German labor force - about 80% - is covered by social security. The data version that I use covers the years 1975-2017. Not only do these data meet the mentioned requirements, but they are highly credible administrative data as well. This is important because, in many non-administrative surveys, there could be the worry that people who move to another location have a high probability of dropping out of the sample. With administrative data, this is not a problem.

Because the SIAB is drawn from social security records, wages are only reported up to the maximum social security contribution level. To circumvent this problem I follow the procedure proposed by Gartner (2005) and impute top-coded wages. I describe the details of this procedure in Section A.1 of the appendix.

The regional variable contained in the SIAB is the county (Landkreis) that a worker's current job is registered in. This means that for workers who are currently unemployed, there is no information about the worker's whereabouts. To deal with this issue I assume that workers who lose their job do not immediately move to another location but look for a new job first and move only when starting a new job in a different county. I then fill the gaps in the location variable using the last valid county reported for each worker.

German counties are relatively fine-grained regional units. It is likely that in some places many workers will live in one county and commute to work in another one. It is therefore questionable if counties are a good measure of local labor markets. To circumvent this potential problem I follow Kuhn et al. (2021) and turn to commuting zones published by the Federal Office for Building and Regional Planning (Bundesamt für Bauwesen und

Raumordnung - *BBSR*). These commuting zones group counties into zones taking into account commuting streams between the counties such that the resulting commuting zones are as self-contained as possible.³ There are 203 such commuting zones in West Germany. I impose a couple of restrictions on the sample. First of all, workers from East Germany are recorded in the SIAB only from 1993, and after decades of socialism labor markets in East Germany are likely to be inherently different along many dimensions. For the purpose of this study, I need as much comparability over time as possible and do not want to deal with these differences. I, therefore, restrict the sample to individuals who have no spells in any East German county. Secondly, because I track cohorts born in the '60s and '70s when female labor supply was much smaller than today, I restrict the sample to men. This also ensures comparability with the related literature that usually imposes the same restriction. Apart from worker data, I need data on unemployment at a regional level. As my primary data source, I use county-level yearly unemployment rates I obtained from the BA. This data reports both unemployment rates and the number of registered unemployed workers. I use this information to calculate unemployment rates at the commuting zone level. While this is the most credible unemployment data source for Germany, county-level data from the BA is unfortunately only available for the years after 1985. Before 1985 the SIAB can be used to directly estimate commuting zone level unemployment. I show that the correlation between SIAB estimates and BA unemployment rates is 92% in 1985 in Figure 19 in the appendix.

It is well known, that price levels can vary greatly between different regions within a country. To measure this, I use a regional CPI at the county level from the *BBSR* for the year 2007 to measure local price levels. Unfortunately, to my knowledge, this is the only comprehensive publicly available regional CPI for Germany.⁴ Of course, price levels can change over time, and calculating real wages using a CPI based on prices measured in 2007 for other years is potentially problematic. I take regional building land prices from the federal office for statistics as well as a local rent index from the *BBSR* to validate this

³The counties reported in the SIAB are counties as of 2017. Because of reforms counties' boundaries do change over time. In West Germany, however, this is relatively rare. There are several versions of these commuting zones from different years. The last version before 2017 is from 2011. Since there are only few differences between the 2011 and earlier versions and county codes in the SIAB are valid in 2017 I use the 2011 commuting zones. I use a county code conversion key, also obtained from the Federal Office for Building and Regional Planning, to translate the 2011 commuting zone key into 2017 county terms.

⁴The *BBSR* did collect regional price data in 1994, but compared only 50 cities across Germany, not taking into account rents. Kosfeld et al. (2009) propose an econometric model to estimate regional CPIs based on these data, but working with the comprehensive CPI from 2009 seems preferable.

CPI for other years. Figures 21 and 20 in the appendix show that the correlations between the regional CPI and the prices in different the available years that are the closest and the furthest apart from 2007 are very similar.

3 Local Unemployment in West-Germany and Definitions

3.1 Local Unemployment

Figure 1 shows unemployment rates across West German commuting zones for the years 1990 and 2005 - a year with comparatively low overall unemployment and a year with infamously high overall unemployment. In both years the heterogeneity in unemployment is well visible with local unemployment rates ranging from 2.8% to 13.06% in 1990 and from 5.5% to as much as 21.5% in 2005.

Figure 1 also hints at another well-established fact. Local unemployment is typically very persistent over time. Figure 2 makes this explicit and plots commuting zone level unemployment rates of 1997 and 2005 as well as 1990 and 2010 against each other. 1997 and 2005 are two years with high overall unemployment while 1990 and 2010 saw much less overall unemployment. In both cases, the green dots that represent the commuting zones align closely along the 45-degree line indicating a very high level of persistence. Given that there are 20 years between 1990 and 2010 it is striking how high the persistence of local unemployment is with a correlation of close to 80% between local unemployment rates in the two years.

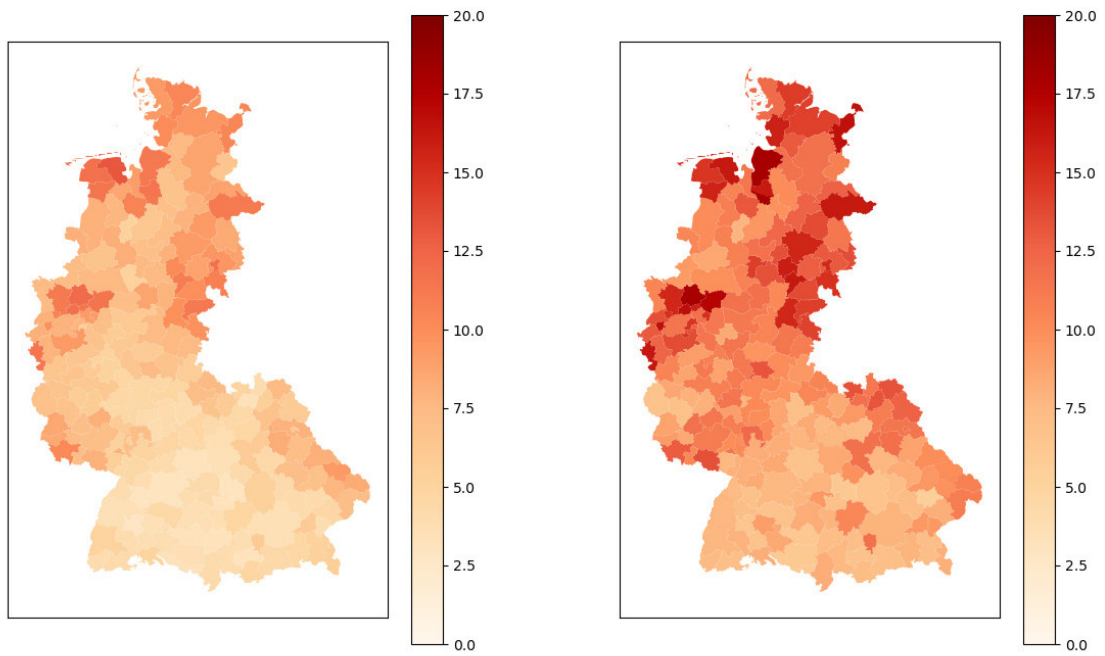
Given this persistence, a natural expectation could be that high and low-unemployment commuting zone are different in some fundamental way. However, in terms of the most obvious candidate observable - urbanization and education - this is not the case. Figure 3 shows the distributions of average unemployment rates over the 1990s and 2000s decades in the three urbanization categories, that the *BBSR* assigns to the commuting zones described in the previous section. Clearly, while it is true that the commuting zones with the most extreme unemployment are urban, it is not true that high unemployment exists only in urban commuting zones. In fact, in both decades, the median unemployment rate in rural areas is higher than in urban areas. Generally, the distributions cover a similar support.

Next, I turn to educational attainment. Figure 4 plots average unemployment rates against

Figure 1: Unemployment across commuting zones

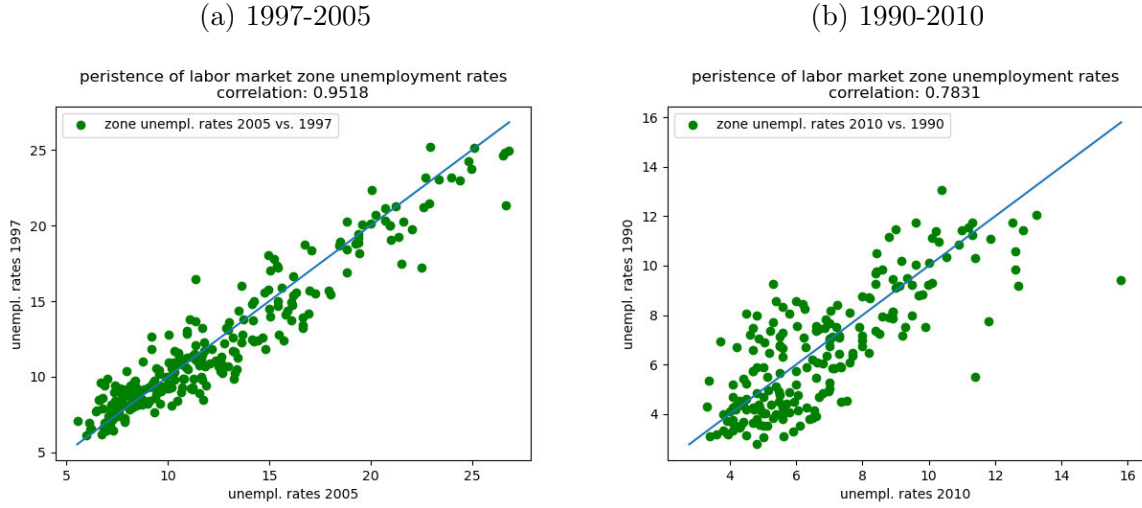
(a) 1990

(b) 2005




Notes: The scale is in %-terms. The maps shown are for the years 1990 (left) and 2005 (right). The regions shown are the commuting zones described in Section 2. Unemployment data comes from the BA.

Figure 2: Persistence of Unemployment




Notes: 1997 unemployment rates against 2005 unemployment rates (left) and 1990 unemployment rates against 2010 unemployment rates (right). Each green dot represents a commuting zone. The scale is in %-terms.

Figure 3: Unemployment and Urbanization



Notes: Boxplots of average unemployment rates over the 1990s (left) and the 2000s (right) in the three urbanization categories for *BBSR* commuting zones. The boxes represent the quartiles of the distribution, the whiskers extend to the unemployment rate the furthest away but within 1.5 inter-quartile ranges from the box. Points beyond that are shown as diamonds. The line through the box represents the median.

Figure 4: Unemployment and Education



Notes: Regression plots of average unemployment rates over the 1990s (left) and the 2000s (right) against the average share of workers with a college degree.

average shares of workers with college degrees in commuting zones in regression plots. To calculate the college share I compute the share of workers in the SIAB who report a tertiary degree.⁵ Figure 4 reveals that by and large, there is no systematic correlation between local unemployment and the college share.

As a consequence, the findings presented in the remainder of this study should not be regarded simply as an artifact of urban and rural differences or agglomeration of educated and uneducated workers. Instead, they are the result of structural differences in the quality of local labor markets.

3.2 Definitions

In order to measure differences between local labor markets with high and low unemployment, I need a notion of what "high" and "low" unemployment mean. In particular, what really matters is not the absolute level of unemployment itself, but how high unemployment in a commuting zone is, relative to the other commuting zones. I decide to work with quintiles and assign the ranks *1 very low* unemployment, *2 low* unemployment, *3 moderate* unemployment, *4 high* unemployment, and *5 very high* unemployment to each commuting zone each year, taking quintiles over commuting zone-level unemployment rates. I refer to these ranks as the *plain ranks* in the remainder of the paper.

⁵I include both, university and university of applied science degrees.

When comparing workers, however, I need one more step because, in principle, workers can move between the commuting zones. When I want to compare workers over their lifecycle I use the most natural option: I simply use the *plain rank* of the commuting zone a worker enters the labor market in for the very first time, in the year the worker starts working there. I refer to this as the *entry rank*.


As Figure 5 demonstrates, depending on the overall level of unemployment in a given year, the *plain ranks* can be assigned to quite different absolute levels of unemployment. While I am interested in relative differences in local unemployment, it is important to keep these considerable differences in overall unemployment between different points in time in mind. Workers who enter the labor market at a time of high overall unemployment are likely to experience different labor market outcomes than a worker who enters under more favorable conditions.⁶ To avoid these differences driving the results I measure, I take a cohort perspective and group the birth years of workers in the SIAB into cohorts. Since I want the cohorts to face similar conditions in the labor market over their lifecycles, I choose cohorts of 2-3 birth years such that the overall level of unemployment in West Germany is similar for all members of a cohort at age 25. Because unemployment was very low in the 1970s and early 80s, I focus on later cohorts who are faced with more substantial unemployment rates during their careers and consider the cohorts 1960-1962, 1963-1965, 1966-1967, 1968-1970, 1971-1973, 1974-1976 and 1977-1979. Many findings are very similar for all cohorts. Instead of showing results for all cohorts in the main text, unless labeled otherwise, I hence show results for the 1963-1965 and the 1971-1973 cohorts in the main text and results for the remaining cohorts in the appendix. I do not consider later cohorts because the SIAB data covers only the years until 2017 and workers born after the 1970s can only be observed until their 30s which makes comparing outcomes over the lifecycle difficult.

4 Worker Mobility

A central question to assess how harmful starting a career in a bad local labor market is is how mobile workers are between local labor markets of different quality. If workers can freely migrate from places with high unemployment to places with better conditions, entering the labor market in a place with high unemployment should not have great effects on long-term labor market outcomes.

⁶see e.g. Kahn (2010b) or Oreopoulos et al. (2012b).

Figure 5: Unemployment rates in the different ranks: 1990 vs. 2005



Notes: Box Plots of unemployment rates by rank in the commuting zones by the rank the commuting zones are assigned. The box plots display the upper and lower quartiles, median, and whiskers extending to the furthest data points within 1.5 times the interquartile range. The left side of the figure corresponds to the plots for the year 1990, while the right side represents the plots for the year 2005.

As a first step, ignoring local unemployment, I study the overall mobility of workers over the course of their entire careers. Table 1 shows the fraction of careers, that take place within a given distance from the commuting zone a worker starts his career in for different cohorts. To ensure that results are not driven by workers who cannot be observed for sufficiently long I restrict the sample to workers who are first observable before or at the age of 25 and who can still be observed at or after the age of 40. As column 1 indicates, about 40% of workers work in the same commuting zone for their entire careers. Column 2 reveals that almost 60% of workers never move further than to a commuting zone, adjacent to their initial commuting zone. Columns 3 and 4 show that about 80% of workers never move further than 100 km and almost 90% of workers never move further than 200 km away from their starting commuting zone.

The high figures presented in table 1 establish the fact that overall worker mobility is low and that most careers in West Germany take place within narrowly confined areas. Nevertheless, about 60% of workers do move at some point in their careers. This raises the question if workers who do move seek out good local labor markets and avoid bad ones. To provide an answer, I collect all moves between commuting zones that happen throughout the whole careers of workers⁷ and check which of the quintile ranks the move is from and which of the ranks the move is to. Then, for each of the ranks of the commuting zones that the moves are from, I calculate the relative frequencies of moves that go to commuting zones with each of the 5 ranks a move can be to:

⁷Results are very similar when looking at moves at different ages separately.

Table 1: Fractions of workers who stay within a given distance from their first commuting zone


Birth Years	Stayers	Neighboring Zone	100 km	200 km
1960-1962	0.42	0.59	0.79	0.87
1963-1965	0.43	0.59	0.79	0.87
1966-1967	0.42	0.59	0.79	0.86
1968-1970	0.42	0.59	0.79	0.87
1971-1973	0.41	0.58	0.78	0.86
1974-1976	0.41	0.56	0.78	0.86
1977-1979	0.42	0.56	0.77	0.86

Notes: The sample is restricted to workers with observations at age 25 and at or after age 40. The first column shows the fraction of workers who stay in the same commuting zone for their whole observable career. The second column shows the fraction of workers that never move further than to a commuting zone adjacent to their initial commuting zone. The third and fourth columns show the fractions of workers that stay within a radius of 100 km and 200 km respectively. Distance between two commuting zones is measured as the minimum distance between the borders of these commuting zones.

$$a_{ij} = \frac{\text{\#moves from rank } i \text{ to rank } j}{\text{\#moves from rank } i}, \quad i, j = 1, 2, 3, 4, 5$$

The resulting 5×5 matrix A is shown in Figure 6 for the 1960-1962 and 1971-1973 cohorts. For example, the top left corner of the left matrix states that 46% of all moves away from a commuting zone with the rank *1 very low* that among workers born between 1963 and 1965 were to another commuting zone with the same rank. By contrast, the top right corner of the same matrix states that only 5.5% of moves away from an *1 very low* commuting zone were to commuting zones with the rank *5 very high*. As the deeper shades of blue towards the diagonals of the two matrices suggest, workers, when they move, tend to move to commuting zones with similar levels of unemployment to the commuting zone they are leaving. Perhaps unsurprisingly, this is particularly pronounced for moves away from very good labor markets, more than 40% of which go to *1 very low*-commuting zones in all cohorts. But even for moves away from high unemployment zones moves tend to be to places with similar degrees of unemployment. In all cohorts over 60% of moves away from commuting zones with rank *5 very high* are to a zone with either the rank *5 very high* or *4 high*. Fewer than 7% of workers who leave commuting zones with the highest two ranks move to the very best labor markets with the lowest rank. While natural expectation could

Figure 6: Fractions of moves by unemployment rank of origin- and destination commuting zones



Notes: Each cell in the matrices above represents the fraction of moves away from commuting zones with an unemployment rank indicated by the row to a commuting zone with an unemployment rank indicated by the column. E.g., the second cell in the first row of the left matrix indicates that 28% of moves away from zones with rank *1 very low* went to zones with rank *2 low*. Workers are ranked based on the unemployment quintile of the commuting zone where they initially enter the labor market. The matrix on the left shows the results for the 1960-1962 cohort, and the matrix on the right shows the results for the 1971-1973 cohort.

be that this pattern could be very different if only the moves of unemployed or employed workers were counted, I show in Figure 26 and Figure 27 in the appendix that this is not the case and that the pattern does not change much even when workers are unemployed just before moving.

Given that as discussed most workers never move over great distances, the question arises whether these findings are simply a product of agglomeration of good and bad local labor markets or whether these patterns persist conditional on the distance over which a move occurs. To formally test whether this is the case, I adopt a gravity view of migration flows and impose that the number of people f_{ijt} that move between zones i and j in year t are a function of the distance d_{ij} between the two zones, size of the two zones s_i and s_j and the numbers of unemployed workers in both zones u_i and u_j , as well as a year effect β_t . Specifically, I assume that

$$f_{ijt} = \beta_t \frac{s_{it}^{\beta_1} s_{jt}^{\beta_2} u_{it}^{\beta_3} u_{jt}^{\beta_4}}{d_{ij}^{-\beta_5}} \quad (1)$$

Taking logs equation 1 can then be estimated as

$$\log f_{ijt} = \log \beta_t + \beta_1 \log s_{it} + \beta_2 \log s_{jt} + \beta_3 \log u_{it} + \beta_4 \log u_{jt} + \beta_5 \log d_{ij} + e_{jit} \quad (2)$$

where e_{ijt} is the error term and the coefficients can be interpreted as moving flow elasticities. Since the SIAB is a 2% sample of the German labor force that is representative at the county level, I estimate the size of each commuting zone by multiplying the number of workers in each commuting zone by 50. f_{ijt} is the total number of moves from i to j for all commuting zones i and j in the SIAB. I use data from the years 1985-2017. Numbers of unemployed workers come from the BA. Distance between commuting zones is measured as the distance between the two centroids of the two commuting zones. Results are reported in column (1) of Table 2.

Table 2: Gravity regression results

	(1)	(2)
log size of origin zone	0.2210 (0.011)	0.2361 (0.010)
log size of destination zone	0.2500 (0.010)	0.2364 (0.010)
log distance	-0.7102 (0.081)	-0.6987 (0.083)
log n. unempl. origin	0.0370 (0.006)	- -
log n. unempl. destination	0.0085 (0.009)	- -
log (n. unempl. or.) \times (n. unempl. dest.)	- -	0.0218 (0.007)
year fixed effects	yes	yes

Notes: The outcome variable is the log of the number of workers who moved from the origin commuting zone to the destination commuting zone. Standard errors are clustered at the year level and reported in parentheses below each coefficient.

Unsurprisingly the size of both the origin and the destination commuting zone are positively associated with the number of moves and the elasticity with respect to distance is negative and sizeable, further emphasizing the point that workers rarely move over long distances. The smaller and positive coefficient for the log number of unemployed workers in the zone of origin is in line with more unemployed workers looking for and eventually finding

jobs in different commuting zones. Crucially though the coefficient for the log number of unemployed workers in the destination zone is neither negative nor statistically significant. This means that even conditional on distance, it is not the case that workers seek out commuting zones with less unemployment.

In order to be able to reproduce the finding from Figure 6 that moreover, workers who move tend to move to similar commuting zones in terms of local unemployment I also estimate a gravity equation that includes an interaction term of the numbers of workers who are unemployed in both origin and destination zone. The gravity equation then becomes

$$f_{ijt} = \beta_t \frac{s_{it}^{\beta_1} s_{jt}^{\beta_2} (u_{it} u_{jt})^{\beta_3}}{d_{ij}^{-\beta_4}} \quad (3)$$

Results from estimating this equation in logs are in column (2) of Table 2. The coefficient for the interaction term is indeed positive and statistically significant, if far smaller than distance. This means that even conditional on distance, workers who move tend to move to commuting zones with similar degrees of unemployment.


The persistence of local unemployment and the now documented lack of mobility of workers over long distances and between good and bad local labor markets together imply that workers spend most of their careers under similar conditions in terms of local unemployment. Figure 7 puts numbers to this fact. It shows the average fraction of days in a career spent in each of the local unemployment quintile ranks, conditional on the *plain rank* of the commuting zone a worker first entered the labor market in at the time of labor market entry. It is worth remembering that while as demonstrated regional mobility of workers is limited and local unemployment is quite persistent over time, in principle, over the course of a career both, locations and a location's unemployment rank can change. Nevertheless, the rank of a worker's initial commuting zone seems to be a good predictor of the degree of unemployment that a worker is faced with throughout his career. Workers born between 1960 and 1962 who start their observable careers in a commuting zone with rank *5 very high* spend 36% of their time in commuting zones that have the same rank. Strikingly, as much as 71% of time is spent in commuting zones with either the rank *5 very high* or with the rank *4 high* on average. At the other extreme, more than half the time of workers who start their careers in the best labor markets with rank *1 very low* is spent in zones with that same rank, and more than 80% of these workers' time is spent in places with the lowest two ranks. Results for the other cohorts are quite similar.

The results presented in this section suggest that unemployment rates are persistent not

Figure 7: Fractions of time spent in the five unemployment ranks

(a) 1960-1962 Cohort

(b) 1971-1973 Cohort



Notes: Fractions of time spent in commuting zones of each of the five ranks by the rank of the initial commuting zone where a worker first entered the labor market. The initial ranks are on the y-axis. For example, the second cell in the first row of the left graph shows that workers who start in a commuting zone with rank *1 very low* spend 32% of their time in commuting zones with rank *2 low* on average. Workers are ranked based on the unemployment quintile of the commuting zone where they initially enter the labor market. The matrix on the left shows results for the 1960–1962 cohort, and the matrix on the right shows results for the 1971–1973 cohort.

only over time across locations, as demonstrated in Section 3. In fact, as a result of this regional persistence and limited worker mobility, they are persistent at the worker level as well. Workers who start their careers in depressed local labor markets are likely to spend most of their time in local labor markets with a high degree of local unemployment.


5 Workers' Labor Market Outcomes

Having established that workers usually spend most of their time in similar local labor markets, I now turn to the question of what this means for their labor market outcomes. The most direct way in which long periods of time could be harmful to workers is leading to more unemployment. I turn my attention to this channel now.

Figure 8: Average accumulated number of job losses over the lifecycle

(a) 1960-1962 Cohort

(b) 1971-1973 Cohort



Notes: The average number of job losses workers from each rank have had at each age. Workers are ranked based on the unemployment quintile of the commuting zone where they initially enter the labor market. Workers who have no observations after an age younger than 55 are assumed to lose no further jobs. Job losses are defined as transitions from an employed spell to an unemployed spell with a gap of fewer than 31 days between them. Birth years 1960-1962 on the left, birth years 1971-1973 on the right.

5.1 Time in Unemployment


Because the SIAB reports observations exactly to the day I can not only detect a job loss in a worker's labor market history but also calculate how long this unemployment spell lasts. A caveat is that not all sources of unemployment spells exist for all years. Before 1997 only spells during which workers receive short-term unemployment benefits are recorded in the data.⁸ From 1997 onwards workers show up as unemployed whenever they are registered as unemployed with the Federal Employment Agency (BA). As a result, in this Section numbers for years before 1997 exclude workers who are long-term unemployed.

Figure 8 shows how often workers have lost their jobs on average at each age. I count any occurrence of an unemployment spell within a month after having been employed as a job loss to ensure that I include job losses where workers fail to register as unemployed immediately which could result in short gaps in the career history. I again use the *entry rank* described in section 3.2 to rank workers.

It is evident from Figure 8 that the higher the unemployment rank, the more job losses workers accumulate on average. This is true for all other cohorts as well, as demonstrated

⁸In Germany, every employed worker contributes to unemployment insurance. Workers who were employed for 12 months prior to becoming unemployed receive unemployment benefits from the insurance for 6-12 months (up to 24 months for workers who are older than 50). Workers who do not meet this criterion or remain unemployed for longer periods of time receive social benefits (ALG II). Information about these unemployment benefits is contained in the SIAB. Unfortunately, there is no information about ALG II recipients before the year 2005 in the SIAB.

Figure 9: Average length of unemployment spells



Notes: The averages presented in this figure represent the average number of days, of unemployment spells for workers aged 25-55. Workers are ranked based on the unemployment quintile of the commuting zone where they initially enter the labor market. The confidence intervals displayed are standard 95% confidence intervals.

in Figure 30 in the appendix. At age 55, workers who started in the rank *1 very low* accumulated about 0.8 job losses on average while their counterparts who started in the rank *5 very high* accumulated about 0.6. At the end of their observable career, (i.e. at age 43) workers with the worst rank have accumulated 0.6 more job losses on average than their peers with the best rank. In all other cohorts, the difference is at a similar order of magnitude - confer Figure 30.


Apart from losing their jobs more often, another channel through which high local unemployment may be harmful is greater difficulty in finding a new job, once unemployed. Figure 9 shows the average length of an unemployment spell in each cohort depending on the rank. Despite the clear trend of unemployment spells becoming shorter on average for younger cohorts, it is clear that workers with higher unemployment ranks have longer unemployment spells. Workers with the highest unemployment rank spend more than 10 days longer in unemployment than their peers with the lowest rank.

Figure 10 shows the result of combining the results presented in Figure 8 and Figure 9 by showing how much time in unemployment workers accumulate on average over the course of their careers by age. As expected, the plots for the different ranks appear in the expected order as workers with higher unemployment ranks lose their jobs more often and have

Figure 10: Average accumulated time spent in unemployment

(a) 1960-1962 Cohort

(b) 1971-1973 Cohort



Notes: The average number of days that workers have accumulated up to each age. Workers who have no observations after an age younger than 55 are assumed to accumulate no further time in unemployment. Workers are ranked based on the unemployment quintile of the commuting zone where they initially enter the labor market. The confidence intervals displayed are standard 95% confidence intervals around the mean at each age. Results for the 1960-1962 cohort are on the left, and results for the 1971-1973 cohort are on the right.


longer unemployment spells on average. At any age workers with higher unemployment ranks accumulate more time in unemployment. At the end of the observable lifecycle workers with the highest unemployment rank accumulate more than a year more time in unemployment on average than their peers with the lowest unemployment rank.

5.2 Wages

Having documented, that workers who enter the labor markets in high unemployment regions spend more time in unemployment, I now turn to wages and ask if workers in depressed local labor markets face lower wages than their peers as well.

Figure 11 shows log wage profiles for the 1960-1962 cohort. Since I only observe daily wages and cannot observe the number of hours a part-time worker works, I restrict the sample to full-time employed workers to ensure comparability. I use the *entry ranks* described in Section 3.2 to rank workers. While the shown log wage profiles all exhibit the familiar hump shape associated with log-wage lifecycle profiles, it is difficult to see the difference between the profiles of the different ranks. For this reason, in the remainder of the paper, I use the log wage profile of the 3 *moderate* rank as a benchmark and subtract it from the other four profiles, when comparing wages over the lifecycle. Figure 12 shows the result

Figure 11: Log-wage lifecycle profiles of workers born 1960-1962



Notes: Workers are ranked based on the unemployment quintile of the commuting zone where they initially enter the labor market. Confidence bands represent standard 95% confidence intervals around the mean log-wage at each age. The sample is restricted to workers who work in full-time employment at each age.

for the 1963-1965 and 1971-1973 cohorts.

For both cohorts, there is a clear difference in levels. At age 25, there is a difference of about 7 log points (or about 7%) for the 1960-1962 cohort and a difference of about 12 log points for the 1971-1973 cohort. A similar difference of about of around 10% exists for all the other cohorts as well.

A second observation is, that for most cohorts not only a difference in level exists, but that a difference in growth exists as well. For the 1971-1973 cohort, this is very pronounced as the wage gap opens up from about 7% at 25 to about 17% at 44. The wage gap for the 1960-1962 cohort opens up to about 15%. However, for all remaining cohorts, the growth level is also visible and the wage gap opens up over the lifecycle (the remaining plots are in Figure 32 in the appendix).


Given that I neither condition the sample on workers who stay in the same labor market nor on workers who are directly affected by unemployment, this relationship between local unemployment at labor market entry and the development of lifecycle profiles is striking. Not only do workers spend a lot of time in labor markets with a similar relative degree of unemployment, but this is clearly reflected in the development of wages over the lifecycle as well, stressing the important role of the local labor markets where workers start their careers.

Even though the wage gaps shown in Figure 12 are sizable, as discussed in Section 4, workers

Figure 12: Log-wage lifecycle profiles relative to rank 3 *moderate*

(a) 1960-1962 Cohort

(b) 1971-1973 Cohort



Notes: Average log-wage profiles of workers by unemployment rank. Workers are ranked based on the unemployment quintile of the commuting zone where they initially enter the labor market. From each of the shown log wage lifecycle profiles the profile of the 3 *moderate* rank has been subtracted. Confidence intervals are constructed at the 95% level, at each age point. The left side of the figure represents birth years 1960-1962, while the right side represents birth years 1971-1973.

do not seem to systematically leave bad labor markets and seek out the commuting zones with lower unemployment and higher average wages. One way of explaining this could be that, while nominal wages are higher in low unemployment labor markets so are local prices. Indeed, some low-unemployment labor markets in Germany such as Munich, are infamous for high prices and rents.

To quantify how much of the result is driven by these price differences I reproduce the same figure using the regional CPI from the *BBSR* described in Section 2. Figure 13 shows the result.


Clearly, a good part of the differences is absorbed by the CPI correction. The 1960-1962 cohort wage profiles exhibit no difference at 25 and the gap opens up to only 10% at 55. The wage gaps in the 1971-1973 cohort, too, shrink somewhat to about 3% instead of 10% at 25, opening up to about 15% instead of about 20% at 43. Nevertheless, a good part of the wage gaps persist and cannot be explained fully by the local price level and for the younger cohorts in particular the wage gap remains quite substantial.

An explanation for spatial wage gaps that is common in the literature is worker and firm sorting where wage differences are attributed to spatial concentration of productive firms or productive workers. To determine whether sorting of worker types is the driver of the observed wage gaps, I make further use of my ability to track workers along their careers in the SIAB and estimate effects on wages of having started in the different quintile ranks conditional on worker fixed effects. I pool the data of all cohorts of age 25-40 together and

Figure 13: Regional CPI-adjusted log-wage lifecycle profiles relative to rank 3 *moderate*

(a) 1960-1962 Cohort

(b) 1971-1973 Cohort



Notes: Average regional CPI-adjusted log-wage profiles of workers by unemployment rank. The CPI adjustment was calculated using the regional CPI from the BBSR. Workers are ranked based on the unemployment quintile of the commuting zone where they initially enter the labor market. From each of the shown log wage lifecycle profiles the profile of the 3 *moderate* rank has been subtracted. Confidence intervals are constructed at the 95% level, at each age point. The left side of the figure represents birth years 1963-1965, while the right side represents birth years 1971-1973.

estimate

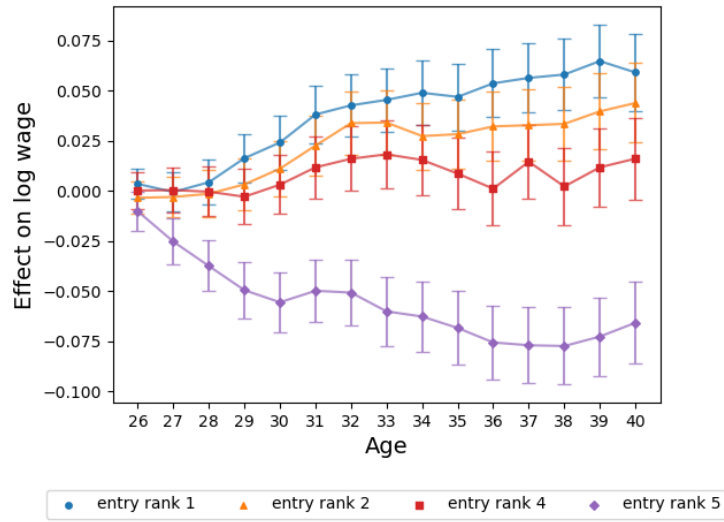
$$\log \text{wage}_{ia} = \sum_{\substack{k=1,2,4,5 \\ j=26,\dots,40}} \mathbb{1}\{a = j\} \mathbb{1}\{\text{entry rank} = k\} \beta_{jk} + a\beta_1 + a^2\beta_2 + \alpha_i + e_{ia} \quad (4)$$

where α_i is a worker fixed effect for worker i , $\beta_{aq(i)}$ is the effect on the log-wage of being a years old and having first entered the labor market in a commuting zone of rank $q(i)$. Since I can only observe daily and not hourly wages, I exclude all workers who ever work part-time.⁹ In order to estimate differential wage growth patterns I control for age and age squared. I take age 25 and the middle rank (3 *moderate*) as benchmarks to avoid multicollinearity. Wages are adjusted using the regional CPI from the BBSR. The resulting fixed effects $\{\delta_{aq(i)}\}_{a=25,\dots,40}$ for the extreme ranks 1 *very low* and 5 *very high* are plotted in Figure 14.

The pattern of differential wage growth from Figures 12 and 13 is preserved even conditional on worker fixed effects. At the beginning of the career at 26, there is no difference in effects on wages for the two entry ranks. At age 40, however, this gap in wage effects has opened up to about 11%. The order of magnitude of this gap is similar to the one exhibited by the raw regional CPI adjusted shown in Figure 13. An important difference is, that wage

⁹I do not keep workers who work part-time only for some time because having worked part-time may have effects on full time wages later in a career.

Figure 14: Fixed effect regression results



Notes: Estimates for $\{\delta_{a1}\}_{a=25,\dots,40}$ from Equation 4 in blue and $\{\delta_{aq5}\}_{a=25,\dots,40}$ in purple. They represent wage effects of age, having started a career in the ranks *1 very low* and *5 very high* respectively, conditional on worker fixed effects. Wages are adjusted using the regional CPI from the *BBSR* before estimating the regression. Confidence intervals are at the 95% confidence level. Standard errors are clustered at the worker level.

effects for the top four ranks are very similar whereas having started a career in the worst rank leads to increasingly less wage growth, compared to the other ranks. Nevertheless, at least from the worker side, the selection of types does not explain the wage gap between the *5 very high* and *1 very low* unemployment regions.

Of course the presented analysis is silent about firm selection as this study focuses on the worker perspective. It is possible that firm sorting as proposed by e.g. Bilal (2023) is indeed important for explaining the wage gap.

5.3 Accumulated Earnings


Section 5.1 documented that workers with higher *entry ranks* accumulate more time in unemployment and in Section 5.2, I documented that workers with higher *entry ranks* earn lower average wages. Both have implications for lifetime earnings. Lower wages directly lead to lower lifetime earnings. Time in unemployment leads to additional wage losses since the unemployment benefits workers receive are lower than the wage they would have earned in employment. On top of that, it is well documented that time in unemployment leads to reduced wages after finding a new job - e.g. through less human capital accumulation or hampered further job ladder climbing.¹⁰ More accumulation of unemployment, therefore, lowers lifetime earnings through both channels. In this section, leveraging the fact, that I can observe daily wages in the SIAB, I calculate accumulated earnings over the lifecycle depending on the *entry ranks*.

Figure 15 shows accumulated lifetime earnings in terms of 2010-Euros by *entry rank*, both unadjusted and corrected for regional price differences using the regional CPI. To make wage levels comparable over cohorts, I CPI-adjust all wages to the level of 2010, using the CPI from the federal statistics office. Since I can only observe workers up to 2017, only the 1960-1962 cohort can be observed for close to a full lifecycle (up to age 55). Results for the other cohorts are therefore only partial lifetime earnings and I will focus on the first cohort in this section. In order to make sure no results are driven by workers dropping out of the sample¹¹, I restrict the sample to workers who are unobservable for no longer than 10% of time between ages 25 and 55 (or the cohort's last observable age).

¹⁰E.g. Jacobson et al. (1993) and Lachowska et al. (2020).

¹¹This can be due to a number of reasons. Workers may genuinely leave the labor force, but they may also become civil servants, and soldiers, become self-employed, or take up a job abroad. Additionally, the data before 1997 has no information on workers under long-term unemployment who receive ALG II unemployment benefits. Unfortunately, there is no way of telling any of these gaps apart and gaps without any information on earnings lower accumulated earnings over the lifecycle.

Figure 15: Average accumulated earnings by Rank



Notes: Accumulated lifetime earnings by unemployment rank of the commuting zone workers started their first job in. The left side of the figure represents birth years 1960-1962, while the right side represents birth years 1971-1973. Earnings are accumulated up to the age of 55 in the left graph and up to the age of 44 on the right. The sample has been restricted to workers who are observable for at least 90% of the time between the age of 25 and the last observable age in each cohort. The price level has been adjusted to the level of 2010 using the CPI from the Federal Statistics Office. The darker bars have additionally been adjusted to the price level of Bonn using the regional CPI from the BBSR.


The gap between the highest and the lowest unemployment *entry rank* for the oldest cohort amounts to 250,000 Euros or about 18% higher accumulated earnings. When correcting for region prices the difference shrinks to 150,000 Bonn-Euros or a difference of about 10%. To put these numbers into perspective, the average price for a new car in Germany in 2010 was about 26,000 Euros. These large numbers emphasize the long-term consequences of a career in a depressed labor market. The same gap for the younger 1971-1972 cohort (shown on the right of Figure 15) in accumulated earnings is about 150,000 nominal Euros or about 90,000 Bonn-Euros at age 44. Since the wage gaps between the unemployment ranks open up with age as documented in Section 5.2, this difference will likely become even larger until workers in this cohort retire.

To determine how much of these gaps in accumulated earnings are individual effects it would be ideal to take out worker fixed effects. This is not feasible, however, since the unemployment rank a worker enters the labor market in and worker fixed effects are collinear. Instead, I project the accumulated, regional-CPI-adjusted earnings reported in Figure 15 on observables. Specifically, I estimate

$$\log \text{accumulated wages}_i = \text{education}_i + \text{industry}_i + \text{occupation}_i + \text{urbanization}_i + e_i \quad (5)$$

I use the highest educational attainment ever reported for each individual i and the oc-

Figure 16: Average residualized accumulated earnings




Notes: Average residualized accumulated earnings until the age of 55 by unemployment rank of the commuting zone workers started their first job in. Residualization is achieved by projecting accumulated earnings on the specification represented by Equation 5. The figure shows results for the birth years 1960-1962.

cupation each individual i spent the most time working in. The SIAB reports 120 levels of occupations based on the German KldB 3-digit classification and 8 levels of industries based on the WZ08 classification. $urbanization_i$ is a 3-level index that classifies the worker i 's first-ever commuting zone into urban, rural, and in-between. It is taken from the *BBSR*. Figure 16 shows average residuals from this regression by unemployment rank for the 1960-1962 cohort.

The difference in the 1960-1962 cohort between the highest and the lowest rank is about 50,000 Euros. Although the difference between the residualized accumulated earnings is much smaller than between the raw accumulated earnings, a sizeable part of lifetime income is not absorbed by the rich observables I project on. This emphasizes the predictive power of initial local labor market conditions for the career outcomes of workers even further and raises the question of why workers do not move away from bad local labor markets.

The results shown in Figure 15 are based solely on labor income. However, unemployed workers receive unemployment insurance benefits. Given the persistence of local unemploy-

Figure 17: Average accumulated earnings and benefits by rank




notes: Accumulated lifetime earnings by unemployment rank of the commuting zone workers started their first job in for the 1960-1962 cohort. Earnings are accumulated up to the age of 55 in the left graph and up to the age of 44 on the right. The sample has been restricted to workers who are observable for at least 90% of the time between the age of 25 and the last observable age in each cohort. The price level has been adjusted to the level of 2010 using the CPI from the Federal Statistics Office. The darker bars have additionally been adjusted to the price level of Bonn using the regional CPI from the BBSR. The numbers include (short-term) unemployment benefits.

ment, unemployment insurance benefits redistribute between regions as every employed worker pays contributions, but only unemployed workers receive benefits. In the remainder of this section, I show that this is the case to a measurable degree, but not enough to change the situation of workers in high unemployment regions meaningfully.

Figure 17 shows accumulated earnings of the 1960-1962 cohort when unemployment benefits are included. The difference between the highest and lowest unemployment rank in accumulated earnings is hardly different when unemployment benefits are included. A likely driver is that wages in low unemployment commuting zones are higher, as demonstrated in Section 5.2, and therefore so are short-term benefits which replace a fixed part of a worker's last wage.

A caveat is, that only short-term unemployment benefits are observable in the SIAB. In Section B.1 in the appendix, I describe a procedure to impute missing unemployment benefits and show that it does not meaningfully change the difference in accumulated income between the unemployment ranks either.

Figure 18: Net unemployment benefits by unemployment ranks over time



Notes: The plot on the left shows contributions to social security subtracted from the sum of benefits as observable in the SIAB by unemployment rank. The plot on the right shows the overall unemployment rate in West Germany for comparison.

Nonetheless, the regional nature of unemployment makes unemployment insurance redistributive between regions on a small scale. Figure 18 plots net benefits (i.e. contributions to social security subtracted from the sum of benefits as observable in the SIAB) and the overall unemployment rate in West Germany. As is evident, particularly in high recessions, net benefits paid in high unemployment regions are higher¹². This is of interest from a policy perspective in particular because local unemployment is so persistent and the regions that benefit the most from unemployment insurance are the same regions for extensive time periods.

6 Conclusion

Starting from the well-known fact that there is sizable and persistent heterogeneity in local unemployment, I documented that workers who are not very mobile between good and bad local labor markets. Instead, I find that driven by a lack of mobility over long distances, the degree of relative unemployment in a worker's initial commuting zone is a very good predictor of the degree of local unemployment the worker will face throughout his working life.

¹²Social security contribution rates are not fixed over time, so the numbers in Figure 18 depend on these rates as well. I show these rates together with more details about the unemployment insurance system in Germany in Section B.1 in the appendix.

I first document that workers who start in worse labor markets also accumulate more time in unemployment, losing their jobs more often and finding new jobs more slowly on average. I then proceed to demonstrate that workers earn different wages, depending on whether they enter the labor market in a good or a bad local labor market. I find wage gaps of about 7-10% at age 25 that open up to about 20% toward the end of the observable lifecycle between workers who started in the highest and workers who started in the lowest local unemployment quintile. I demonstrated that adjusting for the local price level and controlling for worker fixed effects does not do away with these differences and that a difference in wage effects in the order of magnitude of 11% remains.

Controlling for local prices, until the age of 55 there is a difference in lifetime earnings of about 150,000 2010 Euros (adjusted to Bonn prices). Residualizing lifetime earnings with rich observables, a gap of about 50,000 Euros remains.

I conclude that it does seem to matter considerably where a worker starts his working life. The natural way forward for future research would be to establish the mechanism that drives my findings. I view this as a fruitful ground for future research.

References

- T. Allen and C. Arkolakis. Trade and the Topography of the Spatial Economy *. *The Quarterly Journal of Economics*, 129(3):1085–1140, 2014.
- J. G. Altonji, L. B. Kahn, and J. D. Speer. Cashier or consultant? entry labor market conditions, field of study, and career success. *Journal of Labor Economics*, 34(S1):S361–S401, 2016.
- M. Amior and A. Manning. The persistence of local joblessness. *American Economic Review*, 108(7):1942–70, July 2018. doi: 10.1257/aer.20160575. URL <https://www.aeaweb.org/articles?id=10.1257/aer.20160575>.
- N. Baum-Snow and R. Pavan. Understanding the city size wage gap. *The Review of economic studies*, 79(1):88–127, 2012.
- P. Bergman, R. Chetty, S. DeLuca, N. Hendren, L. F. Katz, and C. Palmer. Creating moves to opportunity: Experimental evidence on barriers to neighborhood choice. *American Economic Review*, 114(5):1281–1337, May 2024. doi: 10.1257/aer.20200407. URL <https://www.aeaweb.org/articles?id=10.1257/aer.20200407>.

- A. Bilal. The geography of unemployment*. *The Quarterly Journal of Economics*, 138(3): 1507–1576, 03 2023. ISSN 0033-5533. doi: 10.1093/qje/qjad010. URL <https://doi.org/10.1093/qje/qjad010>.
- O. J. Blanchard, L. F. Katz, R. E. Hall, and B. Eichengreen. Regional evolutions. *Brookings papers on economic activity*, 1992(1):1–75, 1992.
- M. J. Böhm, H.-M. v. Gaudecker, and F. Schran. Occupation growth, skill prices, and wage inequality. *Journal of Labor Economics*, 0:null, 2023 forthcoming. doi: 10.1086/722084. URL <https://doi.org/10.1086/722084>.
- G. Bryan and M. Morten. The aggregate productivity effects of internal migration: Evidence from indonesia. *Journal of Political Economy*, 127(5):2229–2268, 2019.
- T. Bönke, G. Corneo, and H. Lüthen. Lifetime earnings inequality in germany. *Journal of Labor Economics*, 33(1):171 – 208, 2015. URL <https://EconPapers.repec.org/RePEc:ucp:jlabec:doi:10.1086/677559>.
- D. Card, J. Heining, and P. Kline. Workplace Heterogeneity and the Rise of West German Wage Inequality*. *The Quarterly Journal of Economics*, 128(3):967–1015, 05 2013.
- R. Chetty and N. Hendren. The impacts of neighborhoods on intergenerational mobility i: Childhood exposure effects*. *The Quarterly Journal of Economics*, 133(3):1107–1162, 02 2018a. ISSN 0033-5533. doi: 10.1093/qje/qjy007. URL <https://doi.org/10.1093/qje/qjy007>.
- R. Chetty and N. Hendren. The impacts of neighborhoods on intergenerational mobility ii: County-level estimates*. *The Quarterly Journal of Economics*, 133(3):1163–1228, 02 2018b. ISSN 0033-5533. doi: 10.1093/qje/qjy006. URL <https://doi.org/10.1093/qje/qjy006>.
- W. Dauth, S. Findeisen, E. Moretti, and J. Suedekum. Matching in cities. *Journal of the European Economic Association*, 20(4):1478–1521, 2022.
- C. Dustmann, J. Ludsteck, and U. Schönberg. Revisiting the german wage structure. *The Quarterly Journal of Economics*, 124(2):843–881, 2009. ISSN 00335533, 15314650. URL <http://www.jstor.org/stable/40506245>.
- B. Fitzenberger, A. Osikominu, and R. Völter. Imputation rules to improve the education variable in the IAB employment subsample. *FDZ-Methodenreport*, 03/2005, 2005.

- H. Gartner. The imputation of wages above the contribution limit with the German IAB employment sample. *FDZ-Methodenreport*, 02/2005, 2005.
- C. Gaubert. Firm sorting and agglomeration. *American Economic Review*, 108(11):3117–3153, 2018.
- E. L. Glaeser and D. C. Mare. Cities and skills. *Journal of labor economics*, 19(2):316–342, 2001.
- E. D. Gould. Cities, workers, and wages: A structural analysis of the urban wage premium. *The Review of Economic Studies*, 74(2):477–506, 2007.
- J. R. Harris and M. P. Todaro. Migration, unemployment and development: a two-sector analysis. *The American economic review*, pages 126–142, 1970.
- R. Hornbeck and E. Moretti. Estimating who benefits from productivity growth: Local and distant effects of city productivity growth on wages, rents, and inequality. *The Review of Economics and Statistics*, pages 1–49, 06 2022. ISSN 0034-6535. doi: 10.1162/rest_a_01208. URL https://doi.org/10.1162/rest_a_01208.
- K. Huttunen, J. Møen, and K. G. Salvanes. Job loss and regional mobility. *Journal of Labor Economics*, 36(2):479 – 509, 2018. URL <https://EconPapers.repec.org/RePEc:ucp:jlabecon:doi:10.1086/694587>.
- L. S. Jacobson, R. J. LaLonde, and D. G. Sullivan. Earnings losses of displaced workers. *The American economic review*, pages 685–709, 1993.
- L. B. Kahn. The long-term labor market consequences of graduating from college in a bad economy. *Labour Economics*, 17(2):303–316, 2010a. ISSN 0927-5371. doi: <https://doi.org/10.1016/j.labeco.2009.09.002>. URL <https://www.sciencedirect.com/science/article/pii/S0927537109001018>.
- L. B. Kahn. The long-term labor market consequences of graduating from college in a bad economy. *Labour economics*, 17(2):303–316, 2010b.
- P. Kline and E. Moretti. Place based policies with unemployment. *American Economic Review*, 103(3):238–43, 2013.

- R. Kosfeld, H.-F. Eckey, and M. Schüßler. Ökonometrische Messung regionaler Preisniveaus auf der Basis örtlich beschränkter Erhebungen. RatSWD Research Notes 33, German Data Forum (RatSWD), 2009. URL <https://ideas.repec.org/p/rsw/rswrns/rswrns33.html>.
- M. Kuhn, I. Manovskii, and X. Qiu. The geography of job creation and job destruction. Working Paper 29399, National Bureau of Economic Research, October 2021. URL <http://www.nber.org/papers/w29399>.
- M. Lachowska, A. Mas, and S. A. Woodbury. Sources of displaced workers’ long-term earnings losses. *American Economic Review*, 110(10):3231–66, 2020.
- B. L. Levy, T. Mouw, and A. D. Perez. Why did people move during the great recession? the role of economics in migration decisions. *RSF: The Russell Sage Foundation Journal of the Social Sciences*, 3(3):100–125, 2017. ISSN 23778253, 23778261. URL <http://www.jstor.org/stable/10.7758/rsf.2017.3.3.05>.
- I. Marinescu and R. Rathelot. Mismatch unemployment and the geography of job search. *American Economic Journal: Macroeconomics*, 10(3):42–70, 2018.
- E. Moretti. Local labor markets. volume 4B, chapter 14, pages 1237–1313. Elsevier, 1 edition, 2011.
- M. Morten and J. Oliveira. The effects of roads on trade and migration: Evidence from a planned capital city. *NBER Working Paper*, 22158:43, 2016.
- E. Nakamura, J. Sigurdsson, and J. Steinsson. The gift of moving: Intergenerational consequences of a mobility shock. Technical report, National Bureau of Economic Research, 2016.
- OECD. How persistent are regional disparities in employment? the role of geographic mobility. *OECD Employment Outlook 2005 Edition*, 2005.
- P. Oreopoulos, T. Von Wachter, and A. Heisz. The short-and long-term career effects of graduating in a recession. *American Economic Journal: Applied Economics*, 4(1):1–29, 2012a.
- P. Oreopoulos, T. Von Wachter, and A. Heisz. The short-and long-term career effects of graduating in a recession. *American Economic Journal: Applied Economics*, 4(1):1–29, 2012b.

- S. J. Redding and E. Rossi-Hansberg. Quantitative spatial economics. *Annual Review of Economics*, 9(1):21–58, 2017.
- J. Roback. Wages, rents, and the quality of life. *Journal of Political Economy*, 90(6):1257–1278, 1982.
- A. Şahin, J. Song, G. Topa, and G. L. Violante. Mismatch unemployment. *American Economic Review*, 104(11):3529–64, 2014.
- L. A. Sjaastad. The costs and returns of human migration. *Journal of Political Economy*, 70(5):80–93, 1962.
- A. Spitz-Oener. Human capital, job tasks and technology in east germany after reunification. *National Institute Economic Review*, 201:97–106, 2007.
- R. H. Topel. Local labor markets. *Journal of Political economy*, 94(3, Part 2):S111–S143, 1986.
- H. Uhlig. Regional labor markets, network externalities and migration: The case of german reunification. *American Economic Review*, 96(2):383–387, 2006.
- H. Uhlig. The slow decline of east germany. *Journal of Comparative Economics*, 36(4):517–541, 2008.
- W. van den Berge. Bad start, bad match? the early career effects of graduating in a recession for vocational and academic graduates. *Labour Economics*, 53:75–96, 2018. ISSN 0927-5371. doi: <https://doi.org/10.1016/j.labeco.2018.05.011>. URL <https://www.sciencedirect.com/science/article/pii/S0927537118300563>. European Association of Labour Economists 29th annual conference, St.Gallen, Switzerland, 21-23 September 2017.
- A. Young. Inequality, the Urban-Rural Gap, and Migration*. *The Quarterly Journal of Economics*, 128(4):1727–1785, 09 2013.

A Data

A.1 Wage Imputation

Employers in Germany report employers' wages only up to the contribution limit to social security accounts. The wage variable is therefore right-censored. To circumvent this problem I follow the procedure recommended by Gartner (2005) and impute wages in the following way: For every year I estimate the model $\log \text{wage}_{it} = x'_{it}\beta + e_{it}$, where x_{it} is a vector containing *imputed education level*, *labor market status* (full-time employed, part-time employed, marginally employed or apprentice), *the fraction of spells where an individuals wage is censored in other years* and the *mean reported wage from other years* for individual i at time t using TOBIT estimation.


In order not to underestimate the variance of wages above the censoring level I then use the predicted wage $x'_{it}\hat{\beta}_{\text{tobit}}$ and the estimate of the standard deviation $\hat{\sigma}_{\text{tobit}}$ to draw the imputed wage from the truncated normal distribution with mean $x'_{it}\hat{\beta}_{\text{tobit}}$ and standard deviation $\hat{\sigma}_{\text{tobit}}$ truncated to the left at the censoring level.

As the results presented below are mainly based on samples that are restricted to men I impute wages for men and women separately to allow for different wage variances for men and women.

A.2 Education Imputation

The BeH data is taken from reports that employers make about their employees. While this official nature of the data source lends it a great deal of its credibility there is room for interpretation as to what exactly should be reported that may be handled differently by different employers. The education variable in particular suffers from inconsistencies such as reporting lower levels of education after higher levels have been reported for the same individual in previous years. One reason for this may be that employers tend to report only the minimum education level required for a particular job. To deal with these inconsistencies I employ the Education Imputation Procedure 1 proposed by Fitzenberger et al. (2005) and extrapolate reported education such that I am left with 6 consistent education variables: *1 No Degree*, *2 High School*, *3 Vocational Training*, *4 High School and Vocational Training*, *5 Technical College* and *6 University*.

Figure 19: SIAB vs BA unemployment rates (1985)



Notes: BA unemployment rates are on the x-axis and unemployment rates estimated from the SIAB are on the y-axis. Each blue dot represents a commuting zone.

A.3 Spell Error Imputation


In some cases, there are gaps in the reported unemployment history of an individual. Because employers report information about employees in continuous ongoing employment once yearly I follow Böhm et al. (2023 forthcoming) and impute such gaps if there is a one-year gap between two spells with the same employer by filling it with the last valid spell. Other gaps are interpreted as spells during which an individual is not part of the labor force.

B Unemployment Benefit Imputation

B.1 Unemployment Benefit Imputation


The German unemployment benefit system works as follows: Workers who were employed for at least 24 months with the last 5 years before entering unemployment receive 60% of their after tax wages up to the maximum social security contribution level for up to two years. This means that the benefit depends on the tax class and other particularities of the German tax system that are impossible to reconstruct for each individual worker

Figure 20: Regional CPI vs INKAR rent-index



Notes: The INKAR rent-index is provided by the BBSR on the INKAR website. It is based on rents of flats of 40-100 m^2 in average or good neighborhoods advertised on a variety of platforms. It is meant to reflect offers that people who look for a flat online in a given county are typically faced with. It is reported in 1-Euro bins.


Figure 21: Regional CPI vs building land prices



Notes: Building land prices are average prices per square-meter of sold building land within a county. The data comes from Regionaldatenbank Deutschland - the regional data base of the federal and state level statistical offices of Germany.

because they in turn depend on unobserved characteristics such as e.g. marital status and the number of children. I therefore impute these benefits in the following way: Each year, I use the sample of workers for whom I can observe unemployment benefits and run a TOBIT regression, regressing the observed benefit on the last observed wage. I then use the fitted TOBIT model to predict the benefit for unemployed workers with unobservable

Figure 22: Tobit Prediction of Unemployment Benefit 2000



notes: The blue dots plot actually observed benefits on the y-axis against the respective worker's last wage on the x-axis. The red line shows the fitted TOBIT model that I use to impute benefits for unemployed workers without observable benefits.


unemployment benefit. Benefits paid under this scheme are called *Arbeitslosengeld (ALG)*. Until 2005, as a fallback, workers who were employed for at least 12 months within the past 30 months before entering unemployment received lower benefits called *Arbeitslosenhilfe (ALH)*. ALH was based on a similar system as ALG and I impute it in the same way for workers without observable benefits who are eligible for ALH but not for ALG. Figure 22 shows an illustration of the TOBIT predictions.

Until 2005 Long-term unemployment benefits were part of social benefits regulated and paid by local governments. I am unaware of a feasible imputation strategy for these payments. Since 2005 there are nation-wide fixed long-term unemployment benefits. I simply use the official benefit figures for long-term unemployment workers reported through the ASU.

Adding these imputations to the workers observable income results in slightly higher accumulated income than shown in Figure 17. However, the difference between the high and low unemployment ranks doesn't change much. Results for the 1960-1962 cohort are shown in Figure 23.

Unemployment insurance is financed using social security contributions. Employed workers pay a share of their wage to social security. These contribution rates for the years covered by the SIAB are given in Table 3.

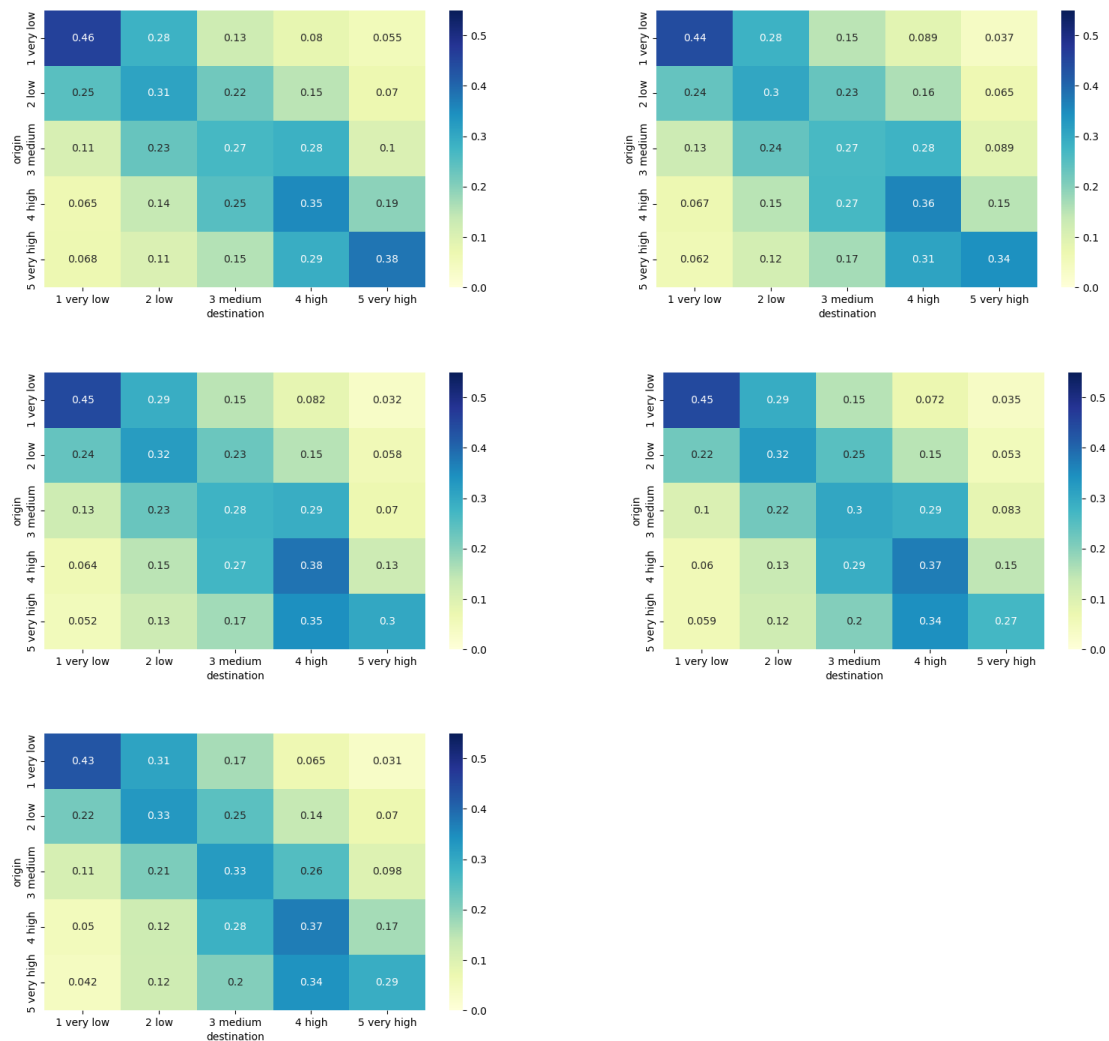
Figure 23: Average accumulated earnings and benefits by rank



notes: Accumulated life-time earnings by unemployment rank of the commuting zone workers started their first job in for the 1960-1962 cohort. Earnings are accumulated up to the age of 55 for in the left graph and up to the age of 44 on the right. The sample has been restricted to workers who are observable for at least 90% of the time between the age of 25 and the last observable age in each cohort. The price level has been adjusted to the level of 2010 using the CPI from the Federal Statistics Office. The darker bars have additionally been adjusted to the price level of Bonn using the regional CPI from the BBSR. The numbers include imputed unemployment benefits.


C Additional Figures

Figure 24: Fractions of moves by unemployment rank of origin- and destination commuting zones



Notes: Each cell in the matrices above represents the fraction of moves away from commuting zones with an unemployment rank indicated by the row to a commuting zone with an unemployment rank indicated by the column. E.g. the second cell in the first row of the left matrix indicates, that 29% of moves away from zones with rank *1 very low* went to zones with rank *2 low*. Workers are ranked based on the unemployment quintile of the commuting zone where they initially enter the labor market. Cohorts are arranged as follows: 1960-1952 - top left, 1966-1967 - top right, 1968-1970 - middle left, 1974-1976 - middle right, 1977-1979 bottom.

Figure 25: Fractions of moves over at least 200 km by unemployment rank of origin- and destination commuting zones




Notes: Each cell in the above matrices represents the fraction of moves away from commuting zones with an unemployment rank indicated by the row to a commuting zone with an unemployment rank indicated by the column. Only moves over a distance of at least 200 km are considered. Distance between two commuting zones is measured as the minimum distance between the borders of these commuting zones. E.g. the second cell in the first row of the left matrix indicates, that 18% of moves away from zones with rank *1 very low* went to zones with rank *2 low*. Workers are ranked based on the unemployment quintile of the commuting zone where they initially enter the labor market. Cohorts are arranged as follows: 1960-1952 - top left, 1966-1967 - top right, 1968-1970 - middle left, 1974-1976 - middle right, 1977-1979 bottom.

Year	Contribution Rate
1975	0.020
1976	0.020
1977	0.020
1978	0.020
1979	0.020
1980	0.030
1981	0.030
1982	0.030
1983	0.030
1984	0.030
1985	0.041
1986	0.041
1987	0.041
1988	0.041
1989	0.041
1990	0.043
1991	0.043
1992	0.043
1993	0.043
1994	0.043
1995	0.065
1996	0.065
1997	0.065
1998	0.065
1999	0.065
2000	0.065
2001	0.065
2002	0.065
2003	0.065
2004	0.065
2005	0.065
2006	0.065
2007	0.042
2008	0.038
2009	0.033
2010	0.028
2011	0.028
2012	0.030
2013	0.030
2014	0.030
2015	0.030
2016	0.030
2017	0.030


Table 3: Social Security Contribution Rates

Figure 26: Fractions of moves by unemployment rank of origin- and destination commuting zones - unemployment workers



Notes: Each cell in the matrices above represents the fraction of moves away from commuting zones with an unemployment rank indicated by the row to a commuting zone with an unemployment rank indicated by the column. Only moves that happen immediately after an unemployment spell are considered. E.g. the second cell in the first row of the left matrix indicates, that 29% of moves away from zones with rank *1 very low* went to zones with rank *2 low*. Workers are ranked based on the unemployment quintile of the commuting zone where they initially enter the labor market. The matrix on the left shows the results for the 1963-1965 cohort, and the matrix on the right shows the results for the 1971-1973 cohort.

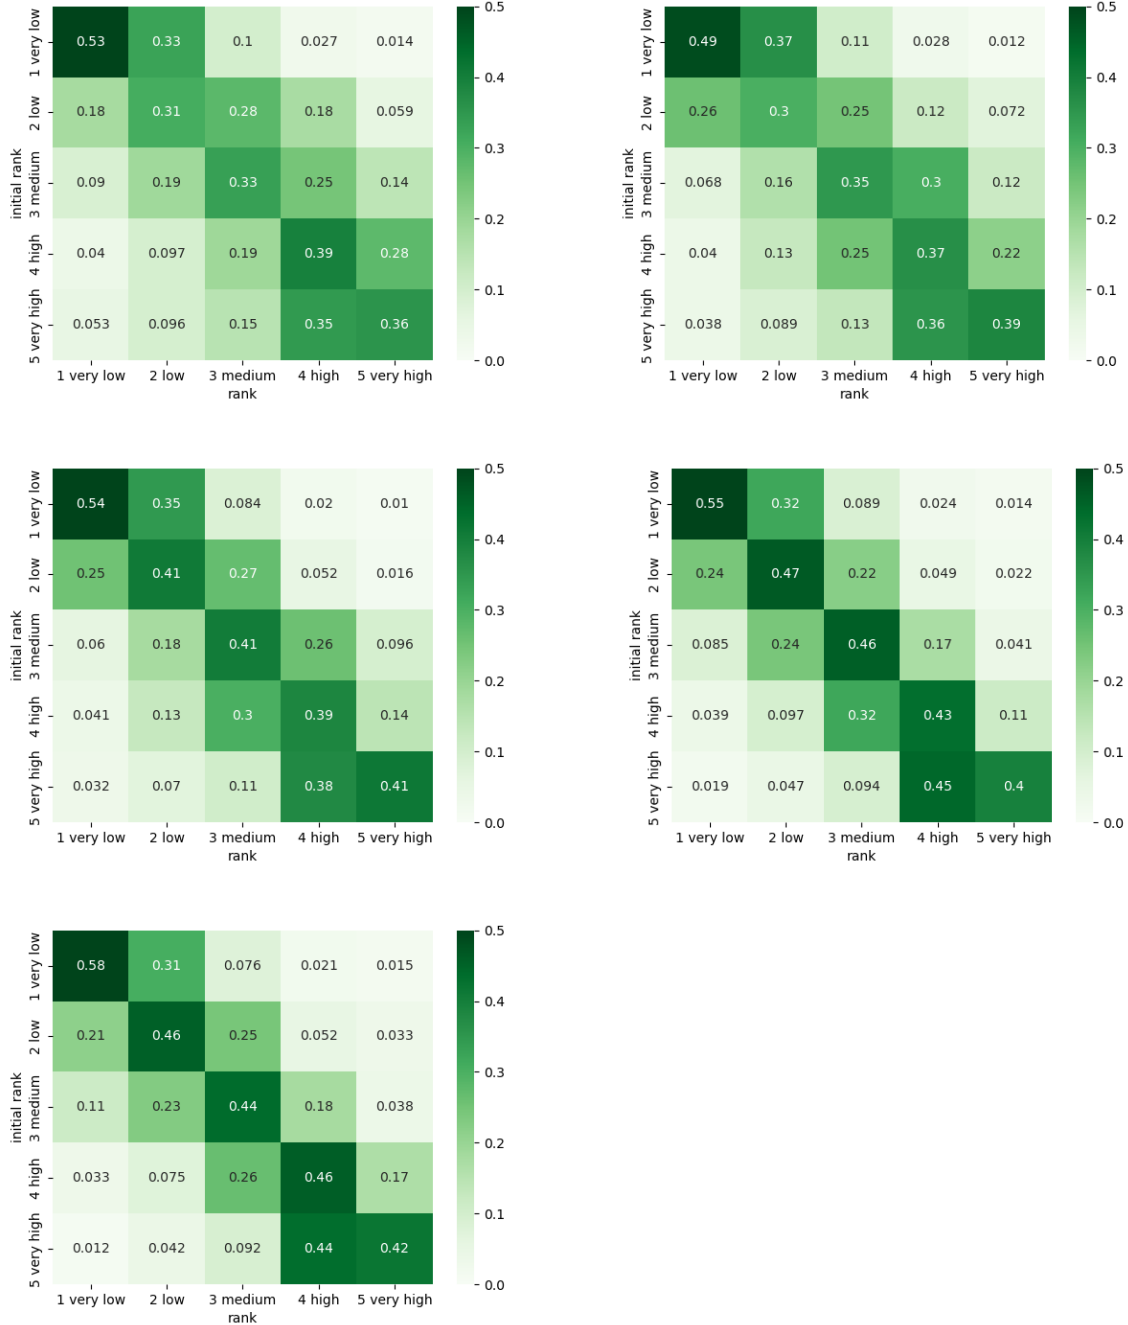
Figure 27: Fractions of moves by unemployment rank of origin- and destination commuting zones - unemployment workers



Notes: Fractions of moves of workers who were employed before moving that occur over the whole life by the rank of origin and rank of destination. The matrix on the left shows the results for the 1963-1965 cohort, the matrix on the right shows the results for the 1971-1973 cohort.


Each cell in the matrices above represents the fraction of moves away from commuting zones with an unemployment rank indicated by the row to a commuting zone with an unemployment rank indicated by the column. Only moves that happen immediately after a spell under employment are considered. E.g. the second cell in the first row of the left matrix indicates, that 29% of moves away from zones with rank *1 very low* went to zones with rank *2 low*. Workers are ranked based on the unemployment quintile of the commuting zone where they initially enter the labor market. The matrix on the left shows the results for the 1963-1965 cohort, and the matrix on the right shows the results for the 1971-1973 cohort.

Figure 28: Fractions of time spent in the five unemployment ranks



Notes: Fractions of time spent in commuting zones of each of the five ranks by the rank of the initial commuting zone of a worker where he first entered the labor market. The initial ranks are on the y-axis. For example, the second cell in the first row of the left graph states that workers who start in a commuting zone with rank *1 very low* spend 32% of their time in commuting zones with rank *2 low* on average. Workers are ranked based on the unemployment quintile of the commuting zone where they initially enter the labor market. Cohorts are arranged as follows: 1960-1952 - top left, 1966-1967 - top right, 1968-1970 - middle left, 1974-1976 - middle right, 1977-1979 bottom.

Figure 29: Fractions of time spent in the five unemployment ranks, apprentices only



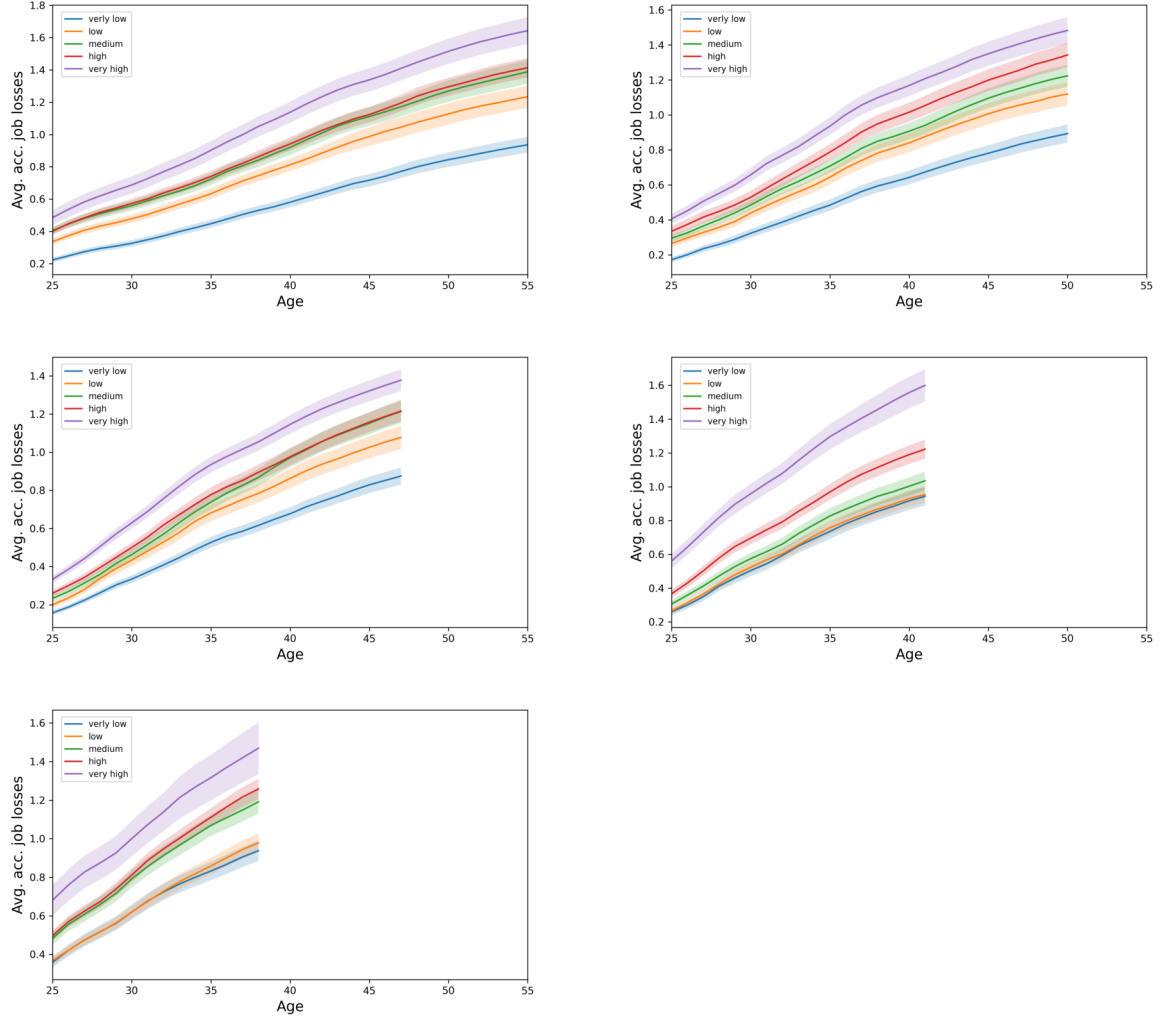
Notes: Fractions of time spent in commuting zones of each of the five ranks by rank of the initial commuting zone of a worker entered the labor market in. The sample is restricted to workers who enter the labor market as apprentices before the age of 20. The initial ranks are on the y-axis. For example, the second cell in the first row of the left graph states that workers who start in a commuting zone with rank *1 very low* spend 32% of their time in commuting zones with rank *2 low* on average. Workers are ranked based on the unemployment quintile of the commuting zone where they initially enter the labor market. The matrix on the left shows the results for the 1963-1965 cohort, and the matrix on the right shows the results for the 1971-1973 cohort.

Table 4: Fractions of workers who start a first job within a given distance from the commuting zone they were apprenticed in

	stayers	neighboring zone	100 km	200km
cohort				
1960-1962	0.886477	0.944741	0.980467	0.989316
1963-1965	0.871097	0.938801	0.978383	0.987640
1966-1967	0.863239	0.932684	0.975736	0.986004
1968-1970	0.866690	0.940537	0.980296	0.988598
1971-1973	0.857920	0.936378	0.977938	0.987726
1974-1976	0.851978	0.932327	0.974109	0.986430
1977-1979	0.845169	0.930024	0.976455	0.987192

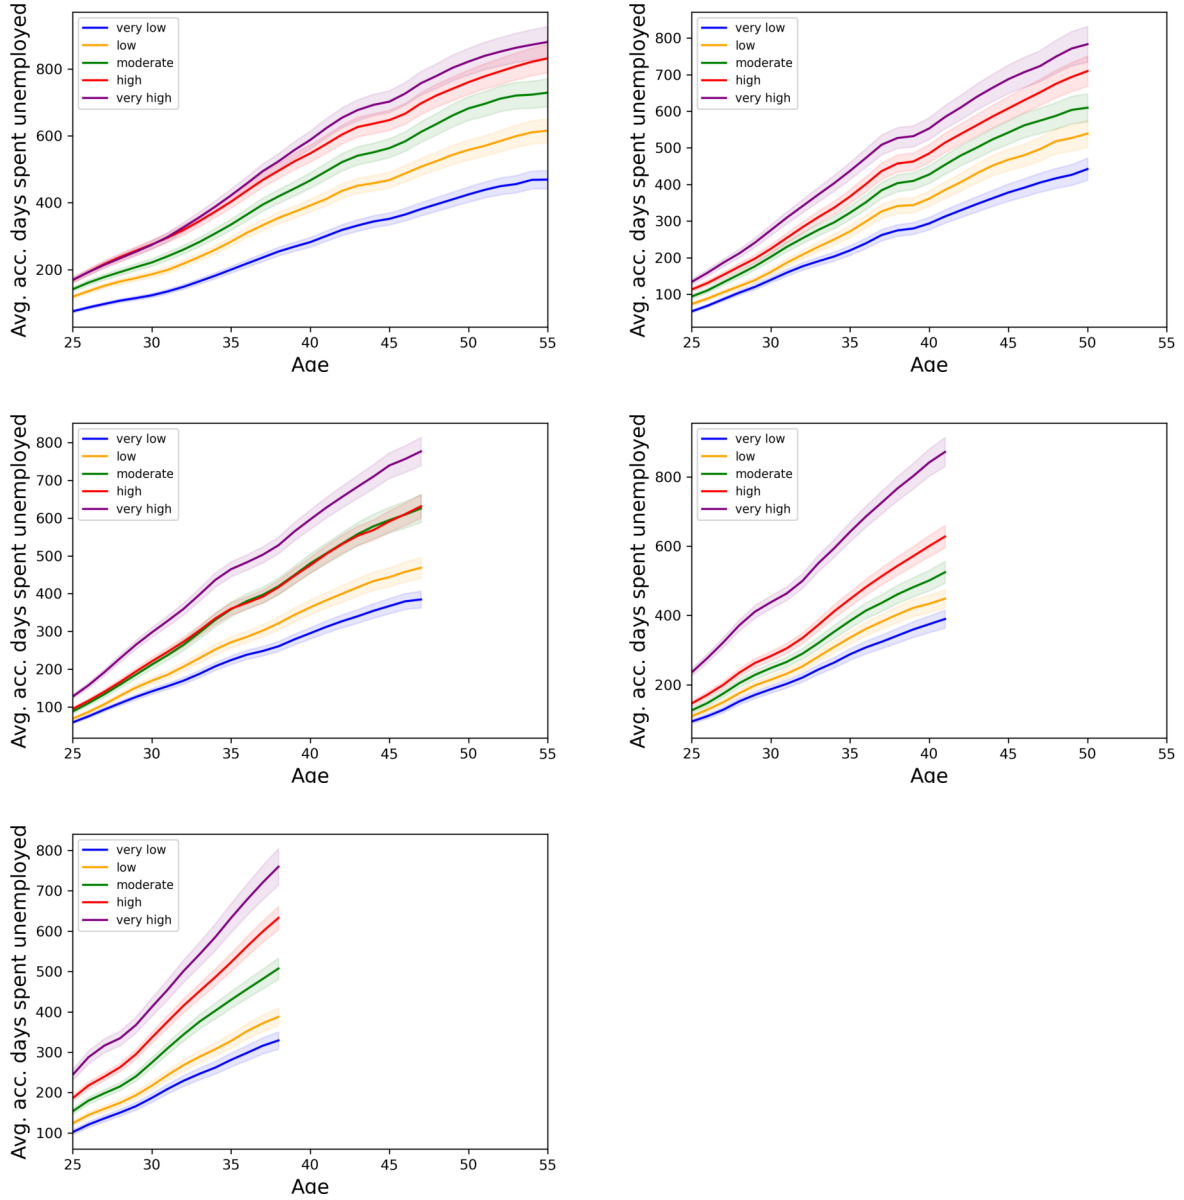
Notes: Sample consists of workers who entered the labor market as apprentices before the age of 20. The first column shows the fraction of such workers who start their first non-apprentice-job in the same commuting zone that they were apprenticed in. The second column shows the fraction of workers that start their first non-apprentice job in the same or a neighboring commuting zone. The third and fourth columns show the fractions of workers that stay within a radius of 100 km and 200 km from where they were apprenticed when they start their first non-apprentice-job respectively. Distance between two commuting zones is measured as the minimum distance between the borders of these commuting zones.

Figure 30: Average accumulated job losses



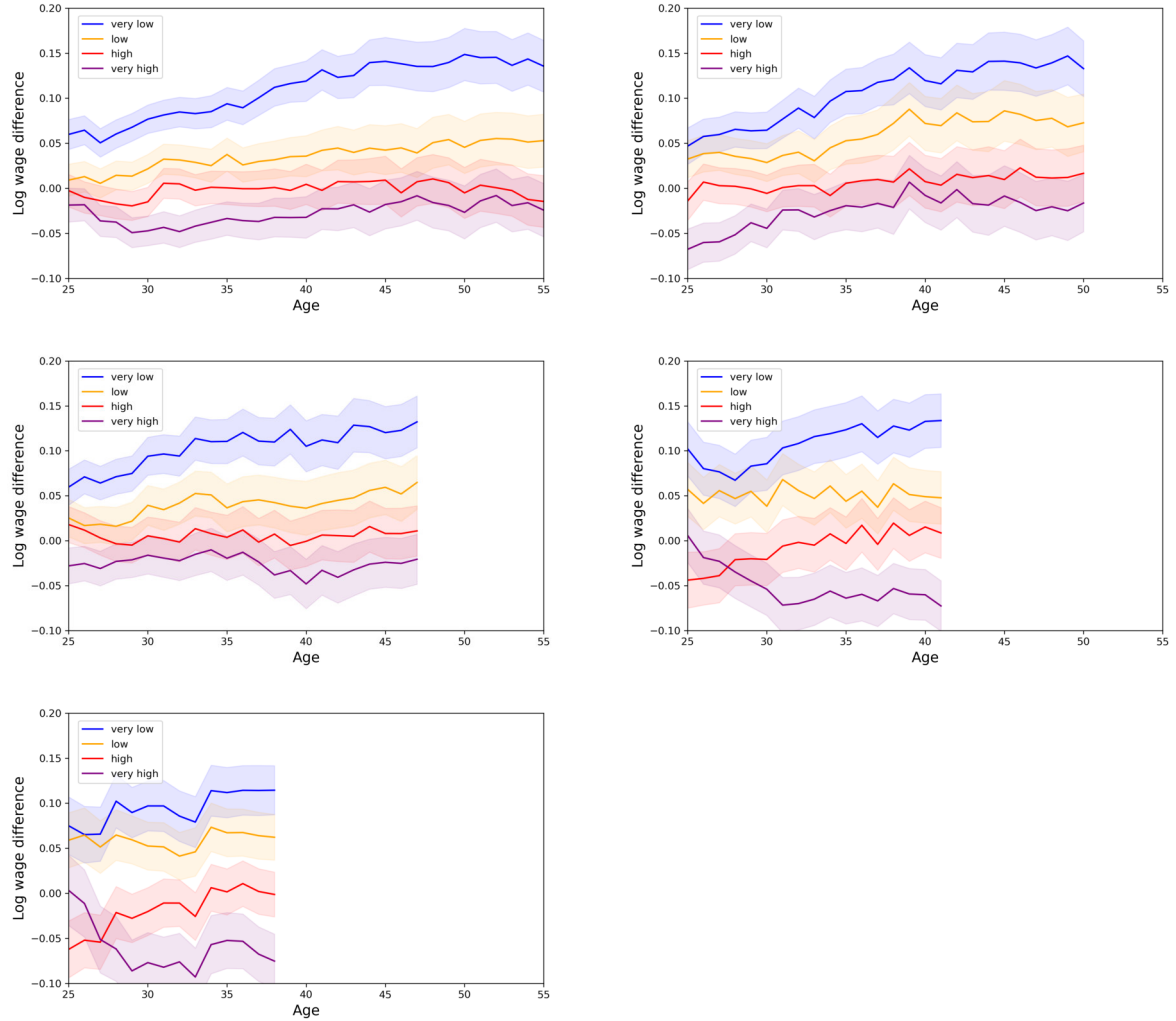
notes: The average number of job losses workers from each rank have had at each age. Workers are ranked based on the unemployment quintile of the commuting zone where they initially enter the labor market. Workers who have no observations after an age younger than 55 are assumed to lose no further jobs. Job losses are defined as transitions from an employed spell to an unemployed spell with a gap of fewer than 31 days between them. Cohorts are arranged as follows: 1960-1952 - top left, 1966-1967 - top right, 1968-1970 - middle left, 1974-1976 - middle right, 1977-1979 bottom.

Figure 31: Average accumulated time spent in unemployment



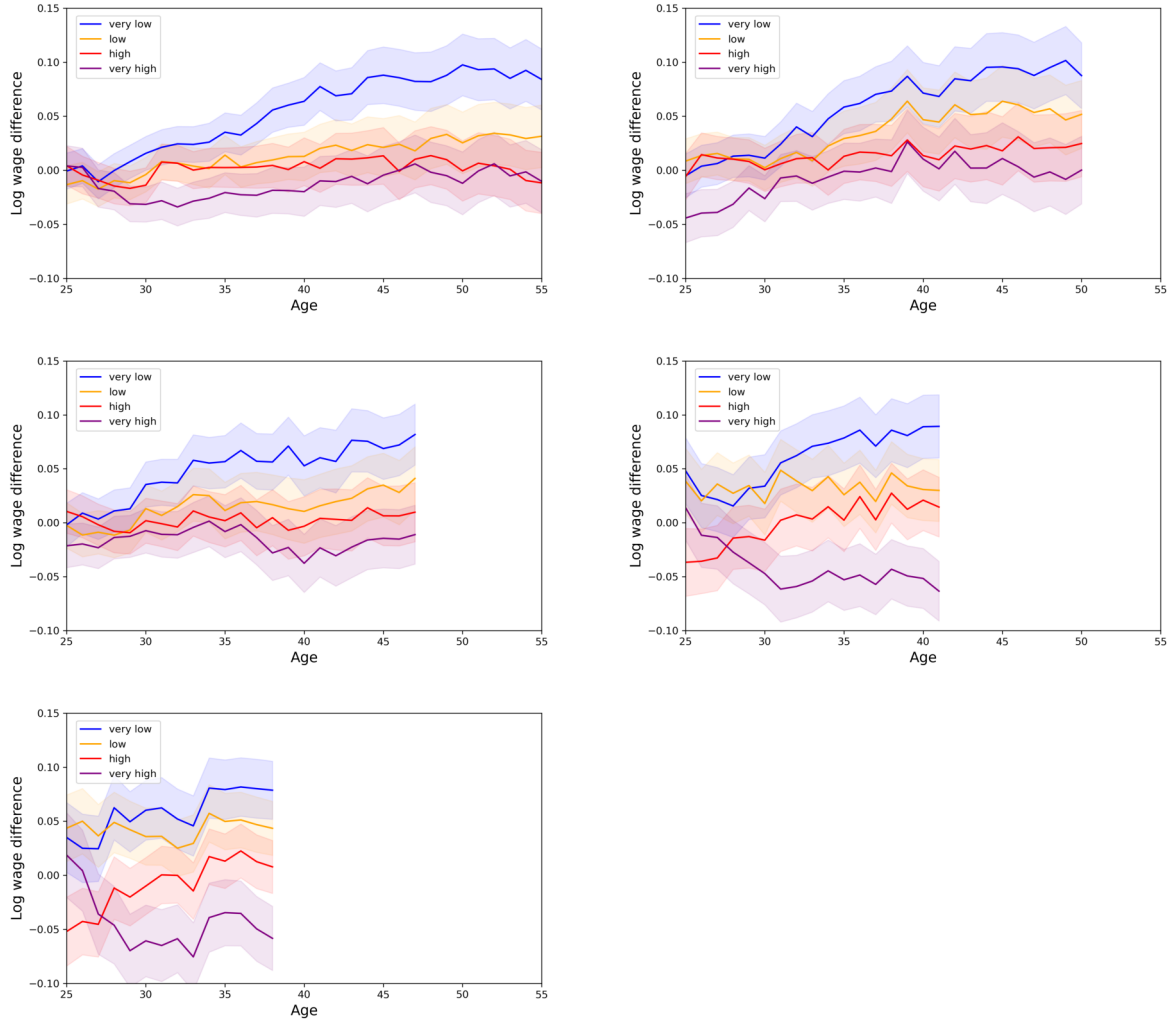
notes: The average number of days that workers have accumulated up to each age. Workers who have no observations after an age younger than 55 are assumed to accumulate no further time in unemployment. Workers are ranked based on the unemployment quintile of the commuting zone where they initially enter the labor market. The confidence intervals displayed are standard 95% confidence intervals around the mean at each age. Cohorts are arranged as follows: 1960-1952 - top left, 1966-1967 - top right, 1968-1970 - middle left, 1974-1976 - middle right, 1977-1979 bottom.

Figure 32: Average Log Wage lifecycle Profiles relative to rank β moderate



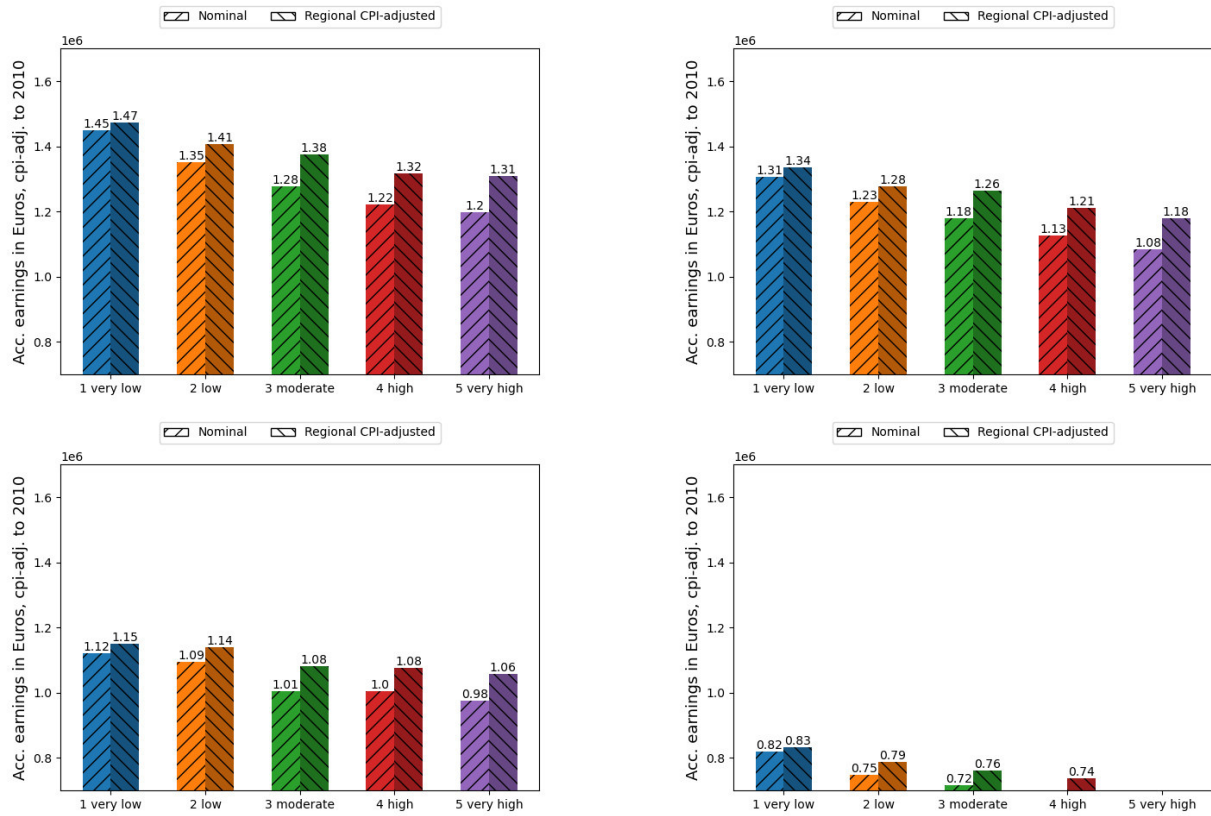
notes: Average log-wage profiles of workers by unemployment rank. Workers are ranked based on the unemployment quintile of the commuting zone where they initially enter the labor market. From each of the shown log wage lifecycle profiles the profile of the β moderate rank has been subtracted. Confidence intervals are computed using two-sided t-tests to assess whether the resulting difference is statistically significant from 0 at the 95% level, at each age. Cohorts are arranged as follows: 1960-1952 - top left, 1966-1967 - top right, 1968-1970 - middle left, 1974-1976 - middle right, 1977-1979 bottom.

Figure 33: Regional CPI-adjusted Log Wage lifecycle Profiles relative to rank 3 *moderate*



Notes: Average regional CPI-adjusted log-wage profiles of workers by unemployment rank. The CPI adjustment was calculated using the regional CPI from the BBSR. Workers are ranked based on the unemployment quintile of the commuting zone where they initially enter the labor market. From each of the shown log wage lifecycle profiles the profile of the 3 *moderate* rank has been subtracted. Confidence intervals are computed using two-sided t-tests to assess whether the resulting difference is statistically significant from 0 at the 95% level, at each age. Cohorts are arranged as follows: 1960-1952 - top left, 1966-1967 - top right, 1968-1970 - middle left, 1974-1976 - middle right, 1977-1979 bottom.

Figure 34: Average accumulated earnings



notes: Average accumulated earnings up to the age of 40. The sample includes workers with gaps in their observable histories that comprise no more than 10% of the time between the ages of 25 and 55. Cohorts are arranged as follows: 1963-1965 - top left, 1966-1967 - top right, 1968-1970 - middle left, 1974-1976 - middle right.