Regional Occupations, Local Rents and Worker Mobility*

Job Market Paper

[Click here for the latest version]

Johannes Weber†

March 27, 2026

Abstract

Despite large regional wage differences, worker mobility is low in Germany and other developed economies. This paper argues that low mobility arises from the interaction between locally accumulated, region-specific human capital and high regional living costs. Using German social security records, I document that workers in high-wage regions have better careers but not higher entry wages, despite higher living costs. I show that large parts of variation in regional wage levels and growth can be accounted for by different regional occupation compositions. I develop a general equilibrium lifecycle model with two regions, local labor and rent markets, and borrowing constraints in which human capital accumulation reflects the skills most in use in a worker's local labor market. I find that in the calibrated model, young workers are kept from moving to the region with larger wage growth by high rent prices. Older workers with more experience have low incentives to move because, on average, their human capital would be mismatched if they did. I show that constructing more housing in the high-wage region leads to more mobility and is welfare-improving for workers in both regions.

Keywords: Worker Mobility, Labor Supply, Lifecycle, Local Labor Markets

JEL: J61, R13, R23, R31

*I thank Moritz Kuhn and Hans-Martin v. Gaudecker for invaluable advice, guidance, and encouragement throughout this project. I also thank Christian Bayer, Michael Peters, Joseph Altonji, Kerstin Holzheu and Brigitte Hochmuth as well as participants of various seminars for their helpful feedback and comments. Support by the Deutsche Forschungsgemeinschaft (DFG) through CRC TR 224 (Project C01) is gratefully acknowledged. The usual disclaimer applies.

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

1 Introduction

There are large differences between regional wage levels in developed economies. Yet, worker mobility remains low. This raises the question of why workers do not move to high-wage regions in large numbers.

In the literature, the most prevalent explanations are high moving costs, sorting of workers into regions by their productivity and high regional living costs. Two important considerations have evaded the spotlight of the discussion, however. Firstly, local labor markets differ not only in productivity and wage levels, but also in the type of work workers perform. Secondly, moving decisions that take into account career opportunities depend crucially on the stage of a worker's career and therefore on age.

In this study, I combine these two points and argue that worker mobility is low because, over their careers, workers accumulate human capital that is tailored to their local labor market and is therefore not easily transferable to other regions. First, I empirically document large differences in real wage growth between regions and show that large parts of these differences can be accounted for by regional occupation compositions. Building on these findings, I develop a general equilibrium model in which the skills workers accumulate through experience reflect their labor market. In the calibrated model, young workers are discouraged from moving to the high-wage region primarily because of high rents, while older workers have little incentive to move since it would render the skills they accumulated less productive. Together, housing market frictions and region-specific skill accumulation are sufficient to lock workers into their initial regions, with only minor direct moving costs. In light of this interaction between local housing and labor markets, I proceed to conduct policy experiments that show that housing construction in high-wage regions can increase mobility, GDP, and welfare, even in a region where no new houses are constructed.

In the first part of the paper, I leverage high-quality administrative data from Germany and document three key facts. Firstly, I show that worker mobility is indeed low in Germany. At the age of 55, about half of all workers have spent their entire career in the very same commuting zone and about 70% have never moved further than to a commuting zone neighboring their very first one.

Secondly, over the life-cycle, the difference between expensive high- and cheap low-wage regions is not a difference in wage levels but in wage growth. Making use of the large correlation between the wage and the rent level, I split the sample into 3 using rent terciles. Taking regional prices into account, workers who begin their career in the lowest rent tercile actually earn about 6% higher real wages than their peers who start in the highest tercile. However, because of greater wage growth, workers from the highest tercile earn more on average in the middle part and especially at the end of their careers. At age 55, workers in the highest tercile earn about 10% higher real wages than their counterparts in the lowest tercile. I show that these growth patterns persist even when controlling for individual and age by occupation fixed effects in a dynamic diff-in-diff regression, casting doubt over the view that pure worker selection causes the wage differences.

Thirdly, occupations are spatially concentrated. Even after taking different sizes of regional workforces into account, each occupation in the data exhibits strictly convex Lorenz curves across commuting zones. The resulting different regional occupational compositions explain a sizable part of regional variation in wage levels and wage growth. Depending on the decade, the composition of occupations explains about 60-80% of variation in the wage level and about 45-55% of variation in wage growth over 25 years at the county level.

In the second part of the paper, I propose a two-region general-equilibrium life-cycle model with local labor and rent markets to explain low worker mobility in light of my empirical findings. Workers draw utility from a final consumption good and housing, and can save but not borrow. Workers can choose to relocate between the regions at a resource and a utility moving cost every year. Human capital is two-dimensional, consisting of two separate skills, and is accumulated stochastically through experience. In line with different occupation compositions in each region, one of the two skills is more productive and can be learned faster, and in each region, a separate intermediate output good is produced. One of the regions has faster wage growth on average because its dominant skill can be accumulated faster.

I calibrate the model to the German data and investigate which forces lead to low worker mobility. I find that workers born in the poorer region with less wage growth face a dilemma. In principle, they would like to move to the region with high wage growth to maximize their expected career outcomes. The strongest force that keeps them from moving there is the high rent price in that region. By contrast, older workers well into their careers have accumulated wealth and can afford high rents more easily, but are kept back by their human capital. On average, it reflects their local labor market and is less applicable in the other region. Taking the dynamic interaction between locally adapted human capital and regional housing costs into account, the implied moving costs—measured as compensating differentials between moving and staying—are much lower than typically estimated in the literature, averaging about €39,000. They are especially small for young workers—only about €8,000 for 20-year-olds—and increase with age as human capital becomes more region-specific. I show that estimating a Mincer regression, which assumes full portability of experience across regions, strongly overestimates potential wage gains of moving to the high wage region.

Finally, I use the model to study housing policy as a lever for regional mobility. I conduct a policy experiment in which the government finances the construction of an additional 3% of housing in the high-rent region with greater wage growth. I find that this policy succeeds in inducing more mobility towards the high-wage region through the resulting rent price reduction. A secondary effect is that wages change in both regions. The larger influx of labor supply leads to lower wages in the region with greater growth and to a wage increase in the low-wage region. The policy is welfare-enhancing for workers in both regions, however. In the high-wage region, the welfare gains from cheaper housing outweigh the welfare loss from the wage decrease. Workers in the low-wage region value the wage increase at home, but crucially, the increased opportunity from easier mobility to the growth region as well. Overall consumption equivalent variation for newborn workers is 0.80%. Because more workers benefit from the faster skill accumulation in the rich region, I register a 0.04% increase in final output.

The remainder of the paper is structured as follows: Section 2 discusses the data I employ for the empirical analysis. Section 3 presents the empirical results. Section 4 discusses model assumptions and calibration. Section 5 explains low mobility through the lens of the model. The policy experiments are described in Section 6. Section 7 concludes.

Related Literature My work relates to several branches of the literature. Going back to Sjaastad (1962), the idea that moving can be seen as an investment has been present in the literature for a long time. In an important recent addition to this view Bilal and Rossi-Hansberg (2021), who stress the short-term cost and long-term pay-offs of location choice, and that moving to a cheap location is unlike financial assets, not subject to borrowing constraints. I document a clear application of long-term investment through location choice in light of regional human capital.

Equally important, there is the literature on moving costs and migration. An early seminal contribution is Roback (1982), who stresses the role of local amenities. In a more recent influential study Kennan and Walker (2011) estimate large moving costs but also substantial effects of expected income in different regions. Other papers replicate large moving cost estimates in other settings (e.g. Tombe and Zhu (2019), Bryan and Morten

(2019), and Ransom (2022)). Central to the mechanism in my model, Bayer and Juessen (2012) show that moving costs are overestimated if dynamic selection is not taken into account. Also related to my work, Kaplan and Schulhofer-Wohl (2017b) use a model with occupations and regions to explain a drop in interstate mobility by better information and less dispersed wages between states within occupations. Other papers focus on spatial frictions such as impeded ability to compete (Schmutz and Sidibé (2018)) or increased search costs (Heise and Porzio (2022), Diaz et al. (2023)). Perhaps closest to my work, Diaz et al. (2023) stress the lifecycle nature of moving and higher rents because of wealthy late-career workers in high-wage regions. Diaz et al. (2023) focus on search frictions, however. While their focus is on search frictions, my paper instead highlights how regional occupation structures and the accumulation of region-specific human capital shape workers' mobility decisions and regional wage dynamics. Moreover, unlike Diaz et al. (2023), my model allows for borrowing constraints and captures general-equilibrium feedback through the production of regional intermediate goods.

I add to this literature in a number of ways. I explicitly take the accumulation of regional human capital because of regional occupation structures over the lifecycle into account. Additionally, I allow workers to accumulate wealth in the presence of borrowing constraints. Finally, I stress the general-equilibrium effects of moving large numbers of workers between regions that produce different intermediate goods.

Further, there is research on local living costs. Moretti (2013) finds that the nominal wage difference between skill groups is, in part, mitigated by high local living costs. Card et al. (2025) find that local wage-premia are offset by local living costs. Relatedly, my work connects to the literature on regional housing supply and its implications. Hsieh and Moretti (2019) find that in a Rosen-Roback style model, restrictions on new housing supply impede GDP growth in the US. Other papers, too, find evidence for regulations driving up housing prices (e.g. Glaeser et al. (2005), Glaeser and Ward (2009), Saiz (2010), Herkenhoff et al. (2017), and Glaeser and Gyourko (2018)). Building on these findings, I take inelastic housing supply as given and add the point that resulting high living costs have career implications by making young workers miss the window of moving towards opportunity.

A key aspect of my theory is the specificity of human capital. Going back to Shaw (1984) and Shaw (1987), this idea has long been present in the literature. More recently, Kambourov and Manovskii (2009) show that there are sizable returns to occupational experience. Sullivan (2010), too, finds that both occupation and industry-specific human capital are key determinants of wages. Another strand of the literature stresses the importance of different tasks performed in different occupations for the accumulation of human capital (e.g. Yamaguchi (2012), Böhm et al. (2024)).

The point that workers experience larger wage growth and learn more or more quickly in some places than in others, which is key to my theory, is well accepted in the literature. Glaeser and Maré (2001), find that workers in cities have higher wage growth and that a wage premium stays with workers who leave cities, suggesting faster human capital accumulation. Roca and Puga (2016), too, find that workers in large cities accumulate more valuable experience. Dauth et al. (2022) stress that workers in cities have better matching opportunities.

Finally, there is an influential branch of the literature that stresses selection and sorting of workers and firms into locations. Regarding firm selections Bilal (2023) explains regional differences in job stability and unemployment through selection of firms by productivity. Lindenlaub et al. (2022) show that firm sorting has an important influence on wage dispersion. Lhullier (2025) explains between and within-city wage inequality through agglomeration of productive firms. Concerning selection from the worker side, the literature is focused on selection of high- versus low-skilled workers. In this spirit, Borjas et al. (1992) stresses the role of differential returns to being high-skilled in different locations. Similarly, Hunt and Mueller (2004) find that workers with high skills migrate to areas with high returns. Dahl (2002) confirms the importance of comparative advantages of high and low education types in Roy-type regional selection models. Behrens et al. (2014) explain the correlation between skills and city population through sorting by talent and productivity. Diamond (2016) attributes increases in sorting to labor demand changes and changes in local amenities. I add to this literature by stressing that skill is multidimensional and that it evolves with experience that depends on the region where it is gained. As a result, in my model, comparative advantage evolves with regional labor market experience.

2 Data

I use the *Sample of Integrated Labor Market Histories* (SIAB) provided by the Institute for Employment Research (IAB) at the German Federal Employment Agency. The SIAB is a 2% random sample of the complete labor market histories of all workers covered by social security in Germany between 1975 and 2017. Approximately 80% of the German workforce is covered, excluding only the self-employed, civil servants, and the military.

The data contain detailed information on daily wages and socio-demographics, including age, gender, and education. A key feature of the SIAB is its regional detail: it reports the county of each workplace, enabling me to track workers' locations over time. Because counties are relatively small administrative units, I measure geographic moves at the level of commuting zones defined by the Federal Office for Building and Regional Planning (BBSR), which better reflect local labor markets. An additional key strength of these data is that they record occupations in 120 levels, based on the German *KldB* classification.

For my analysis, I restrict the sample to male workers born between 1950 and 1964. Female labor force participation was relatively low in these cohorts, and East German observations are only available from 1993; I therefore drop all workers with any East German observations. To focus on long-term career and mobility outcomes, I further exclude workers with very short employment histories, keeping only those whose first observation occurs at age 30 or earlier and whose last observation occurs at age 45 or later.

Although the SIAB is recorded at a daily frequency, I construct an annual panel by retaining only the spell active on June 1 of each year. One limitation of the data is that wages are top-coded at the social security contribution ceiling; about 12% of wage observations are thus right-censored. Following Böhm et al. (2024) and Dustmann et al. (2009), I impute these censored wages with a standard procedure. The imputation procedure is described in detail in Appendix A.1.

In addition to the microdata described above, my analysis requires information on local price levels. To measure the county-level price level, I use a regional consumer price index (CPI) published by the BBSR, based on regional prices in 2007. When measuring wage growth and in the calibration, I deflate all monetary variables to 1995 prices using the national CPI from Jordà et al. (2017). Regional price differences are accounted for using the BBSR regional CPI.

I also require information on local rents. For this, I draw on the rent index from the *Indicators and Maps on Urban and Spatial Development* (INKAR) database of the BBSR. These data measure the average offered rent per square meter for apartments in medium to good residential areas and are available from 2010 onward. Because the rent index is only available from 2010, I use 2010 values to rank counties in earlier years. To validate this, I make use of the high correlation between rents and building land prices. Building land prices are available at the county level since 1995, and for even earlier years for major cities. Appendix A.2 shows that building land prices are extremely persistent between 1985 and 2010, and that after 2010, rents and building land prices are highly correlated. This provides confidence that the relative ranking of local rents is stable over time. Further, in my analysis, I aggregate the rent index into three categories (low, medium, high) to classify counties by relative cost using county rent terciles. As a robustness check, I use the government's regional rent categories for calculating housing allowances for low-income households (*Wohngeld*). These categories are based on rents for relatively low-cost apartments. Rerunning my main empirical analysis with these data does not change my empirical conclusions, as reported in Appendix C.

To calibrate the model, I also use the *Socio-Economic Panel* (SOEP). To estimate initial wealth, I restrict the data to individuals observed between the ages of 16 and 20. I make use of each individual's reported asset values from the personal wealth module (financial assets, business assets, real estate assets, insurance and pension assets, and other assets such as valuable possessions). To estimate household size, I estimate the modified OECD equivalent scale on households with exactly one adult male worker born between 1940 and 1964. Details of this procedure are discussed in Appendix E.2. I also use regional data from the federal statistics office (Destatis) and the statistics offices of the German states on the regional total square meters of housing and the regional number of births. Details on the calibration are discussed in Section 4.3.

3 Empirical Facts

3.1 Local Wages and Worker Mobility

Differences between regional wage levels in Germany are large and persistent. The left panel of Figure 1 shows a map of West Germany with average log daily wages at the commuting zone level in 1990. The difference between the highest and the lowest level is about 0.6 log-points (over 60%)[1]. The right panel of Figure 1 shows the extremely high level of persistence of local wage levels in a scatter plot. The correlation between commuting zone average levels over the 20-year gap from 1990 to 2010 is 0.868.

[1] Figure 1 in the appendix shows that a large part of the variation survives residualization by age and education.

Figure 1: Local Wage Level and Persistence

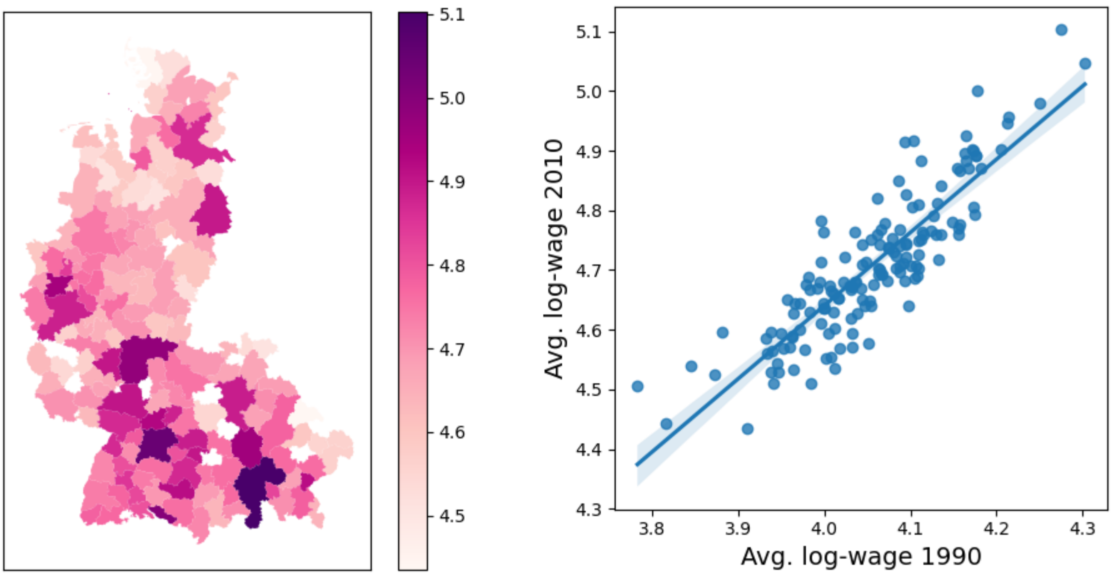

(a) Avg. Log Wage 2010 *(b) Wage Level Persistence*

Notes: The left panel maps commuting zones in West Germany, colored by the average log daily wage in 2010 estimated from the SIAB (darker shades indicate higher values). The right panel shows a scatter of average log daily wages in 2010 (vertical axis) on average log daily wages in 1990 (horizontal axis) with an OLS fit. Each point corresponds to one commuting zone. The sample is restricted to full-time workers.

Despite these large differences in local wage levels, worker mobility during the same time was low. Figure 2 shows the fraction of workers who have never left the commuting zone they initially started their career in against age. By age 55, around half of all workers have never left the commuting zone in which they began their career. The dashed line in Figure 2 reveals that as many as 70% of workers have never worked in a commuting zone that does not neighbor their initial commuting zone. This means that a large fraction of careers in Germany played out in only one commuting zone or close by. In particular, this fact is not simply driven by workers who spend their entire careers in high-wage areas. As shown in Figure 15 in the appendix, the numbers hardly change when reducing the sample to workers who started their career in commuting zones of the lowest wage tercile.

Given the large differences in the local wage level, this is surprising. Should the large and persistent wage level differences not lead to workers moving to high wage regions, leading to high worker mobility? A first candidate answer is that local living costs completely nullify the wage gains. Indeed, Figure 3 shows that a large correlation between the local

Figure 2: Fraction of Never-Movers

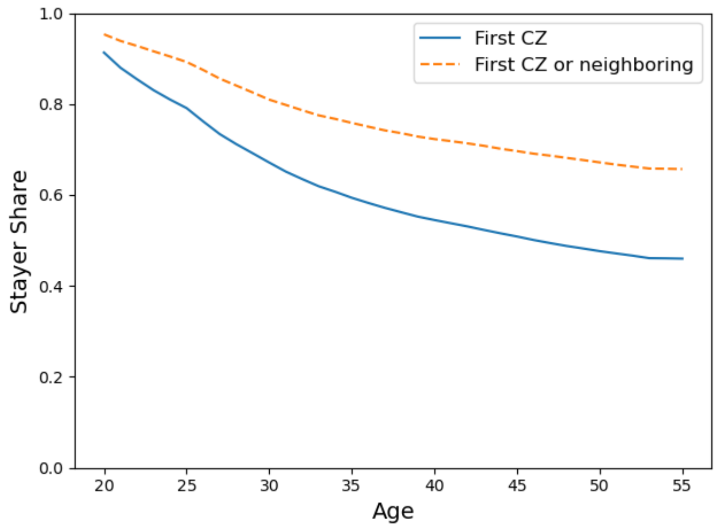

Notes: The solid line plots the share of workers in the SIAB sample who have never left their initial commuting zone, by age. The dashed line plots the share of workers who have either remained in their initial commuting zone or moved only to an adjacent commuting zone.

wage and the local rent level exists. The natural question then becomes whether working in a high-wage/high-rent area even pays off.

To get to the bottom of this, I turn my attention to wage profiles earned over the lifecycle. First, I divide counties into terciles using the 2010 INKAR rent index. Then, I divide workers into three groups by the rent tercile that their very first workplace falls into. I then track these workers over their careers and calculate average log wages of full-time[2] workers at each age, using the BBSR CPI to adjust for regional price differences. The left panel of Figure 4 shows the result for workers born between 1960 and 1964. The right panel shows the same log-wage profiles with the national average log-wage subtracted at each age. This visualizes the log-wage difference between the three rent-terciles at each age. For example, at age 25, the blue line showing the demeaned profile of the lowest tercile is at 0.03 and the green line showing the highest tercile is at -0.03, indicating that workers who started in the lowest tercile earn about 6% higher real wages on average. Results for the remaining birth years are similar and shown in Appendix B.

Figure 4 reveals two key facts. Firstly, after price differences have been accounted for, young workers do not earn higher average real wages in the high-rent/high-wage areas than

[2]Since the SIAB reports whether a worker works part-time but not the number of hours this refers to, I elect to exclude part-time wages from the results presented in Figures 4.

Figure 3: Rent vs. Wage-Level

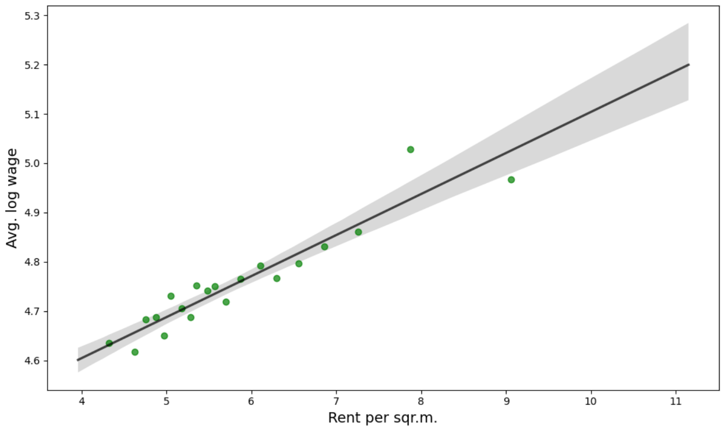

Notes: Binned Scatter of average log-wages (SIAB) against average rents per square meter (INKAR) at the county level with 20 bins. The black line shows an OLS fit from regressing average log wages on average rents per square meter.

their counterparts in the low-rent/low-wage areas. As discussed, in fact, at age 25, real wages are about 6% higher in the cheapest bin than in the most expensive bin. Secondly, looking only at the average wage level of different areas hides the fact that over the lifecycle, the difference between regions is not in the wage level but in wage *growth*. While real wages are higher in the cheap bin for workers under 30, wages in the expensive regions grow faster. At 55, the wage gap between the most and least expensive bins has grown to more than 10% on average. In particular, integrating the differences up over the whole career, it is clear that being in the high-wage/high-rent area does pay off. The late career real wage growth outweighs the lower real wage entry level in terms of real earnings. However, it is also clear that young workers in high-wage areas will tend to be poorer because of high prices. From the point of view of a young worker, a career in a high-wage/high-rent area can be viewed as a long-term investment.

The patterns presented in Figure 4 are striking. In principle, however, they could be an artifact of worker selection — for instance, differences in productivity types or education. I therefore estimate the following regression:

$$\log w_{ij} = \sum_{\substack{k=1,3 \\ a=26,\ldots,55}} \mathbb{1}\{a=j\}\mathbb{1}\{\text{z(i)}=k\}\beta_{jz(i)} + \alpha_i + \delta_j + \gamma_{\text{educ}(i)j} + \mu_{\text{pt.time},ij} + \eta_t + e_{ij} \quad (1)$$

w_{ij} is the wage of worker i at age j. $z(i)$ is the rent bin that worker i started his career in. α_i, δ_j and η_t are worker, age and year fixed effects. $\mu_{\text{pt.-time},ij}$ is a part-time effect

Figure 4: Log Wage Lifecycle Profiles

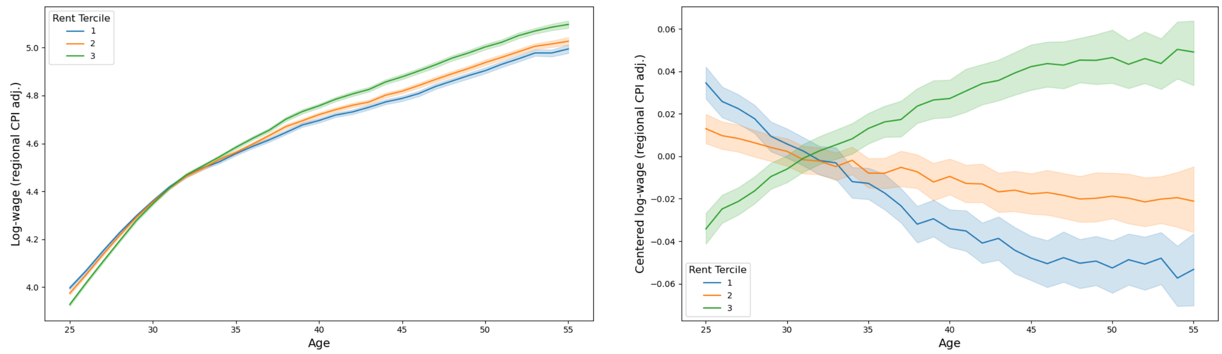

(a) Plain *(b) Mean Subtracted*

Notes: The left panel shows average log wages by age for workers in full-time employment who began their careers in counties belonging to the first, second, and third rent tercile, respectively. Rents are measured using INKAR data from 2010. The green line corresponds to the highest, the blue line to the bottom, and the orange line to the middle rent tercile. Wages are regionally CPI-adjusted using the BBSR CPI. The right panel shows the same wage profiles after subtracting the overall mean at each age. The sample includes workers born between 1960 and 1964.

and $\gamma_{\text{educ}(i)j}$ is an age by education level fixed effect. $\beta_{jz(i)}$ is the coefficient of interest. It measures the effect of having started a career in rent bin $z(i)$ on the log-wage at age j. Note that age 25 and the workers who started their careers in the middle rent tercile are excluded as the baseline.

Equation 1 controls for two key factors. Firstly, it controls for worker fixed effects. If wage differences were to be explained by sorting of plain productivity types, this fixed effect should explain the majority of wage differences. Secondly, it controls for age-by-education fixed effects. If the pattern of differential was to be explained by different education types who have selected into their education group because of their ability to learn and who have faster wage growth precisely because of this better learning ability, this effect should explain the growth differences.

Figure 5 reports the estimated coefficients of interest and standard errors. Neither of the discussed controls explains the difference in wage growth over the career. The difference in wage effects between the top and the bottom rent bin amounts to about 3% at age 26 and continues to grow steadily with age. The difference between the effects of having started the career in the highest vis-à-vis the lowest rent bin on wages at 55 amounts to about 12% - an effect of similar magnitude compared with the raw means presented in Figure 4. This suggests that instead of sorting of plain types or of differently educated types who can learn faster, another mechanism must be at play, producing the observed patterns.

Figure 5: Regression Results

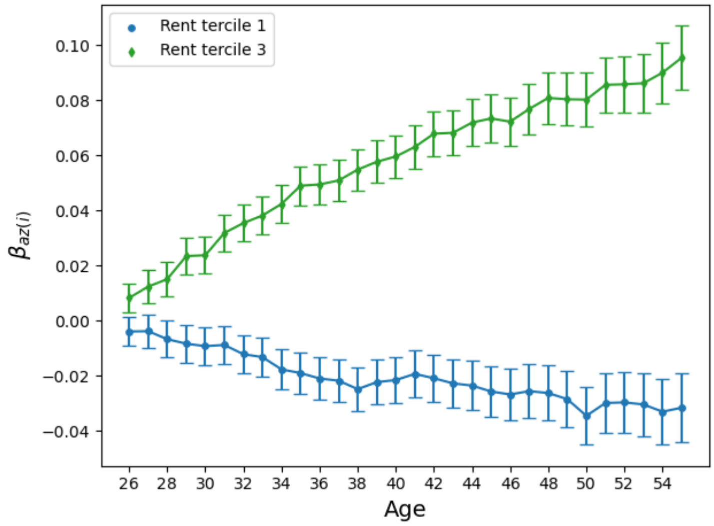

Notes: Regression coefficients $\beta_j z$ for age $j = 26, \ldots, 55$ and rent bins $z = 1, 3$ from estimating Equation 1. The green line corresponds to workers who started their careers in the highest rent tercile, the blue line corresponds to workers who started their careers in the bottom tercile. Age 25 and the middle rent bin $z = 2$ are excluded as the baseline level. β_{jz} measures the effect of having started a career in rent bin $z(i)$ on the log-wage at age j. Standard errors are clustered at the worker level.

Next, I turn my attention to the type of work performed in the different regions - the workers' occupations.

3.2 Regional Occupations

Occupations in Germany are spatially very concentrated. Figure 6 shows Lorenz curves based on the share of all workers working in a given occupation in each commuting zone on the left. Each red line represents one of the 120 *KldB88* occupations reported in the SIAB. The black line represents the hypothetical Lorenz curve for an occupation with equally many workers in each commuting zone. All the Lorenz curves are strictly convex, indicating strong spatial concentration of occupations. At the extreme, the bottom line (corresponding to nautical and aeronautical workers) crosses the point $(0.9, 0.247)$, indicating that as many as 75.3% of all nautical/aeronautical workers work in only 10% of commuting zones. Other occupations represented by lines at the very bottom include insurance professionals, chemists, chemistry lab assistants, publicists, and data specialists. Occupations exhibiting the least spatial concentration include carpenters, scaffolders, plastics-processing technicians, and bricklayers.[3]

Part of the extreme concentration displayed by the red Lorenz curves in the left part of Figure 6 is driven by the different sizes of commuting zones. The right part of Figure 6 shows Lorenz curves in blue that take this into account. Each commuting zone is assigned a section on the x-axis that is proportional to the number of workers contained in the SIAB who work there. Then, commuting zones are sorted by the number of workers per capita in a given occupation. The y-axis still shows accumulated worker shares. The black line now represents the hypothetical Lorenz curve for an occupation with equally many workers per capita in each commuting zone. As expected, the resulting Lorenz curves are closer to the 45-degree line than the uncorrected counterparts. However, all of them are still strictly convex, and occupations still exhibit a large degree of spatial concentration.

Figure 6: Occupation Lorenz Curves

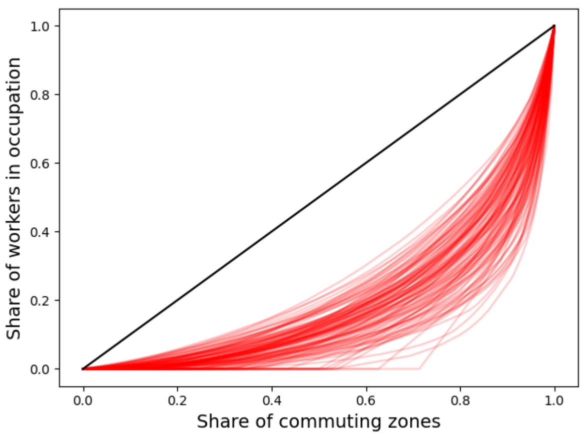

(a) CZ-level

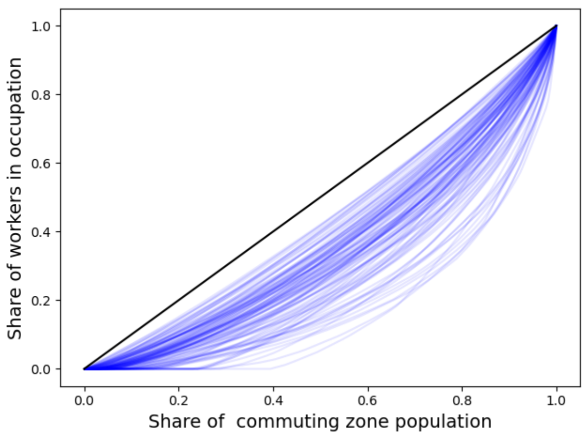

(b) CZ-size-corrected

Notes: Lorenz curves based on shares of all workers working in an occupation who work in each of the BBSR commuting zones. On the left, the commuting zones on the x-axis are sorted from the lowest employment share to the highest employment share separately for each occupation. Each red line represents one of the 120 *KldB88* occupations reported in the SIAB. The black solid line represents the 45-degree line, i.e., the Lorenz curve for an occupation that has equally many workers in each commuting zone. On the right, commuting zones on the x-axis are sorted by employment per capita estimated from the SIAB. Each commuting zone is assigned a segment on the x-axis that is proportional to its size in the SIAB. The black solid line represents the 45-degree line, i.e., the Lorenz curve for an occupation that has equally many workers per capita in each commuting zone. The sample is restricted to the year 2000.

[3] As demonstrated by e.g Dauth et al. (2018), industries are regionally concentrated too. Since industries and occupations are closely connected and the model I propose in Section 4 does not explicitly differentiate between the two notions, and the SIAB is better suited to speak to occupations, I restrict my analysis to occupations.

Figure 7: Observed vs. Counterfactual Log Wage Profiles

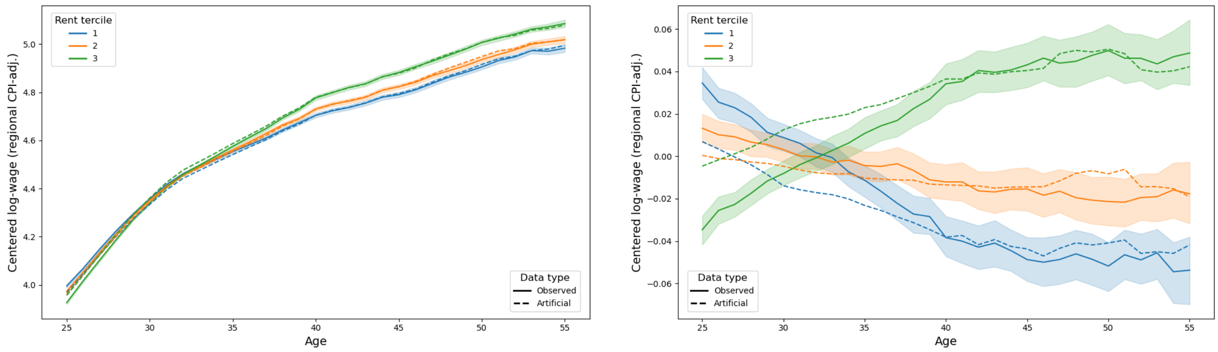

(a) Plain *(b) Mean Subtracted*

Notes: The solid lines show the average log wage profiles for workers in full-time employment who began their careers in counties belonging to the first, second, and third rent tercile, respectively. The dashed lines show counterfactual log wage profile that results from assigning every worker the national average wage, conditional on occupation and age. The green line corresponds to the highest, the blue line to the bottom, and the orange line to the middle rent tercile. Wages are regionally CPI-adjusted to the level of Bonn using the BBSR CPI. The right panel shows the same wage profiles after subtracting the overall mean at each age. Rents are measured using INKAR data from 2010. The sample includes workers born between 1960 and 1964.

The main takeaway from Figure 6 is that workers in different regions are faced with different occupation compositions. This naturally raises the question if these regional occupation structures explain the regional differences in wage levels and wage growth shown in the previous section. To investigate this, I revisit the log-wage profiles shown in Figure 4 and conduct a simple counterfactual exercise that eliminates any regional wage differences *within*-occupation. This isolates the role of occupational composition. If the occupation composition were identical across rent bins, the log-wage profiles should then coincide exactly.

I proceed as follows. As a first step, I calculate average log wages by occupation and age across all regions in West Germany. Each worker is then assigned this national average wage according to their occupation and age, regardless of their region. I then recompute the lifecycle wage profiles of the three regional rent bins using these counterfactual wages.

Figure 7 compares the observed average log-wage profiles already presented in Figure 4 to the resulting counterfactual log-wage profiles. The solid lines are the empirical wage profiles already shown in Figure 4 and the dashed lines are the counterfactual log wage profiles without intra-occupational wage dispersion. The left panel shows average log-wage profiles, while the right panel displays the same profiles after subtracting the national mean at each age. The counterfactual profiles almost exactly replicate the empirical ones. In particular, they closely replicate the differences in wage growth as demonstrated by the right panel.

As discussed, if the occupation composition in the three rent bins were the same, all the log wage profiles would be the same too. In the right panel, this would mean that all three dashed lines should be exactly at 0 at every age. Any difference from 0 must come from the occupation composition behind the three dashed lines. The near-perfect match between empirical and counterfactual profiles implies that the wage differences between workers in the three rent bins can be almost fully attributed to differences in occupational composition, rather than to within-occupation regional wage differentials.

To complement the findings presented in Figure 7 with a simple quantitative measure, I next report how much of the between-county variation in wages can be statistically explained by differences in occupational composition. Specifically, I estimate two straightforward regressions—one for wage levels and one for wage growth—to summarize the explanatory power of regional occupation shares in terms of R^2 values.

I first regress average county log wages on the county's occupational employment shares:

$$\overline{\log w_c} = \sum_{\omega \in \Omega} \beta_\omega s_{c\omega} + e_c \tag{2}$$

where $s_{c\omega}$ denotes the employment share of occupation ω in county c and Ω is the set of 28 two-digit occupations.[4]

Panel A of Table 1 shows that occupation shares explain more than 60% of the between-county variation in wage levels, reaching 80% in the 2000–2010 decade. Thus, most cross-regional wage variation can be attributed to differences in occupational composition.

In Section 3.1, I showed, however, that over the lifecycle, the main difference between high- and low-wage regions is not in the entry level, but in wage growth.[5] To examine whether occupational composition also predicts regional wage *growth*, I estimate a second regression where the dependent variable is the average log-wage difference between ages 50

[4] I work at the county × 2-digit occupation level (rather than BBSR commuting zones × 120 SIAB occupations) to avoid an overparameterized regression. There are 28 2-digit occupations and 269 counties as opposed to 120 occupations and 151 commuting zones. The occupation classification in my SIAB version does not allow differentiation between 2-digit occupations on 3 occasions, leaving only 28 2-digit occupations instead of the full 31 given by the standard KldB classification.

[5] All wage levels are adjusted to the price level of 1995 using the CPI from the Macro History Database (Jordà et al. (2017)).

Table 1: R^2 for Occupation-share Regressions.

Panel A: Wage level by decade		**Panel B: Wage growth by cohort**	
Years	R^2	Cohort	R^2
1980–1990	0.632	1950–1954	0.444
1990–2000	0.708	1955–1959	0.467
2000–2010	0.809	1960–1964	0.543

Notes: Panel A reports R^2-values from estimating Equation 2. The outcome variable is the county-level average log wage. Panel B reports R^2-values from estimating Equation 3. The outcome variable is the difference between the average log wage at 50 and at 25. Wages and occupation shares are measured at the county and 2-digit occupation level.

and 25:

$$\overline{\log w_c^{50}} - \overline{\log w_c^{25}} = \sum_{\omega \in \Omega} \gamma_\omega \tilde{s}_{c\omega} + u_c \tag{3}$$

I estimate Equation 3 separately by cohort. The occupation shares $\tilde{s}_{c\omega}$ are estimated early in each cohort's career - to be precise, 1975-1980 for the 1950-1954 cohort, 1980-1985 for the 1955-1959 cohort, and 1985-1990 for the 1960-1964 cohort. The resulting R^2 values are reported in Panel B of Table 1. Across all three cohorts, the R^2 values are close to 50% of variation explained. This result is particularly striking because wage growth is measured over 25 years, and occupation shares are measured only around the time at which workers begin their careers. Even so, a sizeable part of wage growth variation can be explained by the occupation composition young workers face in a given county.

4 Model

4.1 Model Motivation

In Section 3, I document two key facts. First, using the strong correlation between rent and wage levels, I show that the difference between regions with high wages and high rents and regions with low wages and low rents is not merely a difference in wage levels, but one in wage growth. This is consistent with workers in the high-wage areas accumulating human capital more quickly. Second, the regional occupation composition explains large parts of wage levels and growth, and the different observed wage profiles are wage profiles of workers in different occupations.

In this section, I lay out a general-equilibrium model that reflects these observations and makes use of them to endogenously explain the low levels of worker mobility. The model features two regions with regional labor and rent markets. Regional occupation structure is represented in two ways.

First, in the model, workers accumulate two different skills through experience. Production in each region emphasizes one of the skills, making the accumulation of this skill more likely and more productive. As a result, on average, the human capital of a worker will be more useful in the region in which it was accumulated, making moving less and less attractive as careers progress. This is in line with the literature that documents the accumulation of occupation-specific human capital. Second, I assume that each regional labor market produces a separate intermediate good, reflecting specialization arising from regional occupational structure. These intermediate goods are then aggregated into a final consumption good.

Using the calibrated model, I then proceed to quantify different mechanisms that lead to low mobility. I then conduct a policy exercise in which additional housing is constructed in the region with better career opportunities.

4.2 Model Assumptions

Regions and Production The production function is set up as in the canonical model for production with different skills in Acemoglu and Autor (2011), adding spatial segregation. There are two locations called A and B. In each region, there is a regional representative firm that produces a tradable intermediate good using labor as input. The production function of such a regional firm is

$$F_z(L_z) = L_z$$

where z denotes the region and L_z is the total amount of effective labor supplied to the firm in region z.

The two intermediate goods produced in each region are aggregated to a final consumption good by a national aggregation firm with technology

$$Y = \left(F_A(L_A)^{\frac{\sigma-1}{\sigma}} + F_B(L_B)^{\frac{\sigma-1}{\sigma}}\right)^{\frac{\sigma}{\sigma-1}}$$

where σ is the elasticity of substitution between the two intermediate goods. Labor markets are spatially segmented and competitive. The wage for one unit of labor supplied in region A and the wage for one unit of labor supply $\overline{w}_B$ are given by

$$\overline{w}_A = \frac{\partial Y}{\partial L_A} = \left[1 + \left(\frac{L_B}{L_A}\right)^{\frac{\sigma-1}{\sigma}}\right]^{\frac{1}{\sigma-1}} \qquad \overline{w}_B = \frac{\partial Y}{\partial L_B} = \left[1 + \left(\frac{L_A}{L_B}\right)^{\frac{\sigma-1}{\sigma}}\right]^{\frac{1}{\sigma-1}}.$$

and the wage in Region B takes on the same form. Note that the wage in both regions depends on the ratio of labor supplied in both regions. This is a key consequence of the two regions producing separate intermediate goods. Since both of these goods are needed to produce the final consumption good, wages are not only a function of local labor supply[6], but instead depend on labor supply in the other region too. As a result, worker re-location from one region to the other will - all else equal - raise wages in the region that workers leave and decrease them in the other region.

Each region is endowed with an exogenous stock of housing, denoted by $\overline{\eta_z}$, that is owned by absentee landlords and rented out in spot markets at the endogenous rent price p_z.

Workers - Basic Assumptions Workers are born and draw their birthplace z_0 and an initial endowment of wealth e_0. The worker is born in region A with probability p_A, and the distribution of initial wealth is denoted by F_e. Time is discrete and at annual frequency. Workers work for 40 years and live for 60, spending the last 20 years in retirement. Age is denoted by index j.

Each period before retirement, workers can decide whether to move to the other region at the end of the period. If they decide to move, they pay a resource moving cost κ and incur a utility moving cost $k(j)$, where the latter is increasing with age towards the end of the lifecycle. $k(j)$ captures the emotional and social cost of relocation—leaving behind familiar surroundings, friends, and family. The cost rises late in the life cycle as workers become more settled and socially embedded in their region. Each period before retirement, workers draw an *iid* taste shock for each region ζ_{zj} from an Extreme Value Type 1 distribution.

Workers can invest in an asset a at an exogenous rate R that is determined in interna-

[6]Note that earnings only depending on local TFP and labor supply is the most prevalent assumption in the spatial labor literature - see e.g. Bilal (2023), Hsieh and Moretti (2019), Diaz et al. (2023).

tional financial markets.

Workers - Labor Supply Workers work full-time for the local firm in the region where they live. The number of units of effective labor they supply to the local firm is a function of idiosyncratic productivity ψ_j and two skill levels - h_{aj} and h_{bj}. Idiosyncratic productivity ψ_j follows an $AR(1)$ process in log-space:

$$\log \psi_j = \rho_\psi \log \psi_{j-1} + \varepsilon_j, \quad \varepsilon_j \sim \mathcal{N}(0, \sigma_\psi)$$

Both skill levels can take N discrete values on the grid $\mathcal{H} = \{h_1, \dots, h_N\}$ with $h_1 < \dots, < h_N$. Given productivity and skill levels, the effective units of labor supplied inelastically by each worker are determined by region-specific functions:

$$\ell_A(h_a, h_b, \psi) = h_a^{\alpha_A} \cdot h_b^{1-\alpha_A} \cdot \psi$$
$$\ell_B(h_a, h_b, \psi) = h_a^{\alpha_B} \cdot h_b^{1-\alpha_B} \cdot \psi$$

A worker's labor earnings, given skill levels and productivity, are again region-specific and given by

$$w_A(h_a, h_b, \psi) = \ell_A(h_a, h_b, \psi) \cdot \overline{w}_A$$
$$w_B(h_a, h_b, \psi) = \ell_B(h_a, h_b, \psi) \cdot \overline{w}_B$$

In each region, one of the two skills is more productive at producing the local output good. I assume that h_a is the dominant skill in region A and that h_b is the dominant skill in region B Formally, this means that $\alpha_A > \frac{1}{2}$ and that $\alpha_B < \frac{1}{2}$.

All workers at age $j = 1$ start their careers with skill-levels $h_a = 1$ and $h_b = 1$. Through experience, workers stochastically move up the skill grid. Every year that workers work, there is a probability of the skill index of either or both skills moving to the next grid point if they have not yet reached the h_N. For $i - 1, \dots, N-1$ this probability takes the form

$$P(h'_a = h_{i+1} \mid h_a = h_i, z, j) = \begin{cases} \alpha_A \cdot \overline{p_a} \cdot \delta_a^{j-1}, & z = A, \\ \alpha_B \cdot \overline{p_a} \cdot \delta_a^{j-1}, & z = B, \end{cases}$$

$$P(h'_b = h_{i+1} \mid h_a = h_i, z, j) = \begin{cases} (1-\alpha_A) \cdot \overline{p_b} \cdot \delta_b^{j-1}, & z = A, \\ (1-\alpha_B) \cdot \overline{p_b} \cdot \delta_b^{j-1}, & z = B. \end{cases}$$

The interpretation is straightforward. First, the probability exponentially decays with age j. This means that the older workers become, the slower they learn and the lower the probability of making the next step. Second, the exponential decay is different for the two skills. This means that one of the skills can be learned more quickly, or alternatively, that the probability of a meaningful skill increase is higher simply because there is more to learn. Third, the probability is region-specific in that it is scaled by the Cobb-Douglas exponent from the labor supply function. This means that each skill can be learned more quickly in the region where it is dominant, capturing that skills are learned faster if they are used more through learning by doing, and if there are more people who have this skill, from which it is possible to learn. The Cobb-Douglas exponent captures both - how important the skill is for making the local good - and the exposure to experience, which scales the skill-inherent progress probabilities.

Workers - Value Functions Workers draw utility from consumption of the consumption good c and housing units η. The flow utility function is Stone-Geary and given by

$$u_j(c, \eta) = \frac{1}{1-\lambda}\left(\left(\frac{c}{\nu_j}\right)^{\gamma} \cdot \left(\left(\frac{\eta}{\nu_j}\right) - \nu_j \underline{\eta}\right)^{1-\gamma}\right)^{1-\lambda}.$$

A couple of points warrant discussion. First, there is a subsistence level of housing consumption $\underline{\eta}$, capturing that a worker cannot survive without a minimum amount of shelter. Another realistic effect of this assumption is that poorer workers spend a greater share of their expenditure on housing. After covering expenses for $\bar{\eta}$, workers spend constant shares γ and $1-\gamma$ on consumption and housing, respectively.

Second, there are household size weights $\{\nu_j\}_{j=0}^{T}$. Using the SOEP, I estimate the share of household income in families with one adult male in the cohorts I based my empirical results on to be about 80% (See Appendix E.1 for details). I therefore assume that the labor income earned by any one worker is, in fact, the household income of this worker's household. Consequently, the age-dependent scaling parameter ν_j captures the size of a worker's household over the lifecycle and the fact that the larger a household, the larger the marginal utility of consumption and housing is. Finally, λ is the intertemporal elasticity of substitution.

Each working-age individual solves a dynamic location and consumption–savings problem. At every age j a worker with current state (z, h_a, h_b, a, ψ) decides (i) whether to remain in the current region or move to the other region, and (ii) how much to consume, save, and rent housing for that period. The value function of a working-age worker is therefore

$$V(j, z, h_a, h_b, a, \psi) = \max_{z' \in \{A,B\}} \{V_{z'}(j, z, h_a, h_b, a, \psi) - \mathbb{1}\{z \neq z'\}\, k(j)\, \nu_j + \zeta_{z'}\} \qquad (4)$$

where z' denotes the chosen region for the next period, $V_{z'}(\cdot)$ is the conditional value function given that choice, $k(j)\nu_j$ is the age- and household-size–specific utility cost of moving, and $\zeta_{z'}$ is a region-specific idiosyncratic preference shock drawn from an Extreme Value Type I distribution. This structure yields a standard logit decision rule for regional mobility.

Conditional on choosing region z', the worker optimizes over consumption c, next-period assets a', and housing η:

$$V_{z'}(j, z, h_a, h_b, a, \psi) = \max_{a', c, \eta} \{u_j(c, \eta) + \beta\, \mathbb{E}_j\, [V(j+1, z', h'_a, h'_b, \alpha', \psi')]\}$$

$$\begin{aligned} s.t. \quad & \mathbb{1}\{z \neq z'\}\kappa\nu_j + p_z\eta + c + a' = w_1(\omega(z), h_{\omega(z)}, \psi)(1-\tau) + Ra \\ & a' \geq 0 \end{aligned}$$

The term $\kappa\nu_j$ represents the resource cost of moving, paid only when the worker relocates ($z' \neq z$), and scaled by household size ν_j. Both moving costs are multiplied by ν_j to reflect that larger households face higher financial and utility costs when moving. Expectations are taken over next-period skill levels, idiosyncratic productivity, and regional taste-shocks, conditional on all information available at age j. Workers pay a linear tax τ on labor income.

At the beginning of the lifecycle, there is an initial period $j = 0$. Workers are born with an initial productivity draw ψ_0, wealth endowment e_0, and birthplace z_0. Before entering the first working period ($j = 1$), they decide in which region to start their career. This decision is subject to the usual idiosyncratic taste shocks ζ_A and ζ_B, and takes into account moving costs relative to their birthplace. The value of a newborn worker is

$$V(0, z_0, e_0, \psi_0) = \max_{z' \in \{A,B\}} \left\{V^0_{z'}(z_0, e_0, \psi_0) - \mathbb{1}\{z_0 \neq z'\}\, k(0)\nu_0 + \zeta_{z'}\right\}, \qquad (5)$$

where z' denotes the region chosen for the start of working life. The conditional value function $V^0_{z'}(\cdot)$ is the continuation value from entering region z' at age $j = 1$:

$$V^0_{z'}(z_0, e_0, \psi_0) = \max_{a'} \{\beta \, \mathbb{E}_0 \left[V(1, z', h_{a,1}, h_{b,1}, a', \psi_1)\right]\}$$
$$\text{s.t.} \quad \mathbb{1}\{z_0 \neq z'\}\kappa\nu_0 + a' = e_0, \quad a' \geq 0.$$

Workers do not make consumption or housing choices in this initial period because it is assumed that they still live with their parents and that this choice is made for them.

After the working phase, workers retire in their last workplace and can no longer move. They receive unemployment benefits $b_i = s_b \cdot w_{i40}$, where w_{i40} is the total labor income, worker i earned in the last working-age period and $s_b \in (0, 1)$[7]. The value function of a retired worker for $j < 60$ is

$$V^{\text{ret}}(j, z, b, a) = \max_{c,\eta,a' \geq 0} \left\{u_j(c, \eta) + \beta V^{\text{ret}}(j + 1, z, b, a')\right\} \tag{6}$$
$$s.t. \quad c = Ra + b - a' - p_z\eta$$
$$a' \geq 0$$

In the final period, workers consume all their wealth, i.e., the value function in the final period of a worker's life is

$$V^{\text{final}}(z, b, a) = \max_{c,\eta} \{u_{60}(c, \eta)\} \tag{7}$$
$$s.t. \quad c = Ra + b - p_z\eta.$$

Equilibrium Individual state variables consist of age j, location z, skill levels h_a and h_b, productivity ψ and assets a. The individual state vector is denoted by $x := (j, z, h_a, h_b, \psi, a) \in \mathbb{X}$ where $\mathbb{X}$ is the state space. Let $\mathbb{A} = \{x \in \mathbb{X} \mid z = A\}$ and $\mathbb{B} = \{x \in \mathbb{X} \mid z = B\}$. Let μ_A denote the cumulative distribution of states contained in $\mathbb{A}$ and μ_B denote the cumulative distribution of states contained in $\mathbb{B}$. A competitive steady state equilibrium is a collection of value functions $\{V(x), V^0(x), V^{\text{ret}}(x)\}$[8], worker decision rules $\{z'(x), c(x), \eta(x), a'(x)\}$, equilibrium rent prices p_A, p_B and unit wages $\overline{w}_A, \overline{w}_B$ such that

[7] This means that b is a direct function of the skill levels, productivity, and location in the last period. As such, the value function of the retired worker could be expressed as a function of the same arguments as a working phase value function. To save notation, (h_a, h_b, ψ) are condensed into the state b, however.

[8] Note that I use $V^{\text{ret}}(j, z, b, a) = V^{\text{ret}}(j, z, h_a, h_b, \psi, a) = V^{ret}(x)$ where h_a, h_b and ψ refer to the skill levels and productivity in the last working age period of a worker.

1. Given rent prices p_A, p_B and unit wages $\overline{w}_A, \overline{w}_B$, the decision rules $\{z'(x), c(x), \eta(x), a'(x)\}$ solve the workers' decisions by solving the recursive maximization problems (4)-(7) with the value functions $\{V(x), V^0(x), V^{\text{ret}}(x)\}$.

2. Given the household decisions $\{z'(x), c(x), \eta(x), a'(x)\}$

$$L_A = \int_{\mathbb{A}} \ell_A(h_a, h_b, \psi) d\mu_A(a)$$
$$L_B = \int_{\mathbb{B}} \ell_B(h_a, h_b, \psi) d\mu_B(b)$$

   and the unit wages $\overline{w}_A, \overline{w}_B$ maximize the final consumption good firm's profits given $\frac{L_A}{L_B}$.

3. Rent Markets clear, i.e.

$$\int_{\mathbb{A}} \eta(a) d\mu_A(a) = \bar{\eta}_A$$
$$\int_{\mathbb{B}} \eta(b) d\mu_B(b) = \bar{\eta}_B.$$

The equilibrium is solved for using numerical methods. See Appendix D for details.

4.3 Model Calibration

I calibrate the model to the cohort born between 1960 and 1964. All prices and wages are expressed in 1995 values. Regional price differences are captured endogenously through the model's rent prices and the implied regional CPI. These differences are left untargeted; their fit is discussed in Section 4.4.

Since the model has only two regions, I define the data analogue regions as follows: I define all counties contained in the lower two rent bins (rent tercile 1 and 2 over counties) from Section 3 to be region A and the highest rent bin (tercile 3) to be region B. With this definition, about 56% of observations of the cohort in question in the SIAB are from Region A and the remaining 44% are from B. I interpret model age $j = 1, \ldots, 40$ to represent real-world age $20, \ldots, 59$. In the following, I explain the calibration of parameters.

To translate real-world to model prices, I use the average full-time wage in the SIAB and the average wage in the model.

Utility Parameters The intertemporal elasticity of substitution λ is set to the standard value 2 found, e.g., in Kaplan and Schulhofer-Wohl (2017a). The elasticity of substitution between consumption and housing γ targets 22% of expenditure share for housing by renters reported by Schlattmann (2024), taking into account the subsistence level of housing $\underline{\eta}$. $\underline{\eta}$ itself is set to the model equivalent of $9m^2$ since this is the minimum legal personal living space in most German states.

The discount factor β is set to 0.98.

Finally, the household size weights $\{\nu_j\}_{j=0}^{60}$ are estimated using the modified OECD equivalence scale and SOEP data. I fit a cubic polynomial to the household sizes of 25-70 year old male workers, who are the only adult male in their household born between 1940 and 1964. Before 25 and after 70, I impose that the resulting weight is the same as the nearest fitted value, since there are only relatively few observations of very young and very old households. The estimated values are shown in Figure 24 in Appendix E.

Production Parameters A key parameter is the elasticity of substitution between the intermediate goods produced in each region, as it determines how strong general equilibrium effects of labor relocation will act on local wage levels. Since the production function of the aggregation firm can be interpreted as a representation of trade of sectoral intermediate goods between regions, I choose $\sigma = 4$ in line with the interval of values estimated by the trade literature for Armington elasticities (e.g., Bajzik et al. (2020), Feenstra et al. (2018)).

Region Parameters Regional housing supply is estimated using data from the regional branch of the German Federal Statistics Office. These data report the total square meters of residential buildings at the county level. The total number of square meters in the analogues of regions A and B, denoted by $\bar{\eta}_A^{\text{data}}$ and $\bar{\eta}_B^{\text{data}}$, is obtained by summation across counties. Model housing supplies are then given by

$$\bar{\eta}_A = \bar{\eta}_A^{\text{data}} \cdot \phi, \quad \bar{\eta}_B = \bar{\eta}_B^{\text{data}} \cdot \phi,$$

where ϕ converts total square meters into efficiency units of housing. It is chosen such that the model matches the average rent price per square meter in West Germany in 1995 of €5.22, according to the statistical yearbooks from Destatis.

The probability of being born in region A, denoted by $P(A)$, is estimated at 55% using county-level birth data from Regionalstatistik.

Endowment Distribution I use wealth data from the SOEP for individuals under 20 to estimate the initial wealth distribution F_e. Since 59% of under-20-year-olds report zero wealth, I assume that workers draw an initial endowment of 0 with probability 0.59, and with probability 0.41 draw a positive endowment from a log-normal distribution:

$$\log e_0 \mid e_0 > 0 \sim \mathcal{N}(\mu_e, \sigma_e).$$

The parameters μ_e and σ_e are estimated by maximum likelihood.

Idiosyncratic Productivity The continuous $AR(1)$ process for idiosyncratic productivity is approximated by using Rouwenhorst discretization with five points. The auto-regressive coefficient ρ_ψ is estimated from the SIAB by estimating

$$\log w_{ij} = \alpha_{z(ij)} + \gamma_{\text{edu}(i)} + \delta_{\omega(i)} + t^{\omega}_{ij}\beta_1 + t^{z}_{ij}\beta_2 + j\beta_3 + j^2\beta_4 + u_{ij}$$

where z denotes the sample analogue region of A and B, $\alpha_{z(ij)}$ is a region fixed effect, $\delta_{\omega(i)}$ is an occupation fixed-effect, t^{ω}_{ij} and t^{z}_{ij} are occupational and regional tenure, and $\gamma_{\text{edu}(i)}$ is an education effect. The occupation and region effects proxy the regional wage levels $\overline{w}_A$ and $\overline{w}_B$. The tenure variables, as well as the age profile, capture the part of the wage that the model explains through human capital accumulation of the two skills.

Next, I estimate ρ_ψ as the persistence parameter of the residual component by fitting

$$\hat{u}_{ij} = \rho_\psi \, \hat{u}_{i,j-1} + \varepsilon_{ij},$$

where $\hat{u}_{ij}$ denotes the empirical residuals from the first-stage regression. The variance of the innovation σ^2_ψ is then chosen such that the variance of log wages in the model matches the variance of the empirical log wages.

Moving Costs There are two types of moving costs: the resource cost κ and the utility cost $k(j)$. The resource cost κ is fixed at the model equivalent of €5,000. This value is intended to represent the total out-of-pocket expenses associated with a long-distance move within Germany, including transportation, deposits, and partial furnishing.[9]

The functional form of $k(j)$ is chosen to be approximately constant for most of a worker's career and to rise only toward the end of working life, starting around $j = 26$ (age 45 in the data). This captures that older workers become increasingly reluctant to move as they approach retirement—for example, because they have stronger social and family ties or are more attached to their local environment.

Formally, I specify

$$k(j) = k \cdot \text{scale}(j),$$

where scale(j) is a smooth S-shaped function that remains near one until about age 45 and then increases. The level parameter k is calibrated internally. Details on the exact functional form of scale(j) and the calibrated age profile of $k(j)$ are provided in Appendix E.3.

Human Capital The human capital grid is defined as $\mathcal{H} = \{h_1, \ldots, h_{10}\}$ with $N = 10$ grid points. To anchor its range, I use the empirical 10th and 90th percentiles of log wages in the SIAB, denoted by d_1 and d_{10}. The first point in $\mathcal{H}$ is normalized to one, $h_1 = 1$. The highest grid point h_{10} is pinned down by the empirical 90–10 wage ratio, i.e.

$$\frac{h_{10}}{h_1} = \frac{d_{10}}{d_1}.$$

The remaining points are spaced equidistantly in log space between h_1 and h_{10}.

Retirement Benefits / Taxes I choose the replacement rate s_b such that, on average, workers receive 40% of the average wage they earned as a retirement benefit b. The linear labor income tax τ is adjusted to exactly balance the government budget, i.e., such that the sum over all retirement benefits equals the sum of all collected taxes.

Internal Calibration The above-described procedure leaves 8 parameters that need to be calibrated internally: The skill accumulation probability decay parameters $\bar{p}_a, \bar{p}_b, \delta_a$ and δ_b, the labor skill substitution parameters α_A and α_B, the baseline utility moving cost k, and the scale of the taste shock σ_ζ. These parameters are chosen to match the life cycle

[9] Publicly available estimates for long-distance moves (e.g., between major cities such as Duisburg and Munich) suggest direct costs of about €2,000–2,500 for small flats, excluding deposits and new furniture (Check24, price comparison platform).

wage profiles of both workers who start their career in A and workers who start their career in B, and the fractions of workers who have never left their first region, again both for workers from A and B, all simultaneously.

Table 2 gives an overview of all calibrated parameters.

Table 2: Model Parameters

Parameter	**Description**	**Value**	**Source/Target**
Set			
σ	Armington elast.	4	Trade Lit.
λ	IES	2	standard value
β	Discount factor	0.98	standard value
R	Interest Rate	1.03	standard value
$\underline{\eta}$	Subsistence housing	model eq. of $9\,\mathrm{m}^2$	German Regulation
κ	Moving cost	€5,000	Set
Directly Estimated			
$\{\nu_j\}_{j=0}^{60}$	HH weights	cf. Figure 24	SOEP
$P(A)$	Birth prob.	0.55	Destatis
p_e	Wealth mass	0.41	SOEP
(μ_e, σ_e)	Wealth shape	$(\log 0.136, 1.544)$	SOEP
$\bar{\eta}_A, \bar{\eta}_B$	Housing supply	$(4.482 \cdot 10^8, 2.940 \cdot 10^8)$	Destatis
ρ_ψ	AR(1) coeff.	0.77	SIAB
$\mathcal{H} = \{h_1, \ldots, h_{10}\}$	Skill Grid	$\{1, \ldots, 5.687\}$	SIAB
Internal			
γ	cons.-hs. elast.	0.820	22% hs. exp. share
σ_ψ^2	Shock var.	0.38	$\mathrm{Var}(\log w)$
$\bar{p}_a, \bar{p}_b$	skill acc. prob.	$(0.75, 0.83)$	LC wages/stayer-share
δ_a, δ_b	skill acc. prob. decay	$(0.482, 0.672)$	LC wages/stayer-share
α_A, α_B	Skill subst.	$(0.783, 0.262)$	LC wages/stayer-share
k	Utility moving cost	0.180	LC wages/stayer-share
σ_ζ	Taste scale	0.2	LC wages/stayer-share
ϕ	Housing efficiency	1.335	Avg. rent price p.sq.m
s_b	Replacement rate of b	0.339	40% of avg. wage
τ	Tax rate	0.2025	Balanced budget

4.4 Model Fit

The left panel of Figure 8 shows average log-wage profiles of workers who begin their careers in region A and region B in the model, compared to the same profiles calculated from the SIAB. Both shape and career wage growth in the model fit the data extremely well. Note that neither the model nor the empirical profiles have been adjusted for the regional price level.

Figure 8: Model Fit: Wages

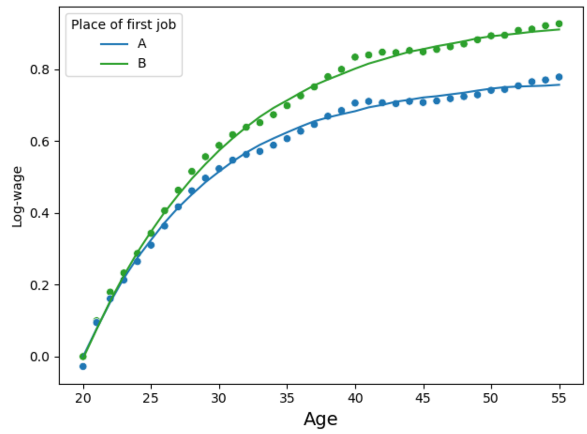

(a) *Plain*

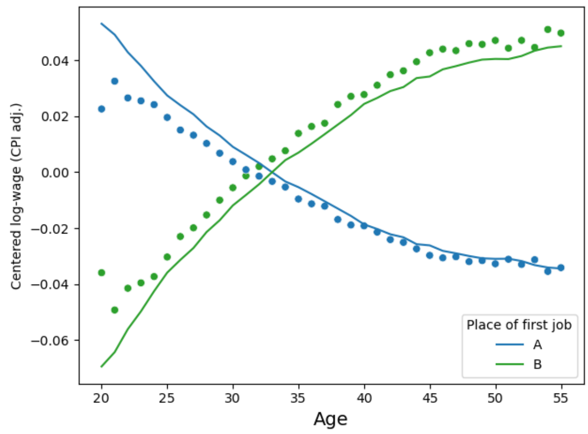

(b) *Mean Subtracted, Regional CPI-adj.*

Notes: The left panel shows average log-wage lifecycles of workers who started their career in A vs B in the model vs the data. The right panel shows the same profiles but with the mean subtracted at every age and adjusted using the regional CPI. For the empirical data, the BBSR CPI is used, the model CPI is calculated using endogenous rent prices and the consumption good as the numeraire. The solid lines represent simulated model outcomes. The dots represent SIAB data estimates.

The right panel, by contrast, shows wage profiles that have been regionally CPI-adjusted. For the empirical data, I use the regional *BBSR* CPI, as in Section 3.1. For the simulated data, I construct an analogous model-internal regional CPI as

$$CPI(A) = \frac{p_A \cdot q_\eta + 1 \cdot q_c}{p_B \cdot q_\eta + 1 \cdot q_c}, \quad CPI(B) = 1$$

where q_η and q_c are the average quantities of housing and consumption consumed in the whole economy, across both regions. I use economy-wide averages since the *BBSR* index also uses the overall German consumption basket, which is used by Destatis. To better illustrate differences between the regions over the lifecycle, the national average log wage is subtracted from each profile in Panel (b).

Note that only the unadjusted wages shown in Panel (a) are targeted directly. The regional real-wage differences in Panel (b), which are implied by the model's endogenous rent and price structure, nonetheless align closely with their empirical counterparts.

Figure 9: Fraction of Never-Movers

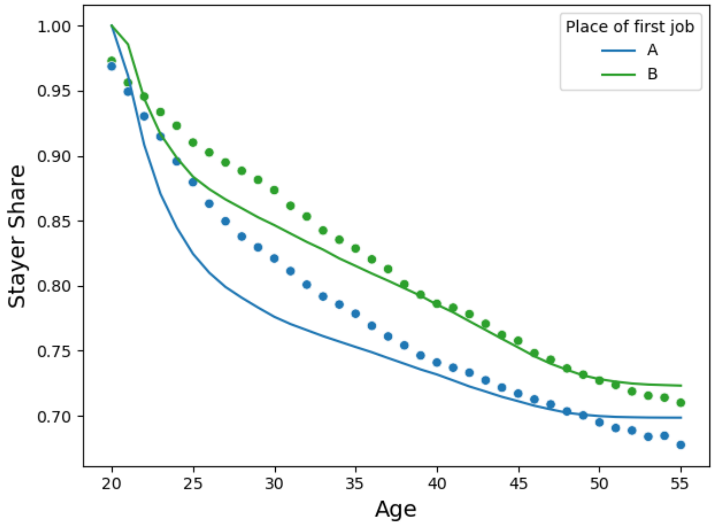

Notes: The solid lines plot the share of workers in the model simulation who have never left their initial region, by their place of labor market entry. The dots plot the same shares in the *SIAB* using the empirical analog of Regions A and B.

Figure 9 shows the fraction of workers who have never left their initial region, by age and region of labor market entry. These empirical profiles are used as calibration targets. The green line represents workers who entered the labor market in region B, the blue line represents workers who entered in region A. The dots show the same statistic in the SIAB using the empirical analog for A and B. The green dots, for example, show the fraction of workers who started their career in a county of the highest rent tercile and have never worked in a county in the lower two terciles. Even though workers move slightly too often in the first 5 years of their lives, the model calibration generates a very realistic mobility level over the lifecycle. In particular, the level difference between workers who began their careers in A vs B is very close to the empirical difference.

Table 3: Model Fit

Moment	**Model**	**Target**
Variance of log-wage	0.50	0.52
Housing expenditure share	22%	22%
Avg. per sq. m. rent	€5.22	€5.22
b-replacement rate of avg. wage	41%	40%

Table 3 reports model outcome values of targets that are not already shown in either Panel (a) of Figure 8 or Figure 9. All targets are met within reasonable tolerance.

Rent prices in the two regions are untargeted. Per-square-meter prices in the model amount to €3.95 in region A and €6.66 in region B. To put these values into perspective, Figure 10 compares them to the distribution of average county rent prices faced by workers in the SIAB in 2010, as measured by the INKAR index. The model rents are reasonably close to the empirical medians and fall well within the support of observed rents.

Figure 10: Rent Price Fit

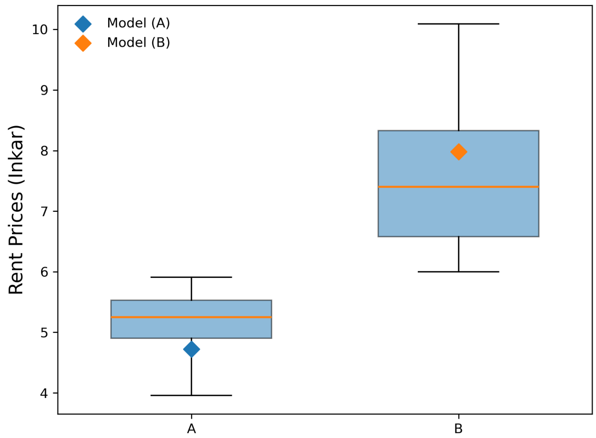

Notes: The boxplots show the distribution of average per-square-meter rents faced by workers in 2010 in counties corresponding to the model's regions A and B, based on the 2010 INKAR rent index. The boxes represent the interquartile range (25th–75th percentile) of county-level average rents, and the whiskers indicate the full observed range. The orange lines mark the medians, and the diamonds show the model's rent levels, rescaled to 2010 using a national rent deflator (FRED), solely for comparability of price levels.

5 Worker Mobility in the Model

5.1 Explaining Low Worker Mobility

As demonstrated in Section 4.4, workers in region B experience substantially stronger wage growth on average. Why, then, do workers not move toward opportunity more frequently? In this section, I quantify different channels that limit mobility.

In the model, four forces limit workers' mobility across regions. Most directly, there are two types of moving costs: the resource cost and the utility cost. Both act as direct frictions, discouraging relocation by reducing the net benefits of moving. In addition, two further mechanisms become important depending on a worker's career stage.

First, there is the high rent price in region B. This holds back young workers in particular. Young workers have neither accumulated assets nor human capital and earn low wages. This means that they have less to spend on consumption and therefore their housing consumption is relatively close to the subsistence level. As a result, the marginal utility from the larger housing they can afford in region A is particularly high, making moving to region B less attractive.

Second, there is human capital mismatch. This, by contrast, affects mostly older workers. Older workers have, on average, accumulated human capital that is more valuable in the region they have spent the most time in. If they move to the other region, their labor earnings drop because their skills are less productive at producing the local output good. On top of that, the older workers, the lower the probabilities of advancing the skill grids become. After a move, older workers are therefore not only mismatched but additionally have lower probabilities of adjusting their skills to the local labor market.

To quantify these effects, at different stages in life, I simulate the moving costs in terms of money, taking all monetary and non-monetary considerations into account. I implement this by calculating how much money each worker in the model would need to be compensated with, in order to choose to move to the other region, abstracting from the *iid* extreme-value taste-shock.[10] The importance of the different channels can then be measured by simulating how the moving costs change if this channel is shut down. To achieve this, both types of moving cost parameters can be set to zero. Note that this includes moving costs for any future moves, given a state. To eliminate the rent price channel, the rent price in region B is set to the level of region A. The remaining residual after all three frictions are removed reflects the contribution of human capital mismatch and the underlying differences in skill accumulation probabilities between regions. Note that when shutting down any of the cost channels, compensating differentials can become negative, meaning that workers would require a payment to stay in their current region.

As a first step, I simulate moving costs for different age groups, presented in Table 4. Overall, average moving costs are about €$39,000$. However, for workers in their 20s, moving costs are only just under €$16,000$ on average. After that, they increase with age.

[10]In the presence of taste shock, it is impossible to pay workers high enough transfers to make them move to the other region with certainty. Instead, I calculate, the minimal payment necessary to move workers to the other region with a probability greater than $\frac{1}{2}$, essentially calculating moving costs in the absence of contemporary taste shocks.

Since it is infeasible to calculate these numbers exactly for each state without extreme use of computational time, I approximate the number by selecting the smallest sufficient payment on an exponentially spaced grid between €−37,500 and €1.5 million with 1,000 points. Workers take prices as given, and no general equilibrium effects are taken into account.

Table 4: Average Moving Costs by Age Group

	Overall	20–29	30–39	40–49	50–59
Mean cost (€)	39,000	15,900	21,600	25,900	90,500

Notes: The table reports average simulated moving costs by age group, expressed in 1995 euros. Moving costs measure the monetary compensation each worker in the model would require to be indifferent between staying and moving to the other region, taking both monetary and non-monetary factors into account. Calculations abstract from idiosyncratic taste shocks and hold equilibrium prices fixed.

For workers approaching retirement, the drastic increase in moving costs can be explained by the exogenous rise in utility moving costs $k(j)$. For workers below 45, however, this happens endogenously, through the discussed channels.

I proceed by conducting the discussed decomposition. First, I turn my attention to young workers. The left panel of Table 5 shows average moving costs by region for 20-year-old workers. Starting with young workers, the baseline in row (1) shows that moving costs are around €8,000 in both regions. When the resource cost κ is set to zero (row 2), moving costs fall sharply, to about €2,000 in both regions—confirming that out-of-pocket expenses are particularly burdensome for young workers who have little wealth. Eliminating the utility cost $k(j)$ (row 3) yields only a modest reduction in moving costs, suggesting that psychic or convenience costs matter less early in life.

Setting rents in region B equal to those in A (row 4) has by far the strongest effect. Without the regional rent difference, moving costs for workers from A drop to €−13,000, on average, indicating that if it wasn't for the high rents in region B, young workers in A had a strong incentive to move. By contrast, workers from B would then require about €14,600 in compensation to move to A—a much higher amount than in the baseline. This highlights that high rents in region B are the key factor deterring young workers in A from moving toward opportunity. At the same time, these high rents make staying in B relatively unattractive for young workers there, offsetting their better wage prospects and resulting in similar baseline moving costs across regions.

When all moving costs and rent differences are eliminated (row 6), workers from A have an even higher willingness to pay to move to B - almost €19,000 on average. Workers from B, by contrast, still require a compensation of €4,500 because moving away slows down their human capital accumulation until they can move back.

In sum, high rent prices in region B are the dominant barrier preventing young workers in A from moving toward better career opportunities. In contrast, workers in B face little incentive to move, as lower rents elsewhere do not compensate for the loss in wage growth potential.

Table 5: Decomposition of Moving Costs by Region and Age (in €)

	Age 20		Age 45	
	Region A	Region B	Region A	Region B
(1) Baseline (all frictions)	8,400	8,100	24,400	27,400
(2) No resource cost	2,100	1,900	18,200	21,200
(3) No utility cost	6,700	6,600	15,400	16,500
(4) Equal rent prices ($p_B = p_A$)	-13,000	14,600	20,300	32,700
(5) No resource or utility cost	500	400	9,200	10,300
(6) No resource or utility costs and $p_B = p_A$	-18,900	4,500	5,500	13,600

Notes: Entries show the average monetary compensation (in euros) required to make a worker move from their current region to the other. Columns report averages for regions A and B at ages 20 and 45. Each counterfactual removes one or more frictions: (2) resource cost κ, (3) utility cost $k(j)$, (4) rent price differential, (5) both costs simultaneously, and (6) all aforementioned frictions combined.

Next, I turn my attention to older workers. The right panel of Table 5 shows average moving costs by region for 45-year-old workers.

Baseline moving costs are much higher at 45 than at 20—about €24,000 in region A and €27,000 in region B. When the resource cost κ is removed (row 2), moving costs fall by roughly one quarter, indicating that out-of-pocket expenses remain relevant but are no longer the main driver. Eliminating the utility cost $k(j)$ (row 3) reduces costs by a similar magnitude.

Equalizing rents between regions (row 4) leads to a modest reduction in moving costs for workers in A but a strong increase for workers in B. At age 45, workers have already accumulated wealth and typically rent larger amounts of housing. As a result, the marginal utility from additional housing space is smaller than for young workers, and higher rents in region B no longer constitute a major barrier to moving.

Setting both the resource and utility costs to zero (row 5) cuts moving costs by about two-thirds relative to the baseline, yet a considerable residual remains. Even when all frictions—including rent differences—are removed (row 6), moving costs remain at roughly €5,500 in A and €13,600 in B. These residual costs capture the effect of human capital mismatch: older workers have accumulated skills that are region-specific, and moving implies a loss in productivity and slower future accumulation in the new local labor market.

It is now clear why workers move so little over the lifecycle. Workers in region A face a dilemma. While they are still young, they would, in principle, like to move to region B to benefit from faster human capital accumulation. But with little wealth and low entry-level wages, high rents make moving to B unattractive. Their only option is to remain in A and save until, with higher wealth levels, they can more easily afford the high rents.

Workers in region B, by contrast, have little incentive to move to region A in search of career opportunities. They endure their early years with limited wealth and high rents but are later compensated through faster wage growth and stronger career progression.

In sum, mobility is not low in the model because moving is extremely costly, but because relatively minor costs at the beginning of young workers' careers lock them into a path. High rents prevent young workers from moving when it would be most beneficial. As they age, the accumulation of region-specific human capital locks workers into their regions.

5.2 Potential Wages

The moving costs calculated in the previous section are quite low, compared to some literature estimates. To illustrate why, I compare potential earnings in the other region in the model with predictions a basic Mincer regression produces when fitted to the model simulations. After all, Mincer regressions are a standard tool in the literature to estimate potential earnings. Through the lens of the model, however, a basic Mincer Regression struggles to produce accurate predictions for potential earnings in the other region.

The most standard Mincer regression with a regional fixed effect takes the form

$$\log \text{wage}_{ij} = \delta + \beta_1 \text{schooling}_i + \beta_2 \text{experience}_{ij} + \beta_3 \text{experience}_{ij}^2 + e_{ij},$$

where $\delta_{z(ij)}$ is the region fixed effect. In the model, schooling does not feature explicitly. However, workers enter the economy with different levels of persistent productivity ψ_{i0}. I therefore use workers' initial productivity to proxy education and let it enter as a fixed effect. Experience is given by age. I therefore estimate

$$\log \text{wage}_{ij} = \alpha_{\psi_{i0}} + \delta_{z(ij)} + \beta_1 \cdot j + \beta_2 \cdot j^2 + e_{ij}. \tag{8}$$

Table 6 shows the mean and median difference between a worker's current log-wage and the log-wage predicted for this worker in the other region by the fitted model given in Equation 8 for 45-year-old workers. For comparison, the mean and median of the actual model counterfactual, given a worker's state, are included as well. The model counterfactual takes into account regional prices by using the model-internal regional CPI discussed in Section 4.3 to adjust wages.

Table 6: Average Log Wage Differences: Mincer Prediction vs. Model Counterfactual

	Region A		Region B	
	Mincer Prediction	Model Counterfactual	Mincer Prediction	Model Counterfactual
Mean	0.104	-0.275	-0.229	-0.455
Median	0.113	-0.310	-0.233	-0.473

Notes: The table reports the mean and median log wage differences between simulated model wages and predicted wages in the other region from the Mincer regression and from the structural model counterfactual, separately by region. Positive values indicate that predicted wages exceed observed wages.

In the case of workers from region A, (Panel (a) Figure 26), the Mincer Regression overestimates potential earnings by a large margin. On average, the regression predicts a sizable wage of over 10%, whereas the actual counterfactual implies a real wage of 0.275 log-points. This is because naively estimating Equation 8 fails to acknowledge that general experience does not capture the level of the skill required in region B. Additionally, it does not take the high rents into account.

Likewise, the regression model underestimates the wage loss of workers from region B, because it fails to take the regional nature of human capital into account.[11]

An econometrician facing the simulated model data who assumes that experience can be fully utilized after a move to another region may estimate a model similar to Equation 8 and conclude that there are large unrealized gains from moving to region B. Such an interpretation would suggest that high moving costs are necessary to explain why workers from region A remain in their home region. Taking regional human capital and living costs into account, however, these large incentives do not exist through the lens of the model, where regional human capital and high living costs are the main drivers of low mobility. In this sense, the model's implications echo the findings of Kambourov and Manovskii (2009), who show that human capital is highly occupation-specific. Acknowledging the regional concentration of occupations, therefore, becomes crucial in explaining low mobility.

[11] Note that this wage loss does not capture that a given level of wealth is worth more in real terms in Region A than in Region B, mitigating any wage losses when moving to Region A.

6 Housing Supply Policy

In Section 5.1, I discussed two main findings. First, moving typically only makes economic sense for young workers, since they have not yet accumulated much region-specific human capital and therefore face little skill mismatch after a move. Second, the main factor preventing young workers in region A from moving to region B is the high rent level. In light of the limited housing supply, older workers who occupy large housing units do not internalize the effect of their housing consumption on the ability of young workers to move to region B. This section presents a policy experiment aimed specifically at mitigating this effect by expanding the housing stock in region B. In particular, I let the government increase the housing stock in region B through public investment.

According to data from the Federal Statistical Office (Destatis),[12] constructing one square meter of housing in 1995 cost €1,254, and a square meter of undeveloped building land cost €58.02. I therefore use the model equivalent of €1,312.02 as the price for building an additional efficiency unit of housing. The construction project is assumed to be fully financed through government debt. Interest payments at the exogenous model rate R are covered by adjustments to the labor income tax τ. The government remains the owner of the newly constructed housing and receives rental income from renting it out at the market rent price.

[12] Specifically, I take the building price from the Destatis report “Bauen und Wohnen, Baugenehmigungen, Baukosten” (Construction and Housing – Building Permits and Construction Costs) and the land price from “Fachserie 17 / Reihe 5 Preise Kaufwerte für Bauland” (Series 17 / Series 5: Prices – Purchase Prices for Building Land).

Figure 11: Fraction of Never-Movers (Counterfactual)

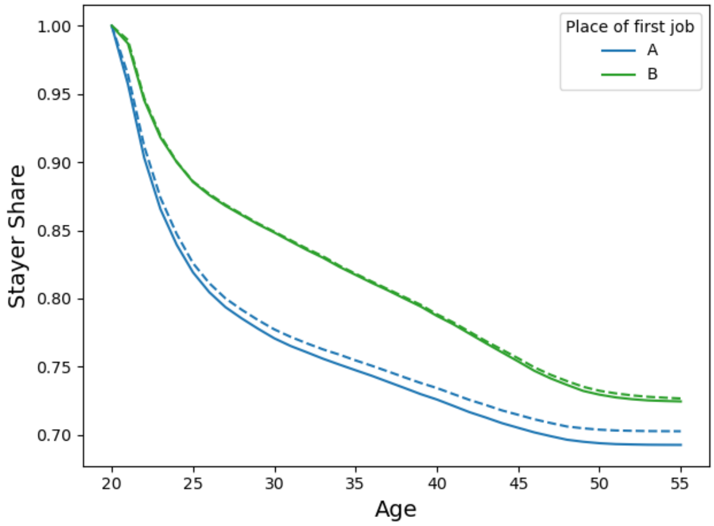

Notes: The solid lines plot the share of workers in the model simulation who have never left their initial region, by their place of labor market entry after a 3% increase in the housing stock in region B. Dashed lines show the same fractions in the baseline economy.

I study a construction project that increases the housing stock in region B by 3%. The magnitude is motivated by large-scale real-world construction projects. For example, an ongoing project in the city of Mannheim[13] is projected to provide housing for about 10,000 people—a roughly 3% increase in the city's population. This policy directly lowers the rent price in region B, which falls to €6.51 compared to €6.66 in the baseline economy. As a result, moving to region B becomes much more attractive, especially for young workers. Figure 11 shows the fraction of workers who have never moved by origin region. The baseline economy is depicted by dashed lines, and the post-policy economy by solid lines. Clearly, mobility out of region A increases. Consequently, rents in region A adjust slightly downward as well—to €3.93 compared to €3.95 in the baseline—reflecting lower local housing demand.

In addition to the rent effects, the policy also influences wage levels. Since the two regions produce separate intermediate goods, reallocating labor from region A to region B raises the price of labor in region A and lowers it in region B. This wage adjustment is amplified by the fact that more workers in region B now accumulate human capital more rapidly. As a result, wages per efficiency unit rise by 0.18% in region A and fall by 0.18% in region B.

[13] Specifically, the Franklin Village project.

Table 7: CEV after Housing Construction

Decomposition	Veil of Ignorance	Conditional on A	Conditional on B
Full Policy	0.80%	0.76%	0.85%
Δp_A	0.07%	0.12%	0.03%
$\Delta \overline{w}_A$	0.15%	0.23%	0.04%
Δp_B	0.41%	0.17%	0.71%
$\Delta \overline{w}_B$	-0.12%	-0.05%	-0.22%
$\Delta \tau$	0.29%	0.29%	0.29%

Notes: Consumption equivalent variation (CEV) of newborn workers with respect to a counterfactual economy where the government increases the housing stock in region B by 3%. The decomposition in rows 2–6 isolates the contribution of each general-equilibrium price change by resetting one variable at a time to its baseline level and recomputing the resulting CEV.

The faster human capital accumulation of a larger number of workers in region B raises the overall effective labor supply in the economy, which increases by about 0.05%. Aggregate output rises as well, but only modestly, since more labor is now devoted to producing the intermediate good from region B, which exhibits diminishing returns in the final output production. Overall, output in the new steady state is 0.03% higher than in the baseline.[14]

Next, I analyze how these changes translate into welfare effects. Table 7 reports the consumption equivalent variation of newborn workers relative to the baseline. Behind the veil of ignorance, newborn workers experience a sizeable welfare gain of 0.8%. Conditional on being born in region B, the welfare gain rises slightly to 0.85%, as these workers directly benefit from the larger housing supply. The CEV for workers born in region A, at 0.76%, is still substantial.

The decomposition in Table 7 reveals that the rent reduction in region B is the dominant welfare channel. Its effect is particularly strong for workers born in B, amounting to 0.41% CEV. This benefit is partially offset by the negative wage effect in region B of about -0.12%. For workers from region A, the changes in local rents and wages both have positive effects (0.12% and 0.23%, respectively). Importantly, however, the lower rent in region B contributes even more—around 0.7% CEV—reflecting the welfare gain from being able to move more easily to the high-opportunity region.

Additionally, the construction project is fiscally profitable. Government rental income

[14] GDP is measured as total output Y, excluding construction activity.

Table 8: CEV by Endowment

Endowment Tercile	Veil of Ignorance	Conditional on A	Conditional on B
1	0.80%	0.79%	0.77%
2	0.76%	0.75%	0.75%
3	0.85%	0.84%	0.79%

Notes: Consumption equivalent variation of newborn workers by endowment tercile with respect to a counterfactual economy where the government increases the housing stock in region B by 3%.

from the new housing exceeds the interest payments on the construction debt, resulting in a further welfare gain of 0.29% CEV through lower taxes.

Finally, Table 8 presents welfare effects by endowment tercile. The gains are largest for workers in the lowest tercile, particularly those born in region B. Since poorer households spend a larger share of income on housing, they benefit disproportionately from lower rents. Although the policy raises welfare for all groups, it is most beneficial to poorer workers.

For comparison, I also consider an alternative policy aimed at reducing housing demand rather than expanding supply. Specifically, I examine a downsizing incentive similar to programs discussed in the public debate, where retired households receive transfers if they move to smaller dwellings. In the model, however, this policy proves largely ineffective. It induces only marginal reductions in housing consumption among retirees, with negligible effects on rents, wages, or mobility. Consequently, the welfare effect is slightly negative, as the fiscal cost of the subsidy outweighs its benefits. The details of the experiment are described in Appendix G.2.

7 Conclusion

In this study, I revisit the question of why workers are not more mobile despite large regional wage differences. Using high-quality administrative data from Germany, I document that regional wage differences manifest themselves not in higher entry-level wages but in larger wage growth over the lifecycle. I show that occupations are regionally concentrated and that large parts of the variation in local wage levels and growth can be accounted for by regional occupation compositions.

Motivated by these findings, I propose a two-region general equilibrium model with local labor and rent markets. In the model, different occupation compositions are represented by different skills that are differentially productive in the two regions and are learned stochastically through experience, as well as different intermediate output goods being made in each region. In one of the regions, workers specialize in a skill that can be learned faster, and experience greater wage growth over their careers. Workers can move between the regions every year, draw utility from consumption and housing, and are subject to a borrowing constraint. I calibrate the model to the data, and I find that for young workers, the high rent prices in the high-wage region are a major reason not to move there. Older workers have, on average, accumulated human capital that is more productive in their current region and have no incentives to move. I find that average moving costs, measured by the amount of compensation needed to move workers, are only about €39,000.

I conduct a policy experiment in which the government pays for the construction of 10% more housing in the high-wage region. I find that this policy leads to more mobility into the high-wage region. In particular, it is welfare-improving for workers from both regions. Workers in the high-wage region value the decreasing rent prices. Workers from the low-wage region experience wage increases and moderate rent decreases. On top of that, they benefit from increased possibilities to move to the high-wage region for a better career.

A similar policy, incentivizing retired workers to downsize their housing consumption, fails to increase welfare.

For future research, there are several ways forward. For once, extending the model to include home-owners is interesting, since home-owners are even more reluctant to move and, in particular, dislike falling house prices that could be induced after construction projects. Another interesting idea could be opening the model economy to international trade. This could alter general equilibrium effects because of labor re-allocation, since international demand for one of the intermediate output goods might be greater. Finally, introducing a notion of differential talent to speak to early life skill mismatch could be interesting to introduce even stronger inequality of opportunity. It could also lead to larger GDP responses after mobility-inducing policies. I view this as fruitful ground for future research.

8 References

D. Acemoglu and D. Autor. Chapter 12 - skills, tasks and technologies: Implications for employment and earnings**we thank amir kermani for outstanding research assistance and melanie wasserman for persistent, meticulous and ingenious work on all aspects of the chapter. we are indebted to arnaud costinot for insightful comments and suggestions. autor acknowledges support from the national science foundation (career award ses-0239538). volume 4 of *Handbook of Labor Economics*, pages 1043–1171. Elsevier, 2011. doi: https://doi.org/10.1016/S0169-7218(11)02410-5. URL https://www.sciencedirect.com/science/article/pii/S0169721811024105.

J. Bajzik, T. Havranek, Z. Irsova, and J. Schwarz. Estimating the armington elasticity: The importance of study design and publication bias. *Journal of International Economics*, 127:103383, 2020. ISSN 0022-1996. doi: https://doi.org/10.1016/j.jinteco.2020.103383. URL https://www.sciencedirect.com/science/article/pii/S0022199620300982.

C. Bayer and F. Juessen. On the dynamics of interstate migration: Migration costs and self-selection. *Review of Economic Dynamics*, 15(3):377–401, 2012. ISSN 1094-2025. doi: https://doi.org/10.1016/j.red.2012.02.002. URL https://www.sciencedirect.com/science/article/pii/S1094202512000075.

K. Behrens, G. Duranton, and F. Robert-Nicoud. Productive cities: Sorting, selection, and agglomeration. *Journal of Political Economy*, 122(3):507–553, 2014. ISSN 00223808, 1537534X. URL http://www.jstor.org/stable/10.1086/675534.

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.

A. Bilal and E. Rossi-Hansberg. Location as an asset. *Econometrica*, 89(5):2459–2495, 2021. doi: https://doi.org/10.3982/ECTA16699. URL https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA16699.

M. J. Böhm, H.-M. von Gaudecker, and F. Schran. Occupation growth, skill prices, and wage inequality. *Journal of Labor Economics*, 42(1):201–243, 2024. doi: 10.1086/722084. URL https://doi.org/10.1086/722084.

G. J. Borjas, S. G. Bronars, and S. J. Trejo. Self-selection and internal migration in the united states. *Journal of Urban Economics*, 32(2):159–185, September 1992. doi: None. URL `https://ideas.repec.org/a/eee/juecon/v32y1992i2p159-185.html`.

G. Bryan and M. Morten. The aggregate productivity effects of internal migration: Evidence from indonesia. *Journal of Political Economy*, 127(5):2229–2268, 2019. doi: 10.1086/701810. URL `https://doi.org/10.1086/701810`.

D. Card, J. Rothstein, and M. Yi. Location, location, location. *American Economic Journal: Applied Economics*, 17(1):297–336, January 2025. doi: 10.1257/app.20220427. URL `https://www.aeaweb.org/articles?id=10.1257/app.20220427`.

C. D. Carroll. The method of endogenous gridpoints for solving dynamic stochastic optimization problems. *Economics Letters*, 91(3):312–320, June 2006. doi: None. URL `https://ideas.repec.org/a/eee/ecolet/v91y2006i3p312-320.html`.

G. B. Dahl. Mobility and the return to education: Testing a roy model with multiple markets. *Econometrica*, 70(6):2367–2420, 2002. doi: https://doi.org/10.1111/j.1468-0262.2002.00443.x. URL `https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1468-0262.2002.00443.x`.

W. Dauth, M. Fuchs, and A. Otto. Long-run processes of geographical concentration and dispersion: Evidence from germany. *Papers in Regional Science*, 97(3):569–593, 2018. doi: https://doi.org/10.1111/pirs.12271. URL `https://onlinelibrary.wiley.com/doi/abs/10.1111/pirs.12271`.

W. Dauth, S. Findeisen, E. Moretti, and J. Suedekum. Matching in cities. *Journal of the European Economic Association*, 20(4):1478–1521, 2022.

R. Diamond. The determinants and welfare implications of us workers' diverging location choices by skill: 1980-2000. *American Economic Review*, 106(3):479–524, March 2016. doi: 10.1257/aer.20131706. URL `https://www.aeaweb.org/articles?id=10.1257/aer.20131706`.

A. Diaz, A. Janez, and F. Wellschmied. Geographic Mobility over the Life-Cycle. IZA Discussion Papers 15896, Institute of Labor Economics (IZA), Jan. 2023. URL `https://ideas.repec.org/p/iza/izadps/dp15896.html`.

T. A. Domencich and D. L. McFadden. *Urban Travel Demand: A Behavioral Analysis*, volume 93 of *Contributions to Economic Analysis*. North-Holland Publishing Company / American Elsevier, Amsterdam / New York, 1975. ISBN 0444108300; 0720431964.

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.

R. C. Feenstra, P. Luck, M. Obstfeld, and K. N. Russ. In search of the armington elasticity. *The Review of Economics and Statistics*, 100(1):135–150, 03 2018. ISSN 0034-6535. doi: 10.1162/REST_a_00696. URL https://doi.org/10.1162/REST_a_00696.

E. Glaeser and J. Gyourko. The economic implications of housing supply. *Journal of Economic Perspectives*, 32(1):3–30, February 2018. doi: 10.1257/jep.32.1.3. URL https://www.aeaweb.org/articles?id=10.1257/jep.32.1.3.

E. Glaeser and D. Maré. Cities and skills. *Journal of Labor Economics*, 19(2):316–342, 2001. ISSN 0734306X, 15375307. URL http://www.jstor.org/stable/10.1086/319563.

E. Glaeser, J. Gyourko, and R. Saks. Why is manhattan so expensive? regulation and the rise in housing prices. *The Journal of Law and Economics*, 48(2):331–369, 2005. doi: 10.1086/429979. URL https://doi.org/10.1086/429979.

E. L. Glaeser and B. A. Ward. The causes and consequences of land use regulation: Evidence from greater boston. *Journal of Urban Economics*, 65(3):265–278, 2009. ISSN 0094-1190. doi: https://doi.org/10.1016/j.jue.2008.06.003. URL https://www.sciencedirect.com/science/article/pii/S0094119008000582.

S. Heise and T. Porzio. Labor misallocation across firms and regions. Working Paper 30298, National Bureau of Economic Research, July 2022. URL http://www.nber.org/papers/w30298.

K. F. Herkenhoff, L. E. Ohanian, and E. C. Prescott. Tarnishing the golden and empire states: Land-use restrictions and the u.s. economic slowdown. Working Paper 23790, National Bureau of Economic Research, September 2017. URL http://www.nber.org/papers/w23790.

C.-T. Hsieh and E. Moretti. Housing constraints and spatial misallocation. *American Economic Journal: Macroeconomics*, 11(2):1–39, April 2019. doi: 10.1257/mac.20170388. URL https://www.aeaweb.org/articles?id=10.1257/mac.20170388.

G. L. Hunt and R. E. Mueller. North american migration: Returns to skill, border effects, and mobility costs. *The Review of Economics and Statistics*, 86(4):988–1007, November 2004. doi: None. URL https://ideas.repec.org/a/tpr/restat/v86y2004i4p988-1007.html.

F. Iskhakov, T. H. Jørgensen, J. Rust, and B. Schjerning. The endogenous grid method for discrete-continuous dynamic choice models with (or without) taste shocks. *Quantitative Economics*, 8(2):317–365, 2017. doi: https://doi.org/10.3982/QE643. URL https://onlinelibrary.wiley.com/doi/abs/10.3982/QE643.

O. Jordà, M. Schularick, and A. M. Taylor. Macrofinancial history and the new business cycle facts. *NBER Macroeconomics Annual*, 31:213–263, 2017. doi: 10.1086/690241. URL https://doi.org/10.1086/690241.

G. Kambourov and I. Manovskii. Occupational specificity of human capital. *International Economic Review*, 50(1):63–115, 2009. doi: https://doi.org/10.1111/j.1468-2354.2008.00524.x. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1468-2354.2008.00524.x.

G. Kaplan and S. Schulhofer-Wohl. Understanding the long-run decline in interstate migration. *International Economic Review*, 58(1):57–94, 2017a. doi: https://doi.org/10.1111/iere.12209. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/iere.12209.

G. Kaplan and S. Schulhofer-Wohl. Understanding the long-run decline in interstate migration. *International Economic Review*, 58(1):57–94, 2017b. doi: https://doi.org/10.1111/iere.12209. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/iere.12209.

J. Kennan and J. R. Walker. The effect of expected income on individual migration decisions. *Econometrica*, 79(1):211–251, 2011.

H. Lhullier. The local root of wage inequality. *University of Chicago, Becker Friedman Institute for Economics Working Paper*, (2025-78), 2025. URL https://ssrn.com/abstract=5285666.

I. Lindenlaub, R. Oh, and M. Peters. Firm sorting and spatial inequality. Technical report, National Bureau of Economic Research, 2022.

E. Moretti. Real wage inequality. *American Economic Journal: Applied Economics*, 5 (1):65–103, January 2013. doi: 10.1257/app.5.1.65. URL https://www.aeaweb.org/articles?id=10.1257/app.5.1.65.

T. Ransom. Labor market frictions and moving costs of the employed and unemployed. *Journal of Human Resources*, 57(S):S137–S166, 2022. ISSN 0022-166X. doi: 10.3368/jhr.monopsony.0219-10013R2. URL https://jhr.uwpress.org/content/57/S/S137.

J. Roback. Wages, rents, and the quality of life. *Journal of Political Economy*, 90(6): 1257–1278, 1982.

J. D. L. Roca and D. Puga. Learning by working in big cities. *The Review of Economic Studies*, 84(1):106–142, 07 2016. ISSN 0034-6527. doi: 10.1093/restud/rdw031. URL https://doi.org/10.1093/restud/rdw031.

A. Saiz. The geographic determinants of housing supply*. *The Quarterly Journal of Economics*, 125(3):1253–1296, 08 2010. ISSN 0033-5533. doi: 10.1162/qjec.2010.125.3.1253. URL https://doi.org/10.1162/qjec.2010.125.3.1253.

L. Schlattmann. Spatial redistribution of carbon taxes. *JMP*, 2024. URL https://lennardschlattmann.github.io/papers/schlattmann_jmp.pdf.

B. Schmutz and M. Sidibé. Frictional Labour Mobility. *The Review of Economic Studies*, 86(4):1779–1826, 09 2018. ISSN 0034-6527. doi: 10.1093/restud/rdy056. URL https://doi.org/10.1093/restud/rdy056.

K. L. Shaw. A formulation of the earnings function using the concept of occupational investment. *The Journal of Human Resources*, 19(3):319–340, 1984. ISSN 0022166X. URL http://www.jstor.org/stable/145876.

K. L. Shaw. Occupational change, employer change, and the transferability of skills. *Southern Economic Journal*, 53(3):702–719, 1987. ISSN 00384038. URL http://www.jstor.org/stable/1058765.

L. A. Sjaastad. The costs and returns of human migration. *Journal of Political Economy*, 70(5):80–93, 1962.

P. Sullivan. Empirical evidence on occupation and industry specific human capital. *Labour Economics*, 17(3):567–580, 2010. ISSN 0927-5371. doi: https://doi.org/10.1016/j.labeco.2009.11.003. URL https://www.sciencedirect.com/science/article/pii/S0927537109001286.

T. Tombe and X. Zhu. Trade, migration, and productivity: A quantitative analysis of china. *American Economic Review*, 109(5):1843–72, May 2019. doi: 10.1257/aer.20150811. URL https://www.aeaweb.org/articles?id=10.1257/aer.20150811.

S. Yamaguchi. Tasks and heterogeneous human capital. *Journal of Labor Economics*, 30 (1):1–53, 2012. doi: 10.1086/662066. URL https://doi.org/10.1086/662066.

A Data

A.1 Wage Imputation

This appendix describes the procedure used to impute wages that are top-coded at the social security contribution ceiling. I follow Böhm et al. (2024), in turn following Dustmann et al. (2009), and implement a Tobit model for (log) daily wages with right-censoring separately for each year. Let $y_i^* = x_i'\beta + \varepsilon_i$ denote the latent log wage with $\varepsilon_i \sim \mathcal{N}(0, \sigma_g^2)$, where g indexes an education $\times$ age-group cell and a three-level bin of the regional rent distribution. Because I later split the sample by the same three rent bins, I perform the imputation separately for these bins. This ensures that the upper-tail fit is not distorted by systematic regional cost differences that correlate with the censoring threshold and wage structure. Importantly, however, my results are almost unchanged if I instead pool regions and omit the rent-bin split at the imputation stage.

The observed outcome is $y_i = \min\{y_i^*, c_t\}$, where c_t is the time-t contribution ceiling mapped to daily real wages. I estimate (β, σ_g) by maximum likelihood using all observations, weighting each record by employment spell length. The covariate vector x_i includes age, the individual's censoring rate, the average wage from other years (to capture persistent individual effects), and education. Allowing the wage variance to differ by education $\times$ age-group cell follows Dustmann et al. (2009) and improves fit for the upper tail.

I then replace censored observations with simulated draws from a truncated normal distribution, using the Tobit prediction as the mean and the estimated Tobit variance as the dispersion parameter.

A.2 INKAR Rent Data

The INKAR rent index is only published from 2010 onwards. To assert that it can still be used to measure rent-terciles in earlier years, I turn to building land prices.

Figure 12 shows the correlation of rents as measured by the INKAR rent index with average building land prices at the county level in a scatter plot. Building land prices have been obtained from the Regional Statistics Database (Regionalstatistik), a database published by the statistics offices of the German Federal States. The correlation is 83%.

Local building land prices can be traced back much further than local rents. At the county level, they are reported from 1995 onwards. Before that, they are published for

Germany's major cities in the annual report of the Federal Statistics Office (Destatis). Panel (a) of Figure 13 shows the auto-correlation of building land prices at the county level between 1995 and 2010. Panel (b) shows the same correlation for the major cities reported in the report of 1980. In both cases, autocorrelation is about 90%.

This provides reassurance that rents are equally persistent over time and the grouping into terciles works. Additionally, I conduct robustness checks with alternative data based on a German rent subsidy scheme reported in Appendix C.

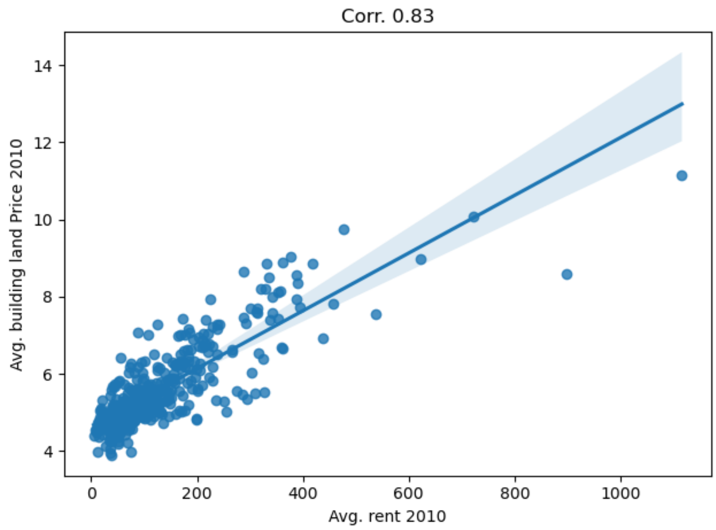

Figure 12: Correlation of County Rents and Building Land Prices in 2010

Notes: Scatter plot of average offered rent per m^2 (INKAR, 2010) against building land prices from the Regional Statistics Database (Regionalstatistik), maintained jointly by the Federal Statistical Office of Germany (Destatis) and the statistical offices of the federal states. Each point is a county. The fitted line shows the OLS regression of log rent on log building land price.

Figure 13: Persistence of Building Land Prices Over Time

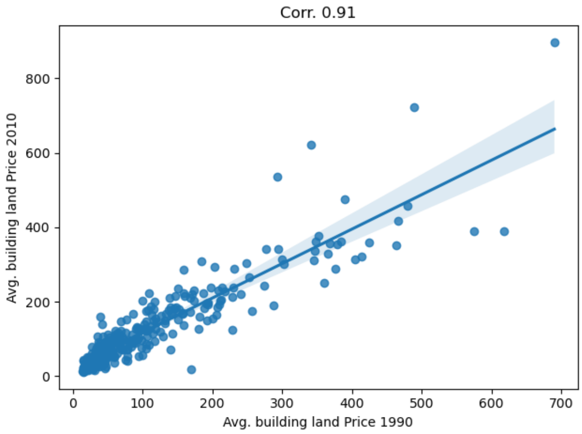

(a) 1995–2010

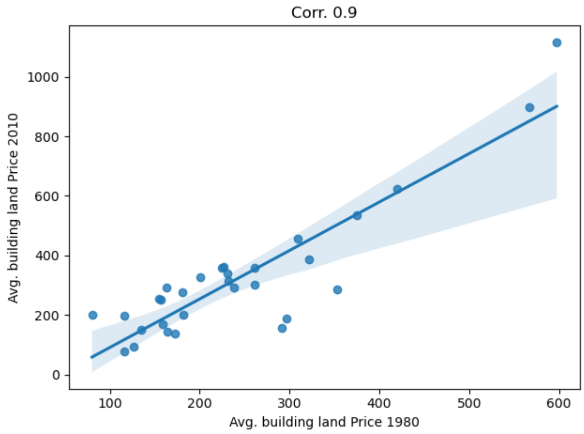

(b) 1980–2010

Notes: Each panel plots the building land price in the initial year on the horizontal axis and the price in 2010 on the vertical axis for all counties. The fitted line is OLS with slope reported in the panel. The prices after 1995 are from the Regional Statistics Database (Regionalstatistik), maintained jointly by the Federal Statistical Office of Germany (Destatis) and the statistical offices of the federal states. The 1980 prices come from the statistical Yearbook (Statistisches Jahrbuch) of the Federal Statistical Office of Germany (Destatis) and include only selected cities.

Figure 14: Local Wage Level and Persistence - Residualized

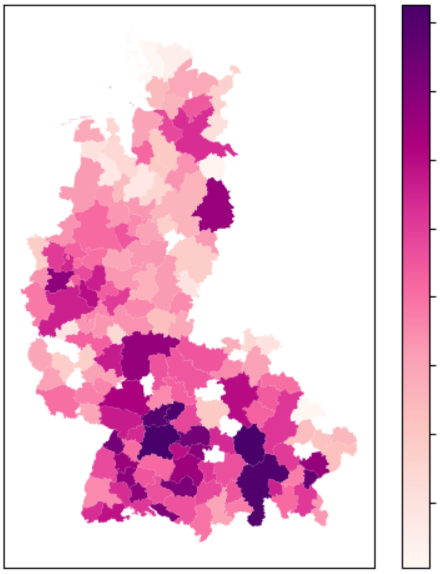

(a) Avg. Log Wage 2010

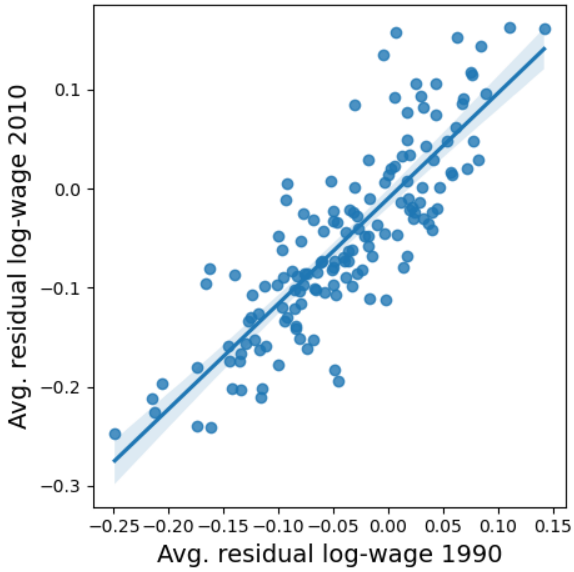

(b) Wage Level Persistence

Notes: The left panel maps commuting zones in West Germany, shaded by the average residual log daily wage in 2010. Residual wages are obtained from regressions of log daily wages on age and education fixed effects (six education levels) estimated using the SIAB. Darker shades indicate higher values. The right panel plots the average residual log daily wage in 2010 (vertical axis) against the average log daily wage in 1990 (horizontal axis) with an OLS fit. Each point represents one commuting zone. The sample is restricted to full-time workers.

B Additional Empirical Results

Figure 14 shows a map of average log wages that have been residualized by regressing on age and six levels of education in the left panel. The right panel shows the auto-correlation between 1990 and 2010 of the same average residuals.

Figure 15 shows the share of workers who never moved to any other commuting zone and the share of workers who never moved further away from their first commuting zone than to a neighboring commuting zone for workers who begin their careers in commuting zones of the lowest wage tercile. Commuting Zone average wages are used to calculate these terciles.

Figure 16 and Figure 17 show the results for the 1960-1964 cohorts in Figure 4 in the main text for the 1950-1954 and the 1955-1959 cohorts, respectively.

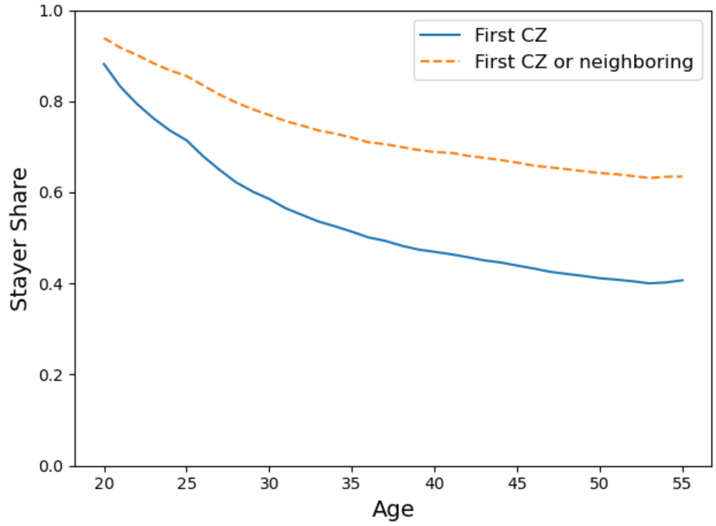

Figure 15: Fraction of Never-Movers

Notes: The sample has been reduced to workers who begin their career in commuting zones of the lowest income tercile. The solid line plots the share of workers in the SIAB sample who have never left their initial commuting zone, by age. The dashed line plots the share of workers who have either remained in their initial commuting zone or moved only to an adjacent commuting zone.

Figure 16: Log Wage Lifecycle Profiles

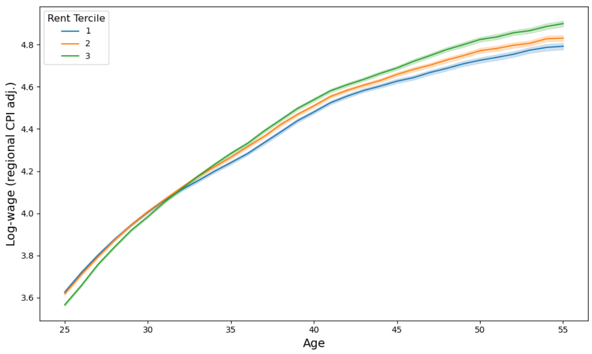

(a) Plain

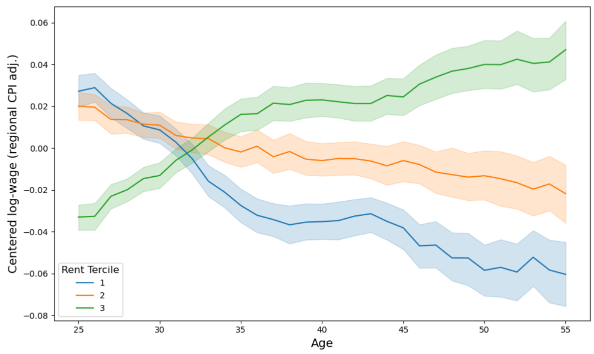

(b) Mean Subtracted

Notes: The left panel shows average log wages by age for workers who began their careers in counties belonging to the first, second, and third rent tercile, respectively. Rents are measured using INKAR data from 2010. The green line corresponds to the highest, the blue line to the bottom, and the orange line to the middle rent tercile. Wages are regionally CPI-adjusted to the level of Bonn using the BBSR CPI. The right panel shows the same wage profiles after subtracting the overall mean at each age. The confidence bands shown are at the 95% level. The sample includes workers born between 1950 and 1954.

Figure 18: Observed vs. Counterfactual Log Wage Profiles

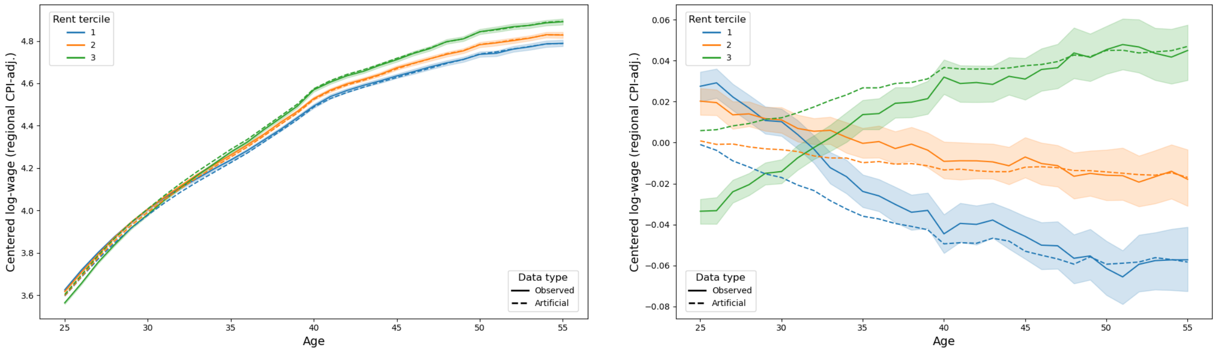

(a) Plain *(b) Mean Subtracted*

Notes: The solid lines show the average log wage profiles for workers in full-time employment who began their careers in counties belonging to the first, second, and third rent tercile, respectively. The dashed lines show counterfactual log wage profile that results from assigning every worker the national average wage, conditional on occupation and age. The green line corresponds to the highest, the blue line to the bottom, and the orange line to the middle rent tercile. Wages are regionally CPI-adjusted to the level of Bonn using the BBSR CPI. The right panel shows the same wage profiles after subtracting the overall mean at each age. Rents are measured using INKAR data from 2010. The sample includes workers born between 1950 and 1954.

Figure 17: Log Wage Lifecycle Profiles

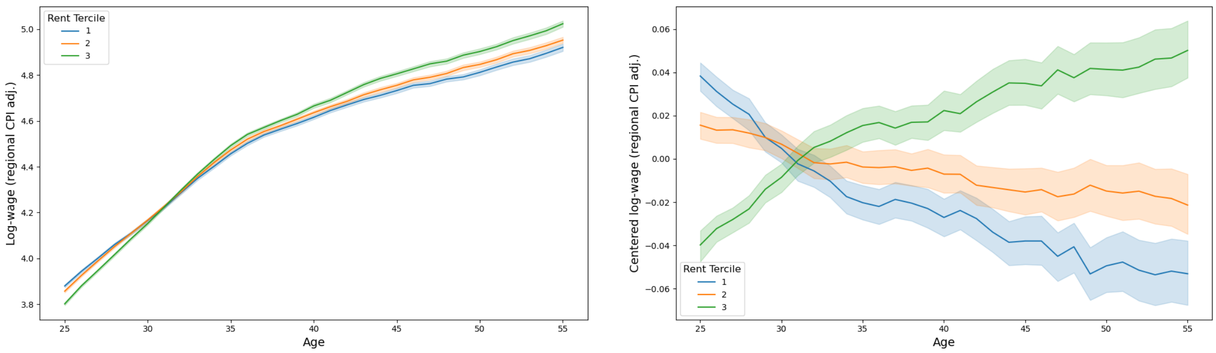

(a) Plain *(b) Mean Subtracted*

Notes: The left panel shows average log wages by age for workers who began their careers in counties belonging to the first, second, and third rent tercile, respectively. Rents are measured using INKAR data from 2010. Wages are regionally CPI-adjusted to the level of Bonn using the BBSR CPI. The right panel shows the same wage profiles after subtracting the overall mean at each age. The sample includes workers born between 1955 and 1959.

Figures 18 and 19 show the results reported in Figure 7 in the main text for workers born 1950-1954 and 1955-1959, respectively.

Figure 20 shows an interesting additional result. The fraction of workers who change

Figure 19: Observed vs. Counterfactual Log Wage Profiles

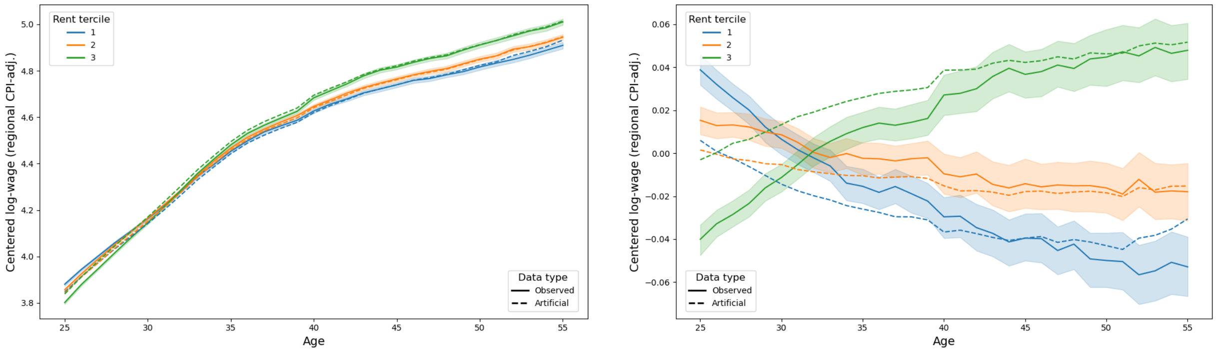

(a) Plain *(b) Mean Subtracted*

Notes: The solid lines show the average log wage profiles for workers in full-time employment who began their careers in counties belonging to the first, second, and third rent tercile, respectively. The dashed lines show counterfactual log wage profile that results from assigning every worker the national average wage, conditional on occupation and age. The green line corresponds to the highest, the blue line to the bottom, and the orange line to the middle rent tercile. Wages are regionally CPI-adjusted to the level of Bonn using the BBSR CPI. The right panel shows the same wage profiles after subtracting the overall mean at each age. Rents are measured using INKAR data from 2010. The sample includes workers born between 1955 and 1959.

their occupation when they move to another commuting zone is quite high - about 50%. This highlights the importance of regional occupational composition.

Figure 20: Occupation Change Shares Among Movers

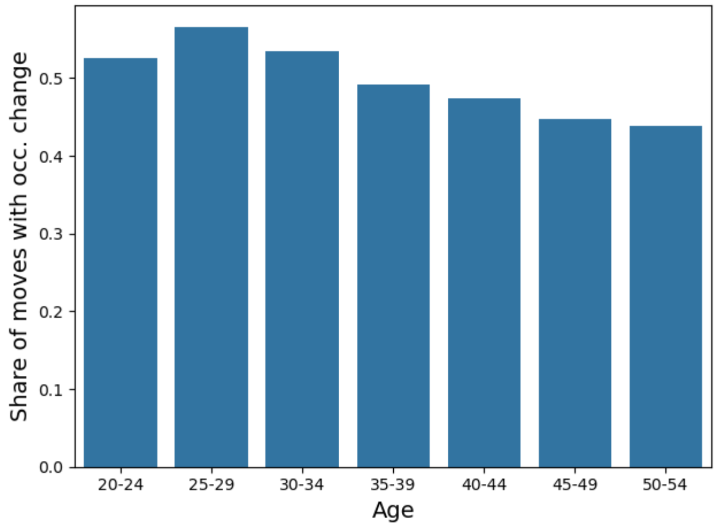

Notes: Fractions of workers who change their occupation among all workers who move to a different commuting zone by age group. Sample pools birth years 1950-1964.

C Alternative Rent Data (Wohngeld)

Figure 21 shows the same results shown in the right panels of Figures 4, 16, and 17 but with a different way of splitting the sample. In Germany, low-income households are eligible for a rent subsidy. Crucially, it depends on the local rent price of non-expensive housing. I obtained these categories from the Federal Statistics Office with the caveat that some counties are missing, leading me to use the INKAR data for the main specifications shown in the main text. To this end, all German counties are grouped into 6 categories. I group counties into three bins of two categories each, and reproduce the log-wage profiles I obtained using rent terciles based on the INKAR rent index. The obtained results are extremely similar, lending confidence that the results are indeed correct.

Figure 21: Log Wage Lifecycle Profiles - Rent Subsidy Category

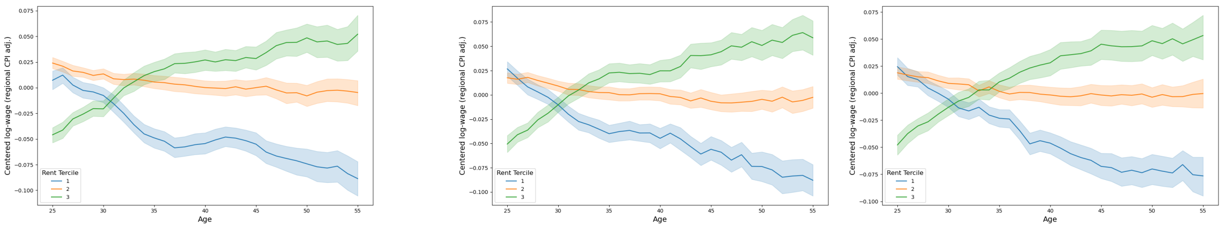

(a) 1950-1954 *(b) 1955-1959* *(c) 1960-1964*

Notes: Log wages by age for workers who began their careers in counties belonging to the three rent bins. Rents are measured using the categories of the subsidiy scheme *Wohngeld* that groups counties into 6 bins ranging from cheap to expensive. The green line corresponds to the highest, the blue line to the bottom, and the orange line to the middle 2 rent bins. Wages are regionally CPI-adjusted to the level of Bonn using the BBSR CPI. All panels show the same wage profiles after subtracting the overall mean at each age. The confidence bands shown are at the 95% level. Each panel shows results for different 5-year cohorts.

Figure 22 reproduces the counterfactual occupation-weighted wage profiles shown in Figure 7 in the main text. Results are extremely similar again.

Figure 22: Observed vs. Counterfactual Log Wage Profiles - Rent Subsidy Category

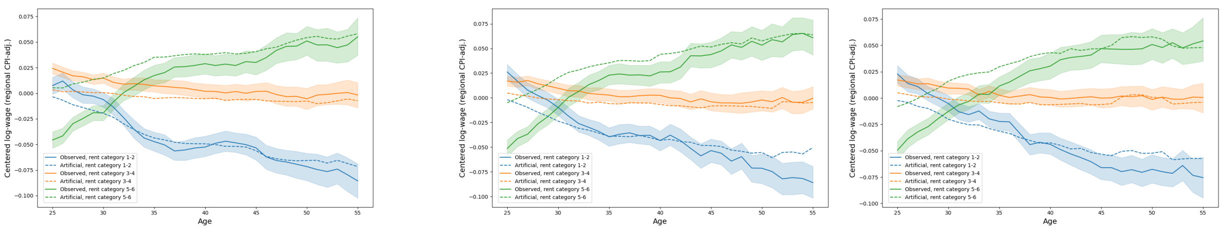

(a) 1950-1954 *(b) 1955-1959* *(c) 1960-1964*

Notes: Log wages by age for workers who began their careers in counties belonging to three rent bins. Rents are measured using the categories of the subsidiy scheme *Wohngeld* that groups counties into 6 bins ranging from cheap to expensive. The dashed lines show counterfactual log wage profile that results from assigning every worker the national average wage, conditional on occupation and age. The green line corresponds to the highest, the blue line to the bottom, and the orange line to the middle 2 rent bins. Wages are regionally CPI-adjusted to the level of Bonn using the BBSR CPI. All panels show the same wage profiles after subtracting the overall mean at each age. The confidence bands shown are at the 95% level. Each panel shows results for different 5-year cohorts.

D Solving the Model

D.1 Solving the Model with EGM

The worker problem is solved by means of backwards induction. This can often be done efficiently using an endogenous grid method (EGM) as in Carroll (2006). However, since the location choice in the model is discrete, I cannot use the standard EGM. Instead, I rely on the endogenous grid method for discrete-continuous dynamic choice models proposed by Iskhakov et al. (2017).

In the last period of the lifecycle, first-order conditions can be used to analytically solve for the worker decisions. In the periods between the working phase and the final period, the inverse Euler equation of retired workers is

$$c = \left[R\beta c'^{-\lambda} \left(\frac{\nu'}{\nu} \right)^{\lambda} \right]^{-1/\lambda}$$

where primes indicate state in the next period. Worker decisions can then be found using standard EGM methods.

During the working phase, the discrete choice needs to be handled. Every period, for every possible state, the worker's problem is solved twice, taking the location choice (i.e., z') as given. Using the worker's inverse Euler equation

$$c = \Bigg[R\beta \int \Bigg(\sum_{h'_a \in \mathcal{H}} \sum_{h'_b \in \mathcal{H}} \sum_{z'' \in \{A,B\}} P(z'' \mid \psi', h'_a, h'_b, z', \psi')\, P(h'_a \mid h_a)\, P(h'_b \mid h_b) \\ c'(z'', h'_a, h'_b, \psi')^{-\lambda} \left(\tfrac{p_z}{p_{z'}} \right)^{(1-\gamma)(1-\lambda)} \left(\frac{\nu'}{\nu} \right)^{\lambda} \Bigg) \mathrm{d}F(\psi' \mid \psi) \Bigg]^{-1/\lambda}.$$

and given a grid of consumption in the next period, implied through next period's already determined consumption policy function by an exogenous grid of assets in the next period, an endogenous grid of consumption in the current period can be calculated. Note that variables with a prime denote state variables in the next period, and z'' denotes the location choice in the period after the next period. Given the functional form of u_j, grid points corresponding to the endogenous contemporary consumption grid of contemporary housing consumption η and subsequently the saving decision, as well as implied conditional value functions V_A and V_B, can be calculated. The borrowing constraint is handled by adding a point with $a = 0$ and implied consumption decisions into the respective grids where workers would like to borrow if they could.

Given that the taste shocks ζ_j are *iid* and $ETV1$ distributed, the probabilities of choosing A or B can be calculated in closed form at every grid point with the famous logit probability formula (Domencich and McFadden (1975))[15]. Policy functions are then interpolated linearly and passed on to the period before the contemporary period.

In the 0-period after birth, the same logit formula can be applied to obtain choice probabilities for each region, taking initial productivity and endowment as given.

This backward induction procedure yields value functions and workers' decisions for given prices $\overline{w}_A, \overline{w}_B, p_A$ and p_B. Since local housing stocks are fixed and $\overline{w}_A$ and $\overline{w}_B$ are closed form function of $\frac{L_A}{L_B}$ solving the model then is a matter of finding $\frac{L_A}{L_B}$ such that the implied wages are profit maximizing for the aggregation firm and result from the workers decision problem given these same wages and rent prices that clear rent markets at the same time.

Aggregate variables $L_A, L_B, \eta_A^d, \eta_B^d$, where the latter two denote local aggregate housing demand, are obtained by simulating 100,000 individuals using the calculated worker decisions.

The fixed-point problem of finding $\frac{L_A}{L_B}, p_A$ and p_B is then solved by updating initial guesses in nested loops. $\frac{L_A}{L_B}$ is updated using Anderson acceleration, and the rent prices are updated with a log-linear updating rule.

D.2 Solving the Model under Downsizing

Under the policy discussed in Section G.2, EGM is no longer applicable for retired workers in region B because there can be corner solutions. Workers who would otherwise have consumed slightly more than $71.02m^2$ may find it optimal to consume exactly the threshold. Similarly, with the transfer, workers who would have rented housing just below this threshold may now find the threshold optimal. I therefore use a standard grid-search method where for every asset grid point and every level of consumption on a grid, the value of choosing freely above the threshold, below the threshold, and the corner, as well as the resulting savings decision solution, are compared. Policy functions are then interpolated on the same asset grid used in the EGM method described in the previous section.

[15] As discussed by Iskhakov et al. (2017), this procedure can produce non-monotonous points in the conditional value grids that need to be eliminated through an upper envelope.

For non-retired workers and workers retired in A, the same EGM algorithm as before is employed.

E Calibration

E.1 Male Income Shares

To document how the within-household income distribution evolves with age, I use the German Socio-Economic Panel (SOEP) to construct the age profile of the male income share shown in Figure 23. I merge the individual-level labor income records with demographic and household identifiers. All monetary amounts are deflated using the German consumer price index from Jordà et al. (2017).

For each household and survey year, I compute individual ages as the difference between survey year and year of birth, dropping implausible values where the implied age is negative. To identify the household head, I restrict attention to males with positive earnings and select, within each household, the oldest such male. As in my main empirical analysis, I restrict the analysis to households whose head was born between 1950 and 1964. I then restrict the sample to households with exactly one adult male earner and a head aged 60 or below.

For each remaining household-year, I aggregate real wages by gender and compute the within-household share of male earnings.

Figure 23 shows the resulting profile. The male income share is close to 100 percent for younger heads and gradually declines with age as female earnings and joint retirement patterns become more prevalent. Only households with a single working-age male are included, so that the pattern captures variation in the relative contribution of female partners rather than compositional shifts in household structure.

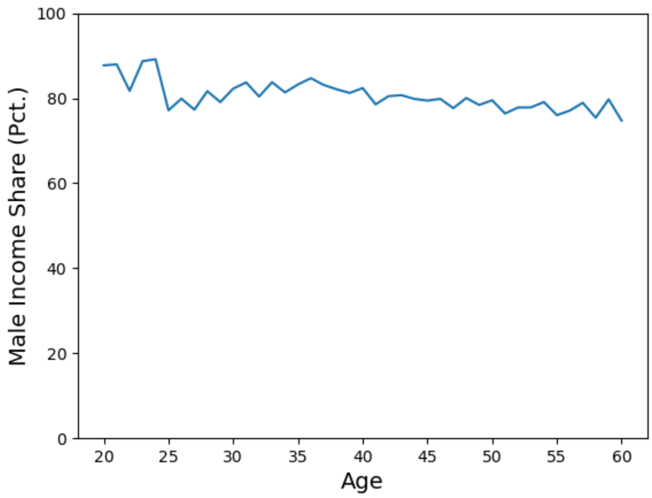

Figure 23: Average Male Income Share by Age of Household Head

Notes: The figure plots the average share of male earnings in total household labor income by the age of the male household head. The sample is restricted to SOEP households with exactly one adult male earner born between 1950 and 1964 and with positive household labor income.

E.2 Household Size Weights

Household consumption is scaled using the modified OECD equivalence scale. To construct the underlying household information, I use the German Socio-Economic Panel (SOEP). I merge the individual-level labor income records with household identifiers and survey years and adjust wages using the consumer price index to obtain real values. For each household, I identify the male head as the oldest working-age male with positive earnings, and restrict the sample to cohorts born between 1940 and 1964.

For these households, I calculate the number of adults and children and then assign an equivalence weight following the OECD method, where additional adults and children increase household needs by fixed fractions relative to the first adult. I then average these weights by the age of the household head, which provides an age profile of equivalence scales.

Because the raw averages are somewhat irregular, I fit a smooth polynomial curve to the values between ages 25 and 70. For ages below 25 and above 70, the curve is extended by holding the fitted values constant at the boundaries. The resulting schedule gives a smooth and age-dependent measure of household equivalence weights that I feed into the model. Figure 24 shows the resulting weights.

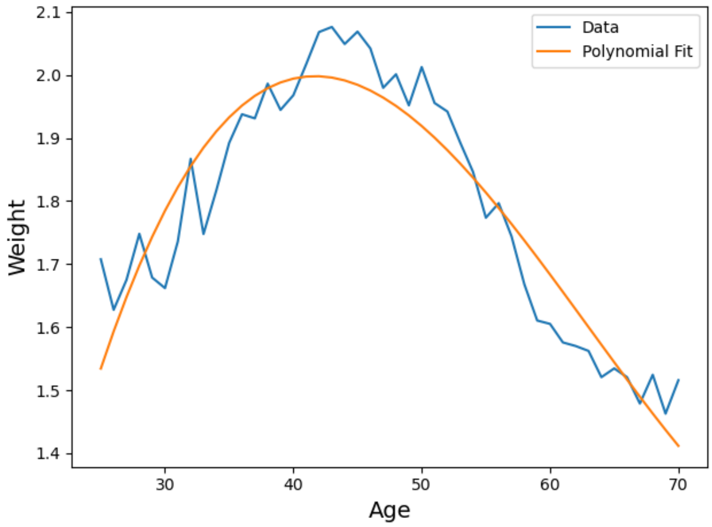

Figure 24: Modified OECD Equivalence Weights

Notes: The blue line shows the modified OECD equivalent weights of households in the SOEP with exactly one male adult born between 1940 and 1964 against age. The orange line is the cubic polynomial fit that is fed into the model.

E.3 Utility Moving Cost

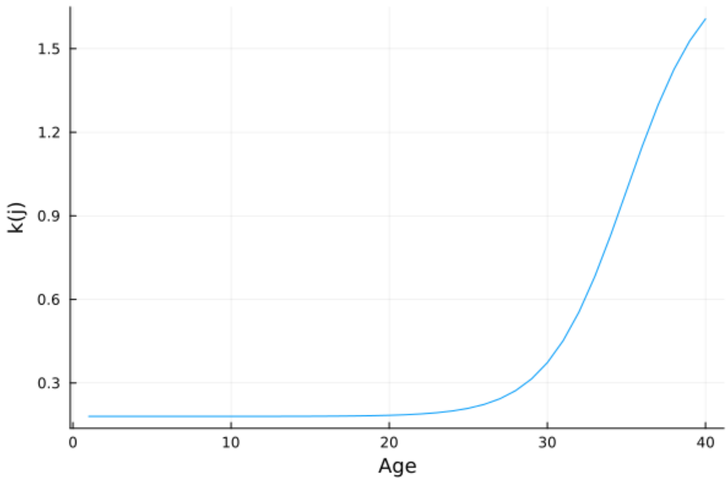

Figure 25: Utility Moving Costs

Notes: Calibrated values of utility moving costs $k(j)$ by age.

The functional form of $k(j)$ is set to

$$k(j) = k\Big[1 + (S-1)\,\sigma\big(2\beta\,(j-\mu)\big)\Big], \qquad s(j) = \sigma\big(2\beta\,(j-\mu)\big).$$

where $\mu = 35$, $\beta = 0.2$ and $S = 10$. The resulting calibrated values are shown in Figure 25.

F Additional Model Results

Figure 26 shows the distribution of predicted and model counterfactual wage changes from working in the other region discussed in Section 5.2.

Figure 26: Predicted Wage Changes

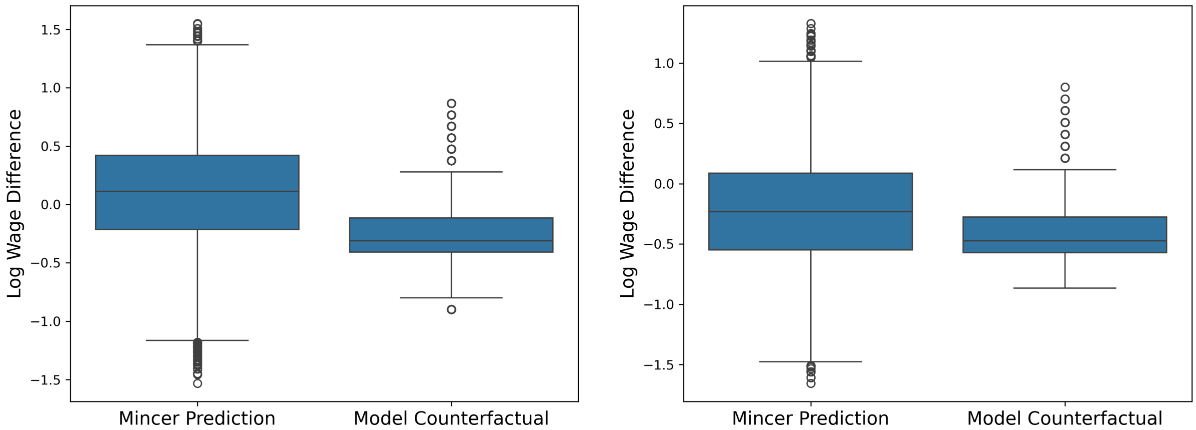

(a) Workers in Region A *(b) Workers in Region B*

Notes: Distribution of log wage differences between observed wages and model-based predictions. The left boxplot in each panel refers to the Mincer regression, while the right boxplot refers to the model counterfactual. In each boxplot, the box represents the interquartile range, the line inside the box the median, and the whiskers extend to 1.5 times the interquartile range. Points beyond the whiskers indicate outliers.

G Additional Policy Results

G.1 CEV of Older Workers

The construction policy presented in Section 6 is beneficial for newborn workers as shown in the main text. However, it is beneficial for older workers too.

At age 45, CEV for workers in region A with median productivity, skill levels, and savings from the construction policy is 0.61%. The number is smaller than the 0.79% CEV for newborn workers. This is because workers at 45 have no more incentives to move to region B because of human capital mismatch and no longer benefit from greater moving opportunities. They do still benefit from the wage increase and the rent price decrease in region A.

By contrast, at 0.84%, median 45-year-old workers in region B have about the same CEV as newborn workers. This is because they stille directly benefit from the rent price reduction.

G.2 Downsizing Policy

Housing construction effectively increases welfare and makes moving to the region with better career opportunities easier, as discussed in the previous section. An alternative policy that is often discussed in public debates is incentivizing households to downsize their housing consumption in order to make room for more people. In the German state of Baden-Württemberg, for example, municipalities can receive a bonus payment of €3000 per household that moves to a flat that is at least 15 square meters smaller than the previous flat. The municipalities may then choose to pay this bonus out to the household directly.

I implement a stylized version of this policy for comparison with the construction project. In the model, there are two reasons that lead to high rents in region B. One is small housing supply as targeted by the construction policy. The other is old workers having accumulated a lot of wealth and consuming large quantities of housing. In the model, in region B, workers under the age of 30 rent a median flat of $58.7m^2$, whereas the median flat of retired workers is as large as $87.05m^2$.

I test the following policy. Workers who are retired and choose to live in a flat that is $15m^2$ smaller than the pre-policy median (i.e. $72.05m^2$) receive a lump sum transfer of

Table 9: CEV after Downsizing Policy

	Veil of ignorance	Conditional on A	Conditional on B
CEV	-0.19%	-0.19%	-0.19%

Notes: Consumption equivalent variation of newborn workers with respect to a counterfactual economy where the government pays lump sum transfers of €250 to retired workers who take up less than $72.05m^2$ of housing.

€250. The net present value of receiving this transfer for every year of retirement is close to €3000. The technical details of the implementation are discussed in Appendix D.2.

The policy is not effective at all. The average housing consumption of retired workers in region B is reduced by only about $0.5m^2$. As a result, none of the general-equilibrium objects change measurably. Table 9 shows the welfare effects on newborn workers. The effect is small and negative - the policy has no meaningful impact on aggregates but still leads to tax increases.

Compared to housing construction, the downsizing policy has two main disadvantages. First, unless incentives are huge (and therefore extremely expensive), it only affects retired workers at the margin, and the housing space that is freed up is therefore limited. Second, houses that are constructed once can be used by subsequent generations, warranting financing the construction through government debt. Incentivizing old households to live in smaller flats needs to be paid for continuously and is therefore expensive.