ECONtribute
MARKETS & PUBLIC POLICY

ECONtribute
Discussion Paper No. 167

Occupation Growth, Skill Prices, and Wage Inequality

Michael Böhm   Hans-Martin von Gaudecker   Felix Schran

May 2022   www.econtribute.de

Funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under
Germany´s Excellence Strategy – EXC 2126/1– 390838866 is gratefully acknowledged.

Occupation Growth, Skill Prices, and Wage Inequality

Michael J. Böhm, Hans-Martin von Gaudecker, and Felix Schran*

May 2022

Abstract

This paper studies the relationship between occupational employment, occupational wages, and rising wage inequality. We document that in all occupations, entrants and leavers earn less than stayers. This suggests selection effects that are negative for growing occupations and positive for shrinking ones. We estimate a model of occupational prices and skills, which includes occupation-specific skill accumulation and endogenous switching across many occupations. Consistent with leading explanations for occupational changes, estimated prices (i.e., selection-corrected wages) and occupational employment growth are positively related. Just over 40% of selection is due to age in the sense that marginal workers have had less time to accumulate skills. The remainder is due to Roy-type selection, i.e., workers reacting to changing prices and shocks unrelated to age. Skill prices establish a long-suspected quantitative connection between occupational changes and the surge in wage inequality.

Keywords: Skill Prices, Selection Effects, Multidimensional Skill Accumulation, Occupational Employment and Wages, Administrative Panel Data, Wage Inequality

JEL codes: J21, J23, J24, J31

*Böhm: University of Bonn and IZA, michael.j.boehm@uni-bonn.de. Gaudecker: University of Bonn and IZA, hmgaudecker@uni-bonn.de. Schran: AXA We would like to thank Matias Cortes, Christian Dustmann, Georg Graetz, David Green, Thomas Lemieux, Uta Schönberg and Shintaro Yamaguchi as well as seminar and conference participants at U-Alberta, U-Bonn, U-British Columbia, UCL, U-Kiel, ESPE 2016, ZEW Mannheim, U-Michigan, EIEF-Rome, USI-Lugano, FU Berlin, SOLE 2017, Canadian Labour Economics Forum 2017, ESEM 2018, the IZA World Labor Conference 2018, the 2019 workshop of DFG Priority Program 1764, the EES Conference at the Swedish Central Bank 2019, and the 4th CRC-TR 224 conference for very valuable comments. Böhm acknowledges support from a Research Fellowship of the Deutsche Forschungsgemeinschaft (DFG, BO 4765/1-2). Gaudecker and Schran gratefully acknowledge financial support by the DFG (GA 1809/2-1, AOBJ: 617258). Gaudecker further acknowledges funding by the DFG under Germany's Excellence Strategy – EXC 2126/1– 390838866 and through CRC TR 224 (Project C01). This study uses the factually anonymous Sample of Integrated Labor Market Biographies (version 7514). Data access was provided via a Scientific Use File supplied by the Research Data Centre (FDZ) of the German Federal Employment Agency (BA) at the Institute for Employment Research (IAB) under contract Number 101357. The Online Appendix is available at https://www.wiwi.uni-bonn.de/gaudecker.

1 Introduction

During the past decades, occupational employment has changed profoundly across Europe and the United States. A burgeoning literature has established fundamental shifts in labor demand as the most important cause of these changes (Autor et al., 2003; Goos and Manning, 2007; Acemoglu and Autor, 2011; Goos et al., 2014, we discuss the literature in detail in the next section). Yet, it remains puzzling that these demand shifts fail to be clearly reflected in occupational wages or wage inequality. First, occupational employment growth has been decoupled from occupational wage growth (Goos and Manning, 2007; Mishel et al., 2013; Green and Sand, 2015; Roys and Taber, 2019; Hsieh et al., 2019). Second, a number of studies suggest that changes in occupation-specific labor demand contributed little to the increase in wage inequality in recent decades (Dustmann et al., 2009; Card et al., 2013; Firpo et al., 2013; Autor, 2019).

Making use of individual-level longitudinal information on wages and occupations from German administrative data, we document three important facts. First, individual workers' wage growth is substantially faster *within* expanding occupations. During our study period 1985–2010, annual individual wage growth is one percentage point higher on average in occupations that double in size compared to those where size is constant. Second, workers who enter an occupation earn 21 percent less than incumbents on average. Similarly, average wages of workers leaving occupations are 16 percent lower than stayers' wages. Third, both of these effects are increasing in net occupation growth. For a doubling in occupation size, the wage gaps between entrants vs incumbents and leavers vs stayers widen by 6 and 3 percentage points, respectively. The raw data thus reveal that net growth of an occupation will attenuate average wages within occupations where selection operates in both directions. This explains why wage and employment growth across occupations are uncorrelated in our data.

In order to interpret these facts, we set up a parsimonious model of occupation choice. Workers have multidimensional occupation-specific skills that evolve heterogeneously across occupations and over the career. The central distinction we make is between *skill prices*—wages paid per constant unit of skill—and average occupational wages. While skill prices are directly affected by shifts in demand, selection may dampen, neutralize, or even overturn their effects on average wages. Less skilled individuals enter growing occupations, which depresses these occupations' average wages. Shrinking occupations also retain the more skilled parts of their workforce, lifting their average wages. Such selection effects imply that between-occupation inequality will underestimate the impact of shifting occupational demand on wage inequality.

We estimate the model using the longitudinal information in our data. Our analysis uncovers three main findings. First, there is a clear positive relation between skill price and employment growth at the level of detailed occupations. This indicates that demand shifts were the dominant drivers of both occupational employment and *skill-constant* wages over the past decades. Characterizing occupations by their task intensities, we find that the patterns are in line with routine-biased technical change (RBTC) as one of the important drivers of occupational demand.[1] More generally, the patterns are consistent with polarization, since employment and skill prices of broad occupation groups with high as well as low wages rose compared to mid-wage occupation groups.

Since employment growth is uncorrelated with average wage growth and positively correlated with skill price growth, skills must deteriorate in growing occupations and increase in shrinking occupations. Our second main finding is that lower-earning workers' net entry into growing occupations and their net exit out of shrinking occupations fully account for the negative relationship between skill changes and employment growth. We term this the *employment growth-selection effect* or simply *growth-selection*. Quantitatively, a doubling of an occupation's workforce translates into a decrease in its average skills by 12 percent. Viewed through the lens of our model, growth-selection stems from both entrants and leavers possessing lower average skills than stayers in any occupation. We exploit the longitudinal dimension of the data to document the nature of growth-selection. One component is that growing occupations draw in and shed younger workers who are equally

[1] Note that this paper does not measure occupational demand or supply shocks directly. We instead infer from the co-movements of quantities and prices that these are consistent with demand shocks. Forces of occupational demand may include RBTC and related technical changes (e.g., Autor et al., 2003), international trade and offshoring (Autor et al., 2013; Goos et al., 2014), transformation of the industry structure (Bárány and Siegel, 2018), changes in consumption patterns (Autor and Dorn, 2013; Mazzolari and Ragusa, 2013), social skills content (Deming, 2017), among others.

skilled as inframarginal workers, but who have had less time to accumulate skills due to their age. We distinguish this *age selection* from *Roy selection*, which refers to differences in skills conditional on experience. Due to the fundamental selection problem, we focus on the current spell in a particular occupation. Aggregating across occupation groups, we find that 42% of employment growth-selection can be attributed to *age selection*.

Our third finding is that changes in occupational skill prices have driven much of the increase in wage inequality over the past decades. Inequality between occupations would have been 70 percent higher at the end of our sample if growth-selection had not counteracted the changing skill prices. We break down the quantiles of the wage distribution using our estimates of skill prices and skills along with workers' occupational histories. Our model tracks the wage distribution well. Demographic composition and faster skill accumulation at the median can account for 70% of the rise in the $50-15$ percentile wage difference. Controlling for other factors, skill prices alone explain 60% of the increase in the $85-50$ percentile difference. These results reconcile the roles of demand shifts and of previous explanations brought forward for the evolution of wage inequality in Germany (e.g., Dustmann et al., 2009) in quantitative terms.

This paper is structured as follows. Next, we describe the data, relate to prior literature, and present our new stylized facts. The third section explains the model and its estimation. Section 4 presents the results on the evolution of skill prices, dissects the growth-selection effect, and reports on extensive robustness checks. In Section 5, we study the impact of skill prices and selection on rising wage inequality. The last section discusses the relationship between our findings and the evolution of labor market institutions.

2 Data, Literature, and Stylized Facts

We use the Sample of Integrated Labor Market Biographies (SIAB) provided by the IAB Institute at the German Federal Employment Agency (Ganzer et al., 2017). The SIAB is a 2% random sample of administrative social security records from 1975 to 2014. It is representative of 80% of the German workforce and includes employees covered by social security, marginal part-time workers, benefit recipients, individuals officially registered as job-seeking, and those participating in active labor market programs. The SIAB excludes the self-employed, civil servants, and individuals performing military service. Most notably, it contains individuals' full employment histories including detailed data on daily wages, industries, and occupations along with socio-demographics such as age, gender, or the level of education. The data is exact to the day as employers need to notify the employment agency upon changes to the employment relationship.

In order to work with a homogeneous sample throughout, we restrict the main sample to German men aged 25 to 54 years who are working full-time (excluding apprentices) in West Germany. See Appendix A.1 for the reasons behind these choices and for details on the wider dataset construction. We will relax all of these restrictions in robustness checks. We transform the spell structure into a yearly panel by using the longest spell in any given year, adjusting wages appropriately for spells that do not last the entire year. Due to a cap on social security contributions, 12% of wages are right-censored at this ceiling; we follow imputation procedures in Dustmann et al. (2009) and Card et al. (2013). We inflate all wages to 2010 prices using the German consumer price index.

A key strength of the SIAB data is that it provides high-quality longitudinal information on workers' occupations. Until 2010, the SIAB Scientific Use File contains a consistent set of 120 occupations; we cannot use subsequent years because the classification changes drastically thereafter. Most of our analyses will be based on the raw 120 occupations. To ease interpretation, we also aggregate them into broader groups following Acemoglu and Autor (2011) and others. These comprise managers, professionals, and technicians (Mgr-Prof-Tech); sales and office workers (Sales-Office); production workers, operators and craftsmen (Prod-Op-Crafts); and workers in services and care occupations (Srvc-Care). See Table A.1 for the mapping of detailed occupations into these groups.

2.1 Wage Inequality and Changes in Occupational Employment

Two of the most important trends in developed countries' labor markets over the past decades have been a strong increase in wage inequality and a substantial reallocation of

Figure 1: Evolution of wage inequality and occupational employment

(a) Wage percentiles (b) Occupations' employment

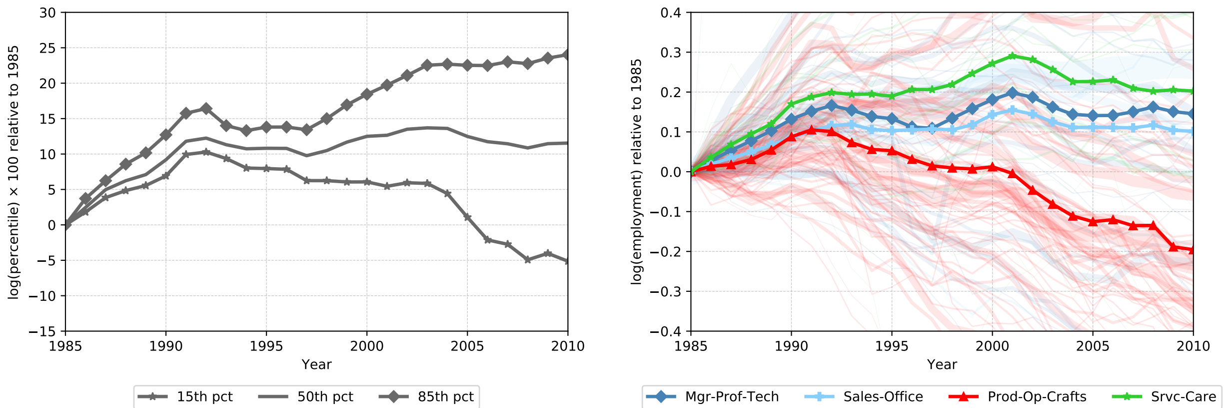

Notes: The vertical axis in Panel 1a shows the 15th, 50th and 85th log wage percentile over time relative to 1985. The vertical axis in Panel 1b shows the log change in the number of employed workers within an occupation over time. Shaded lines in the background represent the 120 detailed occupations in the SIAB SUF. The four groups show an aggregation of these detailed occupations as described in Appendix Table A.1. The thickness of a shaded background line corresponds to the number of employed workers in an occupation averaged across years 1985 until 2010.

employment across occupations broadly characterized by polarization (for a summary see Acemoglu and Autor, 2011). As documented by, e.g., Spitz-Oener (2006), Dustmann et al. (2009), Card et al. (2013), and Goos et al. (2014), Germany is no exception to either phenomenon. Figure 1 reproduces both trends in our dataset.

Figure 1a shows the trends of wage percentiles over the 1985–2010 period normalized to zero in 1985. Inequality increased strongly and steadily both in the upper half, measured by the difference between the 85$^{\text{th}}$ and the 50$^{\text{th}}$ percentile of log wages, and in the lower half ($50 - 15$ difference). Figure 1b plots the trends in the logs of the detailed 120 occupations' employment (shaded lines) and the four aggregated groups (bold lines with markers). Employment in Prod-Op-Crafts occupations declined by more than 20 log points from a baseline share of over 60 percent, whereas the employment share of the other occupation groups increased. This trend has been termed "job" or "employment polarization" because Prod-Op-Crafts workers tend to be located in the middle of the occupational wage distribution (Goos and Manning, 2007). Part of the declining employment in middle-paying occupations appears to be due to changes in technology (affecting codifiable routine-type jobs, see e.g., Autor et al., 2003) as well as international trade and offshoring (affecting manufacturing-type jobs, e.g., Autor et al., 2013).

One may expect to see such shifts of the demand for different types of occupations directly in the wage distribution, not least because the wage and employment trends occurred largely in parallel (Figure 1).[2] There exists, however, surprisingly little quantitative evidence on the role of occupational change for the evolution of wage inequality: holding occupations' wages fixed at their initial levels and reweighting them with employment in subsequent decades, Goos and Manning (2007) show that composition effects due to employment polarization can account for a substantial part of changing wage inequality in the U.K. Autor (2019) finds that in the U.S. a similar exercise explains only small shares of the income growth differentials across five education groups. Further accounting for the degree of urbanization comes close to matching the evolution of real wages of the non-college educated. For the German case, Dustmann et al. (2009) conclude that the rise of lower-half inequality was unlikely to stem from changes in demand. Card et al. (2013) run a set of Mincer regressions and incrementally add occupational identifiers, finding that the role of the latter for rising wage inequality is rather small.

Figure 2a hints at why these analyses tend to have limited explanatory power, plotting changes in employment against changes of average wages over the 1985–2010 period for each occupation. Variation along the horizontal axis shows that employment changes are very substantial. Many occupations grew or shrank by more than 50 log points. Yet, movements of average wages are surprisingly small given the variation in occupation growth and the large increase of inequality. Thus, between-occupation decompositions—such as wage regressions with occupation dummies or reweighting strategies—may attribute little of the trends in wage inequality to factors like changing skill prices and employment structure, and much of its increase to unexplained within-occupation inequality. In fact, the regression line shows that employment and wage changes are uncorrelated. To pick

[2] Both panels of Figure 1 use the same sample. Katz and Murphy (1992) argue for using a sample of workers with high labor market attachment for anything related to wages and a more comprehensive sample for employment. The reason is that compositional shifts in occupational employment may mistakenly be taken as employment growth when using the high-attachment sample only. For example, employment of kindergarten teachers tripled among prime-aged West-German men in our data. In a sample including women and foreigners, employment "only" increased by 57%. Appendix D.1 thus repeats all our analyses where this is possible separating between "wage" and "count" samples, noting important differences along the way. In the main text, we stick to the high-attachment sample in order to avoid switching back and forth depending on whether we need individual-level results along both dimensions or not.

Figure 2: Correlation of changes in employment, average wages, and wage growth

(a) Average wage growth

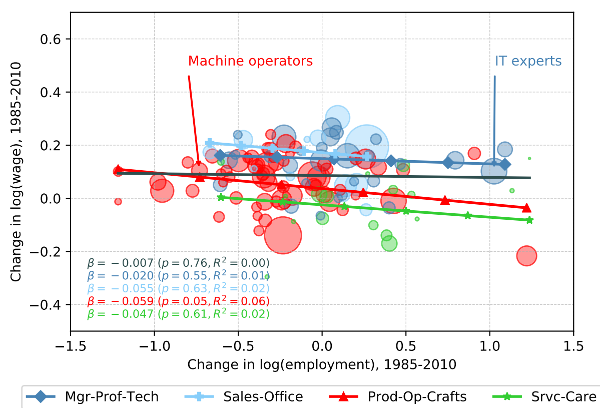

(b) Stayers' wage growth

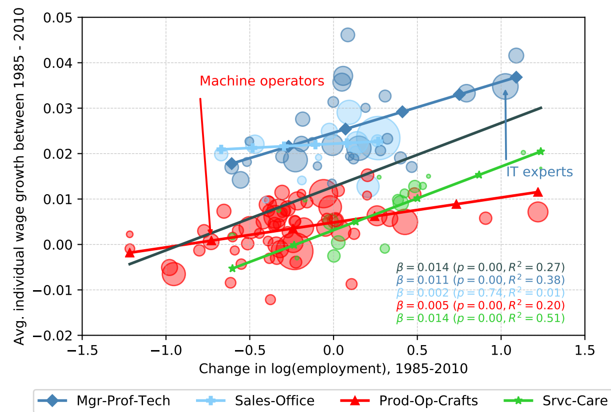

Notes: The vertical axis in Panel 2a shows the change in average wages between 1985 and 2010. The vertical axis in Panel 2b depicts individual wage growth averaged across years 1985 until 2010. The horizontal axis in both panels shows the change of the log number of employed workers within an occupation between 1985 and 2010. One bubble represents one of the 120 detailed occupations in the SIAB SUF. The four groups show an aggregation of these detailed occupations as described in Appendix Table A.1. Bubble size corresponds to the number of employed workers in an occupation averaged across years 1985 until 2010. Regression lines across all occupations (black) and within the four broad groups (colored) are weighted by the number of employed workers.

the two highlighted examples, IT experts' employment rose by 102 log points or 178 %. Their average wages grew by 10 %, just above the overall average. Machine operators—a prototypical occupation to be hit by routine-biased technical change—shrank by 73 log points or 51 %. Yet, average wages grew by the same amount as those of IT experts.

Within the broader groups, the non-correlation between wage and employment growth even turns negative for the lower-earning Prod-Op-Crafts and Srvc-Care occupations. This is consistent with the regressions reported by Dustmann et al. (2009, Section IV.D). They conclude that demand shifts were unlikely to have driven lower-end inequality. Finding little or negative correlation between occupational wage and employment growth is not confined to Germany. Hsieh et al. (2019) and Roys and Taber (2019) document very small correlations between the growth rates of occupational employment and wages in the U.S. Employment in low-skill occupations increased in the U.K. and Canada, while at the same time wages dropped compared to routine occupations (Goos and Manning, 2007; Green and Sand, 2015). Next to the role that occupations have to play for wage inequality, this begets the more fundamental question of whether, on aggregate, shifts in demand versus supply of labor to different occupations were the dominant factor for the changes of the employment structure.[3] We will find that, while the latter may have a role to play, the data strongly suggest that demand changes along the lines of routine-biased technological change or international trade are important.

2.2 Individual-Level Wage Growth and Selection

As a first pass, Figure 2b shows that there is a strong positive correlation between employment and *individual-level* wage growth. The horizontal axis is the same as in Panel a whereas the vertical axis plots the average annual wage growth of workers who stayed in their occupation for any two consecutive years. Wage growth rates within occupations clearly line up with their employment growth. We will control for other important factors, such as occupation- and age-specific returns to experience, below. But the juxtaposition of Figures 2a and 2b already suggests that the key reason for the differences will be selection into occupations. Put differently, demand shifts will indeed be driving the changes of employment and prices paid for skilled labor across occupations. Negative selection of entrants into growing occupations obscures this relation when looking at average occupation-specific wages. The underlying occupational prices would be spreading out more than the average occupational wages, which are captured in the above-discussed decomposition analyses.

Data limitations have prevented a more thorough analysis of such selection effects. In particular, the main sources in the U.S. are repeated cross-sections (CPS, Census) or longitudinal data too small in size for investigating individual-level dynamics across detailed occupations (PSID, NLSY). The SIAB data allow us to track occupational biographies over the entire career. Figure 3 gives more direct evidence on the importance of selection effects by plotting employment changes against the wage differentials between marginal workers who switch and inframarginal workers who stay in their occupations.[4]

[3]E.g., Glitz and Wissmann (2021) argue that a declining supply of medium versus low-skilled young workers in Germany was responsible for part of the rising lower-end inequality depicted in Figure 1a.

[4]McLaughlin and Bils (2001) perform a related exercise with a coarser set of industry sectors in the PSID data, reporting comparable results on wage differences.

Figure 3: Selection into and out of occupations

(a) Entrants' minus incumbents' wages

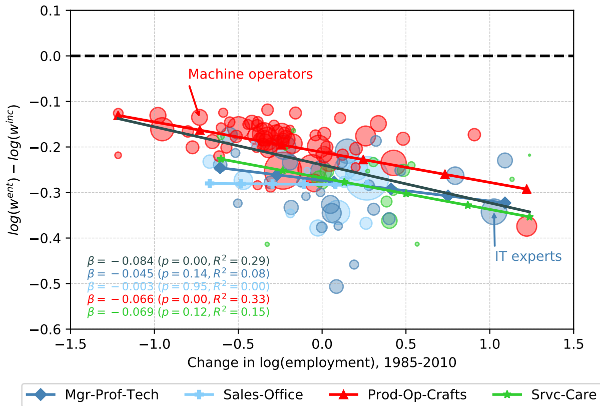

(b) Leavers' minus stayers' wages

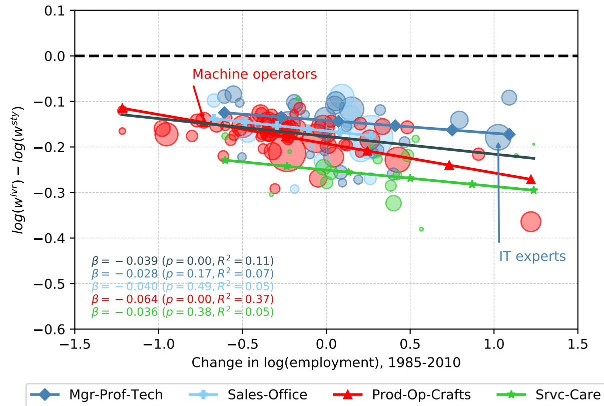

Notes: The vertical axis in Panel 3a shows the average wage of an entrant to an occupation relative to the average wage of incumbents. The average wage is taken across years 1985 until 2010. The vertical axis in Panel 3b shows the average wage of a worker leaving an occupation next period relative to the average wage of stayers. The average is taken across years 1985 until 2009 to avoid all workers being leavers at the sample end. The horizontal axis in both panels shows the change of the log number of employed workers within an occupation between 1985 and 2010. One bubble represents one of the 120 detailed occupations in the SIAB SUF. The four groups show an aggregation of these detailed occupations as described in Appendix Table A.1. Bubble size corresponds to the number of employed workers in an occupation averaged across years 1985 until 2010. Regression lines across all occupations (black) and within the four broad groups (colored) are weighted by the number of employed workers.

The vertical axis of Figure 3a shows the difference between entrants and incumbents. An occupational entrant is someone newly observed in the occupation in the current period. He could be joining the labor force for the first time, switch from a different occupation, or enter from unemployment or outside of the labor force. The wage difference between this group and incumbents is strongly negative. The first row of Table 1 shows that across occupations, entrants earn on average 23.5 log points less than incumbents. Furthermore, the difference is strongly declining in occupation growth. Per 100 log points of occupations' employment growth, it increases by another 8.4 log points. This is consistent with a situation where the skill pool that growing occupations can draw from shrinks with the extent of their expansion (e.g., see the model sketched in Appendix B.1).

In principle, the patterns in Figure 3a could be generated if occupation choice only happened at labor market entry in combination with substantial returns to experience. If this was the sole effect, however, we would expect that the wages of workers leaving their occupations would be higher than the wages of those who stay on. Put differently, in such a scenario individuals dropping out of our sample after age 54 should dominate

Table 1: Selection into and out of occupations: Regression results

Specification	Type	Entrants / Incumbents $\overline{w}\vert X$	Entrants / Incumbents $\hat{\beta}_{\text{occ growth}}$	Leavers / Stayers $\overline{w}\vert X$	Leavers / Stayers $\hat{\beta}_{\text{occ growth}}$
No covariates	Coefficient	-0.235	-0.084	-0.175	-0.039
	S.E.	0.006	0.012	0.005	0.01
Fully stratified by year / education / age	Coefficient	-0.146	-0.069	-0.165	-0.031
	S.E.	0.006	0.011	0.004	0.009
Only occupational switchers	Coefficient	-0.134	-0.026	-0.138	-0.038
	S.E.	0.005	0.011	0.005	0.011
Separate by age group: Fully stratified by year / education / age					
Ages 25–34	Coefficient	-0.11	-0.054	-0.111	-0.027
	S.E.	0.005	0.007	0.004	0.007
Ages 35–44	Coefficient	-0.106	-0.038	-0.112	-0.03
	S.E.	0.004	0.009	0.004	0.009
Ages 45–54	Coefficient	-0.091	-0.021	-0.132	-0.009
	S.E.	0.004	0.008	0.004	0.008

Notes: The numbers are based on regressions of average wage differences between entrants (leavers) and incumbents (stayers) on occupation growth between 1985 and 2010. The average wage is taken across years between 1985 until 2010 and weighted by occupations' employment. In the upper panel, the specification "no covariates" refers to Figure 3. The specification "Fully stratified by year / education / age" controls for the interaction of dummies for calendar year, education (no postsecondary, Abitur or apprenticeship, and university degree), and years of age. The corresponding figure can be found in the Appendix (Figure A.1). "Only occupational switchers" restricts the samples of entrants and leavers to those who directly switch between occupations, i.e., we drop labor market entrants, sample leavers at age 54, and those with an intermittent spell of unemployment or outside the labor force. See Figure A.2 in the Appendix for the graphical representation of the data and regression line. In the bottom panel, we repeat the specification "Fully stratified by year / education / age", splitting the sample into three age groups.

the difference between leavers and stayers. Figure 3b and Table 1 show that the opposite is the case. As for entrants, marginal workers have substantially lower wages than those who stay on. This difference is 17.5 log points on average and increases by 3.9 log points per 100 log points higher employment growth. Put differently, only the lowest-skilled workers leave fast-growing occupations. This suggests that the wage gap is not just due to entrants being at an earlier stage of their career compared to incumbents.

The second specification in Table 1 investigates how much of the wage gaps can be explained with standard observable skill proxies. We first run Mincer regressions of individual workers' wages on full interactions of dummies for calendar year, education in three categories, and years of age. We use the residuals to compute the average wage gaps for entrants/leavers and their slopes with respect to occupation growth as before. Observables account for 38% of the gap between entrants and incumbents; they explain much less of the differences between leavers and stayers. The bottom panel shows that coefficients drop by another 30% on average when splitting the sample into 10-year age groups, thereby making the pools of marginal and incumbent workers more homogenous in that dimension. Qualitatively, results are the same. These patterns underscore that the wage gaps are related to skill selection in education and age. At the same time, the majority of gaps remain unexplained by standard Mincer variables.

The third row of Table 1 investigates the extent to which the gaps persist even among direct occupational switchers (i.e., excluding labor market entrants, sample leavers at age 54, and moves to and from non-employment). Relative to differences in unconditional wages for the entire sample in the first row, the wage gaps decline by 43% and 21%, respectively. This is natural as this sample removes the large skill differences due to entrants/leavers moving between non-employment states or young workers entering the labor market. The average wage gap remains substantial. This shows that direct occupation switchers will measurably contribute to skill selection below and it suggests once more that skills are occupation-specific.

In all specifications contained in Table 1, the slopes of wage gaps with respect to occupation growth are negative. This underscores that, when occupations grow, the composition of workers changes so that average skills become increasingly scarce at all margins – in terms of observables and unobservables, and among workers hired directly from other occupations. The last two columns of Table 1 are related to Groes et al.'s (2014) analysis of the Danish labor market. Also using a sample of stayers and leavers, they find that both low and high earners in an occupation are most likely to switch. In our German data, low earners in an occupation have a much higher probability to leave, dominating the average differences between leavers and stayers (see Appendix A.3).

The prominent models in the literature on occupational changes have difficulty matching the facts in Figure 3 and Table 1 as they feature one-dimensional skills (e.g., Acemoglu and Autor, 2011; Autor and Dorn, 2013). A single dimension leads to a hierarchical ranking of occupations by skill and ensures tractability in general equilibrium. It also implies that upward switchers leave the lower-ranked occupation from above and that downward switchers enter lower-ranked occupations from above (Papageorgiou, 2014). This is hard to square with the finding that even entrants and leavers in low-wage occupations generally earn less than incumbents and stayers, respectively. In principle, the model in Hsieh et al. (2019) could deliver the empirical facts from Figures 2–3 and Table 1, see Appendix B.1 for a stylized theory. In order to make quantitative predictions, it requires time-constant skills and parametric assumptions about individual skills in all sectors. See also Gould (2002) for a related analysis in U.S. data.

The remainder of this paper develops and estimates an empirical model with multidimensional skills that flexibly change over the career. For estimation we only require data on age, occupations chosen, and the corresponding wages. We will use this model to explain and quantify the implications of the new stylized facts shown in this section.

3 Estimating Skill Prices with Longitudinal Data

There are K distinct occupations. At any time t a worker i would earn a potential wage $W_{i,t,k}$ in occupation k. This potential wage is the product of the worker's occupation-specific skill $S_{i,t,k}$ and the occupation-specific price for a unit of skilled labor $\Pi_{t,k}$. We use lowercase letters to denote logs. As in Roy (1951), we assume that workers maximize their incomes by choosing the occupation in which they earn the highest wage:

$$w_{i,t} = \max\{w_{i,t,1}, \ldots, w_{i,t,K}\} = w_{i,t,k(i,t)} = \pi_{t,k(i,t)} + s_{i,t,k(i,t)}. \tag{1}$$

The occupation subscript's argument (i, t) indicates that k is i's choice at time t.

Our goal is to estimate the evolution of skill prices $\pi_{t,k}$ over time. Based on (1), we first employ an approximation that allows us to disentangle prices from skills based on observed data. In the empirical formulation, prices evolve at the aggregate level. Skills will grow and be subject to idiosyncratic shocks over an individual's career; we describe how we deal with the challenges that arise from this in a second step.

The Roy model is hard to estimate in its general form (1) and requires strong restrictions. Böhm (2020) derives a computationally simple approximation, which he uses in repeated cross-sections data under the assumption that observable and time invariant proxies for each $s_{i,t,k}$ are available. We adapt Böhm's result to a panel data setting.

For a worker who switches occupations, i.e., $k(i,t-1) \neq k(i,t)$, note from (1) that $w_{i,t-1,k(i,t)} - w_{i,t-1,k(i,t-1)} \leq 0$ and $w_{i,t,k(i,t)} - w_{i,t,k(i,t-1)} > 0$. We assume:

Assumption 1. *For workers who switch occupations, one can approximate the wage at indifference, where $w_{i,\tau,k(i,t)} = w_{i,\tau,k(i,t-1)}$ with $\tau \in [t-1,t)$, to be in the middle of the wage interval $[(w_{i,t-1,k(i,t)} - w_{i,t-1,k(i,t-1)}), (w_{i,t,k(i,t)} - w_{i,t,k(i,t-1)})]$. That is, we assume the switch point to be halfway between the previous and the current relative wage:*

$$\underbrace{w_{i,t-1,k(i,t)} - w_{i,t-1,k(i,t-1)}}_{\textit{relative wage in period } t-1} + \underbrace{w_{i,t,k(i,t)} - w_{i,t,k(i,t-1)}}_{\textit{relative wage in period } t} = 0 \tag{2}$$

For occupation switchers, the Roy model implies negative relative wages in $t-1$ and positive relative wages in t. Assumption 1 posits that the absolute values of these relative wages (i.e., the wage rents from the chosen occupation or the wage distance to indifference) are the same in both periods. From this, we can express realized wage changes as a function of potential wage changes in the occupations worker i chooses in $t-1$ and t:

Result 1. *Under Assumption 1, worker i's observed wage growth between $t-1$ and t can be written in terms of the potential wage growth in the occupation(s) he chooses in those periods:*

$$\Delta w_{i,t} = \frac{1}{2}\big(\Delta w_{i,t,k(i,t)} + \Delta w_{i,t,k(i,t-1)}\big), \tag{3}$$

where $\Delta w_{i,t,k(i,t)} \equiv w_{i,t,k(i,t)} - w_{i,t-1,k(i,t)}$ and $\Delta w_{i,t,k(i,t-1)} \equiv w_{i,t,k(i,t-1)} - w_{i,t-1,k(i,t-1)}$.

See Appendix B.2 for the derivation. Result 1 is economically attractive as it includes workers' *endogenous* switches between occupations *due to* changing potential wages. Naturally, if a worker stays in an occupation during adjacent periods, his realized wage change is equal to the change in his potential wage in the chosen occupation. If the worker decides to switch occupations and $k(i,t-1) \neq k(i,t)$, half of his wage gain stems from the wage change he would have experienced had he stayed in his previous occupation ($\Delta w_{i,t,k(i,t-1)}$). The other half is the wage change had he been in the current occupation all along ($\Delta w_{i,t,k(i,t)}$).[5] Importantly, only potential wages in the previous and current occupations matter for i's observed wage change; potential wages in all other occupations are irrelevant. This stems from the fact that the worker has comparative advantage in both of these occupations and that he chooses accordingly.

Using Equation (1), we can decompose Result 1 into changes of prices and skills:

$$\Delta w_{i,t} = \frac{1}{2}(\Delta \pi_{t,k(i,t)} + \Delta s_{i,t,k(i,t)}) + \frac{1}{2}(\Delta \pi_{t,k(i,t-1)} + \Delta s_{i,t,k(i,t-1)}). \tag{4}$$

In the next section, we show how panel data on wages and occupation choices allow us to estimate the evolution of prices and skills imposing minimal yet informative structure on $\Delta s_{i,t,k}$. In Section 3.2, we develop a straightforward estimation strategy that is feasible even for the 120 occupations $\times$ 35 years in our application.

3.1 Identification of Skill Price Changes

The empirical strength of Result 1 is that both individuals' realized wage changes $\Delta w_{i,t}$ and occupation choices $k(i,t-1), k(i,t)$ are observed in the SIAB data. This allows linking potential and observed wage changes. However, we need to put more structure on Equation (4) in order to separate changes in skill prices from changes in skills. Our framework relies on differences. We thus do not place restrictions on either object's initial levels.

We model the skill accumulation process as learning-by-doing on the job. Its speed is occupation-specific and depends on age; working in one occupation k impacts subsequent skills in all other occupations. In particular, we assume:

Assumption 2. *Individuals' occupation-specific skill changes are time invariant in expectation. For all $k \in \{1, \ldots, K\}$ and $i \in \{1, \ldots, N\}$*

$$E\left[\Delta s_{i,t,k} \mid a(i,t-1), k(i,t-1), k\right] = \Gamma_{a(i,t-1),k(i,t-1),k} \tag{5}$$

with $a(i,t-1)$ denoting the age group of individual i in period $t-1$.

[5] Böhm and von Gaudecker (2021) reports on a large set of Monte Carlo experiments related to our framework. Using Approximation 1 instead of actual switchpoints is not a cause of concern.

Assumption 2 restricts reduced-form parameters to be time-invariant. It can be interpreted as consisting of two restrictions. The first restriction is time invariance of the structural skill accumulation parameters introduced in Section 3.2. This means that ex ante skill accumulation within $a(i, t-1) \times k(i, t-1)$ cells for all potential k has not changed over time.[6] This is consistent with the literature studying occupational changes (e.g., Acemoglu and Autor, 2011; Bárány and Siegel, 2018; Böhm, 2020; Cortes, 2016; Firpo et al., 2013; Gottschalk et al., 2015; Yamaguchi, 2018; Young, 2014) and more generally with the idea that differences in returns to worker characteristics over time are due to changes in the returns to skills rather than changes in skill endowments.

The second restriction assumes that changing prices do not affect ex post observed skill accumulation among workers of age $a(i, t-1)$ with realized choices $k(i, t-1)$ and $k(i, t)$. This is at odds with the idea of self-selection in the Roy model. The latter would imply, for example, that workers who stay in an occupation during a period of declining prices will be positively selected on skill changes compared to staying workers in periods with constant prices. Therefore, on the one hand, our approach makes progress compared to the literature by incorporating workers' endogenous choices based on changing potential wages (Result 1). On the other hand, it retains the limitation that skill accumulation parameters are not adjusted for changing self-selection over time.

To examine the plausibility of Assumption 2 empirically, Figure 4 plots the year-to-year wage growth of 25–34 year-olds (Panel a) and 35–44 year-olds (Panel b) minus the wage growth of 45–54 year-olds. We subtract the overall mean everywhere. Under Assumption 2, the ratio of wage growth across different demographic cells should remain constant over time; the eight lines should thus be flat at zero. In Panel b, all lines come

[6]The cells need not only be conditioned on age $a(i, t-1)$. In robustness checks, we add a full set of interactions with education levels. This does not make a material difference. We also include a full set of dummies for each year of potential experience (potexp$(i, t-1)$, not interacted with $k(i, t-1), k$), to capture any further non-linearities within the age groups, particularly in the early phase of the career. Conditional on these observables, the restriction does however imply that the speed of learning on the job has not changed over time. For example, a car mechanic in 2010 may well spend more time replacing electronic components than his counterpart in 1975. A secretary will send e-mails rather than typing letters. But there is no temporal change in the speed at which these people get better at their jobs from one year to the next.

Figure 4: Individual wage growth relative to 45–54 year olds

(a) 25–34 year olds (b) 35–44 year olds

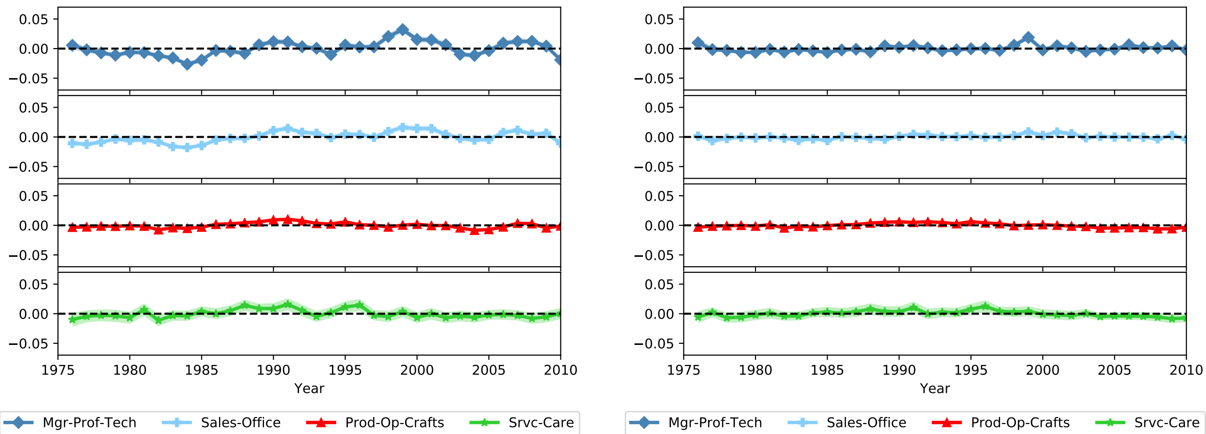

Notes: The lines show average individual wage growth from $t-1$ to t by year of 25–34 (Figure 4a) and 35–44 (Figure 4b) year olds minus average wage growth of 45–54 year olds. Results are centered at zero to show trends over time. The shaded areas around the four lines are 95% confidence intervals. The four groups are based on an aggregation of detailed occupations in the SIAB SUF as described in Appendix Table A.1.

close to it. The left panel is somewhat noisier, but there are no systematic trends.[7] These results suggest that, within occupations, relative wage growth of different age groups has not substantively changed over time.

Assumption 2 identifies changes in skill prices from the difference between wage changes and expected skill changes up to a constant. We thus need one further normalization, which we apply to skill price growth during part of our sample.

Assumption 3. *Skill prices are constant during a base period, that is,*

$$\Delta\pi_{t,k} = 0 \;\forall\; k \in \{1, \dots K\}, t \in \{1, \dots, T_{base}\}.$$

Under Assumption 3, Γ is identified from the base period due to its time-invariant nature. Accordingly, the estimates of $\Delta\pi_{t,k}$ can be interpreted as actual changes of skill

[7] The noise is largest in Mgr-Prof-Tech occupations, which is not too surprising given that many in this group enter the labor market at ages 25–34 and initial wage growth should be more susceptible to business cycles. For example, the largest changes of wage growth across age groups can be found during the dotcom bubble in the late 1990s. As an alternative for the 120 occupations, we split the sample at midterm and plot the change in log employment against the change in wage growth of young (age 25–34 or 35–44) minus old (45–54) workers in the resulting two periods. There is more variation than for the four broad occupation groups but most of the occupations have modest changes in relative wage growth rates and we cannot detect substantive patterns among them. See Appendix A.4.

prices for $t > T_{\text{base}}$. Distinguishing between price and skill growth is a general challenge of estimation based on panel data; it is often done rather implicitly.[8] Using the whole decade 1975–1984 as the base period, we will abstract from short-term fluctuations. This period covers the entire business cycle.[9]

In case Assumption 3 does not hold, we can still identify *accelerations or decelerations of skill price changes* relative to the base period. Proceeding with Assumption 3 as a normalization and focusing on occupation stayers for notational simplicity, we identify $\Gamma_{a(i,t-1),k,k} + \overline{\Delta\pi_{k,base}}$. Accordingly, the skill price estimates $\Delta\pi_{t,k} - \overline{\Delta\pi_{k,base}}$ for $t > T_{\text{base}}$. If $\overline{\Delta\pi_{k,base}} \neq 0$, the estimated skill price changes in subsequent years are accelerations or decelerations relative to their (unknown) trends during the base period. In our discussion, we mainly stick with the easier literal interpretation of the parameter estimates under Assumption 3.

Another identification strategy is to assume that seasoned workers' average skill growth is zero ("flat spot identification", Heckman et al., 1998). Under this assumption, occupation-specific wage growth over the decades is identical to price growth for these ages. Wage growth still incorporates endogenous choices as in Result 1. We will explore this as a robustness check, finding results similar to our main specification.

3.2 Estimation and Interpretation of Skill Accumulation Parameters

We combine the equations for wage growth (4) and skill accumulation (5) to obtain our baseline estimation equation:

$$\begin{aligned}\Delta w_{i,t} = \Delta\pi_{t,k(i,t-1)} \cdot \frac{I_{k(i,t-1)}}{2} + \Delta\pi_{t,k(i,t)} \cdot \frac{I_{k(i,t)}}{2} \\ + \Gamma_{a(i,t-1),k(i,t-1),k(i,t-1)} \cdot \frac{I_{a(i,t-1)} \cdot I_{k(i,t-1)}}{2} \\ + \Gamma_{a(i,t-1),k(i,t-1),k(i,t)} \cdot \frac{I_{a(i,t-1)} \cdot I_{k(i,t-1)} \cdot I_{k(i,t)}}{2} + \varepsilon_{i,t}\end{aligned} \tag{6}$$

[8] As an alternative, Cortes (2016) and Cavaglia and Etheridge (2020) do not use a base period and thus implicitly set one of the skill accumulation parameters to zero (details in Böhm and von Gaudecker, 2021).

[9] See Figure A.3 in the Appendix for employment and wage trends over the full 1975–2010 period. A useful side-effect of choosing the first decade of data as the base period is that the resulting analysis period of 1985–2010 is the same as in Card et al. (2013), easing comparisons.

with $I_{a(i,t)}$ and $I_{k(i,t)}$ denoting indicator variables for i's age group and choice of occupation, respectively. In line with Assumption 3, we set $\Delta\pi_{k,t} = 0 \ \forall \ k \in \{1,\ldots K\}, t \in \{1,\ldots,T_{\text{base}}\}$ and estimate (6) for the whole period $t \in \{1,\ldots,T\}$ by OLS.

Regression (6) is perfectly saturated in age groups, previous and current occupations. The regression error $\varepsilon_{i,t}$ reflects individuals' skill shocks or idiosyncratic deviations that are different from the average worker in the respective cell spanned by these dummies. That is, we can write $\varepsilon_{i,t} = \frac{1}{2}(u_{i,t,k(i,t-1)} + u_{i,t,k(i,t)})$ where $u_{i,t,k} \equiv \Delta s_{i,t,k} - \Gamma_{a(i,t-1),k(i,t-1),k}$. By Assumption 2, we have $E\left[u_{i,t,k} \,|\, a(i,t-1), k(i,t-1), k\right] = 0$ and the error term in equation (6) is uncorrelated with the regressors.

Result 2. *Under Assumptions 1–3, OLS estimation of equation* (6) *consistently identifies price changes* $\{\Delta\pi_{t,k} | t > T_{base}\}$ *and average skill changes* $\{\Gamma_{a(i,t-1),k(i,t-1),k}\}$*. Without Assumption 3, OLS identifies these parameters up to the normalization of the base period.*

As discussed with Assumption 2, average occupation-specific skill changes Γ will not coincide with structural skill accumulation parameters in general. Suppose there exists a learning-by-doing production function, such that *ex ante* a worker accumulates skills according to $\Gamma^*_{a(i,t-1),k(i,t-1),k} = E\left[\Delta s_{i,t,k} \,|\, a(i,t-1), k(i,t-1)\right]$ with idiosyncratic innovations $u^*_{i,t,k}$. What we identify in regression (6) is *ex post* accumulation in the sense that

$$\Gamma_{a(i,t-1),k(i,t-1),k(i,t)} = \Gamma^*_{a(i,t-1),k(i,t-1),k(i,t)} + E\left[u^*_{i,t,k(i,t)} \,|\, a(i,t-1), k(i,t-1), k(i,t)\right] \quad (7)$$

is conditional on the current choice $k(i,t)$ and $u_{i,t,k} = u^*_{i,t,k} - E\left[u^*_{i,t,k(i,t)} \,|\, a(i,t-1), k(i,t-1), k(i,t)\right]$ is the deviation of the structural shock from its ex post mean. By the nature of the data, $\Gamma_{a(i,t-1),k(i,t-1),k}$ with $k(i,t-1) \neq k$ will be identified from switchers. For these, often $E\left[u^*_{i,t,k} \,|\, a(i,t-1), k(i,t-1), k\right] > E\left[u^*_{i,t,k} \,|\, a(i,t-1), k(i,t-1)\right]$. That is, the off-diagonal accumulation parameters tend to be upward-biased because of a classic self-selection problem. We thus cannot interpret them as what is learned for $k(i,t)$ while working in $k(i,t-1)$ in the entire population. More subtly, stayers' $\Gamma_{a(i,t-1),k(i,t-1),k(i,t)}$ with $k(i,t-1) = k(i,t)$ will tend to be overestimated relative to Γ^*, too. By self-selection, stayers are likely to have experienced more positive shocks compared to leavers.

Nevertheless, regression (6) succeeds in estimating the skill price changes as long as Assumption 2 holds. Böhm and von Gaudecker (2021) bring the model to its limits using Monte Carlo simulation and a range of possible extensions (switching costs, trends in amenities, employer learning, different magnitudes of structural shocks $\mathrm{var}(u^*)$). The price changes are well identified using our method for plausible parameter ranges. It becomes clear, however, that the estimated Γ matrix needs to be interpreted in the above *ex post* sense. For our decomposition purposes and analysis of skills in occupations, this will be sufficient. The caveats are that we can neither make statements about workers' skill (changes) in occupations that they do not choose, nor about the population distribution of the K-dimensional skill vector.

4 Skill Prices and Skill Selection

This section first presents the estimation results. These include the evolution of skill prices, the accumulation of skills over the career, and the relationship of prices and occupations' average skills with employment growth. We then dig deeper into the nature of the implied selection effects, showing that the skill differences between marginal workers and those who remain in their occupations drive the strongly negative association of employment growth with average skill changes in an occupation. We further decompose these in Section 4.3, before reporting on various robustness exercises in the final part. One of those is an alternative way to think about unemployment and spells outside the labor force, which are not fully obvious in our setup (in the main estimation, we treat such spells as missing data).

4.1 Estimated Skill Price Changes and Average Skill Accumulation

Figure 5a depicts the evolution of skill prices, normalizing them to zero in 1985 and cumulating the yearly changes until 2010. In the broad occupation groups, skill prices increased strongly among Mgr-Prof-Tech occupations, modestly among Sales-Office and Srvc-Care, and they decreased among Prod-Op-Crafts. The thin lines in the background

Figure 5: Estimated skill prices and stayers' average skill accumulation

(a) Skill prices over time (b) Skill accumulation over the career

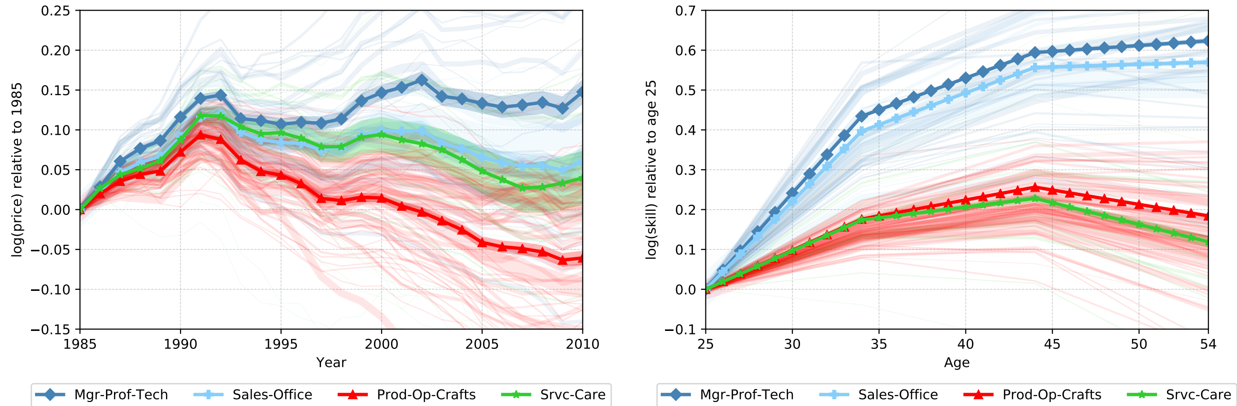

Notes: The figure shows estimated skill price changes over time and occupation stayers' skill accumulation over the life cycle. OLS estimates as described by Equation (6). Shaded lines in the background represent the 120 detailed occupations in the SIAB SUF. The four groups show an aggregation of these detailed occupations as described in Appendix Table A.1. The thickness of a shaded background line corresponds to the number of employed workers in an occupation averaged across years 1985 until 2010. The shaded areas around the four lines are 95% confidence intervals.

show that these broad estimates mask substantial heterogeneity among the 120 detailed occupations. We will explore this in greater detail below.

Several distinct periods are noticeable. All prices increased in the favorable economic conditions between 1985 and 1991, although this was already less pronounced for the Prod-Op-Crafts occupations. These have experienced a continuous decline thereafter to the point that prices in 2010 were more than five percent below their initial value in 1985. For the other occupations, there was a drop during the 1992–93 recession as well; prices then stayed constant until they rebounded before the turn of the century. This rebound was most pronounced for Mgr-Prof-Tech occupations; prices in this group did not change much for the remainder of our sample period. Skill prices fell by about 5 percentage points for Sales-Office and Srvc-Care occupations between 2000 and 2010. These broad patterns are consistent with the job polarization of Figure 1b above.

Figure 5b graphs the estimates for the diagonal elements of Γ, i.e., the average skill accumulation for stayers. See Table C.1 in the Appendix for the full set of estimated coefficients in Γ. Again, bold lines are for the four broad groups and thin lines for the 120 detailed occupations. Skill growth in the early years of the career is steep. Absent changes in skill prices, it implies a 20% growth in wages between age 25 and age 34 for Prod-

Op-Crafts or Srvc-Care occupations and 50% or more for the other two. It slows down mid-career and flattens out or turns negative toward the end of prime age. This reflects the well-established concavity of life-cycle wage profiles (e.g., Lagakos et al., 2018). Skill growth differs substantially by occupation. It is initially fast in high-earning Mgr-Prof-Tech and Sales-Office occupation groups and never completely ceases. Growth is flatter and eventually peters out or turns negative in the Prod-Op-Crafts and Srvc-Care groups, i.e., occupations that often require physical labor.[10] Again, the broad groups mask substantial heterogeneity across the detailed occupations. At the same time, it is the case that the "blue" occupations on the one hand and the "red and green" occupations on the other hand are almost separate; there are hardly any to be found in the other block. This shows that life-cycle wage profiles are decidedly different across occupations and accounting for this is critical in producing reasonable estimates of prices and skills.

We now hone in on our key finding of this section, namely that employment growth and skill price growth go hand in hand. Figure 2a has shown that the left-hand side of:

$$\underbrace{\mathrm{E}\left[w_{i,t,k(i,t)}\right] - \mathrm{E}\left[w_{i,t-1,k(i,t-1)}\right]}_{\text{Mean wage change}} = \underbrace{\Delta\pi_{t,k}}_{\text{Price change}} + \underbrace{\mathrm{E}\left[s_{i,t,k(i,t)}\right] - \mathrm{E}\left[s_{i,t-1,k(i,t-1)}\right]}_{\text{Mean skill change}}, \tag{8}$$

i.e., the average wage change of professions over the 1985–2010 period, is unrelated to employment growth at the level of the 120 detailed occupations. Figure 6 plots its two components on the right-hand side separately.

Figure 6a shows that detailed occupations' log employment changes are positively related to cumulative skill price changes over the same period. The black regression line summarizes this strong relationship for the 120 occupations, which is in marked contrast to the zero correlation for wages in Figure 2a. The relationship holds within occupation groups, showing that the result is more general than a particular demand shifter that predominantly impacts broad occupations. Figure 6a indicates that shifts in demand were the dominant driver of occupational employment changes. Pure shifts in supply would

[10] Compared to other countries Sales-Office is quite high-earning in Germany (Cavaglia and Etheridge, 2020, has a direct comparison with the U.K.). Its average wages for men are about halfway between Mgr-Prof-Tech and Prod-Op-Crafts, employment is not declining over time.

Figure 6: Correlation of changes in employment, skill prices, and skills

(a) Skill prices

(b) Skills

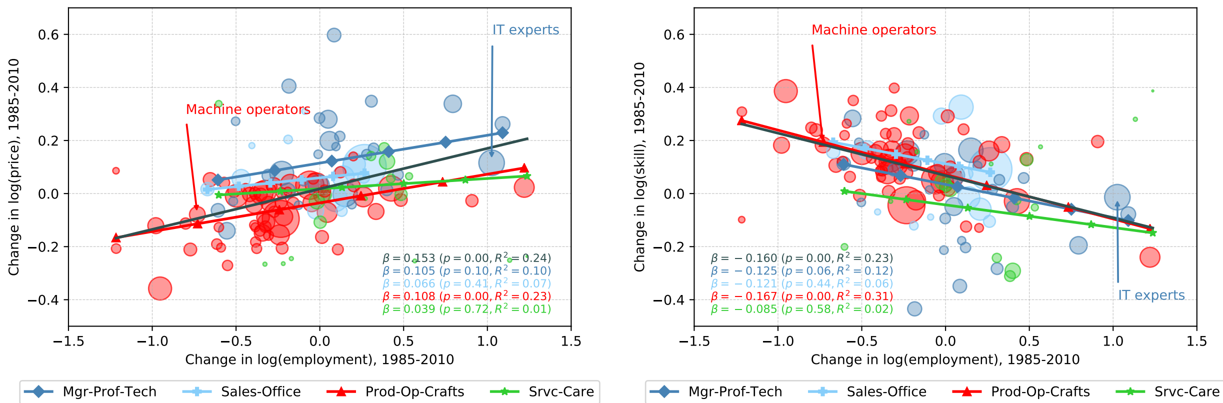

Notes: The vertical axis in Panel 6a shows the change in skill prices between 1985 and 2010 estimated as detailed in Section 3.2. The vertical axis in Panel 6b depicts the change in skills between 1985 and 2010 estimated as the residual between price and wage changes as shown in Equation (8). The horizontal axis in both panels shows the change of the log number of employed workers within an occupation between 1985 and 2010. One bubble represents one of the 120 detailed occupations in the SIAB SUF. The four groups show an aggregation of these detailed occupations as described in Appendix Table A.1. Bubble size corresponds to the number of employed workers in an occupation averaged across years 1985 until 2010. Regression lines across all occupations (black) and within the four broad groups (colored) are weighted by the number of employed workers.

have induced a negative correlation of employment with prices, all else equal. At the same time, there exists heterogeneity of occupations' employment and price growth around the regression line(s). This may, among other things, be due to contemporaneous shifts in supply but also due to variation in the elasticity of skill supply (inelastic occupations would have small employment and large price growth).

Figure 6b shows that implied skill changes are the flipside of the skill price estimates. The regression line indicates that skills of the occupations with the fastest growth declined by 35 log points on average compared to those that shrank the most. These are large effects; we devote the next subsections to examining their components and plausibility.

4.2 Defining and Estimating Employment Growth-Selection

We have documented in Section 2.2 that entering (leaving) workers' skills on average appear decidedly below those of incumbents (stayers) and that faster-growing occupations draw even less skilled entrants (leavers). Given that growing sectors by definition experience net entry, this could substantially drag down growing occupations' average wages despite rising demand and increasing skill prices. Here we formalize and quantify this effect in the context of our model, showing that it indeed drives the systematic part of the relationship between employment growth and skills.

The change in average skills of an occupation in Equation (8) is determined by three mutually exclusive groups of workers: Those who leave the occupation after period $t-1$; those who stay on after period $t-1$ and are thus incumbent in period t; and those who enter in period t. Denoting the share of leavers after $t-1$ by $p^{lvr}_{t-1,k}$ and the share of period-t entrants by $p^{ent}_{t,k}$, we can decompose the change of skills in occupation k into three terms:

$$\begin{aligned}
\underbrace{\mathrm{E}\left[s_{i,t,k(i,t)}\right]-\mathrm{E}\left[s_{i,t-1,k(i,t-1)}\right]}_{\text{Mean skill change}} \quad = \quad & \underbrace{\left(1-\frac{p^{lvr}_{t-1,k}+p^{ent}_{t,k}}{2}\right)\cdot \mathrm{E}\left[\Delta s^{incumb}_{i,t,k}\right]}_{\text{1. Skill accumulation of } t-1 \text{ stayers}} \qquad (9)\\
& + \underbrace{\frac{p^{lvr}_{t-1,k}+p^{ent}_{t,k}}{2}\cdot\left(\mathrm{E}\left[s^{ent}_{i,t,k}\right]-\mathrm{E}\left[s^{lvr}_{i,t-1,k}\right]\right)}_{\text{2. Churning: difference entrants in } t\text{, leavers after } t-1}\\
& + \underbrace{\frac{p^{ent}_{t,k}-p^{lvr}_{t-1,k}}{2}\cdot\left(\mathrm{E}\left[s^{ent}_{i,t,k}\right]-\mathrm{E}\left[s^{incumb}_{i,t,k}\right]+\mathrm{E}\left[s^{lvr}_{i,t-1,k}\right]-\mathrm{E}\left[s^{sty}_{i,t-1,k}\right]\right)}_{\text{3. Employment growth-selection}}.
\end{aligned}$$

See Section B.4 of the Appendix for the steps of the derivation. Summing the year-to-year changes in (9) over the period 1985–2010 gives the components of skill changes we analyze in the following. Note that while $t-1$ stayers are the same individuals as period-t incumbents, $\mathrm{E}\left[s^{sty}_{i,t-1,k}\right] \neq \mathrm{E}\left[s^{incumb}_{i,t,k}\right]$ because their skills change over time. The first term of Equation (9) reflects just this, i.e., the skill accumulation of workers who remain in the occupation. Its impact on occupational skill changes is high if turnover is small and skill accumulation of staying workers is high. Figure 7a shows that this term is positive for the vast majority of occupations and larger among higher-accumulation occupations.

The second term of Equation (9) is churning, which is composed of average turnover multiplied with the skill differences between period-t entrants and $t-1$ leavers. Churning tends to impact aggregate skills negatively since leavers will have accumulated more

Figure 7: Employment growth vs. the components of skill changes

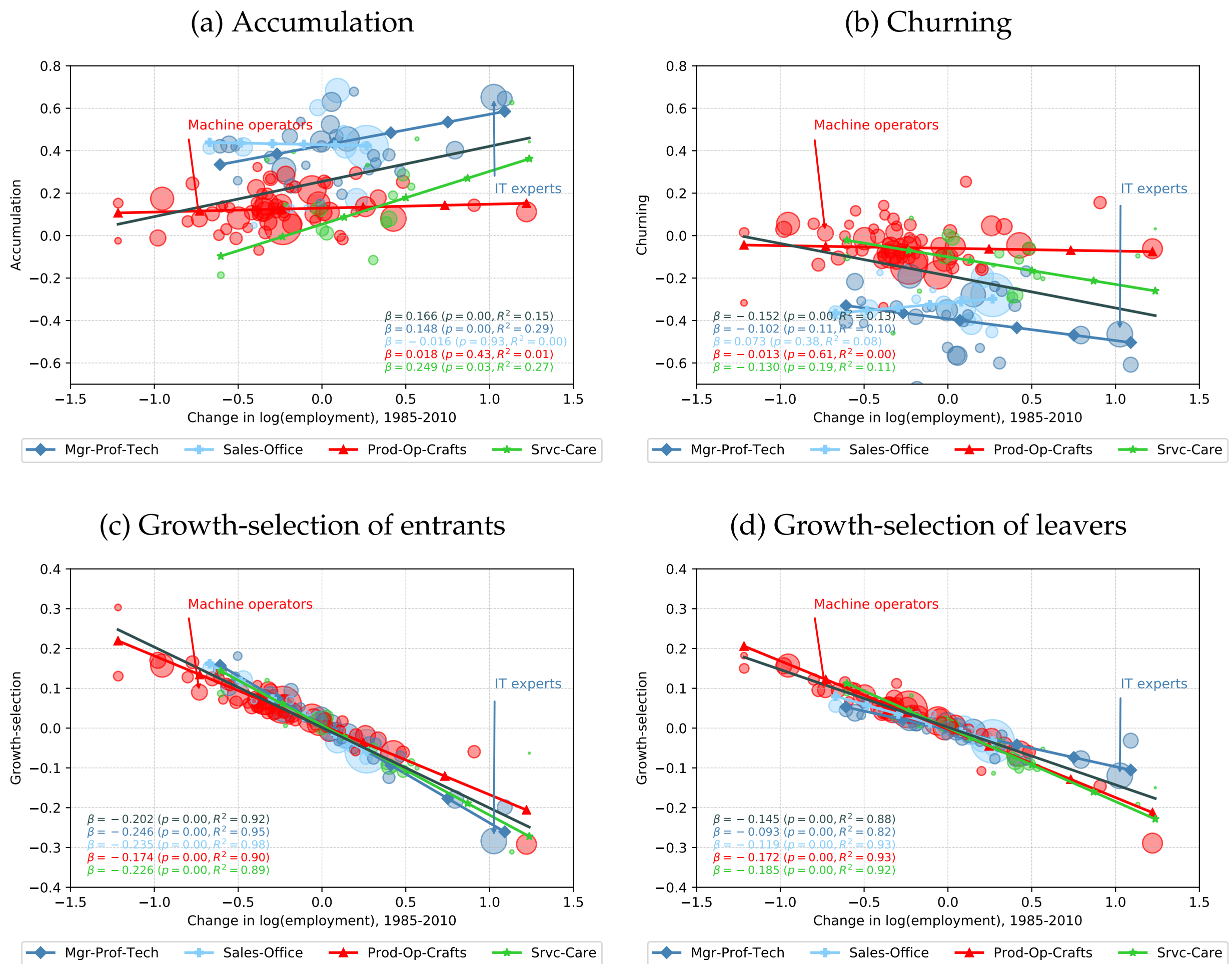

Notes: Results correspond to sample averages following Equation (9). The horizontal axis in both panels shows the change of the log number of employed workers within an occupation between 1985 and 2010. One bubble represents one of the 120 detailed occupations in the SIAB SUF. The four groups show an aggregation of these detailed occupations as described in Appendix Table A.1. Bubble size corresponds to the number of employed workers in an occupation averaged across years 1985 until 2010. Regression lines across all occupations (black) and within the four broad groups (colored) are weighted by the number of employed workers.

skills than entrants. It becomes more negative for high turnover occupations and for occupations with high levels of skill accumulation. Hence, the accumulation and churning effects will often act in opposite directions. Figure 7b shows this to be the case.

Occupation growth does not have a first order effect on the total of accumulation and churning. By inducing variation in turnover $\left(p^{lvr}_{t-1,k} + p^{ent}_{t,k}\right)$, occupation growth pushes terms 1. and 2. in opposite directions. It is thus no surprise that the slopes in Figures 7a and 7b have opposite signs. Indeed, we fail to find a systematic relation between their sum and occupation growth (see Appendix Figure C.1).

Occupation growth directly enters the third term of Equation (9), which we call the *growth-selection effect*. It is the product of net entry and the skill difference between marginal and inframarginal workers. Because that difference is negative for both entrants and leavers everywhere (see Figure 3), occupation growth will determine its sign. The growth-selection effect is negative for growing occupations; it is positive for shrinking occupations; and it is zero when there is no change in size.[11] Growth-selection thus formalizes and quantifies the intuition developed in Section 2, whereby the more an occupation grows, the more net entry of less skilled workers it experiences.

The bottom panels of Figure 7 plot employment growth-selection. Both for entrants (Panel c) and leavers (Panel d) the relationship between employment growth and growth-selection is strongly negative. Both contribute to the declining skills of growing occupations. Growth-selection determines the entire negative slope between changes in employment and changes in skills from Figure 6b (the location of the regression line and most of the variation around it stem from accumulation and churning). Therefore, the systematic part of the large selection effects we found in Section 4.1 can be traced back to changes in occupations' sizes multiplied with the skill gaps between marginal and inframarginal workers. An interesting observation is that in the model sketched in Appendix B.1, based on Hsieh et al. (2019), rising skill prices attract workers of exactly offsetting lower quality to an occupation. Without making that assumption in the estimation, it comes out of our analysis. It also holds in the U.S.-CPS data analyzed in Hsieh et al. (2019).

4.3 Contributors to Employment Growth-Selection

Having established its quantitative importance, we now investigate further the nature of the employment growth-selection effect. We decompose (9) in a way that allows us to gauge the relative contributions of age and of different worker types. That is, is growth-selection primarily determined by entry of young workers who simply have not had a

[11] If an occupation is stable in the sense that its size and skill composition are constant, accumulation and churning must cancel each other out. The growth-selection effect will be zero because of constant employment and the left-hand-side of (9) will be zero because of constant skills. Skill accumulation of staying workers must equal the churning effect due to the difference in skills between entrants and leavers.

chance to accumulate a lot of skills? Or is it that entrants into growing occupation have lower skills than incumbents had at the same point in their life-cycle? Toward the end of this section, we alternatively consider the contributions of different types of switchers, i.e., other occupations, unemployment, etc.

Disentangling the roles played by age and worker type would be simple if we observed the distribution of skills (or potential wages) of workers in all professions and age groups. We could define a worker type—a natural measure would be his skills at labor market entry—and consider how choices over the life-cycle have changed his skills. Due to the fundamental selection problem, we will only be able to proxy this ideal. Nevertheless, it is useful to clearly define the forms of selection that we shall try to tease apart:

1. *Roy selection* is sorting based on workers' skills, holding experience constant.

2. *Age selection* is sorting of workers who are equally skilled holding experience constant, but whose skill levels are different because young workers have had less time to accumulate human capital.

In our framework, skills change dynamically depending on the chosen occupation. Our estimates show that these effects are substantial in size. Moreover, we cannot hope to interpret our estimates of Γ's off-diagonal elements as skill changes in occupations not chosen by workers (see the discussion around Assumption 2 in Section 3.2 and the final paragraph of Section 3.2). We thus focus on the current occupation spell.

We decompose the growth selection effect into contributions by different age groups and into the sources of skills relevant for the current occupation spell – endowment at entry, expected skills accumulated during the spell, and shocks relative to this expectation. We map these components to Roy and age selection, respectively. We do so separately for entrants versus incumbents and leavers versus stayers.

In deriving the decomposition formulas, we focus on growth-selection of leavers, i.e., $\frac{1}{2}\left(p_{k,t}^{ent} - p_{k,t-1}^{lvr}\right) \cdot \left(\mathrm{E}\left[s_{i,t-1,k(i,t-1)}^{lvr}\right] - \mathrm{E}\left[s_{i,t-1,k(i,t-1)}^{sty}\right]\right)$. Growth-selection of entrants works accordingly. First, we split the overall skill difference between leavers and stayers into components that are due to G different groups of leavers:

$$\mathrm{E}\left[s^{lvr}_{i,t-1,k(i,t-1)}\right] - \mathrm{E}\left[s^{sty}_{i,t-1,k(i,t-1)}\right] = \sum_{g=1}^{G} p^{lvr,g}_{k,t-1}\left(\mathrm{E}\left[s^{lvr,g}_{i,t-1,k(i,t-1)}\right] - \mathrm{E}\left[s^{sty}_{i,t-1,k(i,t-1)}\right]\right) \quad (10)$$

where $\sum_{g=1}^{G} p^{lvr,g}_{k,t-1} = 1$ and the notation $p^{lvr,g}_{k,t-1}$ indicates the share of group g among workers leaving k after $t-1$. Stayers' skills do not depend on the group. Both the relative size of each group ($p^{lvr,g}_{k,t-1}$) and its average skills ($\mathrm{E}\left[s^{lvr,g}_{i,t-1,k(i,t-1)}\right] - \mathrm{E}\left[s^{sty}_{i,t-1,k(i,t-1)}\right]$) matter for its contribution to growth-selection. Net entry into the occupation ($p^{ent}_{k,t} - p^{lvr}_{k,t-1}$) is a constant that everything is multiplied with. As skill prices are the same for leavers and stayers, we can infer the components of (10) directly from wage gaps between leavers and stayers.

Decomposition (10) by itself would yield the contributions of different age groups to employment growth selection. To the extent that Roy selection matters, however, the decomposition is not informative with respect to age selection in the above sense. In order to shed light on this, we further decompose (10) by the source of skills within the current occupation spell. Focusing on the skill differences between a particular group g of leavers and workers who stay put, we obtain:

$$\begin{aligned} \mathrm{E}\left[s^{lvr,g}_{i,t-1,k(i,t-1)}\right] - \mathrm{E}\left[s^{sty}_{i,t-1,k(i,t-1)}\right] \quad = \quad & \left(\mathrm{E}\left[s^{lvr,g,endow}_{i,t-1,k(i,t-1)}\right] - \mathrm{E}\left[s^{sty,endow}_{i,t-1,k(i,t-1)}\right]\right) \\ & + \left(\mathrm{E}\left[s^{lvr,g,\Gamma}_{i,t-1,k(i,t-1)}\right] - \mathrm{E}\left[s^{sty,\Gamma}_{i,t-1,k(i,t-1)}\right]\right) \\ & + \left(\mathrm{E}\left[s^{lvr,g,shocks}_{i,t-1,k(i,t-1)}\right] - \mathrm{E}\left[s^{sty,shocks}_{i,t-1,k(i,t-1)}\right]\right). \end{aligned} \quad (11)$$

Decomposition (11) relies on our estimates of prices and the diagonal elements of the accumulation matrix Γ. We require the longitudinal information in the data to recover the skills at each individual's occupation entry. See details of our procedure in Appendix C.2.

The first summand on the right represents skill differences in *endowments* at the most recent entry into occupation k. When we write about endowments below, they always concern skills in the current occupation at the start of the current spell. We attribute these endowments to Roy selection because they measure skill differences at the time when leavers and stayers respectively entered the occupation. This interpretation would be questionable if experience of marginal workers at entry were lower than that of inframarginal workers at the time when they entered. However, age at entry of marginal workers is actually about three years higher than age at entry of inframarginal workers. Magnitudes vary, but the difference is positive for all occupation groups and for both entrants and leavers. This pattern lends support to interpreting skill differences at the start of an occupation spell as a form of Roy selection.

The second summand with superscript Γ are skills that were on average accumulated since the entry into an occupation. We interpret these as age selection. With the covariates in our main specification (6), the summand only depends on the occupation, age at entry into the occupation, and the time that has elapsed since then. For leavers, the sign of the difference between marginal and inframarginal workers is not clear ex ante. On the one hand, sample leavers at age 54 will contribute to longer spells for marginal workers. On the other hand, young workers tend to have higher occupational mobility and this will lead to shorter spells among leavers relative to stayers.

The last summand in (11) are individual deviations (*shocks*) from the average accumulation in the second summand. These are a dynamic form of Roy selection – over time, we expect workers who stay put to experience more favorable shocks in an occupation than workers who eventually leave, contributing to skill differences unrelated to age.

For entrants, we decompose growth-selection analogously to (11). Entrants arrive in period t with their endowment of skills in occupation $k(i,t)$. This endowment has been shaped by previous experience. However, by definition, they have not had a chance yet to accumulate skills during this particular spell, whether through predicted accumulation or shocks. Hence, the equivalent to the right hand side of (11) simplifies:

$$\mathrm{E}\left[s_{i,t,k(i,t)}^{ent,g,endow}\right] - \mathrm{E}\left[s_{i,t,k(i,t)}^{incumb,endow}\right] - \mathrm{E}\left[s_{i,t,k(i,t)}^{incumb,\Gamma}\right] - \mathrm{E}\left[s_{i,t,k(i,t)}^{incumb,shocks}\right].$$

The results of the decomposition for the broad professions are shown in Table 2. We combine the older two age groups to make the table easier to parse. Results with our usual three groups are in the Appendix Table C.2. The upper panel of Table 2 shows relative contributions to growth selection; columns sum to one. The bottom panel shows the growth selection effect and the fractions of entrants and leavers by age group.

Growth-selection is negative for the three growing professions and positive for the shrinking Prod-Op-Crafts. Among Mgr-Prof-Tech occupations, average skills declined by 3% between 1985 and 2010 because of employment growth-selection. Of this 3% decline, 19.6% were due to entrants aged 25–34 whose skill endowments were below those of incumbents. 15.8% were due to leavers in that same age group whose skills at the time of their entry were below the endowments of those staying in the occupation. The endowment effects for 35–54 olds (-5.3% for entrants and -9.8% for leavers) work in the opposite direction. Hence, summing up the first four rows, endowments make up 20% of growth-selection among Mgr-Prof-Tech occupations. This overall number is similar in Sales-Office (23%). The corresponding percentages are substantially higher in the professions with lower wages and low rates of skill accumulation. Beginning-of-spell endowments contribute 42% (Prod-Op-Crafts) and 64% (Srvc-Care) to growth selection. In both cases, initial skills of middle-aged entrants and leavers also contribute positively to growth selection, i.e., these are lower on average than endowments of inframarginal workers at the beginning of their spells.

Expected skill accumulation, shown in the second panel of Table 2, matters more in occupations with high wages and high rates of skill growth. Among Mgr-Prof-Tech occupations, 25.7% and 13.8% come from skills accumulated by incumbents during their spell relative to young and old entrants, respectively. Workers leaving a profession have accumulated less skills since the beginning of their spells than the workers who stay on. This effect contributes 4.3% to growth-selection for young and 12.5% for older workers. That older leavers contribute more stems from their higher share among all leavers, shown in the bottom of the table. The sum of these four numbers gives the total contribution of differences in predicted skill accumulation between marginal and inframarginal workers for Mgr-Prof-Tech occupations. It amounts to 56% and is similar for Sales-Office (63% overall) while it is much lower in Prod-Op-Crafts (33%) and Srvc-Care (14%) occupations. These

Table 2: Contributions to growth-selection by source of skills

Source	Type	Age	Mgr-Prof-Tech	Sales-Office	Prod-Op-Crafts	Srvc-Care
Endowment at the most recent entry into the occupation group	Entrants	25–34	0.196	0.188	0.096	0.238
		35–54	-0.053	-0.029	0.117	0.088
	Leavers	25–34	0.158	0.158	0.084	0.218
		35–54	-0.098	-0.083	0.125	0.098
Predicted **skill accumulation** since the most recent entry	Entrants	25–34	0.257	0.284	0.141	0.047
		35–54	0.138	0.128	0.090	0.028
	Leavers	25–34	0.043	0.074	0.045	0.010
		35–54	0.125	0.141	0.049	0.057
Deviation of skills from the prediction since the most recent entry	Entrants	25–34	0.064	0.047	0.082	0.055
		35–54	0.034	0.019	0.053	0.032
	Leavers	25–34	0.044	0.037	0.045	0.048
		35–54	0.091	0.036	0.072	0.081
Background						
Growth Selection	Total		-0.030	-0.021	0.059	-0.043
Fractions	Entrants	25–34	0.620	0.636	0.667	0.574
		35–54	0.380	0.364	0.333	0.426
	Leavers	25–34	0.234	0.299	0.300	0.375
		35–54	0.766	0.701	0.700	0.625

Notes: Numbers in the first panel represent relative contributions to the growth-selection effect. Columns sum to one. The first row in the second panel shows the growth-selection effect within each broad occupation group during 1985–2010. The last four rows show $p_{k,t}^{ent,g}$ and $p_{k,t-1}^{lvr,g}$, averaged over the entire period.

are the contributions of age selection; the large heterogeneity underlines the importance of flexibly modeling skill accumulation by occupation.

The final contributor to growth-selection in Table 2 are deviations from average skill accumulation. Again, shocks since the beginning of the spell are zero by construction for entrants. Among Mgr-Prof-Tech occupations, $6.4\% + 3.4\% = 9.8\%$ of growth selection is due to incumbents having experienced positive skill shocks relative to the estimated $\Gamma_{a,k,k}$ during their tenure. Inframarginal workers also experience better shocks than those who decide to leave. These differences contribute another $4.4\% + 9.1\% = 13.5\%$. Hence, just under a quarter (23%) of growth selection is accounted for by more favorable shocks among inframarginal workers in Mgr-Prof-Tech occupations. In all other professions the contributions of both entrants' and leavers' shocks are positive, too. In total, they make up 14% of growth-selection for Sales-Office, 25% of Prod-Op-Crafts, and 22% of Srvc-Care. Staying in an occupation is thus endogenous in the sense that only workers who receive sufficiently favorable skill shocks will decide to remain in it. We interpret this as a dynamic form of Roy selection, which is quantitatively substantial and supports the self-selection model underlying the estimation method.

We now take another look at the original question of age selection versus Roy selection. First, to calculate the aggregates over all occupation groups mentioned in the introduction, we weight the numbers for age selection (56% for Mgr-Prof-Tech, 63% for Sales-Office, 33% for Prod-Op-Crafts, 14% for Srvc-Care) with the average employment shares over the 1985–2010 period (22%, 17%, 56%, and 5%, respectively). We thus obtain that 42% of employment growth-selection can be attributed to age selection, or marginal workers not having had the same amount of time to accumulate skills in an occupation as inframarginal workers. The remaining 58% are due to static and dynamic Roy selection, or skill differences between marginal and inframarginal workers within occupations, holding age selection constant. Of these 58%, static Roy selection (endowments) contributes 35 percentage points and dynamic Roy selection 23 percentage points.

Second, we study the contributions to growth selection by young workers only. We thus go back to (10), decomposing growth selection by age group. For young workers, this amounts to summing all the contributions of entrants and leavers aged 25–34 in Table 2. This occupational sorting of young workers makes up 76% of growth-selection for Mgr-Prof-Tech occupations, 79% for Sales-Office, 49% for Prod-Op-Crafts, 62% for Srvc-Care. It contains all sources of skill differences with inframarginal workers.[12] In particular, 9–45 percent of it are due to average skill accumulation, or young workers having had less time to acquire skills – age selection of young workers. Another 37–74 percent are due to lower endowments – static Roy selection of young workers. 11–26 percent stem from skill shocks

[12]That young workers are more mobile and have below-average occupation skills is consistent with Neal (1999) and Pavan (2011). Their models emphasize finding out about match quality through search while we focus on specific skill acquisition. An alternative interpretation of shocks in our model is as learning about skills (see Böhm and von Gaudecker, 2021), which would be closer to Neal (1999) and Pavan (2011).

– dynamic Roy selection of young workers.[13] Hence, even within the youngest age group, which had little time to accumulate skills, Roy selection effects dominate quantitatively.

In Appendix Table C.3 we explore an alternative break-down of growth-selection. In that version, we work directly with decomposition (10), considering 12 groups. These are made up of the 2 age groups and 6 types of origin/destination states of entrants/leavers. We consider switches between the broad professions, unemployment, temporary spells outside the labor force, and initial labor market entry/sample exit. We find interesting differences between the three growing occupation groups on the one hand and Prod-Op-Crafts on the other. For the former, the single largest contributor are young labor market entrants, making up 25–40% of growth-selection. For Prod-Op-Crafts, switches to and from unemployment contribute almost half of the total. Direct switches between occupations are non-negligible and almost always positively contribute to growth-selection. That is, the fact that both entrants and leavers are less skilled than inframarginal workers holds independently of the origin/destination occupation. The one exception are middle-aged switchers involving Mgr-Prof-Tech occupations. On average, workers who leave this group have higher skills than incumbents in their destination. Switchers into Mgr-Prof-Tech are higher-skilled than stayers in their origin occupations. The broader patterns, however, once again indicate that inframarginal workers have strong specific skills compared to marginal workers, which is hard to reconcile with a ranking of occupations by a single skill. We can actually reject the one-dimensional model, see Appendix B.3.

Finally, we do the same decompositions for individual occupations. The results are in Appendix Tables C.5 and C.6 for a selected set of occupations experiencing large changes in employment. In absolute value, growth-selection varies between 0.09 and 0.29 among these, compared to 0.02–0.06 among the broader professions. The basic patterns from before hold up: Skill accumulation and selection at labor market entry are particularly important for occupations with high rates of skill growth (e.g., growing consultants and tax advisors, or shrinking accountants and valuers). Endowments and deviations from

[13] For example, of young workers' 79% contribution to growth-selection in Sales-Office, $\frac{18.8\%+15.8\%}{78.8\%} =$ 44% are due to endowments, $\frac{28.4\%+7.4\%}{78.8\%} = 45\%$ to accumulation, and $\frac{4.7\%+3.7\%}{78.8\%} = 11\%$ due to shocks.

expected skill changes are quantitatively large throughout; moves involving unemployment contribute substantial shares among Prod-Op-Crafts occupations.

4.4 Robustness of Results

Appendix D examines the robustness of our empirical results in alternative samples and estimation specifications. In the following, we briefly summarize these robustness checks. Finally, we connect to prior literature by describing occupations via their task content.

Comprehensive sample and different demographic groups: In Appendix D.1 we redraw our key graphs with occupational growth on the x-axes among a comprehensive "employment sample". In that sample, we include women and foreigners while expanding the age range to 20–60. Consistent with Katz and Murphy (1992), we keep prime age West German males for the wages, skill prices, and skills on the y-axes, since these are better measured in a high-attachment "wage sample". We do this because changes in the demand or supply of labor to different occupations should affect overall employment beyond prime age males, and there can be direct or feedback effects on occupational prices from this. Figure 6 and all other relationships we show turn out similar with the comprehensive sample. Hence, accounting for potential equilibrium effects via other demographic groups does not change our results. In Appendix D.2, we also change the composition of the entire sample (employment and wages) by varying the inclusion of East Germans, foreigners, women, and age ranges. All key results are similar throughout.

Exogenous versus endogenous job separations: Modeling direct switches between occupations as endogenous choices is a key feature of our approach compared to the standard exogenous mobility assumption. In Appendix D.3, we further expand this notion and construct a sample which also endogenizes non-employment transitions, adding all intermittent non-employment spells by imputing workers' wages and their occupation choices. We do this by comparing wages before and after the non-employment spell, and assign workers the lower of those two adjusted for inflation as well as the corresponding occupation's prices. That is, we assume that workers could have worked in the lower paying occupation all along but decided to become unemployed or exit the labor force for some time instead. Re-running the analysis on the sample constructed this way, we find that skill accumulation estimates are generally smaller and turn negative in cases that one would expect to be "downward" switches (e.g., from Mgr-Prof-Tech to Prod-Op-Crafts). Yet the other results are similar to before: The correlation between wage and employment growth is approximately zero. Price and employment growth have a strong positive relation, which is only slightly flatter compared to the main sample. The implied skill changes are thus negative and they continue to be closely related to growth-selection.

Different estimation specifications: Appendix D.4 reports on different model specifications mentioned in Section 3. First, we employ Heckman et al.'s (1998) flat spot approach, which assumes that mature workers' skill growth is flat. We thus set skill accumulation to zero for the 45–54 year old subsample. Our estimated coefficients depicted in Figure 5b lend support to this assumption. At the same time, possibly forward-looking choices—i.e., via picking occupations with high skill accumulation—should be less of a concern in this age group. This obtains a similarly steep relationship between skill price and employment changes as in the baseline estimation. Growth-selection also works in the same direction as the implied skill changes, although it is somewhat flatter. We then return to our main sample and modify the skill accumulation function. Whether we restrict it by making it age-independent, enrich it to be specific to education $\times$ age $\times$ occupation cells, or add potential experience to the latter, task price estimates are very similar (Figures D.12–D.14). Our tentative conclusion is that the functional form assumptions do not seem to significantly impact our results. Similar findings hold for extensions to the basic model, such as non-pecuniary benefits, considered in Böhm and von Gaudecker (2021).

Connecting to the task-based approach: A large literature has investigated occupational changes with the task-based approach (e.g., Autor et al., 2003; Firpo et al., 2013; Yamaguchi, 2018). For Germany, several authors have used the Qualifications and Career Surveys (QCS) to measure routine versus non-routine task content (e.g., Spitz-Oener, 2006; Antonczyk et al., 2009; Gathmann and Schönberg, 2010). In Appendix D.5, we also employ the QCS to construct routine as well as analytical, interactive, and manual task content for our 120 occupations. We then relate these measures to employment growth, wage changes, and our estimates for prices and skills. We find that occupations intensive in analytical (Mgr-Prof-Tech) and interactive (Mgr-Prof-Tech and Sales-Office) tasks indeed grew quite strongly, whereas employment in routine-intensive (Prod-Op-Crafts) occupations declined. High analytical and interactive task content predict rising wages, but the relation with skill prices is even steeper. Conversely, implied skills deteriorate in analytical and interactive task content. The correlation between routine task intensity and changing average wages is zero; this is composed of falling prices and rising skills. All this is consistent with the impact of RBTC on these occupations and with our finding that skill price changes are counteracted by selection effects. These results demonstrate the usefulness of task measures as a dimension-reduction device (e.g., Gathmann and Schönberg, 2010; Yamaguchi, 2012), particularly when working with more limited datasets.

5 Skill Prices and Wage Inequality

We have argued that selection effects are the reason why occupational wages and employment growth are uncorrelated over the period under study. By a similar token, selection may shroud the relation of demand shifts with wage inequality, particularly between occupations. In this section, we thus examine to what extent selection may be responsible for the result that occupations exhibit limited explanatory power in the increase of wage inequality. In Mincer-regressions, Card et al. (2013) obtain only small decreases of residual inequality when adding occupation dummies. Dustmann et al. (2009) find that in the lower half of the distribution, occupational demand is not a first-order factor driving wage differences. We first use only the estimated skill prices and selection to quantify the forces driving between-occupation inequality. We then employ the full model to disentangle components that affect various percentiles of the wage distribution,[14] paying particular attention to entry wages, demographics, skill accumulation, and prices.

[14]We study wage inequality as opposed to, e.g., overall inequality. It is thus worth noting that we do not see a clear trend in labor force participation rates of German men over most of the study period. There was a decline between 1975 and 1989 but rates stabilized around 93-94% thereafter. This is in contrast to U.S. men, where rates dropped from almost 94% to below 90% between 1989 and 2010.

5.1 The Attenuating Effect of Selection On Inequality

Over the period of our study, the variance of log wages multiplied with 100 went up by 12.4 points from a baseline of 14.3. The component due to differences between occupations started at a value of 5 in 1985. It more than doubled and reached almost 40% of the overall inequality in 2010. A substantial share of the increase thus occurred between occupations, consistent with occupational demand (e.g., due to routine-biased technical change and offshoring as in Acemoglu and Autor, 2011) driving inequality.

There are competing explanations. Several papers before us have asked to what extent demographic change has contributed to the increase in inequality. The first column of Table 3 reports on a counterfactual analysis similar to that of Figure 16 in Autor (2019).[15] Holding wages at their 1985 level, we reweight observations to match the distribution of age, foreigner status, and education in 2010.[16] This exercise answers the question: what if choices conditional on these observables and wages were constant at their 1985 levels, but the demographic structure had shifted to that of 2010 (due to population aging, increased immigration, and rising educational achievements of younger cohorts)? Quantitatively, the answer is similar to that of Autor (2019). The effects make up roughly a fifth of the total increase and a third of the increase in between-occupation inequality.

With the help of our model we can gain further insights into the components that have driven changes of between inequality. Denoting average wages (skills) in an occupation by $\bar{w}_k$ ($\bar{s}_k$), we can write the change between 1985 and 2010 as:

$$\Delta\sigma^2(\bar{w}_{k,t}) = \underbrace{2 \cdot \sigma(\Delta\pi_{k,t}, \bar{w}_{k,1985}) + 2 \cdot \sigma(\Delta\bar{s}_{k,t}, \bar{w}_{k,1985})}_{2\cdot\sigma(\Delta\bar{w}_{k,t},\bar{w}_{k,1985})} + \underbrace{\sigma^2(\Delta\pi_{k,t}) + \sigma^2(\Delta\bar{s}_{k,t}) + 2 \cdot \sigma(\Delta\pi_{k,t}, \Delta\bar{s}_{k,t})}_{\sigma^2(\Delta\bar{w}_{k,t})} \tag{12}$$

[15]The closest to his analysis adds occupational choice to the reweighting. The results from this are similar to Table 3 and in Appendix E.1. We prefer the specification here because age, foreigner status, and education are arguably all factors that mostly contribute to occupational supply as opposed to demand.

[16]To compute the weights, we follow DiNardo et al. (1996) using a logit model with 30 dummies for the ages between 25–54, a dummy for being permanently German, and three dummies for education status.

Table 3: Decomposition of the between-variance of wages, data and counterfactuals

		Counterfactuals			Actual
		Rewgt. age, foreign, educ.	Prices only	Prices + rewgt. age, foreign, educ.	Data + price estimates
Overall $\Delta\sigma^2(w_{i,t})$		2.41	5.13	9.56	12.41
Between $\Delta\sigma^2(\bar{w}_{k,t})$		1.74	5.13	8.88	5.25
$2 \cdot \sigma(\Delta\bar{w}_{k,t}, \bar{w}_{k,1985})$	$2 \cdot \sigma(\Delta\pi_{k,t}, \bar{w}_{k,1985})$	0.00	3.23	3.23	3.23
	$2 \cdot \sigma(\Delta\bar{s}_{k,t}, \bar{w}_{k,1985})$	0.53	0.00	0.53	-0.53
$\sigma^2(\Delta\bar{w}_{k,t})$	$\sigma^2(\Delta\pi_{k,t})$	0.00	1.89	2.36	1.76
	$\sigma^2(\Delta\bar{s}_{k,t})$	1.21	0.00	1.21	3.02
	$2 \cdot \sigma(\Delta\pi_{k,t}, \Delta\bar{s}_{k,t})$	0.00	0.00	1.55	-2.24

Notes: All values are multiplied with 100. The levels in 1985 are 14.3 (overall) and 5.0 (between). Based on specification with 120 occupations. $\bar{w}_{k,t}$ refers to the average wage in occupation k in year t. The counterfactual experiments are: *Rewgt. age, foreign, educ.*: take observations in 1985 and reweight them to match the 2010 distribution of these characteristics with weights computed following DiNardo et al. (1996) in order to obtain 2010 wages. *Prices only*: take individual wages in 1985 and add our estimated price changes to obtain 2010 wages. *Rewgt. age, foreign, educ. + prices*: Combine both experiments.

See Appendix E.1 for the detailed derivation. First consider the terms underneath the braces, which involve only wages. If between-occupation wage inequality was constant $\Delta\sigma^2(\bar{w}_{k,t}) = 0$ and the wage structure across occupations changed $\sigma^2(\Delta\bar{w}_{k,t}) > 0$, it must be that, on average, occupations at the bottom of the distribution experienced higher wage growth than those at the top $\sigma(\Delta\bar{w}_{k,t}, \bar{w}_{k,1985}) < 0$. The main terms in (12) break down these components into changes of prices and changes of skills. We apply this decomposition to counterfactual experiments to better understand the mechanisms at work.

The first column of Table 3 shows that the reweighting procedure only affects skills; all terms involving price changes are zero. The covariance of skill changes with baseline wage levels is positive. This reflects that the population grew older and more educated together with high-wage, high-education occupations (Mgr-Prof-Tech, Sales-Office) featuring faster skill accumulation over the life cycle. Quantatively, these changes in the demographic structure play a limited role for explaining between-occupation inequality.

The second column of Table 3 reports on the opposite experiment, which isolates the effect of price changes. Holding constant the 1985 demographics and occupation choices, we add the changes of occupational skill prices until 2010 to individuals' wages. This alone generates almost the entire increase of between-occupation inequality. The bulk of the effect stems from the covariance between price changes and initial wage levels. Prices rose in Mgr-Prof-Tech, Sales-Office, and Srvc-Care occupations. The size of the first two is by far larger and they featured high wages already in 1985. Our interpretation of this term is that it reflects the nature of demand shifts: during the period under study, they happened to benefit high-wage occupations on average.

In the third column, we turn on both experiments. The variance of price changes rises somewhat due to the different weights; all other effects from the separate experiments remain the same. The covariance between price and skill changes is substantial and positive. Overall, this counterfactual overestimates the rise of between-occupation inequality by two thirds.[17] Looking at the first line only, one may even be tempted to think that this exercise explains a large share of the total rise in inequality (77 %).

However, a comparison with the last column—the actual between-occupation variance and its components—reveals that this large "explained" share is far off, since there are important dampening effects of selection on wage inequality. In particular, the economic mechanism described at length in the previous section—a deterioration of skills in occupations where prices rose—has a strong impact on inequality. This strong negative covariance is not mechanical. If we interpret price changes as mainly driven by demand shifts and the demographic changes captured by the reweighting as shifts in supply to occupations, the third column of Table 3 suggests that these shifts covaried positively. The impact of -2.24 points in the actual data as well as the negative covariance of skill changes and initial wage levels are therefore important attenuating effects related to growth-selection. As a result, the actual contribution of skill changes to between-occupation inequality is negligible whereas in the counterfactual it is $+3.29$ points overall.

[17]We use skill prices estimated from the observed market equilibrium, which involves workers endogenously selecting into occupations based on these prices. In the counterfactual of column 3, labor supply is inelastic beyond the effect of observables. The partial equilibrium analysis will thus be biased, where the direction of bias depends on nature of the labor market. In competitive markets, without attenuating worker movements, prices should diverge even more strongly and counterfactual between-occupation inequality would rise even further. If occupations acted like monopsonists and only raised wages when supply is sufficiently elastic to actually attract new workers, the bias would be opposite.

What we just described is consistent with theoretical results by Heckman and Honoré (1990) for a two-sector Roy economy. Heckman and Honoré showed that, if the population distribution of skills is log concave, self-selection in the Roy model will generally lead to more equal wages compared to random assignment into occupations. In the particular case at hand, when the correlation of skills in the different occupations is sufficiently low, average skills in the occupation with declining prices will unambiguously improve and they will deteriorate in the occupation with rising prices.

These results show why decompositions based on observables alone have difficulties generating meaningful increases of inequality. So long as average wages across occupations hardly vary with employment growth (Figure 2a), changing demographics and even large shifts of employment across occupations exert limited impact. The reason is that underlying skill prices and supply changes, which would have raised between-occupation inequality further than what is observed, are counteracted by strong selection effects.

5.2 Factors Contributing to Wage Inequality

While the economic forces under scrutiny in this paper are most important for inequality between occupations, our model and the longitudinal SIAB data can be employed to gain a better understanding of the overall development of the wage distribution, too. In the following, we use our estimates to disentangle the factors that contributed to differences between the quantiles of the wage distribution.

Figure 8 plots the evolution of the percentiles of the wage distribution in the data and in various scenarios based on our model. Figure 8a just repeats Figure 1a for ease of comparison; it shows the widening of the German wage distribution (Dustmann et al., 2009; Card et al., 2013). Figure 8b plots the individual-level predictions from our model. To obtain an individual's predicted wage in a particular year, we start from the initial wage observed in the data[18] and follow his occupational choices over the life-cycle, adding the relevant skill accumulation parameters and price estimates along the way. The predictions

[18]For workers who may have entered the labor market before the sample starts—those born before 1950 observed to be working in 1975—we use the skill accumulation estimates to impute their initial wages at age 25, assuming they stayed in the same occupation all along.

track the data closely, both qualitatively and quantitatively. Note that all percentiles in all panels are normalized to zero in 1985; Table E.2 in the Appendix shows that the model is close to its targets also for the levels of these percentiles and the variance.

The remaining panels of Figure 8 investigate the drivers of the model prediction by starting with the most basic version and turning on features one after the other. Panel c reports how the three percentiles would have evolved if workers had kept their initial wages for the entire working life. Many variants of supply changes would be reflected in this scenario. For example, the expansion of tertiary education may lead to higher entry wages for the additional university graduates, raising the upper percentiles. The results show that the median and 85$^{\text{th}}$ percentile rose somewhat. Quantitatively, this is not very important, making up between one fifth (median) and one eighth (85$^{\text{th}}$ percentile) of the total increase. All three percentiles evolve smoothly; the distinct temporal patterns in Panels a and b are thus not driven by changing conditions at labor market entry. After a small initial increase, the 15$^{\text{th}}$ percentile exhibits a pronounced decline starting in the mid-nineties, which is so strong that it could explain the drop of that percentile between 1985 and 2010. This large drop seems due to temporary workers and naturalized citizens, who are often the same.[19] Excluding workers ever coded as foreigners from our sample reduces the fifteenth percentile drop by more than two thirds both in Panels b and c (see Appendix Figure E.1). This is consistent with Dustmann et al.'s (2009) hypothesis that, from the 1990s, low-skilled immigrants and ethnic Germans from Eastern Europe flowed into the West Germany, worsening the composition of employment at the lower end.

In Figure 8d, we continue to assign workers to their initial occupation, now adding the skill accumulation coefficients. There is hardly any change for the fifteenth percentile compared to Panel c, but the median and 85$^{\text{th}}$ percentile rise strongly. The incremental changes are 4 points at the median and 5 points at the 85$^{\text{th}}$ percentile, amounting to one half (one third) of the overall changes between 1985 and 2010. Again, all changes happen

[19]We identify temporary workers from the detailed occupation "assistant laborers", which mostly appears in the industry group "Credit and insurance intermediation, land and housing, rentals". This industry group contains the subgroup "labor recruitment and provision of personnel" where temporary agencies are listed. Temporary work has increased a lot in Germany (Eichhorst and Tobsch, 2013).

Figure 8: Wage inequality scenarios

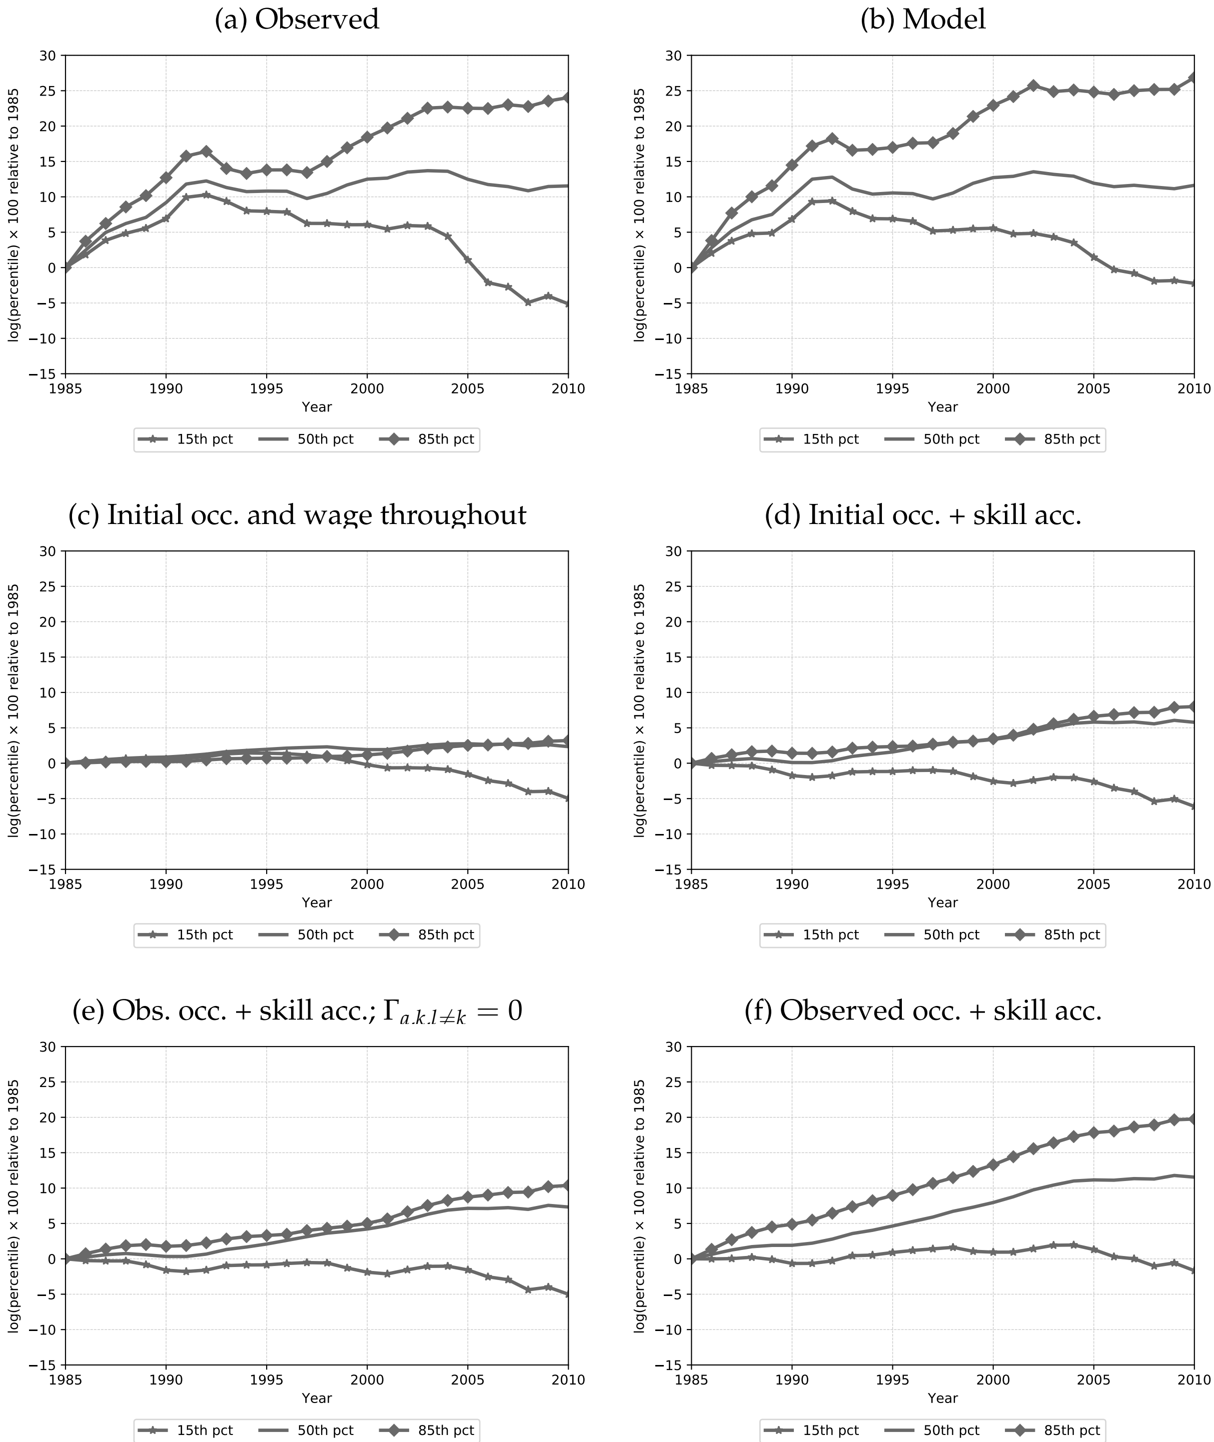

Notes: Panel a: Observed wages. *Panel b*: Simulated life-cycle trajectories based on our full model: Starting from the initial wage and occupational choice, add all skill accumulation and price change estimates using occupational choices observed in the data. *Panel c*: Workers keep their initial wage throughout the life cycle. *Panel d*: Workers stay in their initial job throughout the life-cycle; in each period, we add the skills they would have accumulated in that job (i.e., $\Gamma_{a(i,t_{i,0}),k(i,t_{i,0}),k(i,t_{i,0})}$). *Panel e*: Use observed switches, setting direct gains from switching to zero, i.e., $\Gamma_{a(i,t-1),k(i,t-1),k} = 0 \; \forall \; k \neq k(i,t-1)$. Price changes are zero as well, so the difference to Panel d comes purely from differential skill accumulation in occupations. *Panel f*: As in Panel e, but adding the direct gains from switching. The only difference to the full model in Panel b are the price changes, which continue to be zero. In all scenarios, we treat unemployment or out-of-the-labor force spells as follows: When such a spell is observed in the data, simulated workers do not enter the inequality statistics. Furthermore, we assume no depreciation and upon re-entry into paid work add—where relevant—$\Gamma_{a(i,\tilde{t}),k(i,\tilde{t}),k}$ with $\tilde{t}$ the period before the spell.

rather smoothly. The scenario shows that the demographic and occupational composition has a quantitatively strong impact on the rise of the upper half of the wage distribution.

We add the observed switches to careers in Figure 8e, but do not turn on the direct gains from switching, i.e., we set the off-diagonal elements of Γ to zero. This drives up the median and 85$^{\text{th}}$ percentile by an additional three points; it hardly affects the fifteenth percentile. The results show that switches from occupations with flatter age profiles to those with steeper profiles matter even if one ignores the oftentimes large jumps associated with switches. Yet, these skill accumulation differentials are not large enough to drive a majority of inequality. Part of this may be due to timing: For switches after age 35, accumulation differentials between occupation groups are smaller than early in careers. Still, the rise in the median and 85$^{\text{th}}$ percentile is large in Panels d and e compared to Panel c. In Appendix E.3, we show that this is due to the aging of the workforce with more middle-aged workers at the median in 2010 compared to 1985. Demographics were thus partly responsible for the increase of lower half inequality. In contrast, since the 85$^{\text{th}}$ percentile rose similarly to the median, upper half inequality did not increase much due to demographic changes or skill accumulation within occupations.

Adding the gains associated with occupation changes in Figure 8f raises all statistics and disproportionately affects the 85$^{\text{th}}$ percentile. This is not surprising given the large coefficient estimates for switches into Mgr-Prof-Tech and Sales-Office occupations. Comparing the end points of our sample period, this scenario explains three quarters of the increase in the 85$^{\text{th}}$ percentile and the entire increase in the median; we are too optimistic about the evolution of the fifteenth percentile by 3 points. There are two things to note. First, the temporal pattern is smooth and we do not track the intermittent evolution well. Second, there is a stronger decline in the 15$^{\text{th}}$ percentile when we make unemployment or exiting the labor force a choice by filling such spells with the lowest adjacent wage (Appendix Figure E.2). This suggests that careers at the lower end became more fragmented and our main way of treating non-employment spells hides parts of this.

Comparing Figures 8f and 8b shows that skill prices explain most of the remaining differences with the actual wage distribution. Changing prices raise the 85$^{\text{th}}$ percentile and upper half inequality by an additional seven log points. As for between-occupation inequality, they have a strong impact. In the specification with unemployment a choice, price changes hurt the median and the 15th percentile, again highlighting that we overestimate the gains from switching at the lower end because occupation changes involving wage losses often go via unemployment spells. Finally, adding the price changes allows us to track the evolution of all quantiles. Skill prices are not only aligned with employment across occupations; they also align the temporal patterns of the wage distribution.

In sum, we show that demographic factors and occupational skill accumulation account for most of the increase in lower half inequality. Alternative specifications suggest that more unstable employment biographies and adverse price developments have some role to play, too. This is consistent with the hypothesized effects in Dustmann et al. (2009), which is an important finding overall because it has previously been hard to rationalize polarizing demand for occupations together with wage inequality that increased across the board in most countries and time periods (Goos and Manning, 2007; Mishel et al., 2013; Green and Sand, 2015; Naticchioni et al., 2014). Occupational switches and changing skill prices have a particularly important role to play in the upper half of the wage distribution, driving almost all of the additional wedge that opened up between the 85th percentile and the median over the period 1985–2010.[20]

6 Discussion and Conclusion

This paper studies how occupational employment growth relates to occupational wages and overall wage inequality. We first establish stylized facts that indicate strong negative (positive) selection effects into growing (shrinking) occupations. We then estimate a model of occupation choice based on Roy (1951), which remains empirically tractable for many occupations and accommodates heterogeneous skill changes over the career. Res-

[20]The results in this section are robust to the more general acceleration/deceleration interpretation of skill price changes discussed in Section 3.1. First, wages in the full model, which include both skill accumulation and skill prices, are unaffected by this interpretation. Second, before the estimated prices are included, one would still like to add the average rates of price changes in the base period to the skill accumulation in order to obtain scenarios where “only” entry wages, initial occupations, or occupational switching changed. This is effectively what we do in Panels c–f of Figure 8.

ults indicate that skill-constant occupational wages *(skill prices)* evolved in a way that is consistent with occupational demand shifts. Skill selection of workers completely masks this relationship in raw occupational wages, the development of which is unrelated to employment changes. The systematic part of the skill price-employment growth nexus is due to what we term the *growth-selection effect*; net entry into an occupation multiplied with the skill differences between occupation entrants/leavers versus incumbents/stayers.

The selection effects we uncover are richer than those considered in classic Roy models, where workers' skills across occupations are fixed over time. The closest analogue to this classic Roy selection—skill endowments at the start of an occupation spell—contribute between 20% and 64% to growth-selection depending on the profession. Between 14% and 63% can be attributed to age selection, i.e., that marginal workers in growing occupations have had less time to accumulate (specific) skills. Finally, a dynamic version of Roy selection—positive deviations from average skill accumulation during tenure in an occupation, which raise the likelihood of staying put—contributes another 14–25%. The strong variation of relative magnitudes across occupations reflects the fact that occupational life-cycle wage profiles are very heterogeneous.

Similar lines of reasoning carry over to wage inequality, where we establish a quantitative connection to demand shifts and occupational employment changes. Relative to a situation where workers in the 1980s were given the skill prices of later decades while holding their occupational choices fixed, selection leads to lower wage inequality between occupations. Selection thus makes it appear that occupational changes were not very important. Using our model to understand the trends in overall wage inequality, we instead find that differentially evolving skill prices and heterogeneous skill changes across occupations are the most important drivers of upper-half inequality. Initial occupation choice (low-skilled entrants reducing the 15^{th} percentile) and demographic changes (more skills accumulated at the median) are the main drivers of lower-half inequality.

Our explanation seems consistent with other accounts of rising wage inequality in Germany. One of the most prominent is based on de-unionization and a decentralization of the wage bargaining process (Dustmann et al., 2009, 2014). These phenomena have the strongest impact in manufacturing, i.e., the industry sector that is most important for the declining Prod-Op-Crafts occupations. We deem it plausible that demand shifts are a deeper cause for this, as unions and works councils understand their deteriorating bargaining position due to the threats of substitution by machines or foreign workers (see Baumgarten and Lehwald, 2019, for direct evidence of the latter). Our findings are also consistent with work showing that German firms tend to upgrade labor through investment in skills (Battisti et al., 2019; Dauth et al., 2021).[21] These responses may reflect the institutional environment, which results in relatively cooperative labor relations in Germany. For example, unions and works councils are represented on boards of large companies and thereby involved in managerial decisions (Jäger et al., 2021).

References

ACEMOGLU, D. AND D. H. AUTOR (2011): "Skills, Tasks and Technologies: Implications for Employment and Earnings," in *Handbook of Labor Economics, vol. 4 Part B*, ed. by D. Card and O. Ashenfelter, Elsevier, chap. 12, 1043–1171.

ANTONCZYK, D., B. FITZENBERGER, AND U. LEUSCHNER (2009): "Can a Task-Based Approach Explain the Recent Changes in the German Wage Structure?" *Jahrbücher für Nationalökonomie und Statistik*, 229, 214–238.

AUTOR, D. (2019): "Work of the Past, Work of the Future," in *Richard T. Ely Lecture given at Annual Meeting of American Economics Association, Atlanta, GA, January*, vol. 4.

AUTOR, D. H. AND D. DORN (2013): "The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market," *American Economic Review*, 103, 1553–1597.

AUTOR, D. H., D. DORN, AND G. H. HANSON (2013): "The China Syndrome: Local Labor Market Effects of Import Competition in the United States," *American Economic Review*, 103, 2121–2168.

AUTOR, D. H., F. LEVY, AND R. J. MURNANE (2003): "The Skill Content of Recent Technological Change: An Empirical Exploration," *Quarterly Journal of Economics*, 118, 1279–1333.

BÁRÁNY, Z. AND C. SIEGEL (2018): "Job Polarization and Structural Change," *American Economic Journal: Macroeconomics*, 10, 59–89.

BATTISTI, M., C. DUSTMANN, AND U. SCHÖNBERG (2019): "Technological and Organizational Change and the Careers of Workers," Working Paper.

[21]We find few switches of occupations to be systematically associated with large losses, even if we fill intermittent spells of non-employment.

BAUMGARTEN, D. AND S. LEHWALD (2019): "Trade Exposure and the Decline in Collective Bargaining: Evidence from Germany," Working Paper.

BÖHM, M. J. (2020): "The price of polarization: Estimating task prices under routine-biased technical change," *Quantitative Economics*, 11, 761–799.

BÖHM, M. J. AND H.-M. VON GAUDECKER (2021): "The Performance of Recent Methods for Estimating Skill Prices in Panel Data," arXiv preprint, available at https://arxiv.org/abs/2111.12459.

CARD, D., J. HEINING, AND P. KLINE (2013): "Workplace Heterogeneity and the Rise of West German Wage Inequality," *Quarterly Journal of Economics*, 128, 967–1015.

CAVAGLIA, C. AND B. ETHERIDGE (2020): "Job polarization and the declining quality of knowledge workers: Evidence from the UK and Germany," *Labour Economics*, 66, 101884.

CORTES, G. M. (2016): "Where Have the Middle-Wage Workers Gone? A Study of Polarization Using Panel Data," *Journal of Labor Economics*, 34, 63–105.

DAUTH, W., S. FINDEISEN, J. SUEDEKUM, AND N. WOESSNER (2021): "The Adjustment of Labor Markets to Robots," *Journal of the European Economic Association*, 19, 3104–3153.

DEMING, D. J. (2017): "The Growing Importance of Social Skills in the Labor Market," *Quarterly Journal of Economics*, 132, 1593–1640.

DINARDO, J., N. M. FORTIN, AND T. LEMIEUX (1996): "Labor Market Institutions and the Distribution of Wages, 1973-1992: A Semiparametric Approach," *Econometrica*, 64, 1001–1044.

DUSTMANN, C., B. FITZENBERGER, U. SCHÖNBERG, AND A. SPITZ-OENER (2014): "From Sick Man of Europe to Economic Superstar: Germany's Resurgent Economy," *Journal of Economic Perspectives*, 28, 167–188.

DUSTMANN, C., J. LUDSTECK, AND U. SCHÖNBERG (2009): "Revisiting the German Wage Structure," *Quarterly Journal of Economics*, 124, 843–881.

EICHHORST, W. AND V. TOBSCH (2013): "Has Atypical Work Become Typical in Germany?" Working Paper.

FIRPO, S., N. M. FORTIN, AND T. LEMIEUX (2013): "Occupational Tasks and Changes in the Wage Structure," *Working Paper*.

GANZER, A., A. SCHMUCKER, P. VOM BERGE, AND A. WURDACK (2017): "Sample of Integrated Labour Market Biographies Regional File 1975-2014 (SIAB-R 7514)," Tech. Rep. 1, FDZ–Datenreport.

GATHMANN, C. AND U. SCHÖNBERG (2010): "How General is Human Capital? A Task-Based Approach," *Journal of Labor Economics*, 28, 1–49.

GLITZ, A. AND D. WISSMANN (2021): "Skill premiums and the supply of young workers in Germany," *Labour Economics*, 72, 102034.

GOOS, M. AND A. MANNING (2007): "Lousy and Lovely Jobs: The Rising Polarization of Work in Britain," *Review of Economics and Statistics*, 89, 118–133.

GOOS, M., A. MANNING, AND A. SALOMONS (2014): "Explaining Job Polarization: Routine-Biased Technological Change and Offshoring," *American Economic Review*, 104, 2509–2526.

GOTTSCHALK, P., D. A. GREEN, AND B. M. SAND (2015): "Taking Selection to Task: Bounds on Trends in Occupational Task Prices for the U.S., 1984-2013," Working Paper.

GOULD, E. D. (2002): "Rising wage inequality, comparative advantage, and the growing importance of general skills in the United States," *Journal of Labor Economics*, 20, 105–147.

GREEN, D. A. AND B. M. SAND (2015): "Has the Canadian Labour Market Polarized?" *Canadian Journal of Economics*, 48, 612–646.

GROES, F., P. KIRCHER, AND I. MANOVSKII (2014): "The U-Shapes of Occupational Mobility," *Review of Economic Studies*, 82, 659–692.

HECKMAN, J. J. AND B. E. HONORÉ (1990): "The Empirical Content of the Roy Model," *Econometrica*, 58, 1121–1149.

HECKMAN, J. J., L. LOCHNER, AND C. TABER (1998): "Explaining Rising Wage Inequality: Explorations with a Dynamic General Equilibrium Model of Labor Earnings with Heterogeneous Agents," *Review of Economic Dynamics*, 1, 1–58.

HSIEH, C.-T., E. HURST, C. I. JONES, AND P. J. KLENOW (2019): "The allocation of talent and U.S. economic growth," *Econometrica*, 87, 1439–1474.

JÄGER, S., B. SCHOEFER, AND J. HEINING (2021): "Labor in the Boardroom," *Quarterly Journal of Economics*, 136, 669–725.

KATZ, L. F. AND K. M. MURPHY (1992): "Changes in relative wages, 1963–1987: supply and demand factors," *The quarterly journal of economics*, 107, 35–78.

LAGAKOS, D., B. MOLL, T. PORZIO, N. QIAN, AND T. SCHOELLMAN (2018): "Life-Cycle Wage Growth Across Countries," *Journal of Political Economy*.

MAZZOLARI, F. AND G. RAGUSA (2013): "Spillovers from High-Skill Consumption to Low-Skill Labor Markets," *Review of Economics and Statistics*, 95, 74–86.

MCLAUGHLIN, K. J. AND M. BILS (2001): "Interindustry Mobility and the Cyclical Upgrading of Labor," *Journal of Labor Economics*, 19, 94–135.

MISHEL, L., H. SHIERHOLZ, AND J. SCHMITT (2013): "Don't Blame the Robots: Assessing the Job Polarization Explanation of Growing Wage Inequality," Working Paper.

NATICCHIONI, P., G. RAGUSA, AND R. MASSARI (2014): "Unconditional and Conditional Wage Polarization in Europe," Working Paper.

NEAL, D. (1999): "The complexity of job mobility among young men," *Journal of labor Economics*, 17, 237–261.

PAPAGEORGIOU, T. (2014): "Learning Your Comparative Advantages," *Review of Economic Studies*, 81, 1263–1295.

PAVAN, R. (2011): "Career Choice and Wage Growth," *Journal of Labor Economics*, 29, 549–587.

ROY, A. D. (1951): "Some Thoughts on the Distribution of Earnings," *Oxford Economic Papers*, 3, 135–146.

ROYS, N. AND C. TABER (2019): "Skill Prices, Occupations and Changes in the Wage Structure for Low Skilled Men," Working Paper.

SPITZ-OENER, A. (2006): "Technical Change, Job Tasks, and Rising Educational Demands: Looking outside the Wage Structure," *Journal of Labor Economics*, 24, 235–270.

YAMAGUCHI, S. (2012): "Tasks and Heterogeneous Human Capital," *Journal of Labor Economics*, 30, 1–53.

——— (2018): "Changes in Returns to Task-Specific Skills and Gender Wage Gap," *Journal of Human Resources*, 53, 32–70.

YOUNG, A. (2014): "Structural Transformation, the Mismeasurement of Productivity Growth, and the Cost Disease of Services," *American Economic Review*, 104, 3635–3667.