

Understanding occupational wage growth

Adrian Adermon

Simon Ek

Georg Graetz

Yaroslav Yakymovych

The Institute for Evaluation of Labour Market and Education Policy (IFAU) is a research institute under the Swedish Ministry of Employment, situated in Uppsala.

IFAU's objective is to promote, support and carry out scientific evaluations. The assignment includes: the effects of labour market and educational policies, studies of the functioning of the labour market and the labour market effects of social insurance policies. IFAU shall also disseminate its results so that they become accessible to different interested parties in Sweden and abroad.

Papers published in the Working Paper Series should, according to the IFAU policy, have been discussed at seminars held at IFAU and at least one other academic forum, and have been read by one external and one internal referee. They need not, however, have undergone the standard scrutiny for publication in a scientific journal. The purpose of the Working Paper Series is to provide a factual basis for public policy and the public policy discussion.

More information about IFAU and the institute's publications can be found on the website www.ifau.se

ISSN 1651-1166

Understanding occupational wage growth*

Adrian Adermon[†] Simon Ek Georg Graetz
Yaroslav Yakymovych

June 2, 2026

Abstract

We jointly estimate growth in occupational wage premia as well as time-varying occupation-specific life-cycle profiles for Swedish workers 1996–2013. Our novel identification strategy is based on re-centering of life-cycle profiles around their flat spot. We document a substantial increase in between-occupation wage inequality due to differential growth in premia, and show that changes in worker composition partly counteracted this trend. The association of wage premium growth and employment growth is positive, suggesting that premium growth is predominantly driven by demand-side factors. We also find that wage growth due to occupation-specific skill acquisition was more dispersed in the early years of the sample period. Our results are robust to varying the assumed flat spot over a reasonable range, as well as to allowing for occupation-level changes in returns to cognitive and psycho-social skills. The results suggest that Swedish wage setting institutions have not prevented wages and quantities from adjusting to technological change or consumer demand shifts.

Keywords: Wage Growth, Inequality, Occupational Mobility

JEL Classification: C23, J24, J31, J62

*We thank Michael Böhm, Stefan Eriksson, Peter Fredriksson, Lena Hensvik, Lisa Laun, Martin Nybom, Oskar Nordström Skans, as well as seminar participants at the Uppsala Center for Labor Studies and the Future of Labor workshop in Berlin for helpful suggestions. Adermon & Ek: Institute for Evaluation of Labor Market and Education Policy (IFAU). Graetz: University of Edinburgh. Yakymovych: Institute for Housing and Urban Research (IBF), Uppsala University. Yakymovych gratefully acknowledges funding from Handelsbankens forskningsstiftelser (Grant No. W23-0026).

[†]Corresponding author. Email: adrian.adermon@ifau.uu.se.

1 Introduction

The past four decades have seen systematic shifts in occupational employment across industrialized countries, with high- and low-paying occupations gaining at the expense of the middle. This is commonly interpreted as reflecting labor demand shifts induced by technological change, consumer demand, or offshoring. However, the impact of such occupation-level demand shifts on the wage structure is far from clear. First, occupations appear to play a minor role in driving changes in wage inequality, at least in terms of descriptive decomposition exercises. Second, occupational employment and wage growth typically do not feature a strong positive correlation. Finally, wage inequality trends differ substantially across countries, while occupational employment shifts are highly similar.¹

In this paper, we shed light on these puzzles by studying occupational wage growth in Sweden 1996–2013. Swedish employment shifts are similar to those elsewhere (Adermon and Gustavsson, 2015), but the wage structure is dramatically compressed compared to most other industrialized countries, and growth in inequality has been moderate and episodic (Graetz, 2020). We show that, as elsewhere, occupations do not appear to play an important role in basic decompositions of changes in wage inequality.

However, as has long been recognized, any analysis of occupational demand shifts and wages must address selection problems arising from workers’ systematic sorting into occupations (see for instance Roy, 1951; Acemoglu and Autor, 2011; Böhm, 2020). For example, a positive demand shock to computer programmers may manifest itself as an increase in the price paid for a unit of programming output. At the same time, this increased *occupational wage premium* draws in workers from other occupations, who may be less productive than incumbents, thus leaving *observed* wages approximately unchanged. We address selection by drawing on existing work focusing on occupation stayers (Cortes, 2016; Böhm et al., 2024), whose wage growth comes closer to the growth in premia as time-invariant skills are differenced out and the composition of workers is left unchanged.

Our main contribution is to propose a new method for addressing another major challenge when interpreting occupational wage growth: differences in the returns to experience across occupations (Deming, 2021) and over time (Deming and Noray, 2020). Such differences give rise to an identification problem whereby it is impossible to distinguish premium growth from returns to experience when comparing within-individual wage growth

¹The polarization of occupational employment in the US and elsewhere has been documented by Wright and Dwyer (2003); Goos and Manning (2007); Autor et al. (2006); and Goos et al. (2014). See in particular Adermon and Gustavsson (2015) for the Swedish case. Goos et al. (2014) provide evidence in favor of a technological explanation. Barany and Siegel (2018) emphasize structural change and consumer demand instead. In a decomposition exercise, Hoffmann et al. (2020) find only a minor role for occupations in driving rising wage inequality. Roys and Taber (2019), Böhm (2020), and Böhm et al. (2024) highlight the lack of a strong positive correlation between occupational employment and wage growth in the US and Germany.

across occupations. A theoretically motivated restriction that has been suggested as a solution to this identification problem is the concept of a “flat spot”, a point in the life-cycle when the derivative of human capital with respect to experience is zero (Heckman et al., 1998; Bowlus and Robinson, 2012). We propose a novel approach for implementing this restriction, namely to re-center the experience profiles around the flat spot. This leaves us with greater statistical power as we are not forced to restrict the sample to workers near the flat spot. More importantly, it allows us to estimate experience profiles separately for each occupation and point in time.

Not only does our method capture changes to experience profiles due to new technologies—for instance, technology may render workers’ skills obsolescent at uneven rates across occupations. The method also accounts for changes in the rate of occupation-specific human capital accumulation that result from worker selection. As highlighted by Böhm et al. (2024), when workers switch to fast-growing occupations, a higher fraction of entrants in those occupations means that on average, accumulation of occupation-specific skills proceeds at a faster rate. This is because the entrants start with a low level of occupation-specific skills and because there are diminishing returns to skill accumulation. While Böhm et al. (2024) use a pre-period to estimate time-invariant experience profiles and then use those to predict workers’ occupation-specific skills, our method instead estimates these profiles directly, allowing them to vary over time as well as by occupation.

A further contribution of our paper is to explore to what extent premium growth is driven by changes in occupation-specific skill returns. A growing literature documents changing skill returns in the aggregate, and suggests that occupations may be important drivers of such trends (Deming, 2017; Edin et al., 2022). Given the availability of cognitive and psycho-social skill measures from the Swedish military enlistment, we are able to control for differential changes in skill returns in our estimation.

Our findings are as follows. First, premium growth is positively correlated with employment growth (and more strongly so than is raw wage growth). Second, premium growth is also positively correlated with initial wages. These two findings together imply our third finding, namely that in the absence of compositional changes in the workforce, between-occupation wage inequality would have increased more than it actually has. Fourth, experience profiles vary strongly across occupations at any given point in time, and while they are stable in some occupations, in others they show large changes. Fifth, life-cycle profiles seem more dispersed and a greater contributor to growing inequality in the early part of our sample period, the late 1990s.

The positive association between premium growth and employment shifts suggests that variation in premium growth is mostly due to labor demand factors. At the same time, our results point toward an important life-cycle component to shifts in the occupational wage structure. While the former builds up gradually over our sample period, the latter appears

especially relevant in the late 1990s, when wage coordination within the Swedish collective bargaining system was relatively weak. This suggests that wage setting institutions may be more important for life-cycle wage growth than for common occupational wage premia, although alternative explanations related to labor market imperfections are possible. We revisit this issue in the conclusion.

Our results are robust to allowing for changes in specific returns to cognitive and psycho-social skills. They are also robust to alternative, reasonable choices of the flat spot. By construction, our method identifies premium growth with wage growth at a level of potential experience equal to the flat spot. Given that wage growth generally declines in potential experience, a higher flat spot implies a lower estimate of premium growth. Our baseline choice is 30 years. Indeed, setting the flat spot at 25 years suggests a doubling of the contribution of the variance in premium growth to the overall increase in the variance of log wages. Re-assuringly, however, setting the flat spot to 35 years only reduces the contribution by about 25 percent. Moreover, choosing occupation-specific flat spots slightly increases the contribution.²

Our method highlights the need to account for—and re-center—experience profiles. Using average wage growth among all occupation stayers to estimate premium growth, thus ignoring differences in returns to experience across occupations, yields an estimate of premium-driven growth in wage inequality roughly three times larger than that implied by our baseline specification. Even more strikingly, controlling for experience profiles using only higher-order polynomial terms without re-centering, which amounts to setting the flat spot at zero, results in an estimated contribution of premium growth to the growth in wage inequality that is around 25 times that of our baseline model.

Our findings are consistent with a recent and growing literature documenting the importance of compositional changes in counteracting occupation-level demand shifts (Cortes, 2016; Böhm, 2020; Cavaglia and Etheridge, 2020; Böhm et al., 2024). Our contributions compared to these studies are first, to allow for—and to estimate—time-varying occupation-specific experience profiles when estimating wage premium growth. Second, to provide estimates of premium growth that are robust to time-varying occupation-level skill returns. And third, to provide evidence on occupational demand shifts for the Swedish economy, which at first glance features a very different wage structure and wage setting institutions than the economies studied in the literature (Bhuller et al., 2022; Mogstad et al., 2025).

We also conduct a complementary analysis using data from the U.S. Current Population Survey in Section 4.5. Unfortunately, these estimates are too imprecise to allow any strong

²We set occupation-specific flat spots in a data-driven way by minimizing the sensitivity of variance components to small changes in the flat spot. The procedure requires strictly concave experience profiles (in log wages) apart from possible flat regions.

conclusions.

The remainder of the paper is structured as follows. Section 2 presents our theoretical framework, discusses identification challenges as well as our proposed solutions. We describe our data in Section 3. Section 4 presents our results along with a decomposition and counterfactual scenarios, and Section 5 concludes.

2 Theoretical framework and empirical strategy

The theoretical motivation for our empirical exercise is the standard Roy model in which workers sort into occupations based on comparative advantage. Rather than estimating a fully specified model, our point of departure is an assumption about the data-generating process for potential wages. In this section we explore how key parameters of the wage equation can be identified under different assumptions about occupational choice. In Section 4.3, we show how changes in overall wage inequality can be attributed to occupation-level driving forces, and develop counterfactual scenarios based on our estimated wage equation.

2.1 Identifying the parameters of the wage function

Suppose that individual worker i 's log wage in occupation k and year t , w_{ikt} , is given by

$$w_{ikt} = \pi_{kt} + \alpha_{ik} + \mathbf{s}_i' \boldsymbol{\beta}_k + g_k(x_{ikt} - x^*) + \varepsilon_{ikt}, \quad (1)$$

where π_{kt} is a potentially time-varying occupation-specific wage premium; α_{ik} is an unobserved worker-occupation fixed effect; \mathbf{s}_i is a vector of observable skills with its associated occupation-specific returns $\boldsymbol{\beta}_k$; x_{ikt} is the worker's experience in the occupation measured in years and centered around x^* , to be discussed below; g_k is an occupation-specific experience profile; and ε_{ikt} is an i.i.d. shock. Our main goal is to estimate π_{kt} for each occupation, or at least its *change relative to a reference occupation*.

For the moment, let us assume that workers choose the occupation in which they earn the highest wage in each period, abstracting from dynamic considerations. Furthermore, let us assume for now that the shock ε_{ikt} is realized *after* workers have made their choice. These assumptions are the same as in Cortes (2016). This leaves us with two potential threats to identification: Selection on unobserved time-invariant characteristics, and occupation-specific experience profiles. We address these in turn.

2.2 Selection on time-invariant characteristics

Consider the first difference of equation (1),

$$\Delta w_{ik} = \Delta \pi_k + g_k(x_{ikt} - x^*) - g_k(x_{ikt} - 1 - x^*) + \Delta \varepsilon_{ik}, \quad (2)$$

where Δ is the first difference operator, so that $\Delta X \equiv X_t - X_{t-1}$. If we estimate equation (2) using the sample of occupation stayers, we can be sure that selection on time-invariant skills α_{ik} and s_i is accounted for, since these terms are differenced out. An alternative method accomplishing this is of course to estimate equation (1) in levels and to include worker-by-occupation fixed effects, as in Cortes (2016). We prefer the first difference specification for two reasons. First, it allows us to run separate regressions for each year, and thus work with datasets of manageable size. Second, our data on wages and occupations come from repeated cross-sectional samples, so that it is difficult to construct long panels of individuals workers, and to accurately capture longer occupational spells (see Section 3).

2.3 Occupation-specific experience profiles

For concreteness, we approximate the profile by a polynomial of order M , $g_k(x) = \sum_{m=1}^M \gamma_{km}(x - x^*)^m$. Under this assumption, the component of wage growth due to experience—among occupation stayers—now becomes

$$g_k(x_{ikt} - x^*) - g_k(x_{ikt} - 1 - x^*) = \gamma_{k1} + \sum_{m=2}^M \gamma_{km} \{(x_{ikt} - x^*)^m - (x_{ikt} - 1 - x^*)^m\}.$$

The wage growth equation to be estimated is thus

$$\Delta w_{ik} = \Delta \pi_k + \gamma_{k1} + \sum_{m=2}^M \gamma_{km} \{(x_{ikt} - x^*)^m - (x_{ikt} - 1 - x^*)^m\} + \Delta \varepsilon_{ik}. \quad (3)$$

Estimation of equation (3) for a given occupation yields a constant term $\theta_{kt} = \Delta \pi_k + \gamma_{k1}$. Thus, the challenge is to separate out changes in premia from the constant term of the experience profile. Note that γ_{k1} is the effect of additional experience at the point $x_{it} = x^*$. Human capital theory (Ben-Porath, 1967; Heckman et al., 1998) suggests that there comes a point in a worker's life cycle when human capital accumulation stops, or even reverses due to depreciation—a so-called flat spot where the marginal effect of experience on wages is zero. Thus, if x^* is set to be at the flat spot, then $\gamma_{k1} = 0$, solving the identification problem as we now have $\theta_{kt} = \Delta \pi_k$.³

³Our approach is related to Fosse and Winship (2019), who address the identification problem arising in the presence of age, cohort, and time effects. They highlight that it is only linear effects that are unidentified, and explain how one can bound these. However, a single restriction is often sufficient for

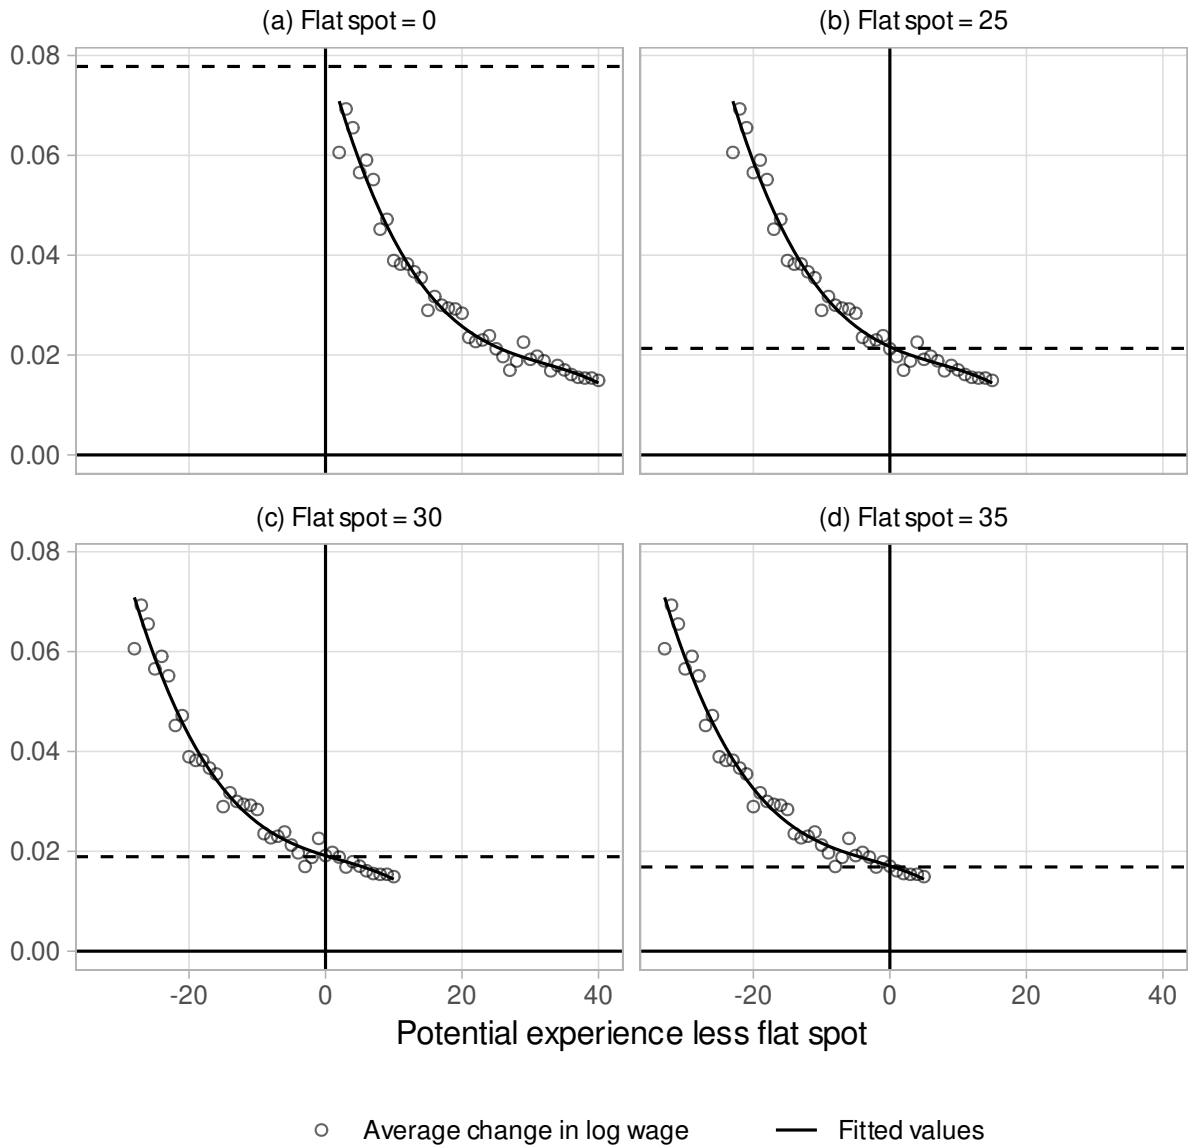


Figure 1: Illustration of flat spot identification

Notes: Dashed lines mark the constant term from estimating the experience profiles, equal to wage premium growth under the respective flat spot assumptions. The data include all individuals who worked as physical and engineering science technicians in 2005 and 2006. See Section 3 for further details on sample selection.

We illustrate this strategy using a concrete example: the wage growth of *physical and engineering science technicians* from 2005–06. Figure 1 plots changes in log wages, together with the fitted polynomial, against potential experience re-centered around different values—the assumed locations of the flat spot. The fitted polynomial comes from estimating equation (3) choosing $m = 4$. Grey dashed lines mark the constant term estimated by the regressions, equal to premium growth under the assumption $\gamma_{k1} = 0$.⁴ The data reveal a strong downward trend in wage growth, consistent with faster skill accumulation among inexperienced workers, as well as a flattening of this relationship at higher levels of potential experience. The top-left panel does not re-center the data, thus yielding a large estimated premium growth of around 8 percent. But an assumption of zero skill accumulation for labor market entrants is of course highly implausible. Assuming flat spots at higher values such as 25, 30, or 35 all yield estimated premium growth around 2 percent, as shown in the remaining panels.

Figure 1 illustrates that choosing the flat spot means picking a point on the fitted first-differenced experience profile and attributing all wage growth at that point to growth in the premium. Relying on a parametric prediction for the profile yields greater statistical power compared to simply using average wage growth at the flat spot.

Figure 1 also raises the question whether the flat spot can be determined in a data-driven way. In general, the answer is no. We illustrate the issue in Figure A.1, which plots three simulated experience profiles in levels and changes. While the flat spot is readily observable in levels, we only observe the first-differenced data. As the figure shows, looking for flat regions in changes will not work—while these *might* correspond to the vicinity of the flat spot (column (a)), they could also arise from a locally linear section of the experience profile (column (b)), or not arise at all (column (c)).

Given that the flat spot cannot be identified without further assumptions, our approach is to set it at 30 for all occupations, while also reporting results for alternative values. In a further robustness check, we estimate flat spots under the additional assumption that the true profiles are strictly concave except for possible flat sections (that is, linear segments with non-zero slope, as in the middle column of Figure A.1, are prohibited). In this case, the second derivative of the profile will be maximized (closest to zero) at the flat spot, so that any statistic of interest should change by the least amount—in absolute value—at the true flat spot. See Appendix B for further details.

A key advantage of our method is that it allows us to jointly estimate experience

point identification, as is the case in our context.

⁴To be precise, the flat spot assumption says that $g'_k(x^*) = 0$. In the polynomial case, $g'_k(x^*) = \sum_{m=1}^M \gamma_{km} m(x - x^*)^{m-1} = \gamma_{k1}$. The flat spot assumption $\gamma_{k1} = 0$ concerns the first derivative, not the first difference. Therefore, the assumption does not imply that $\Delta w_{ik} |_{x=x^*} = \Delta \pi_k$ exactly, which requires $\sum_{m=2}^M \gamma_{km} \{-(-1)^m\} = 0$. However, in practice these equations will hold approximately, as is the case in Figure 1. This justifies the heuristic outlined in the text.

profiles and premium growth. Moreover, as we estimate separate models for each year, we essentially estimate time-varying experience profiles.⁵ A third advantage is that we retain greater statistical power than existing approaches in the literature which implement the flat spot idea using only data on workers near the flat spot (Bowlus and Robinson, 2012; Cavaglia and Etheridge, 2020).⁶

To quantify the gain in statistical precision from our estimation procedure, we estimate premium growth non-parametrically by the average wage growth of individuals with experience levels in the vicinity of the flat spot (29.5–30.5 years of experience, to be precise). The average standard errors are around 2.5 times as large as those from our estimator (see Appendix C for details).

We note that the interpretation of our estimated profiles is affected by the measurement of occupation-specific experience. With a panel that is relatively short (20 years) relative to the typical length of working lives, it is not possible to construct complete occupational histories for each worker.⁷ In our baseline specification we therefore use potential overall labor market experience, based on age and years of schooling. Given this, the occupation-specificity of the γ_{km} 's means that experience is differently valued across occupations, but it does not matter in which occupation this experience was gained. Alternatively, one can simply interpret the estimated profiles as describing the wage growth in a given occupation and year as a function of potential overall labor market experience. This function will depend not only on deep structural parameters, but also on the characteristics—such as occupational histories—of the workers staying in that occupation in that year (and the year before).

2.4 Time-varying skill returns

A key finding in recent research on inequality is that wage returns to various skills have evolved differently over time, with occupations seemingly playing an important role (Deming, 2017; Edin et al., 2022). While this is interesting in its own right, here we are mainly concerned with the impact of such changes on our ability to estimate changes in occupational wage premia. Suppose, then, that returns to portable skills vary over time,

$$w_{ikt} = \pi_{kt} + \alpha_{ik} + \mathbf{s}'_i \boldsymbol{\beta}_{kt} + g_k(x_{ikt} - x^*) + \varepsilon_{ikt},$$

⁵Strictly speaking, the experience profile in *levels* must be constant across the two adjacent years. This would not matter if we had specified the profile in changes in the first place. However, starting with a levels specification is arguably more natural given the Roy framework.

⁶It is also possible to implement flat spot identification via an iterative procedure (Lagakos et al., 2017).

⁷Another challenge is that information on occupation and wage rates of Swedish private sector workers is only available for large cross-sectional samples, as we discuss in the data section.

so that wage growth now becomes

$$\Delta w_{ik} = \Delta \pi_k + \mathbf{s}'_i (\Delta \boldsymbol{\beta}_{kt}) + \sum_{m=2}^M \gamma_{km} \{(x_{ikt} - x^*)^m - (x_{ikt} - 1 - x^*)^m\} + \Delta \varepsilon_{ik}. \quad (4)$$

For selected cohorts of Swedish men we actually have at our disposal the skill measures for which changing wage returns have been documented. We can thus assess whether our baseline estimates of $\Delta \pi_k$ are robust to controlling for these measures, by estimating equation (4) where the vector \mathbf{s}_i contains cognitive and psycho-social skills, as described further in the data section.

2.5 Selection on idiosyncratic shocks

Let us now allow for selection on the idiosyncratic shock ε_{ikt} , as well. The constant term from estimating equation (3), imposing the flat spot assumption $\gamma_{k1} = 0$, now becomes $\theta_{kt} = \Delta \pi_k + E[\Delta \varepsilon_{ik} \mid k_{it} = k_{i,t-1} = k]$. The second term no longer equals zero, due to selection. Other things equal, occupations experiencing relatively fast premium growth will retain more workers with a bad realization of the shock, while occupations in which premia decline only retain those workers with very good realizations. Therefore, selection on idiosyncratic shocks biases downward the between-occupation variance in premium growth. This bias is more severe the larger is the variance of ε_{ikt} . A method to correct for this bias, developed by Böhm et al. (2024), is to include occupation switchers in a regression of wage growth on workers' average choices. We implement this method as a robustness check.

2.6 Remaining issues

There are a number of issues which are beyond the scope of this paper. These include forward-looking occupational choice, amenities, search frictions, and long-term wage contracts. We believe that addressing any one of these requires estimation of a fully specified structural model (for recent examples, see Roys and Taber, 2019; Traiberman, 2019).

3 Data description

3.1 Data sources

We obtain demographic information (year of birth, sex, municipality of residence, education, immigration status) for the population of Swedish residents in 1985–2013 from Statistics

Sweden’s LISA database.⁸ LISA also contains information on employment status in November of each year, annual labor income, as well as industry and municipality of the individual’s workplace.

Information regarding weeks and hours worked, as well as occupation, is obtained from the Swedish Wage Structure Statistics (WSS). Occupations are coded according to the *SSYK96* classification.⁹

We also use information on contractual monthly wage rates from the WSS. This in combination with annual labor income allows us to determine annual labor supply. Most importantly, these contractual wage rates are the main outcome of interest for our analysis, since we are interested in the price of labor.

A drawback of WSS is that outside the public sector, we do not observe the entire population of workers but only a large sample. Sampling is stratified by firms, with large firms being more likely to be drawn. This does not pose any problems for cross-sectional analyses, as sampling weights are provided. It does introduce some difficulties when analyzing dynamic phenomena such as occupational mobility. We discuss this issue further in Section 3.2 below.

Table A.1 shows summary statistics for the first and last pairs of years in our data. We show statistics both for the raw sample, consisting of all individuals with non-missing wages in the WSS, and for our final analysis sample. Individual characteristics are largely unchanged by our sample restrictions, the exception being the occupational shares which increase moderately for Managerial-Professional-Technical occupations and shrink for the others.

Table A.2 shows how each sample restriction affects our final sample size. The most impactful restrictions are the matching of individuals across adjacent years and the restriction to occupation stayers.

3.2 Sample selection and construction of variables

Our population of interest includes all Swedish employees aged 18–64 during the years 1996–2013. Employees are individuals who are employed in November and whose annual labor earnings correspond to at least one “price base amount”.¹⁰ We calculate individual wage growth for all adjacent years, dropping anyone with wage growth below the first or

⁸We end the analysis in 2013 due to a change in the occupational classification in 2014, which generates breaks in employment trends even at higher levels of aggregation.

⁹SSYK (short for *Standard för Svensk Yrkesklassificering*), the *Swedish Standard Classification of Occupations*, is a national version of the International Standard Classification of Occupations (ISCO).

¹⁰The price base amount is used as a base for calculating a large number of social security, taxation and legal payments in Sweden and is indexed to CPI inflation. During 1996–2013, a price base amount corresponded to the 3^d–4th percentile of the annual earnings distribution among employed individuals with positive earnings.

above the 99th percentile for each pair of years.¹¹

We calculate potential labor market experience as years elapsed since year of graduation, based on the highest attained level of education and a school starting age of six. To reduce noise, we drop observations with potential experience below two and greater than 40 years. Due to the limited length of the panel as well as due to sampling, we are unable to construct actual occupation-specific experience.

We use sampling weights to adjust for stratification. The raw weights supplied in WSS feature some extremely large values, and this may introduce noise, especially when multiplying the weights for a first-difference analysis using a two-year panel. Whenever we work with individual, two-year panel data, we therefore trim the weights following the procedure of Potter (1990).¹² However, when computing aggregate moments, we use the original weights.

For our baseline analysis we use the 3-digit-level *SSYK96* occupational classification, which includes 101 occupations. However, we sometimes use a coarser classification for descriptive and other purposes.¹³

4 Results

4.1 Raw wages, wage premia, and employment

To set the stage, we document the relationship between growth in average wages and growth in employment as well as initial wages, across occupations for the period 1996–2013. Panel (a) of Figure 2 plots the long difference in log wages against the long difference in the log of employment, with each marker representing one occupation. First, by moving along the horizontal axis, we see much variation in employment growth. Production, operators, and craft occupations tend to see low (often negative) employment growth, while on average, employment growth appears highest among managers, professionals, and technicians. Clerical and services occupations fall somewhere in between. However, there is much variation even within these broad categories. Turning to wage growth, there is a positive but rather weak relationship with employment growth. Panel (a) of Figure 3 reveals an even weaker relationship of average wage growth with initial (1996) average wages.

¹¹Extreme values of wage growth—five or more standard deviations away from the mean—may occur because individuals enter into and exit from executive positions (Skans et al., 2009). We drop extreme values as these can have a large impact on the results.

¹²The procedure is as follows. We first fit a $\text{Beta}(\alpha, \beta)$ distribution to the weights. Second, weights whose estimated cumulative probability is above 99 percent are trimmed to the estimated 99th percentile. Third, weights are re-scaled such that their sum is unchanged. This procedure is repeated ten times.

¹³This consists of four broad occupation groups: Managers, professionals and technicians; Clerical occupations; Production, operators and craft occupations; and Service and elementary occupations.

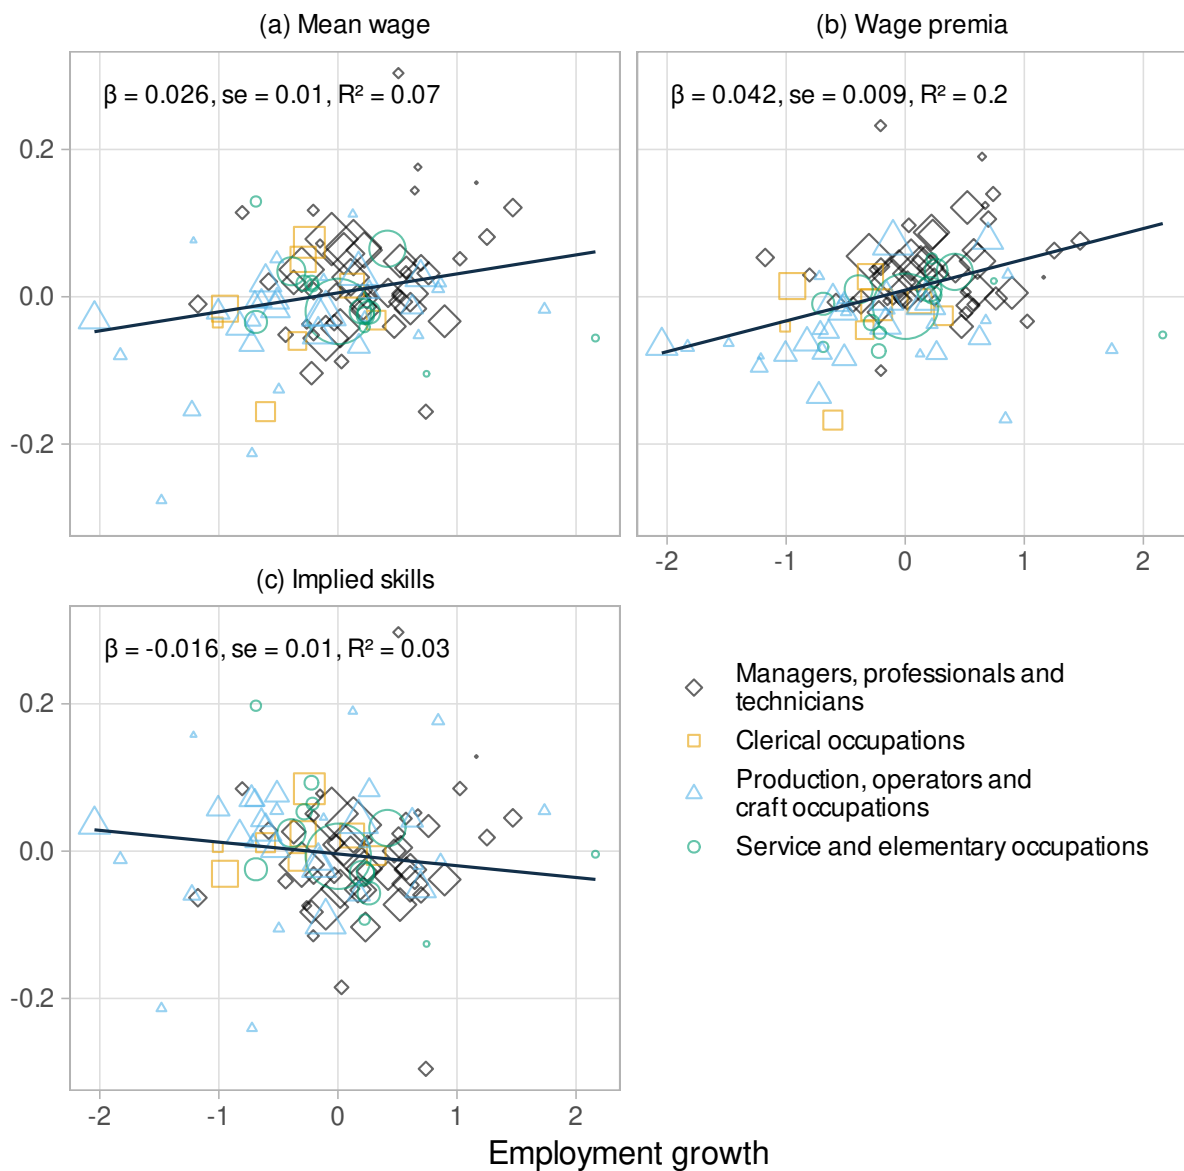


Figure 2: Growth in wages, premia, and skills against employment growth, 1996–2013

Notes: The figure plots the growth mean log wages, cumulative estimated wage premia, as well as the implied change in mean skills, against the change in log employment. Wage premia are estimated according to our baseline specification equation (5). Each marker represents one of 101 occupations. The size of each marker is determined by the employment share in the first year and the regression line is weighted accordingly. We use original survey weights when calculating occupation size and mean log wage.

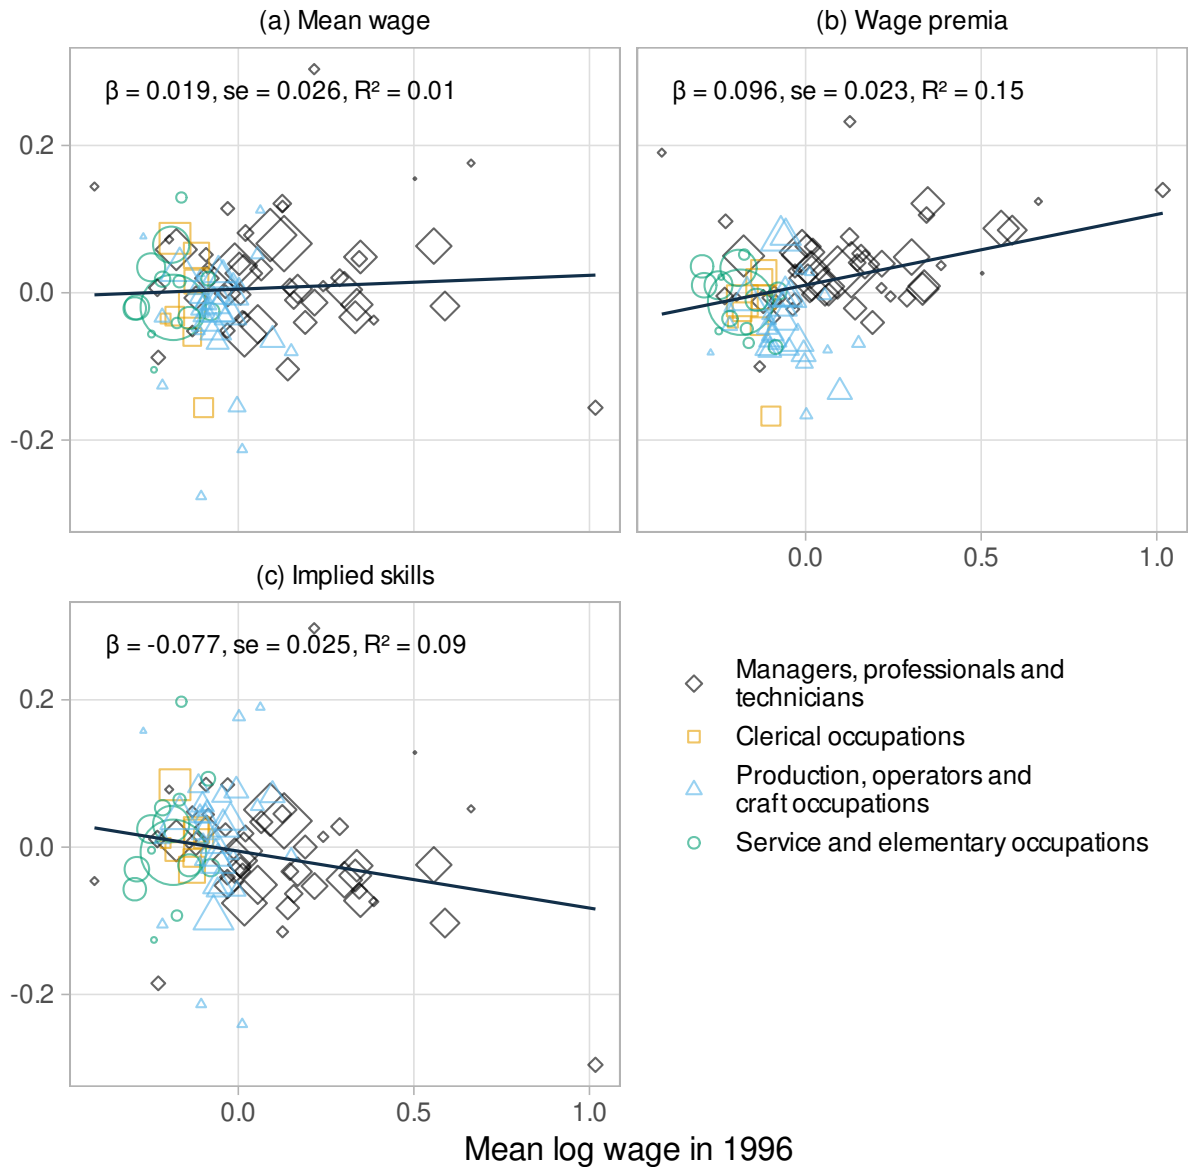


Figure 3: Growth in wages, premia, and skills against initial wages, 1996–2013

Notes: The figure plots the growth mean log wages, cumulative estimated wage premia, as well as the implied change in mean skills, against initial mean log wages. Wage premia are estimated according to our baseline specification equation (5). Each marker represents one of 101 occupations. The size of each marker is determined by the employment share in the first year and the regression line is weighted accordingly. We use original survey weights when calculating occupation size and mean log wage.

However, as discussed in Section 2, average wage growth captures both changes in occupational wage premia and changes in worker composition and hence average skills. In order to isolate changes in wage premia, we operationalize equation (2) by estimating separate regressions of year-on-year changes in individual log wages on occupation fixed effects and a polynomial in potential experience:

$$\Delta w_{it} = \varphi_{kt} + \sum_{m=2}^4 \gamma_{km} \{ \tilde{x}_{it}^m - (\tilde{x}_{it} - 1)^m \} + u_{it}, \quad (5)$$

where φ_{kt} are occupation-specific fixed effects; $\tilde{x}_{it} \equiv x_{it} - x^*$ is potential experience re-centered around the assumed flat spot in the experience profile; and γ_{km} are polynomial coefficients allowed to vary by occupation. In our main specification, we use a fourth-order polynomial and re-center potential experience at 30 years. We report robustness checks with respect to these choices below. We estimate separate regressions for each pair of adjacent years in our sample. In order to control for changes in worker composition, we use only individuals who remained in the same occupation across both years, $k_{it} = k_{i,t-1}$.

Under the assumption that there is no selection on idiosyncratic shocks and that $\gamma_{k1} = 0$ (the flat spot assumption), the fixed effects φ_{kt} consistently estimate premium growth $\Delta\pi_k$ for an adjacent pair of years. We estimate premium growth over the full period by simply accumulating the estimated year-on-year changes, $\widehat{\Delta}_{1997}^{2013} \pi_k = \sum_{t=1997}^{2013} \widehat{\varphi}_{kt}$.

Our premium growth estimates are plotted against employment growth in Panel (b) of Figure 2. The relation between premium growth and employment growth is stronger than that of mean wage growth—the slope is steeper, and R^2 almost triples. This pattern implies that while demand factors were pushing up wage premia during this period, changes in the skill composition of workers acted as a counteracting force, resulting in the tempered trend we see in average wage growth. This is consistent with a situation where growing labor demand in certain occupations attracts new workers with lower productivity than the incumbents—and conversely, occupations with falling labor demand might let their lower-productivity workers go first. The implied change in skill composition can be backed out from our estimates by simply subtracting the estimated changes in premia from the observed changes in average wages. This is shown in Panel (c) of Figure 2. As expected, faster growing occupations have seen falling implied skill levels in their workforce, although this relationship is not very strong.

Panel (b) of Figure 3 shows that premium growth is strongly positively associated with initial wages. Given equation (7), this suggests that premium growth would cause an increase in between-occupation wage inequality in the absence of compositional changes. However, panel (c) of Figure 3 already gives an idea of how strong these compositional changes might be—growth in average skills are strongly negatively related to initial wages. We explore these issues in detail in Section 4.3.2.

One way to assess the plausibility of the estimated growth in skills is to check its association with changes in years of schooling. Panel (c) of Figure A.2 shows that there is indeed a positive relationship, with a fairly high R^2 of 0.2. On the other hand, panel (b) of the same figure shows a negative association between premium growth and changes in years of schooling, consistent with lower educated workers moving into occupations experiencing positive demand changes.¹⁴

4.2 Changes in occupational experience profiles

A key advantage of our empirical approach is that we are able to estimate occupational experience profiles that vary over time. We estimate profiles for 101 occupations and each pair of years from 1996–2013. For the sake of conciseness and clarity, we only show estimated profiles for the largest (in terms of average employment 1996–2013) 3-digit occupation in each of nine main categories, for the years 1997, 2002, 2008, and 2013.

The estimated profiles are shown in Figure 4. There are several noteworthy findings. First, in all occupations wage growth is fastest for inexperienced workers, but this pattern is much more pronounced in some occupations (finance & sales professionals, building frame workers) than in others (personal care workers). Second, while in some occupations the profiles are stable (building frame workers), in others they show large changes over time (computing professionals). Third, profiles are steepest in the late 1990s in several cases, but this not a universal pattern.

To further investigate changes over time, we plot the median as well as quartiles of two measures capturing the steepness of the profiles, namely, the value of the profile at ten years of potential experience as well as the maximum value (both in levels and relative to the value at zero experience). Panel (a) of Figure A.4 reveals that, by both measures, profiles were indeed somewhat steeper in the late 1990s. But even more striking is that the steepness of the profiles was much more dispersed in that period.

Finally, we explore if there is a systematic relationship between dispersion in wage-experience profiles and dispersion in wage premium growth. Panel (b) of Figure A.4 plots the variances of the two steepness measures along with the variance of premium growth against time. It appears that years with higher dispersion in profiles also tend to see higher dispersion in premium growth.

4.3 Occupational drivers of changes in wage inequality

A key objective of this paper is to assess the quantitative importance of occupations for the evolution in wage inequality. Therefore, we need to formally characterize how changes

¹⁴For completeness, Figure A.3 displays the respective bi-variate associations of wage growth, premium growth, and implied skill growth, showing positive correlations between wage growth and premium growth, and wage growth and skill growth, and a negative correlation between premium growth and skill growth.

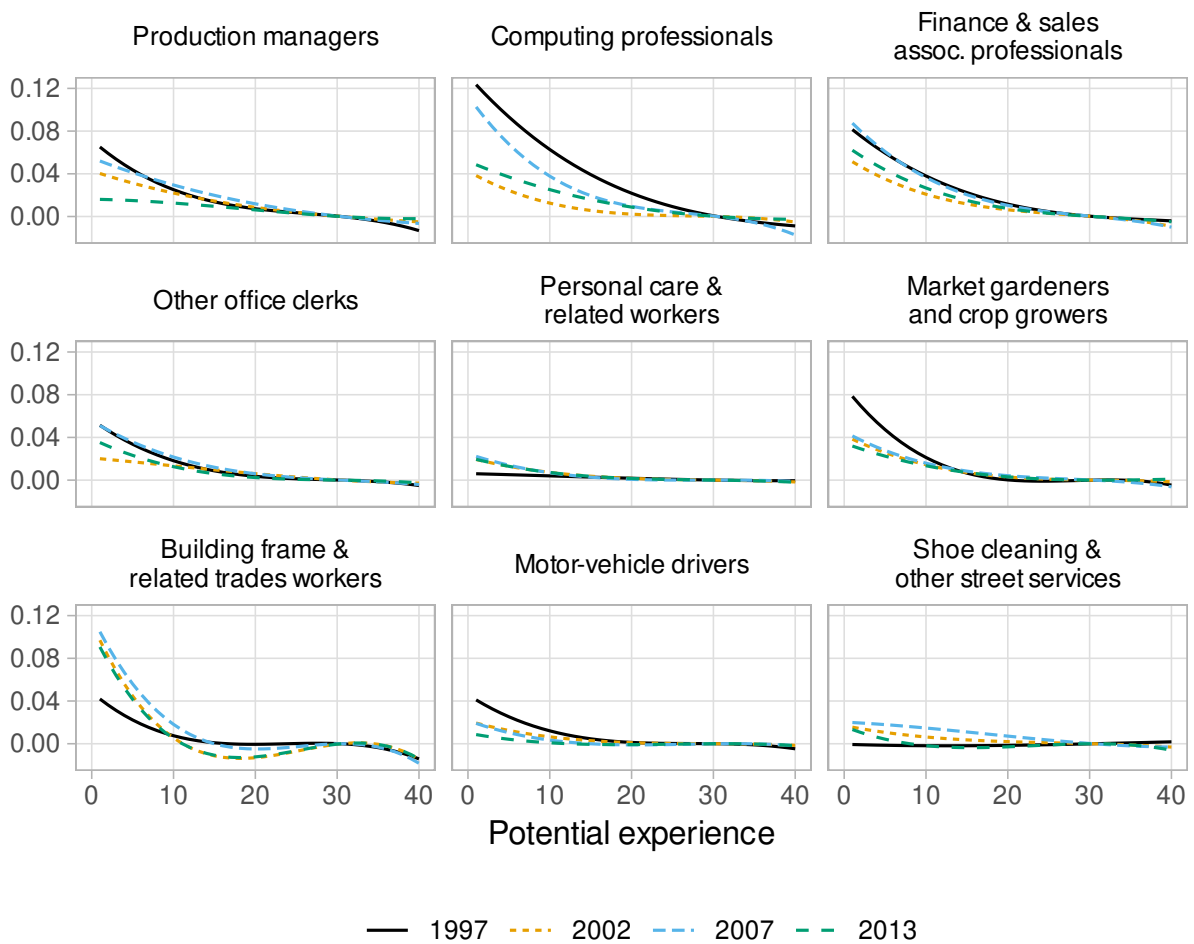


Figure 4: Estimated occupational experience profiles for selected occupations and years
Notes: The figure plots the estimated experience profiles from equation (5) for the indicated occupations and years.

in inequality relate to occupation-level changes such as differential premium growth and worker re-allocation. We closely follow Böhm et al. (2024).

4.3.1 Decomposing changes in between-occupation wage inequality

First, by the Law of Total Variance, $\text{Var}(w_{it}) = \text{E}[\text{Var}(w_{it} | k)] + \text{Var}(\text{E}[w_{it} | k])$. That is, overall wage inequality can be decomposed into a within-occupation and a between-occupation component. Without specifying the distribution of skills, it is difficult to say much about how changes in premia affect the within component, so we focus on the between component.

To simplify notation, we define $w_{kt} \equiv \text{E}[w_{it} | k]$ and $\Delta w_k \equiv \Delta \text{E}[w_i | k]$, and similarly for other variables. The difference operator $\Delta X \equiv X_1 - X_0$ denotes changes between two points in time 0 and 1, not necessarily adjacent years.

Note, to integrate out the conditioning variable—occupational choice—we must specify a distribution of occupational employment. When decomposing the variance at a given point in time, the obvious choice is to use the distribution at that point. But when considering changes over time, we need to be explicit about the distribution. We use subscripts on the variance and covariance operators to do so.

The change in between-occupation wage inequality can be written as

$$\text{Var}_1(w_{k1}) - \text{Var}_0(w_{k0}) = \underbrace{\text{Var}_0(w_{k1}) - \text{Var}_0(w_{k0})}_{\text{change at initial employment}} + \underbrace{\text{Var}_1(w_{k1}) - \text{Var}_0(w_{k1})}_{\text{change due to occupation sizes}}. \quad (6)$$

Define $y_{kt} \equiv w_{kt} - \pi_{kt}$, which captures workers' skills in the broadest sense—all parts of log wages not determined by the occupation premium. The first component on the right-hand side of equation (6) can be broken down as

$$\begin{aligned} \underbrace{\text{Var}_0(w_{k1}) - \text{Var}_0(w_{k0})}_{\text{change at initial employment}} &= \text{Var}_0(\Delta w_k) + 2 \text{Cov}_0(w_{k0}, \Delta w_k) \\ &= \underbrace{\text{Var}_0(\Delta \pi_k) + 2 \text{Cov}_0(w_{k0}, \Delta \pi_k)}_{\text{direct premia effect}} \\ &\quad + \underbrace{\text{Var}_0(\Delta y_k) + 2 \text{Cov}_0(\Delta \pi_k, \Delta y_k) + 2 \text{Cov}_0(w_{k0}, \Delta y_k)}_{\text{effect of changes in skill composition}}. \end{aligned} \quad (7)$$

From equation (7), we see how differential changes in premia may affect changes in wage inequality, and at the same time, how their effects may be offset by opposing forces. In particular, all the components of the decomposition involving changes in average skills Δy_k (labelled *effect of changes in skill composition*), as well as the occupation size term from equation (6), can be seen as potentially countervailing effects due to workers' re-sorting.

In contrast, all terms only involving $\Delta \pi_k$ and initial mean wages w_{k0} (labelled *direct pre-*

mia effect) can be interpreted as giving the counterfactual increase in between-occupation wage inequality in the absence of re-sorting. The latter is a key object of interest in our analysis. Equation (7) shows that, holding worker composition constant, changes in wage premia have a large effect on wage inequality if they are very dispersed, or if they are positively correlated with initial mean wages

4.3.2 Decomposition results

To quantify the role of differential premium growth for changes in between-occupation inequality in Sweden, we use our estimates to calculate the counterfactual scenarios developed in Section 4.3. We first focus on the long difference 1996–2013 and then examine changes at annual frequency.

Table 1: Decomposition of changes in between-occupation wage inequality

	Baseline	Common flat spot		Occ.-spec.
		25	35	flat spot
	(1)	(2)	(3)	(4)
Total				
$\Delta\text{Var}(w_{ik})$	2.57			
Between				
$\Delta\text{Var}(w_k)$	1.31			
$\Delta\text{Var}_0(w_k)$.39			
Components				
$\text{Var}_0(\Delta\pi_k) + 2\text{Cov}_0(w_{k0}, \Delta\pi_k)$.94	2.03	.29	.81
$\text{Var}_0(\Delta\pi_k)$.23	.43	.17	.28
$2\text{Cov}_0(w_{k0}, \Delta\pi_k)$.71	1.59	.12	.54
$\text{Var}_0(\Delta y_k)$.26	.37	.24	.25
$2\text{Cov}_0(w_{k0}, \Delta y_k)$	-.57	-1.45	.02	-.4
$2\text{Cov}_0(\Delta\pi_k, \Delta y_k)$	-.24	-.55	-.16	-.27

Notes: The table reports results from a decomposition of the change in the between-occupation variance in wages between 1996 and 2013 for different flat spot levels. See equation (7) for the formal statement of the decomposition. Column (1) uses a common flat spot for all occupations, at 30 years of potential experience, when estimating growth in wage premia. Columns (2)–(4) vary this common flat spot as indicated. Column (5) estimates a flat spot for each occupation using the procedure described in Appendix B. All figures have been multiplied by 100 for readability.

The first three rows in column (1) of Table 1 show the change in the observed variance of log wages, the change in between-occupation variance, as well as the change in between-occupation variance holding occupational employment fixed at 1996. The variance of log wages increased by 0.026 between 1996 and 2013, from 0.073 in 1996—an increase of around 36 percent. To avoid excessive decimal places, we multiply the variance and its

components by 100 from here on. Although the wage distribution in Sweden is still highly compressed compared to other countries (Graetz, 2020), this increase is large in relative terms.

Between-occupation wage inequality accounts for just over half of the increase in overall variance. But this is allowing for the employment weights in the calculation of variance to change over time. If employment shifts from middle- to both high- and low-paying occupations, we should expect between-occupation inequality to increase even if wage premia do not grow differentially. The phenomenon of job polarization has been extensively documented in the literature (Goos et al., 2014; Adermon and Gustavsson, 2015), and Figure A.5 confirms that it is present also in our sample period.

Our main interest, however, is in occupation-level drivers of wage inequality that are due to differential changes in compensation for a fixed set of workers. The third row in Table 1 shows that holding employment fixed at 1996 levels, the contribution of between-occupation variance shrinks by more than two thirds. But, as discussed above, changes in observed wages at the occupation level may mask changes in composition. To assess the role of differential growth in occupational wage premia, we perform the decomposition given by equation (7).

Column (1) of Table 1 presents our baseline results, with the flat spot set at 30 for all occupations. Holding worker composition constant, the increase in between-occupation variance would have been 0.94 based on our decomposition. This is more than twice the increase in the between-occupation variance of raw wages (at constant employment), almost 40 percent of the increase in the overall variance of log wages, and 72 percent of the increase in observed between-occupation variance. Most of this effect is due to a positive covariance between initial wages and premium growth, while the variance in premium growth plays a relatively minor role. The last two rows in column (1) of Table 1 show the attenuating forces: Changes in worker skills are negatively correlated with both initial wages and growth in wage premia.

Figure 5 shows the evolution of the variance components year-on-year. Panel (a) shows overall growth in wage inequality holding different components constant, while panel (b) shows the detailed decomposition of between-occupation variance at initial employment (i.e., holding occupation sizes constant). Interestingly, during the period 1996–2001, which saw the fastest growth in wage inequality, the attenuating forces of a changing skill composition are absent, and the counterfactual absent skill changes closely tracks between-occupation inequality in raw wages (at constant employment). The attenuating forces emerge only after 2001. Panel (b) shows that this attenuating effect is not due to the variance in skill changes, which closely follow premium growth throughout the period, but rather due to the negative covariances between skill changes and premium growth as

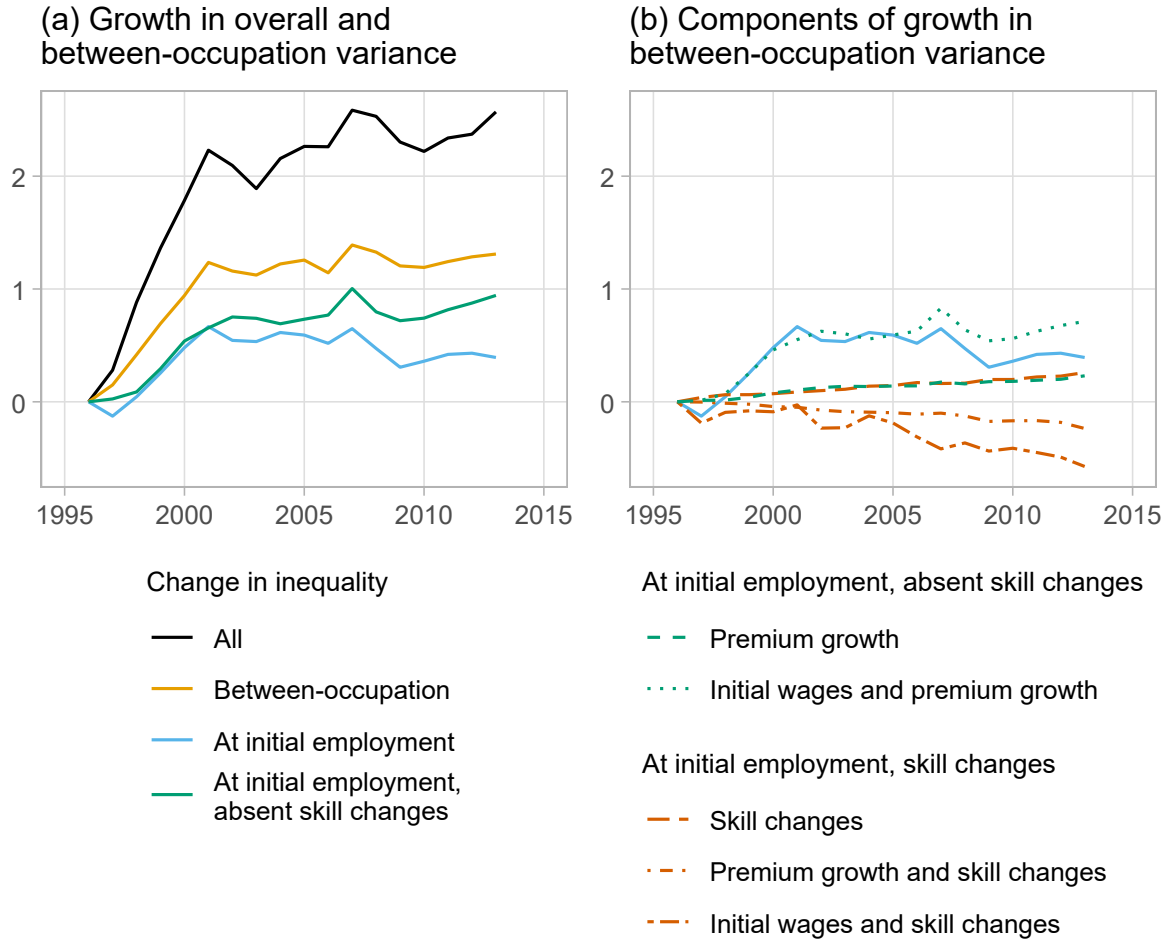


Figure 5: Decomposition of changes in between-occupation inequality 1996-2013

Notes: The figure plots the results from the decomposition given by equation (7) for every year pair $\{1996, t\}$ $\forall t \in \{1996, \dots, 2013\}$. The labels correspond to the components of equation (7) as follows: All: $\Delta \text{Var}(w_{ik})$, Between-occupation: $\Delta \text{Var}(w_k)$, At initial employment: $\Delta \text{Var}_0(w_k)$, At initial employment, absent skill changes: $\text{Var}_0(\Delta\pi_k) + 2 \text{Cov}_0(w_{k0}, \Delta\pi_k)$, Premium growth: $\text{Var}_0(\Delta\pi_k)$, Skill changes: $\text{Var}_0(\Delta y_k)$, Premium growth and skill changes: $2 \text{Cov}_0(\Delta\pi_k, \Delta y_k)$, Initial wages and premium growth: $2 \text{Cov}_0(w_{k0}, \Delta\pi_k)$, Initial wages and skill changes: $2 \text{Cov}_0(w_{k0}, \Delta y_k)$.

well as initial wages.¹⁵

4.4 Robustness checks

We conduct a number of robustness checks for the results that depend on the estimation of wage premium growth. First, we vary the location of the flat spot. As expected given the shape of wage-experience profiles and the above discussion of Figure 1, the decomposition results are sensitive to the choice of flat spot, as seen in columns (3) and (4) of Table 1. The sensitivity varies by component: The variance of premium growth appears more stable than the covariance of premium growth and initial wages.¹⁶

However, when we attempt to determine occupation-specific flat spots in a data-driven way—based on the assumption of strictly concave wage-experience profiles (except for possible flat regions) as discussed in Appendix B—we obtain results quite similar to our baseline specification (column (4)). Note also that setting the flat spot at zero, which would be implied if we simply added higher-order terms of potential experience without re-centering them, yields clearly unreasonable results (column (2) of Table A.4).

Table A.4 displays the decomposition components for a set of further robustness checks. These include changing the order of the polynomial in potential experience; adjusting for endogenous mobility using the method of Böhm et al. (2024); allowing for differential growth in wage premia at the level of regions and industries; pooling the data to estimate time-invariant experience profiles; restricting the data to men with non-missing enlistment scores; controlling for time-varying returns to cognitive and non-cognitive skills within this restricted sample; and dividing the data by gender.¹⁷ In the majority of cases, the no-sorting counterfactual is of similar or even larger magnitude compared to the baseline.¹⁸

Finally, we probe the robustness of the associations of premium growth and implied skills with employment growth, initial wages, and years of schooling. The results are shown in Figures A.8 to A.10, and once again are largely similar to the baseline specification.

¹⁵Column (1) in panel B of Table A.3 displays the decomposition results for the sub-period 2001–2013. Figures A.6 and A.7 display the relationships between growth in wages, premia, and implied skills on the one hand, and employment growth and initial wages on the other, for 2001–2013. While overall inequality changed little during this time, the pattern of premium growth and compositional changes is qualitatively very similar to that for the whole sample period.

¹⁶As premium growth and skill growth are strongly negatively correlated, this difference in sensitivity is mirrored by the other components.

¹⁷Enlistment scores were collected during military enlistment in the last decades of the 20th century, after which conscription was gradually phased out. Among birth cohorts 1952–1981, more than 90 percent of Swedish-born men are covered by these data. We use a combined measure of cognitive skills based on four different standardized tests of inductive, verbal and spatial skills, and technical comprehension, and a measure of psycho-social skills (sometimes called “non-cognitive skills”) based on a half-hour, semi-structured interview with a certified psychologist. See Lindqvist and Vestman (2011) and Fredriksson et al. (2018) for more details on these data.

¹⁸Note that using a polynomial of order one or forcing the experience profiles to be constant over time are more restrictive and thus inferior to our baseline specification.

4.5 Complementary analyses for U.S. data

In addition to our main analysis using Swedish data, we also conduct a complementary set of empirical analyses using data from the Merged Outgoing Rotation Group of the U.S. Current Population Survey. A detailed description of the U.S. sample is provided in Appendix D.1, and the corresponding results are presented in Appendix D.2. The aim of these analyses is twofold: to assess the extent to which our findings generalize across countries, and to further validate the usefulness of our identification strategy in contexts involving relatively small samples.

Consistent with prior research, we find clear evidence of job polarization in the U.S. labor market over the period 1979–2019. Notably, the relationship between employment growth and growth in average occupational wages appears to be considerably stronger in the U.S. than in Sweden. We also observe a strong, positive association between employment growth and the evolution of wage premia. However, in contrast to our findings for Sweden, we do not find that shifts in average implied skill levels systematically offset the effects of changing wage premia on average wages. Instead, there is a positive correlation between employment growth and growth in implied skills. Unfortunately, these estimates are too imprecise to allow any strong conclusions.

5 Conclusion

We contribute to the literature on shifts in the wage structure by jointly estimating growth in occupational wage premia and occupation-specific life-cycle wage profiles. We document substantial changes in occupations' relative premia in Sweden in recent decades, which are masked somewhat in the raw wage data due to worker sorting. There is a positive association between premium growth and employment growth, suggesting that workers have been responsive to changes in occupational demand. The relative premia changes are estimated to have substantially contributed to the increase in overall wage inequality. We also document large heterogeneity in life-cycle profiles across occupations, as well as substantial shifts of the profiles over time. Allowing for occupation-level changes in returns to cognitive and psycho-social skills has little effect on the results.

Despite the differences in approach and setting, our findings are remarkably similar to the ones for Germany reported by Böhm et al. (2024) when expressing them in terms of the overall increase in wage inequality. In our data, between-occupation inequality accounts for 51 percent of the overall increase, compared to 42 in Germany. The importance of differential growth in occupational wage premia for changes in between-occupation wage inequality is quite similar across the two countries: 72 percent for Sweden, and 95 percent for Germany. This suggests that wage setting institutions do not pose a hard constraint for re-allocation of employment in line with technological developments.

Our results suggest that although the overall wage structure in Sweden is highly compressed, forces related to technological change do influence the wage structure and drive workers' occupational choices. An open question is why the increase in Swedish wage inequality was concentrated in the late 1990s. We estimate a contribution of wage premium growth which builds gradually over time. In contrast, there is more dispersion of life-cycle wage growth in the late 1990s. That time period saw a temporary rise in the flexibility of collective bargaining. Our results could thus be interpreted as disruptions to wage setting institutions affecting workers differently across the life-cycle. Alternatively, technology or consumer demand evolve according to a less gradual process, but due to labor market imperfections this is not reflected in common occupational wage premia. If so, our results will be a lower bound on the importance of occupation-level labor demand shifts for growth in inequality. We leave these issues to future research.

The method we propose to estimate changes in occupational wage premia may fruitfully be applied to other settings, especially those in which experience profiles appear to change over time, and in cases where only short (two-year) panels of workers are available.

References

- ACEMOGLU, D. AND D. AUTOR (2011): “Skills, Tasks and Technologies: Implications for Employment and Earnings,” in *Handbook of Labor Economics*, ed. by O. Ashenfelter and D. Card, Elsevier, vol. 4, chap. 12, 1043–1171.
- ADERMON, A. AND M. GUSTAVSSON (2015): “Job Polarization and Task-Biased Technological Change: Evidence from Sweden, 1975–2005,” *The Scandinavian Journal of Economics*, 117, 878–917.
- 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., L. F. KATZ, AND M. S. KEARNEY (2006): “The Polarization of the U.S. Labor Market,” *American Economic Review*, 96, 189–194.
- (2008): “Trends in US wage inequality: Revising the revisionists,” *The Review of economics and statistics*, 90, 300–323.
- BARANY, Z. L. AND C. SIEGEL (2018): “Job Polarization and Structural Change,” *American Economic Journal: Macroeconomics*, 10, 57–89.
- BEN-PORATH, Y. (1967): “The Production of Human Capital and the Life Cycle of Earnings,” *Journal of Political Economy*, 75, 352–365.
- BHULLER, M., K. O. MOENE, M. MOGSTAD, AND O. L. VESTAD (2022): “Facts and Fantasies about Wage Setting and Collective Bargaining,” *Journal of Economic Perspectives*, 36, 29–52.
- BÖHM, M. J., H.-M. VON GAUDECKER, AND F. SCHRAN (2024): “Occupation Growth, Skill Prices, and Wage Inequality,” *Journal of Labor Economics*, 42, 201–243.
- BOWLUS, A. J. AND C. ROBINSON (2012): “Human Capital Prices, Productivity, and Growth,” *American Economic Review*, 102, 3483–3515.
- BURNETTE, J. D. (2017): “Inequality in the labor market for native American Women and the great recession,” *American Economic Review*, 107, 425–429.
- BÖHM, M. J. (2020): “The Price of Polarization: Estimating Task Prices under Routine-Biased Technical Change,” *Quantitative Economics*, 11, 761–799.
- 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.
- CENGIZ, D., A. DUBE, A. LINDNER, AND B. ZIPPERER (2019): “The effect of minimum wages on low-wage jobs,” *The Quarterly Journal of Economics*, 134, 1405–1454.
- 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.

- DEMING, D. J. (2017): “The Growing Importance of Social Skills in the Labor Market,” *The Quarterly Journal of Economics*, 132, 1593–1640.
- (2021): “The Growing Importance of Decision-Making on the Job,” Working Paper 28733, National Bureau of Economic Research.
- DEMING, D. J. AND K. NORAY (2020): “Earnings Dynamics, Changing Job Skills, and STEM Careers,” *The Quarterly Journal of Economics*, 135, 1965–2005.
- EDIN, P.-A., P. FREDRIKSSON, M. NYBOM, AND B. ÖCKERT (2022): “The Rising Return to Noncognitive Skill,” *American Economic Journal: Applied Economics*, 14, 78–100.
- FOSSE, E. AND C. WINSHIP (2019): “Bounding Analyses of Age-Period-Cohort Effects,” *Demography*, 56, 1975–2004.
- FREDRIKSSON, P., L. HENSVIK, AND O. N. SKANS (2018): “Mismatch of Talent: Evidence on Match Quality, Entry Wages, and Job Mobility,” *American Economic Review*, 108, 3303–38.
- GOOS, M. AND A. MANNING (2007): “Lousy and Lovely Jobs: The Rising Polarization of Work in Britain,” *The 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, 2509–26.
- GRAETZ, G. (2020): “Technological Change and the Swedish Labor Market,” Working Paper 2020:19, IFAU.
- 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.
- HOFFMANN, F., D. S. LEE, AND T. LEMIEUX (2020): “Growing Income Inequality in the United States and Other Advanced Economies,” *Journal of Economic Perspectives*, 34, 52–78.
- LAGAKOS, D., B. MOLL, T. PORZIO, N. QIAN, AND T. SCHOELLMAN (2017): “Life Cycle Wage Growth across Countries,” *Journal of Political Economy*, 126, 797–849.
- LINDQVIST, E. AND R. VESTMAN (2011): “The Labor Market Returns to Cognitive and Noncognitive Ability: Evidence from the Swedish Enlistment,” *American Economic Journal: Applied Economics*, 3, 101–128.
- LIU, D. C. AND J. NOCEDAL (1989): “On the limited memory BFGS method for large scale optimization,” *Mathematical Programming*, 45, 503–528.
- MOGSTAD, M., K. G. SALVANES, AND G. TORSVIK (2025): “Income Equality in the Nordic Countries: Myths, Facts, and Lessons,” *Journal of Economic Literature*, 63, 791–839.

- POTTER, F. J. (1990): "A study of procedures to identify and trim extreme sampling weights," in *Proceedings of the American Statistical Association, Section on Survey Research Methods*, American Statistical Association Washington, DC, vol. 225230.
- ROY, A. D. (1951): "Some Thoughts on the Distribution of Earnings," *Oxford Economic Papers*, 3, 135–146.
- ROYS, N. A. AND C. R. TABER (2019): "Skill Prices, Occupations, and Changes in the Wage Structure for Low Skilled Men," Working Paper 26453, National Bureau of Economic Research.
- SKANS, O. N., P.-A. EDIN, AND B. HOLMLUND (2009): "Wage Dispersion Between and Within Plants: Sweden 1985-2000," in *The Structure of Wages: An International Comparison*, National Bureau of Economic Research, Inc, NBER Chapters, 217–260.
- TRAIBERMAN, S. (2019): "Occupations and Import Competition: Evidence from Denmark," *American Economic Review*, 109, 4260–4301.
- VOM LEHN, C., C. ELLSWORTH, AND Z. KROFF (2022): "Reconciling occupational mobility in the current population survey," *Journal of Labor Economics*, 40, 1005–1051.
- WRIGHT, E. O. AND R. E. DWYER (2003): "The patterns of job expansions in the USA: a comparison of the 1960s and 1990s," *Socio-Economic Review*, 1, 289–325.

A Additional figures and tables




Figure A.1: Simulated wage-experience profiles

Notes: The top row shows various simulated experience profiles. The bottom row plots first differences of the profiles above, assuming wage premium growth of two percent.

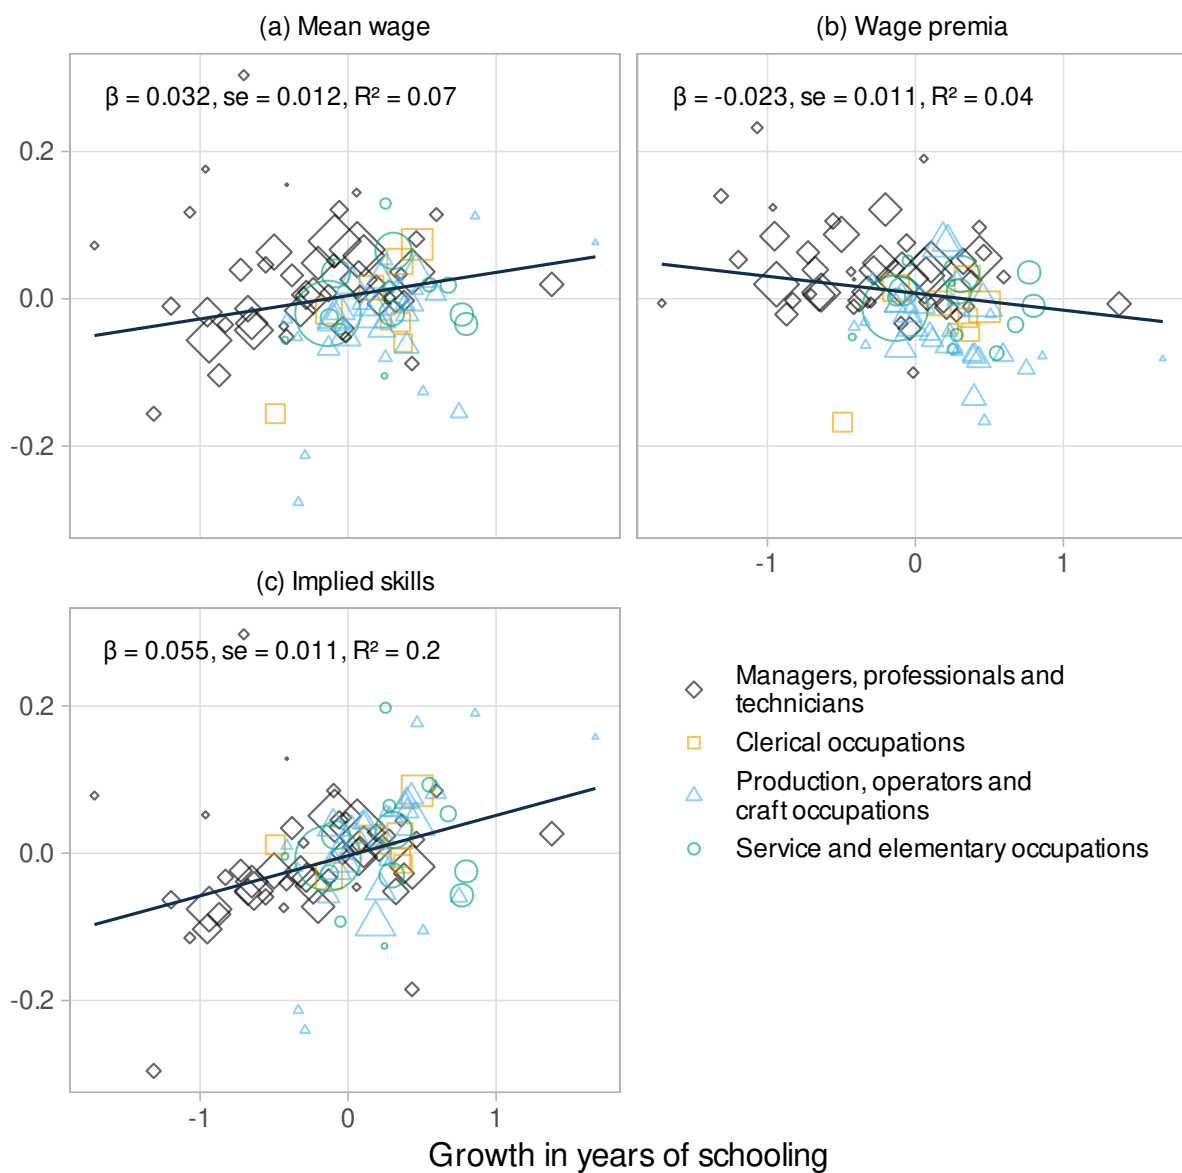


Figure A.2: Growth in wages, premia, and skills against growth in schooling, 1996–2013

Notes: The figure plots the growth mean log wages, cumulative estimated wage premia, as well as the implied change in mean skills, against initial the growth in average years of schooling. Wage premia are estimated according to our baseline specification equation (5). Each marker represents one of 101 occupations. The size of each marker is determined by the employment share in the first year and the regression line is weighted accordingly. We use original survey weights when calculating occupation size and mean log wage.

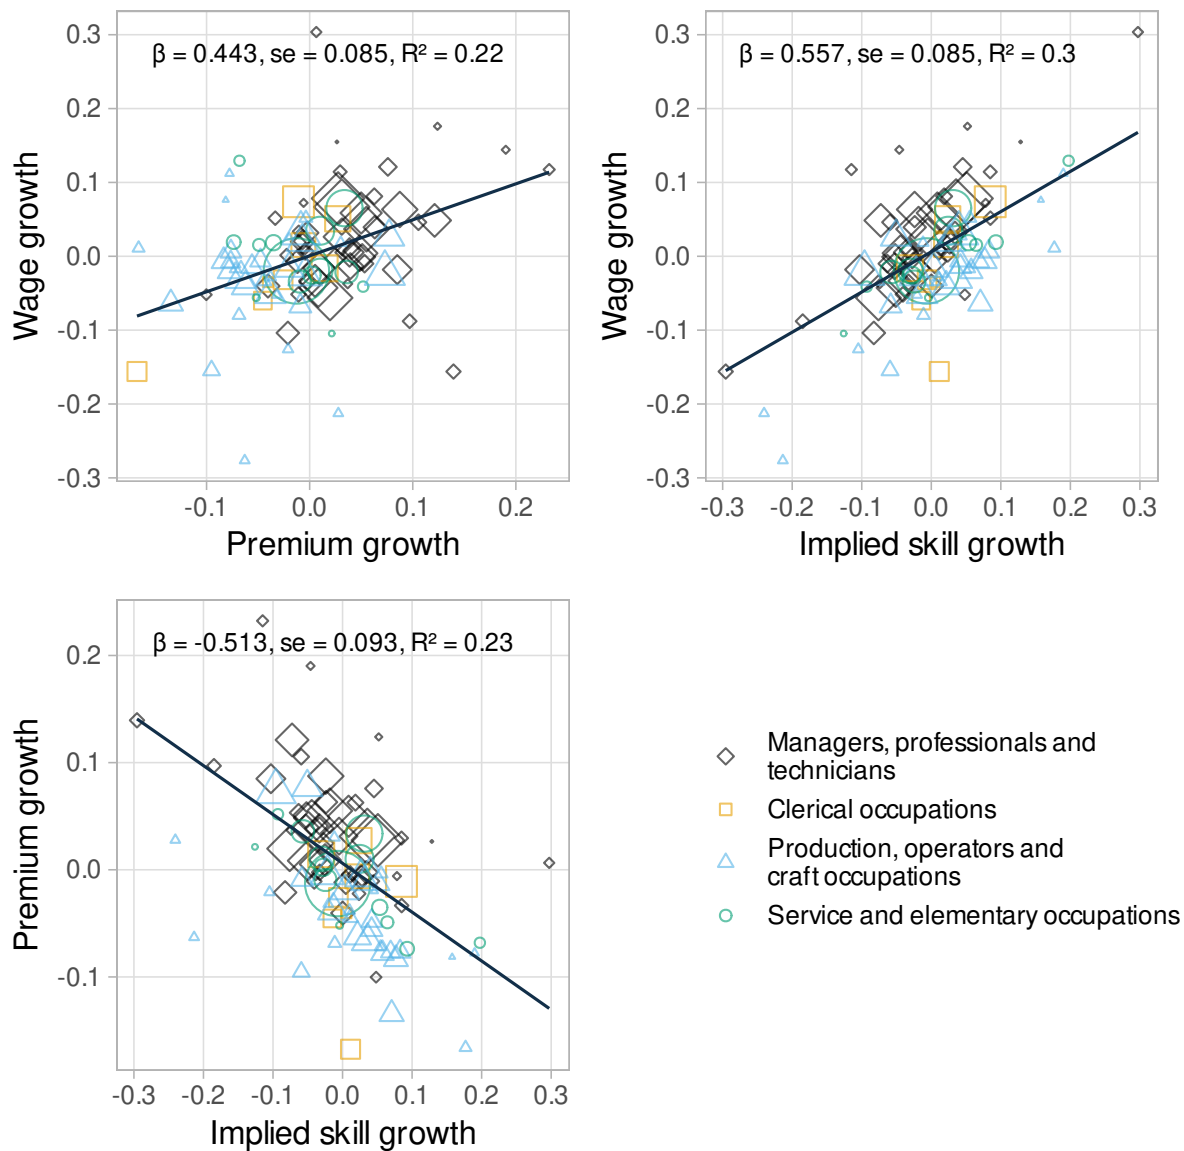
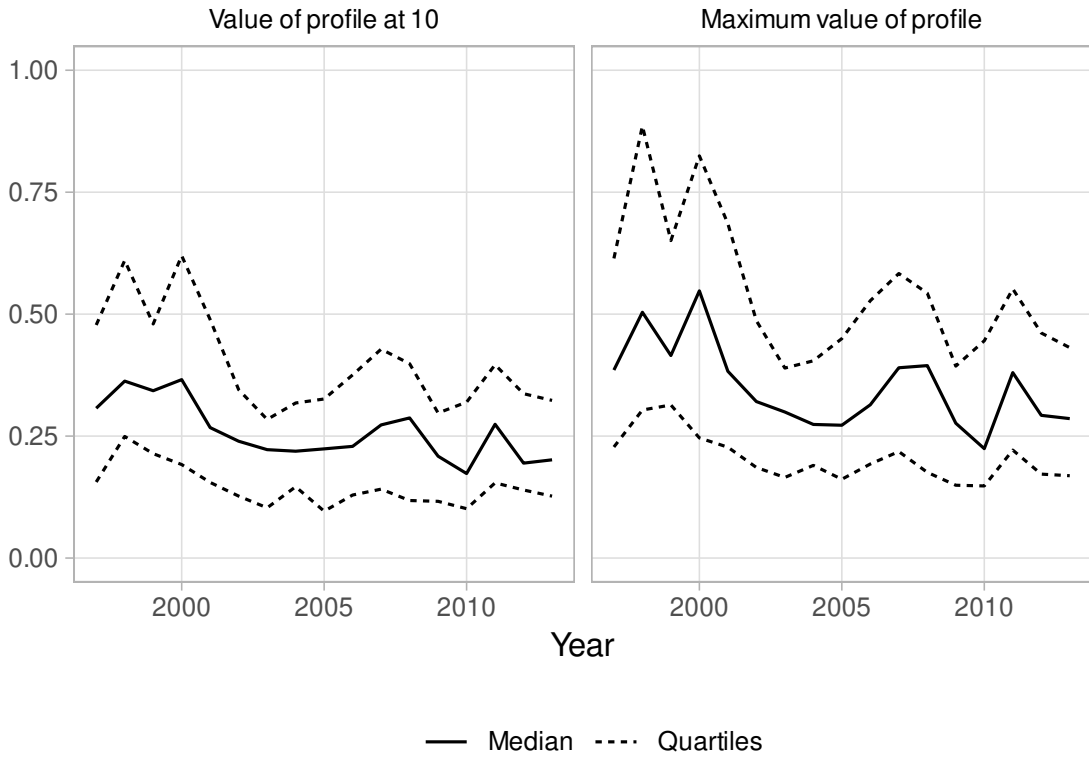


Figure A.3: Relations between growth rates

Notes: The figure plots the bivariate relationships between the growth in mean log wages, cumulative estimated wage premia, and the implied change in mean skills. Each marker represents one of 101 occupations. The size of each marker is determined by the employment share in the first year and the regression lines are weighted accordingly. We use original survey weights when calculating occupation size and mean log wage.

(a) Steepness of wage-experience profiles



(b) Variation in profiles and premium growth




Figure A.4: Wage-experience profiles over time

Notes: The figure characterizes the distribution of the experience profiles estimated by equation (5) over time (panel (a)) and shows how variance in selected characteristics of experience profiles is related to variance in premium growth (panel (b)).




Figure A.5: Job polarization

Notes: The figure plots the growth in log employment against mean log wages in 1996. In Panel (b), log wages have been percentile-ranked, weighted by initial employment. Each marker represents one of 101 occupations. The size of each marker is determined by the employment share in the first year and the regression lines are weighted accordingly. We use original survey weights when calculating occupation size and mean log wage.

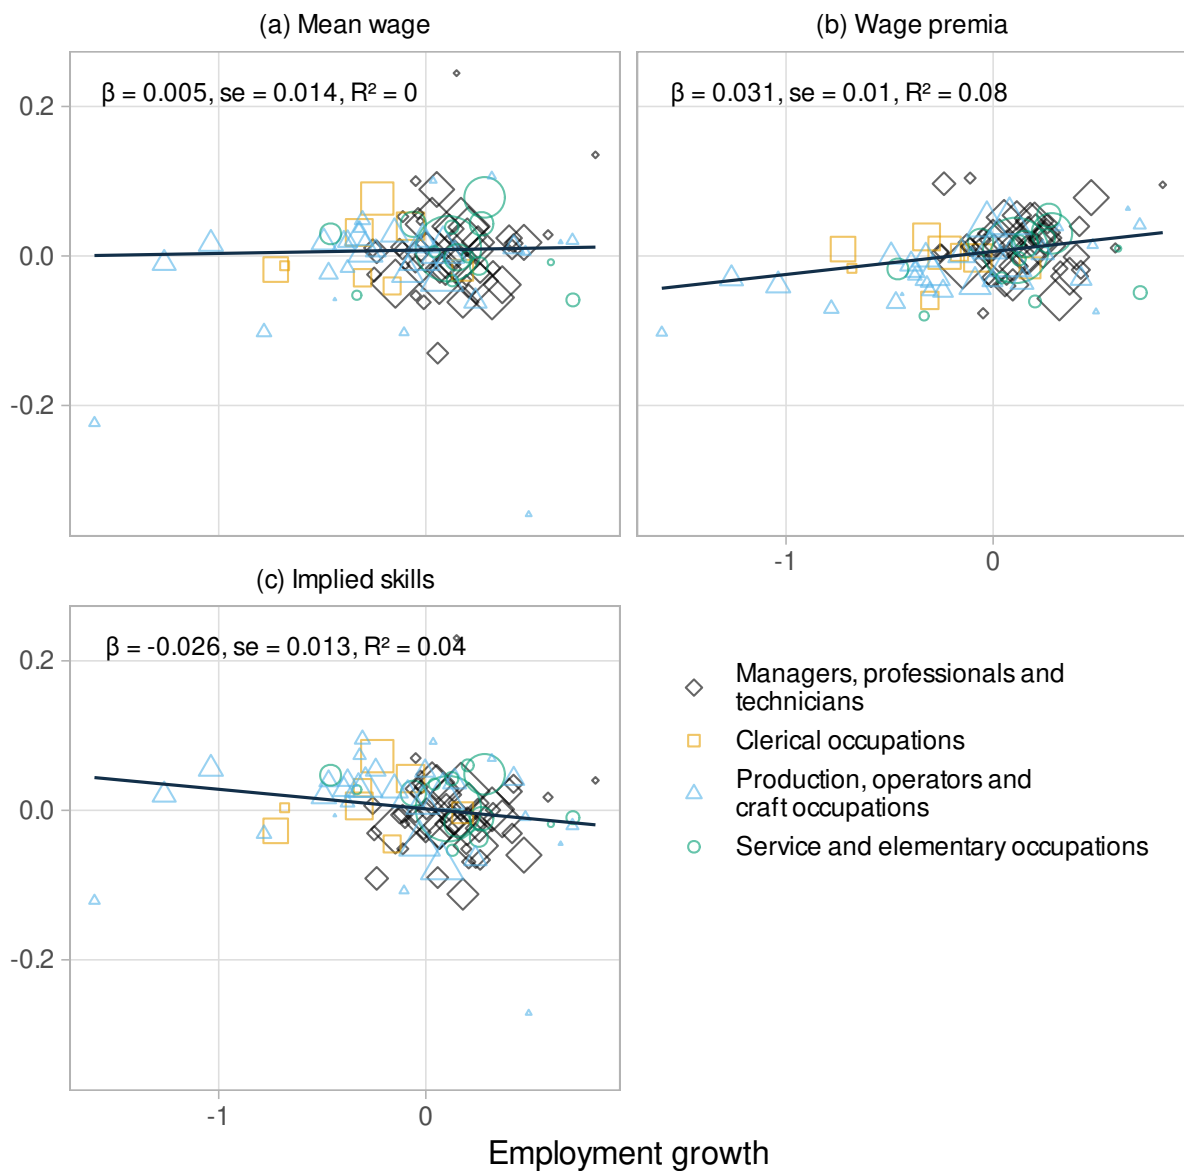


Figure A.6: Growth in wages, premia, and skills against employment growth, 2001–2013

Notes: The figure plots the growth mean log wages, cumulative estimated wage premia, as well as the implied change in mean skills, against the change in log employment. Wage premia are estimated according to our baseline specification equation (5). Each marker represents one of 101 occupations. The size of each marker is determined by the employment share in the first year and the regression line is weighted accordingly. We use original survey weights when calculating occupation size and mean log wage.

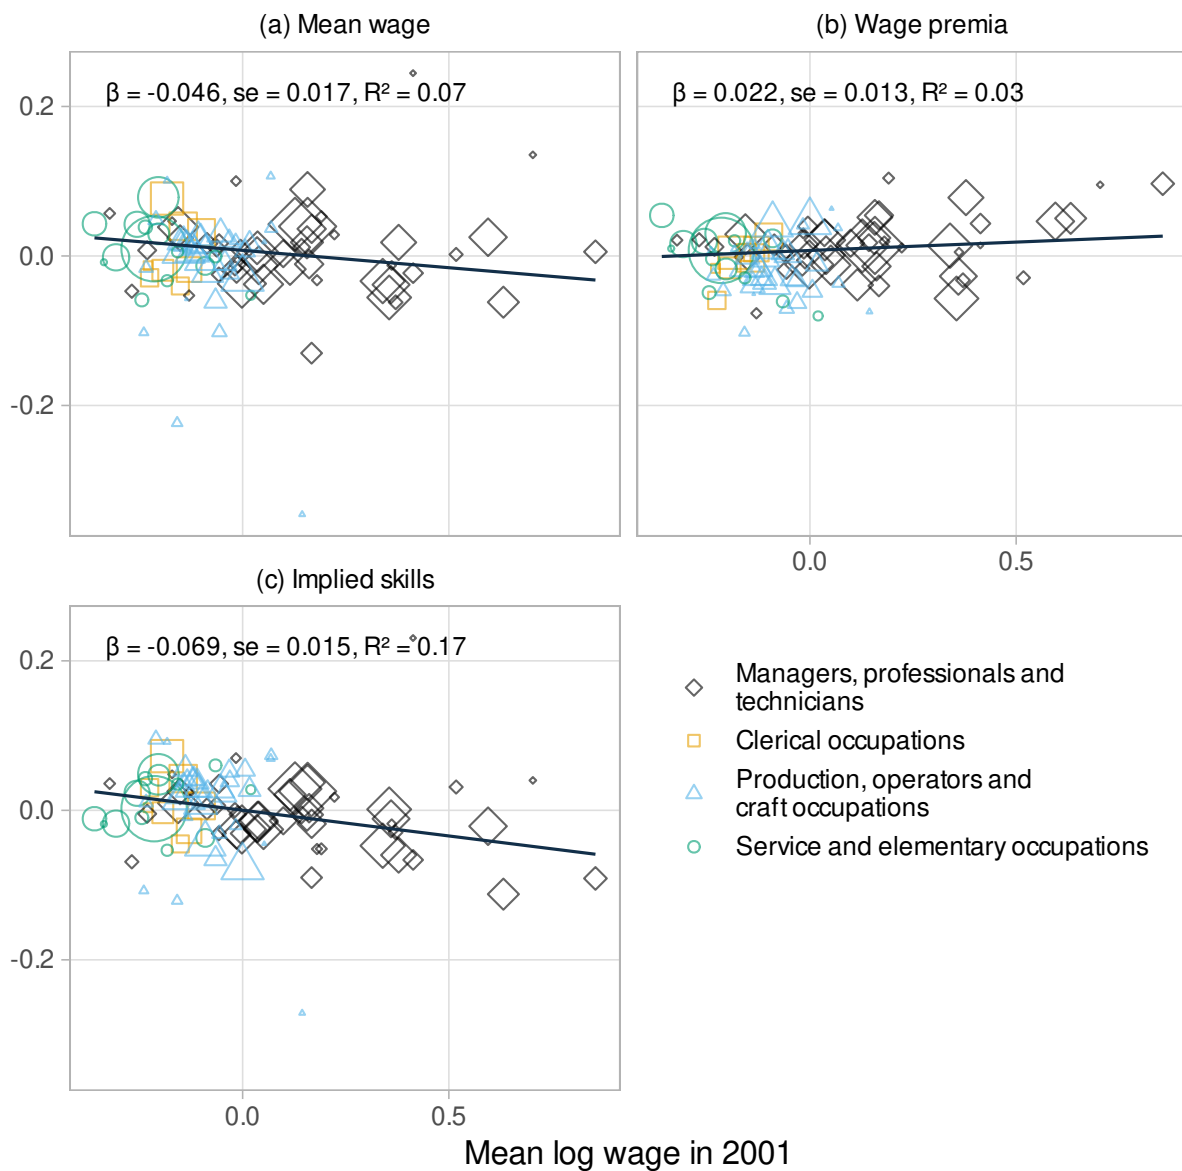


Figure A.7: Growth in wages, premia, and skills against initial wages, 2001–2013

Notes: The figure plots the growth mean log wages, cumulative estimated wage premia, as well as the implied change in mean skills, against initial mean log wages. Wage premia are estimated according to our baseline specification equation (5). Each marker represents one of 101 occupations. The size of each marker is determined by the employment share in the first year and the regression line is weighted accordingly. We use original survey weights when calculating occupation size and mean log wage.




Figure A.8: Premia, skills, and employment growth—robustness checks

Notes: The table reports the coefficients, standard errors, and coefficients of determination from separately regressing cumulative estimated wage premia and the implied change in mean skills (growth in average wage minus premium growth) against the change in log employment at the occupation level for different sets of premia estimates. See the text for descriptions of how these estimates are produced. The weight assigned to each occupation is determined by the employment share in the first year. We use original survey weights when calculating occupation size and mean log wage.




Figure A.9: Premia, skills, and initial wages—robustness checks

Notes: The figure reports the coefficients, standard errors, and coefficients of determination from separately regressing cumulative estimated wage premia and the implied change in mean skills (growth in average wage minus premium growth) against initial mean log wage at the occupation level for different sets of premia estimates. See also the notes to Figure A.8.

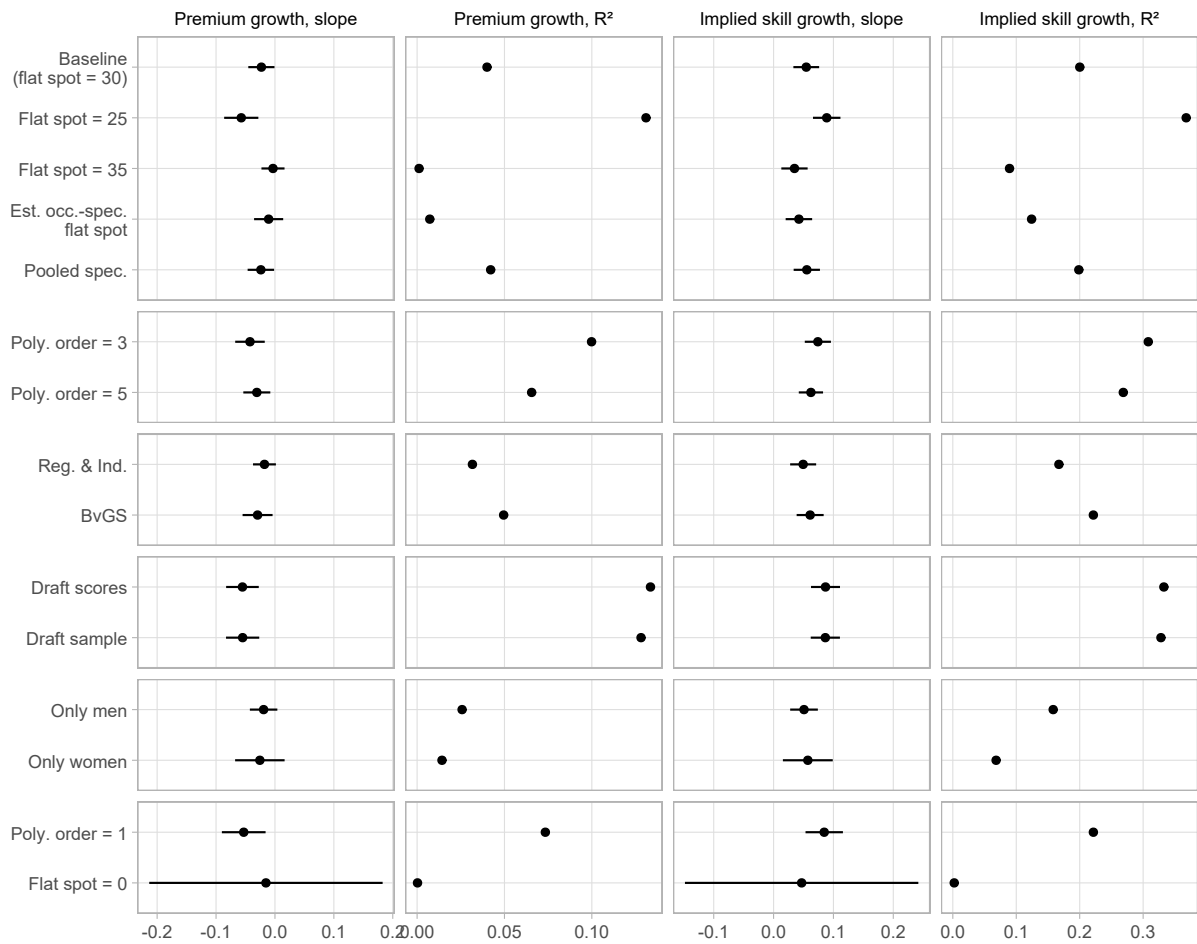


Figure A.10: Premia, skills, and schooling—robustness checks

Notes: The figure reports the coefficients, standard errors, and coefficients of determination from separately regressing cumulative estimated wage premia and the implied change in mean skills (growth in average wage minus premium growth) against growth in average years of schooling at the occupation level for different sets of premia estimates. See also the note to Figure A.8.

Table A.1: Descriptive statistics for raw and final sample

	1996		1997		2012		2013	
	Raw sample (1)	Final sample (2)	Raw sample (3)	Final sample (4)	Raw sample (5)	Final sample (6)	Raw sample (7)	Final sample (8)
Un-weighted no. of obs.	2,124,596	1,309,910	2,159,931	1,309,910	2,346,712	1,461,698	2,354,000	1,461,698
Weighted no. of obs.	3,130,314	1,673,285	3,291,612	1,664,835	3,729,016	1,989,294	3,754,542	2,009,409
Mean wage	21,133 (7,503)	21,670 (7,398)	21,993 (8,218)	22,443 (7,633)	30,055 (13,256)	30,959 (13,193)	30,868 (14,228)	32,049 (13,659)
Mean log wage	9.92 (0.27)	9.94 (0.26)	9.95 (0.28)	9.98 (0.26)	10.25 (0.31)	10.29 (0.30)	10.28 (0.31)	10.32 (0.30)
Mean log wage growth			0.04 (0.10)	0.04 (0.07)			0.04 (0.10)	0.03 (0.06)
Mean age	40.95 (11.58)	41.09 (9.77)	41.23 (11.56)	42.12 (9.76)	42.03 (12.46)	41.38 (10.30)	41.99 (12.45)	42.38 (10.29)
Share female	0.51	0.53	0.50	0.53	0.50	0.51	0.50	0.51
Share foreign-born	0.09	0.08	0.09	0.08	0.14	0.13	0.14	0.13
Mean years of education	11.57 (2.69)	11.96 (2.61)	11.63 (2.67)	11.96 (2.62)	12.65 (2.43)	13.00 (2.38)	12.72 (2.42)	13.00 (2.39)
Mean experience	22.92 (11.78)	22.80 (9.79)	23.15 (11.78)	23.82 (9.79)	23.25 (12.81)	22.28 (10.50)	23.15 (12.79)	23.27 (10.50)
Occupation shares								
Managerial-Professional-Technical	0.38	0.45	0.41	0.45	0.45	0.51	0.46	0.51
Clerks	0.11	0.10	0.11	0.10	0.08	0.07	0.08	0.07
Production-Operators-Crafts	0.22	0.19	0.22	0.19	0.19	0.17	0.18	0.17
Service-Elementary	0.28	0.26	0.27	0.26	0.28	0.25	0.28	0.25

Notes: The raw sample consists of all observations with nonmissing wages in the Wage Structure Statistics (there can be more than one observation per individual if the individual is sampled at several workplaces). The final sample is the sample used when estimating equation (5). All statistics (other than number of observations) weighted by unadjusted cross-sectional sampling weights from the Wage Structure Statistics (our regressions use longitudinal weights). Standard deviations in parentheses. Wages in 2016 SEK.

Table A.2: Share of raw Wage Structure Statistics sample remaining as restrictions imposed

	1996		1997		2012		2013	
	Un-weighted obs. (%)	Weighted obs. (%)	Un-weighted obs. (%)	Weighted obs. (%)	Un-weighted obs. (%)	Weighted obs. (%)	Un-weighted obs. (%)	Weighted obs. (%)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Raw sample	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0
Removing individual-year duplicates	96.9	97.1	97.1	97.5	97.6	97.8	97.5	97.7
Matched with LOUISE	96.7	96.9	96.9	97.3	97.4	97.6	97.4	97.5
Matched across years	82.4	74.2	81.0	70.4	81.9	71.8	81.6	72.0
Employment restrictions	79.7	71.5	78.4	67.9	80.2	70.1	79.9	70.3
Nonmissing occupations	74.8	67.4	73.5	64.1	79.0	69.2	78.8	69.3
Experience restrictions	68.5	61.7	67.4	58.7	68.8	60.7	68.6	60.8
Occupation stayers	62.8	54.6	61.8	51.7	63.4	54.6	63.2	54.8
Excluding extreme wage growth (final sample)	61.7	53.5	60.6	50.6	62.3	53.3	62.1	53.5

Notes: The raw sample consists of all observations with nonmissing wages in the Wage Structure Statistics (there can be more than one observation per individual if the individual is sampled at several workplaces). Weighted columns use cross-sectional weights (our regressions use longitudinal weights). Excluding extreme wage growth involves dropping the top and bottom percentiles of wage growth among occupational stayers (weighted by longitudinal weights).

Table A.3: Decomposition results—sub-periods

	Baseline		Common flat spot		Occ.-spec. flat spot (4)
	(1)	(2)	25	35	
				(3)	
Panel A: 1996–2001					
Total	2.23				
$\Delta \text{Var}(w_{i,k})$			1.13	.36	.55
Between	1.24		.17	.08	.1
$\Delta \text{Var}(w_k)$.67		.96	.28	.45
$\Delta \text{Var}_0(w_k)$.1	.09	.08
Components			-.44	.25	.08
$\text{Var}_0(\Delta\pi_k) + 2\text{Cov}_0(w_{k0}, \Delta\pi_k)$			-.13	-.03	-.05
$\text{Var}_0(\Delta\pi_k)$					
$2\text{Cov}_0(w_{k0}, \Delta\pi_k)$					
$\text{Var}_0(\Delta y_k)$					
$2\text{Cov}_0(w_{k0}, \Delta y_k)$					
$2\text{Cov}_0(\Delta\pi_k, \Delta y_k)$					
Panel B: 2001–2013					
Total	.34				
$\Delta \text{Var}(w_{i,k})$.93	-.08	.26
Between	.07		.13	.08	.1
$\Delta \text{Var}(w_k)$	-.31		.79	-.16	.16
$\Delta \text{Var}_0(w_k)$.18	.12	.15
Components			-1.25	-.3	-.62
$\text{Var}_0(\Delta\pi_k) + 2\text{Cov}_0(w_{k0}, \Delta\pi_k)$			-.17	-.06	-.11
$\text{Var}_0(\Delta\pi_k)$					
$2\text{Cov}_0(w_{k0}, \Delta\pi_k)$					
$\text{Var}_0(\Delta y_k)$					
$2\text{Cov}_0(w_{k0}, \Delta y_k)$					
$2\text{Cov}_0(\Delta\pi_k, \Delta y_k)$					

Notes: The table reports results from a decomposition of the change in the between-occupation variance in wages for different flat spot levels and periods. See equation (7) for the formal statement of the decomposition. Column (1) uses a common flat spot for all occupations, at 30 years of potential experience, when estimating growth in wage premia. Columns (2)–(4) vary this common flat spot as indicated. Column (5) estimates a flat spot for each occupation using the procedure described in Appendix B. All figures have been multiplied by 100 for readability.

Table A.4: Decomposition results—further specifications

	Baseline	Flat spot = 0					BvGS	Reg. & ind.	Pooled spec.	Draft sample	Draft scores	Men	Women
		1	3	5	5	5							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	
Total													
$\Delta \text{Var}(w_{i,k})$	2.57												
Between													
$\Delta \text{Var}(w_k)$	1.31												
$\Delta \text{Var}_0(w_k)$.39												
Components													
$\text{Var}_0(\Delta \pi_k) + 2\text{Cov}_0(w_{k0}, \Delta \pi_k)$.94	24.49	2.57	1.21	1.07	1.28	.66	.97	1.64	1.59	.84	2.02	
$\text{Var}_0(\Delta \pi_k)$.23	17.6	.67	.31	.25	.31	.17	.24	.41	.4	.25	.8	
$2\text{Cov}_0(w_{k0}, \Delta \pi_k)$.71	6.89	1.9	.89	.82	.97	.49	.73	1.23	1.19	.59	1.22	
$\text{Var}_0(\Delta y_k)$.26	17.09	.56	.31	.25	.29	.25	.27	.4	.39	.28	.83	
$2\text{Cov}_0(w_{k0}, \Delta y_k)$	-.57	-6.75	-1.77	-.75	-.68	-.83	-.35	-.59	-1.09	-1.05	-.45	-1.08	
$2\text{Cov}_0(\Delta \pi_k, \Delta y_k)$	-.24	-34.44	-.98	-.37	-.25	-.35	-.18	-.25	-.55	-.54	-.28	-1.38	

Notes: The table reports results from a decomposition of the change in the between-occupation variance in wages for different specifications and periods. See equation (7) for the formal statement of the decomposition and the text for details on the different specifications. All figures have been multiplied by 100 for readability.

B Procedure for estimating occupation-specific flat spots

Suppose that experience profiles are strictly concave except for possible flat regions. That is, linear segments with non-zero slope, as in the middle column of Figure A.1, are prohibited. Formally, $g''(x) \leq 0$, $g''(x) = 0 \Rightarrow g'(x) = 0$. This implies that the second derivative of the profile will be maximized (closest to zero) at the flat spot, so that any statistic of interest should change by the least amount—in absolute value—at the true flat spot. We use this insight to pin down the flat spot in a data-driven way.

Recall from Section 4.3 that the change in between-occupation variance of log wages, at constant employment, can be decomposed as

$$\begin{aligned} \text{Var}_0(w_{k1}) - \text{Var}_0(w_{k0}) &= \text{Var}_0(\Delta w_k) + 2 \text{Cov}_0(w_{k0}, \Delta w_k) \\ &= \text{Var}_0(\Delta \pi_k) + \text{Var}_0(\Delta y_k) + 2 \text{Cov}_0(\Delta \pi_k, \Delta y_k) \\ &\quad + 2 \text{Cov}_0(w_{k0}, \Delta \pi_k) + 2 \text{Cov}_0(w_{k0}, \Delta y_k). \end{aligned} \quad (\text{B.1})$$

Denote by μ the components of the decomposition,

$$\mu \in \mathcal{M} \equiv \{\text{Var}_0(\Delta \pi_k), \text{Var}_0(\Delta y_k), 2 \text{Cov}_0(w_{k0}, \Delta \pi_k), 2 \text{Cov}_0(\Delta \pi_k, \Delta y_k)\}.$$

Each of the elements of \mathcal{M} depends on the change in the premia $\Delta \pi_k$, which in turn depend on the chosen flat spots. However, the sum of all components on the right-hand side of equation (B.1) is constant, so we exclude $2 \text{Cov}_0(w_{k0}, \Delta y_k)$ from the set \mathcal{M} .

Let ϖ denote the vector of changes in premia, and let $\tilde{\mathbf{x}}$ denote the vector of candidate flat spots. Both vectors contain K elements, where K is the total number of occupations, indexed by k . We denote the above-mentioned functional dependence by $\mu \equiv \mu(\varpi(\tilde{\mathbf{x}}))$. Using the chain rule, we define the sensitivity of μ to changing the flat spot, in absolute terms, as

$$|d\mu(\varpi(\tilde{\mathbf{x}}))| \equiv \left| \sum_k \frac{\partial \mu}{\partial (\Delta \pi_k)} \times \sum_{k'} \frac{\partial (\Delta \pi_{k'})}{\partial \tilde{x}_{k'}} \times d\tilde{x}_{k'} \right|.$$

Under strictly concave experience profiles, we conjecture that $|d\mu(\varpi(\tilde{\mathbf{x}}))|$ attains its minimum at or near the vector of true flat spots \mathbf{x}^* , and similarly for the sum over $|d\mu(\varpi(\tilde{\mathbf{x}}))|$,

$$\mathbf{x}^* = \arg \min_{\tilde{\mathbf{x}}} \sum_{\mu \in \mathcal{M}} |d\mu(\varpi(\tilde{\mathbf{x}}))|. \quad (\text{B.2})$$

We implement the optimization problem given by equation (B.2) in practice by solving

$$\mathbf{x}^* = \arg \min_{\tilde{\mathbf{x}}} S \times \sum_{\hat{\mu} \in \mathcal{M}} \left[(\hat{\mu}(\tilde{\mathbf{x}} + \tau) - \hat{\mu}(\tilde{\mathbf{x}}))^2 + (\hat{\mu}(\tilde{\mathbf{x}} - \tau) - \hat{\mu}(\tilde{\mathbf{x}}))^2 \right],$$

where $\hat{\mu}$ denote the estimated moments, τ is size- k vector with constant elements representing step size, and S is a scaling factor chosen for numerical stability. We set the elements of τ to equal 0.01 and $S = 1e+7$. We use the L-BFGS-B method (Liu and Nocedal, 1989) implemented by the `optim` package in R. We impose $\tilde{x}_k \in [25, 40] \forall k$. As the procedure appears to be sensitive to initial values, we draw initial values at random from the continuous uniform distribution $U(26, 39)$ for each \tilde{x}_k . This process is repeated 100 times. We then choose the \mathbf{x}^* with the lowest associated loss.

Note that in principle, given strictly concave profiles one should be able to find the flat spots by minimizing the sensitivity of the $\Delta\pi_k$'s instead of a moment that is a function of them. However, approximating the experience profiles by a polynomial does not guarantee that the estimated profiles are actually strictly concave. Alternatively, one could impose a functional form on the profiles that does guarantee strict concavity. We attempted to do this, but the estimation turned out to be highly unstable.

C Comparing the statistical precision of our estimation method to a non-parametric estimator

Assume that the general assumptions regarding the flat spot discussed in section 2 hold and, in addition, that there is a flat region of sufficient width around this spot. It is then possible to define a local, non-parametric estimator that is simply based on the average wage growth of individuals with experience levels in the direct vicinity of the flat spot.

We use this simple estimator to quantify the gain in statistical precision from our method of jointly estimates the wage premia and experience profiles. In practice, we include experience levels in the range of 29.5 to 30.5, with the assumed flat spot at 30. We then compare the coefficients and standard errors for all premia estimates at the occupation-year-level to our baseline estimates.

The non-parametric estimates are associated with an average standard error of 0.0063356. This is 2.55 times larger than the average for our baseline estimates, 0.0028302. The average ratio of all the estimates is instead 2.82.

Figure C.11 reports the average coefficients and standard errors from the two procedures separately by decile groups of either the coefficients or standard errors from our main method. The average coefficients are closely aligned, meaning that our method on average yields very similar coefficients to the non-parametric estimator. The absolute difference in statistical precision, unsurprisingly, is increasing in the baseline standard errors. Thus, we gain more in absolute terms when estimates are generally imprecise (in turn, primarily due to relatively few observations).

Taken together, the results show that our parametric method can help improve statistical precision, especially when the number of observations are relatively few, compared to using observations only in the direct vicinity of a hypothesized flat spot.

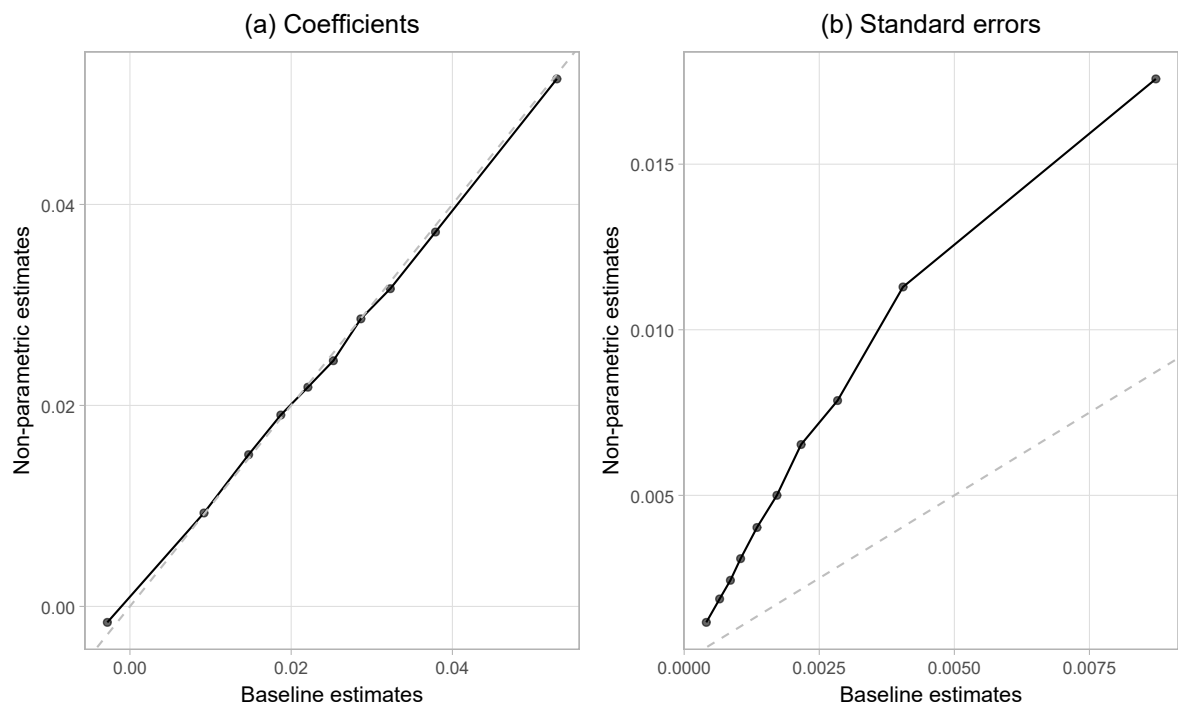


Figure C.11: Comparison of baseline and non-parametric estimates

Notes: The figure reports the average coefficients (subfigure (a)) and standard errors (subfigure (b)) from our baseline method and the non-parametric procedure described above separately by decile groups of either the coefficients or standard errors from our main method.

D US analysis

D.1 Description of US data

We use data from the Current Population Survey Merged Outgoing Rotation Group (CPS MORG), which includes weights so as to be representative of the US population. In the CPS MORG, households are interviewed about a large set of questions, including occupations and wages, at two occasions one year apart. This provides a panel structure which makes it possible to estimate changes in premia. We use data for the years 1979–2019, with gaps in 1984–1985 and 1994–1995 when linkages across survey waves are unavailable.

The CPS interviews are targeted at those who live at a certain address, rather than individuals; everyone in the household is interviewed. If households move, the new inhabitants at the address are interviewed instead. We avoid such false matches through standard procedures (see e.g. Autor et al., 2008; Burnette, 2017; Cengiz et al., 2019) which involve dropping those with different sex, race, age or education across survey waves (the restriction on education is standard in the literature, but also follows from our focus on employed occupation stayers).

We study real log wages of employed for-pay non-military workers aged 18–64 who report that they worked at least 35 hours the week before the survey. We retain only those with non-imputed wages and hours worked. In line with the restrictions for the Swedish sample, we keep only those who have between 1 and 39 years of labor market experience in the first of the two survey waves. Labor market experience is measured as age minus education minus six. Education is measured based on the estimates of Autor et al. (2008), who provide sex- and race-specific estimates of education for respondents with different answers to the years of education question in the CPS.

Occupations are translated to the 1990 census occupational classification based on Autor and Dorn (2013) and vom Lehn et al. (2022). Two-digit occupations are used in the analysis. Four small categories (Private household, Farm operators and managers, Farm except managerial and Extractive) are excluded. We keep occupation stayers, i.e. those who are in the same two-digit occupation in both survey waves.

In the sample of those who satisfy these restrictions, and whom we are able to match between survey waves, we exclude those with wage growth below the 1st or above the 99th percentile to avoid outlier effects.

D.2 US results

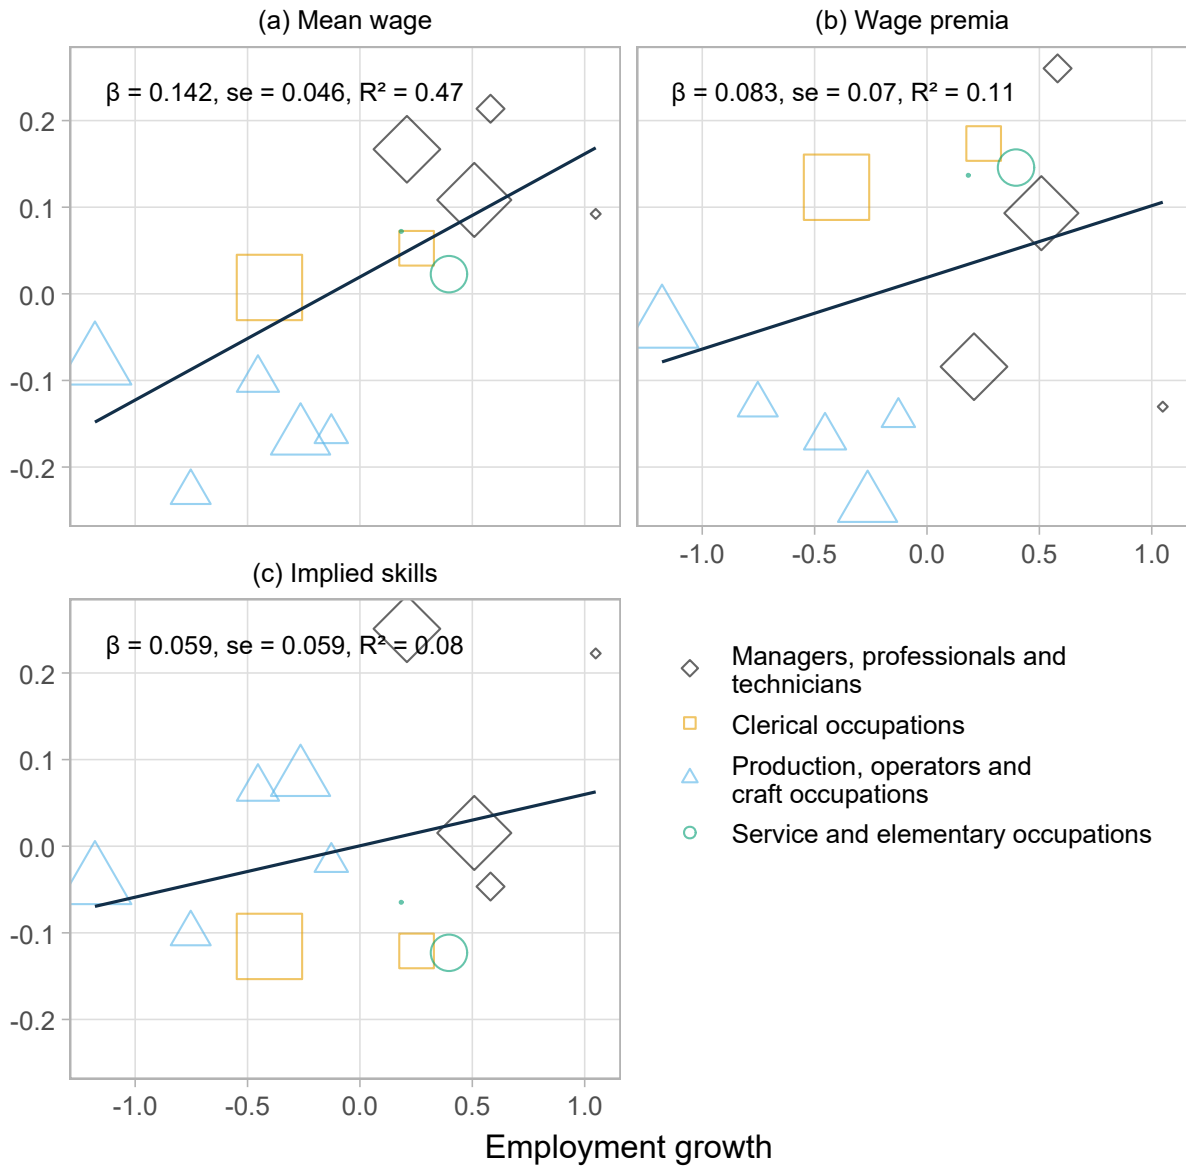


Figure D.12: Growth in wages, premia, and skills against employment growth, 1979–2019, US

Notes: The figure plots the growth mean log wages, cumulative estimated wage premia, as well as the implied change in mean skills, against the change in log employment. Wage premia are estimated according to our baseline specification equation (5). Each marker represents one of 13 occupations. The size of each marker is determined by the employment share in the first year and the regression line is weighted accordingly.

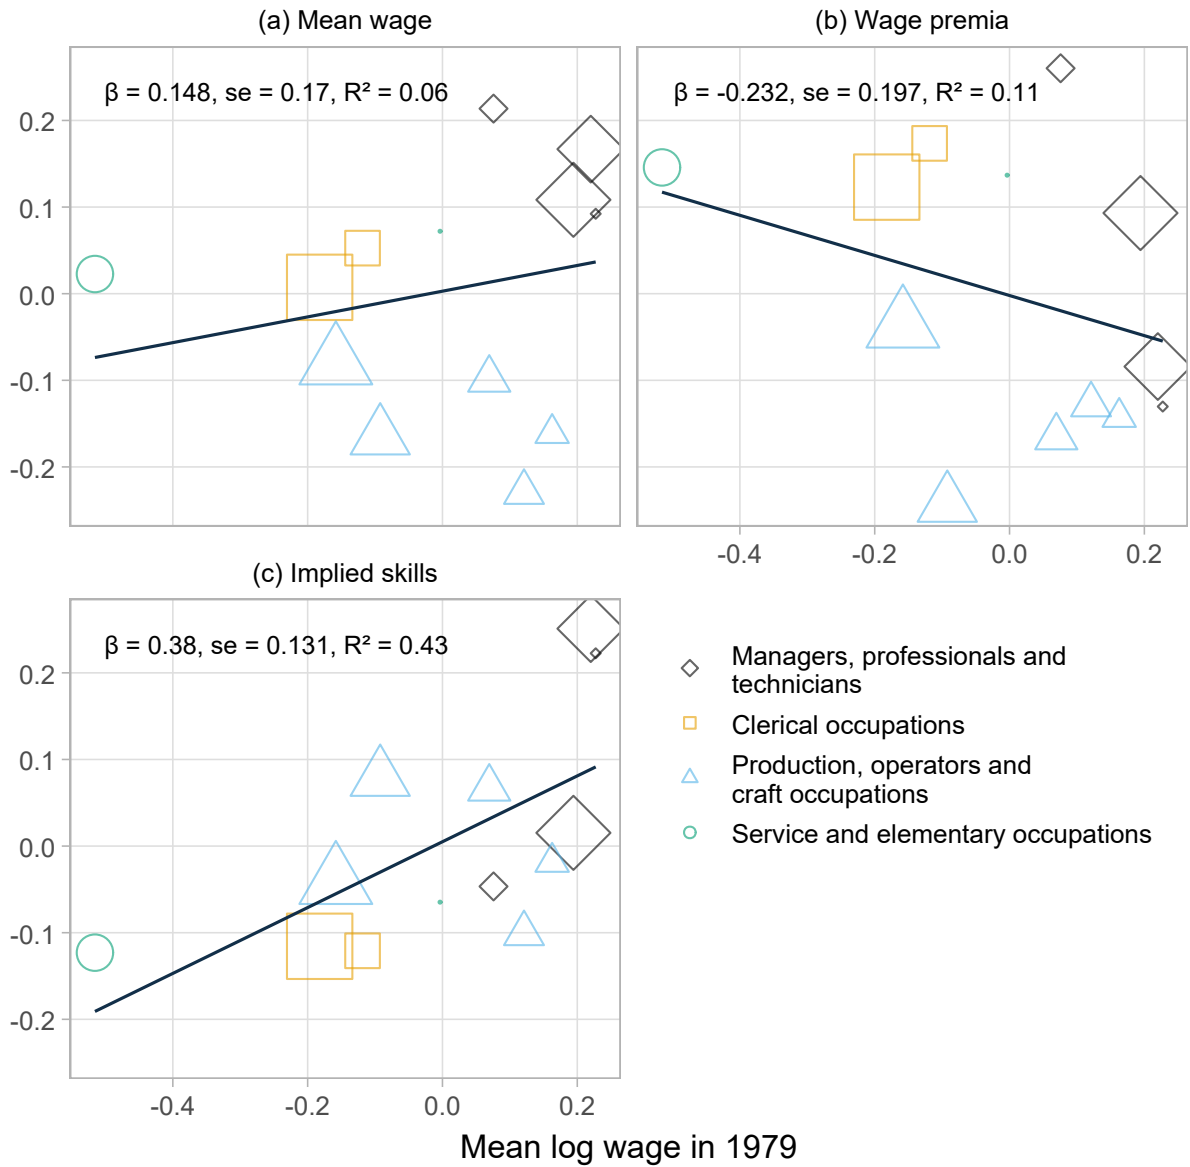


Figure D.13: Growth in wages, premia, and skills against initial wages, 1979–2019, US

Notes: The figure plots the growth mean log wages, cumulative estimated wage premia, as well as the implied change in mean skills, against initial mean log wages. Wage premia are estimated according to our baseline specification equation (5). Each marker represents one of 13 occupations. The size of each marker is determined by the employment share in the first year and the regression line is weighted accordingly.

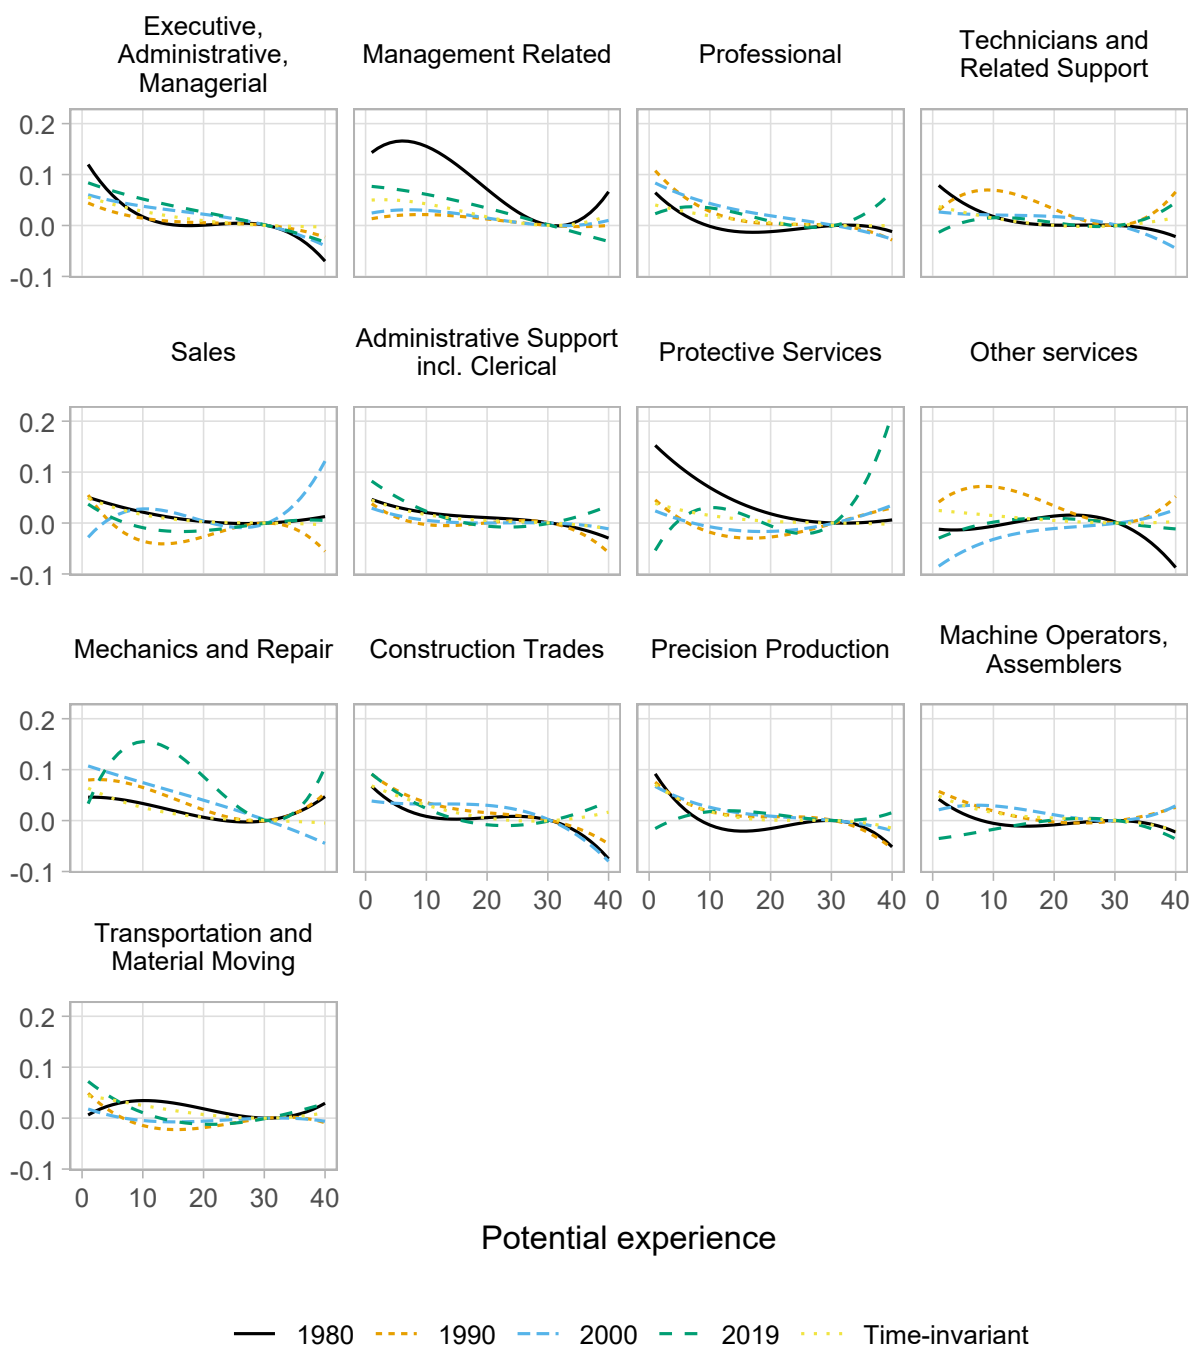


Figure D.14: Estimated occupational experience profiles for selected occupations and years, US

Notes: The figure plots the estimated experience profiles from equation (5) for the indicated occupations and years.

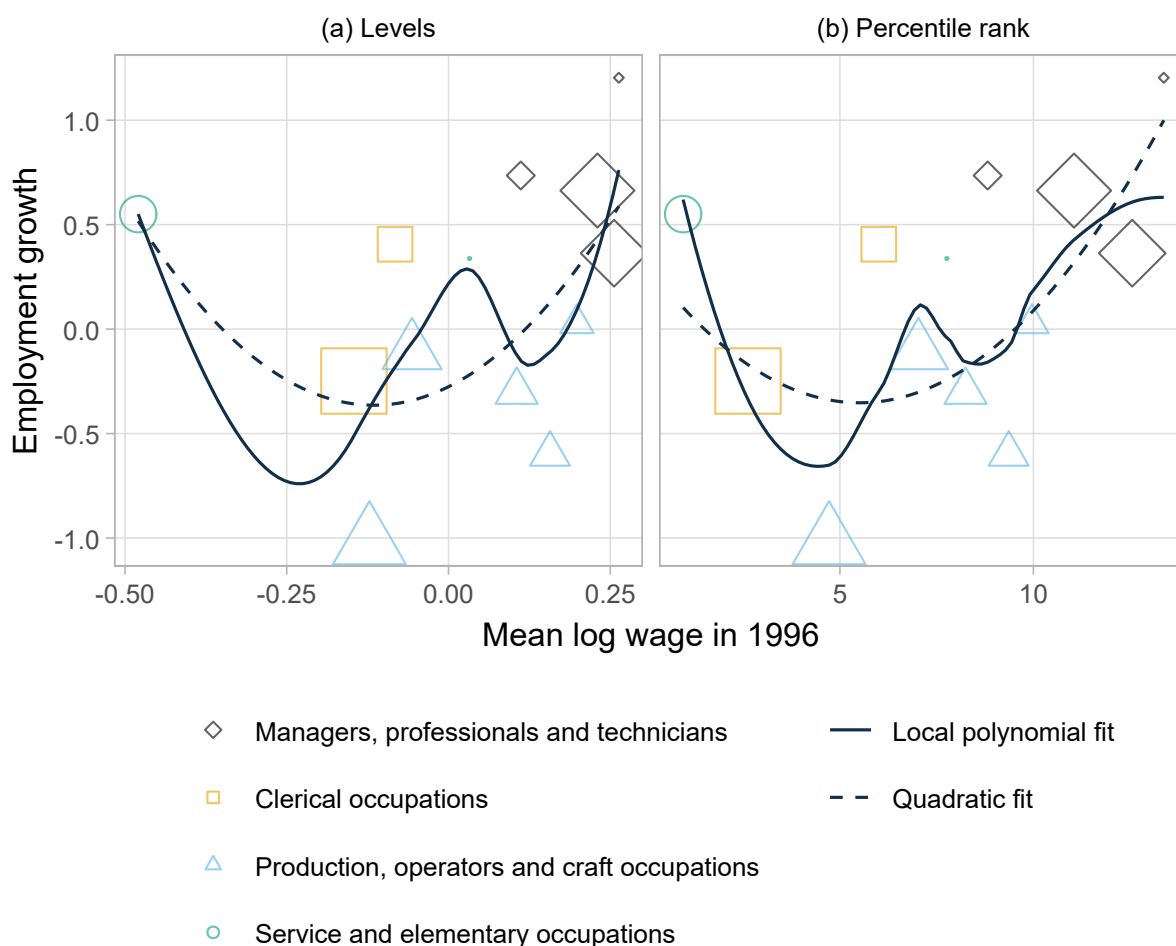


Figure D.15: Job polarization, US data

Notes: The figure plots the growth in log employment against mean log wages in 1979. In Panel (b), log wages have been percentile-ranked, weighted by initial employment. Each marker represents one of 13 occupations. The size of each marker is determined by the employment share in the first year and the regression lines are weighted accordingly.




Figure D.16: Relations between growth rates, US

Notes: The figure plots the bivariate relationships between the growth in mean log wages, cumulative estimated wage premia, and the implied change in mean skills. Each marker represents one of 13 occupations. The size of each marker is determined by the employment share in the first year and the regression lines are weighted accordingly.

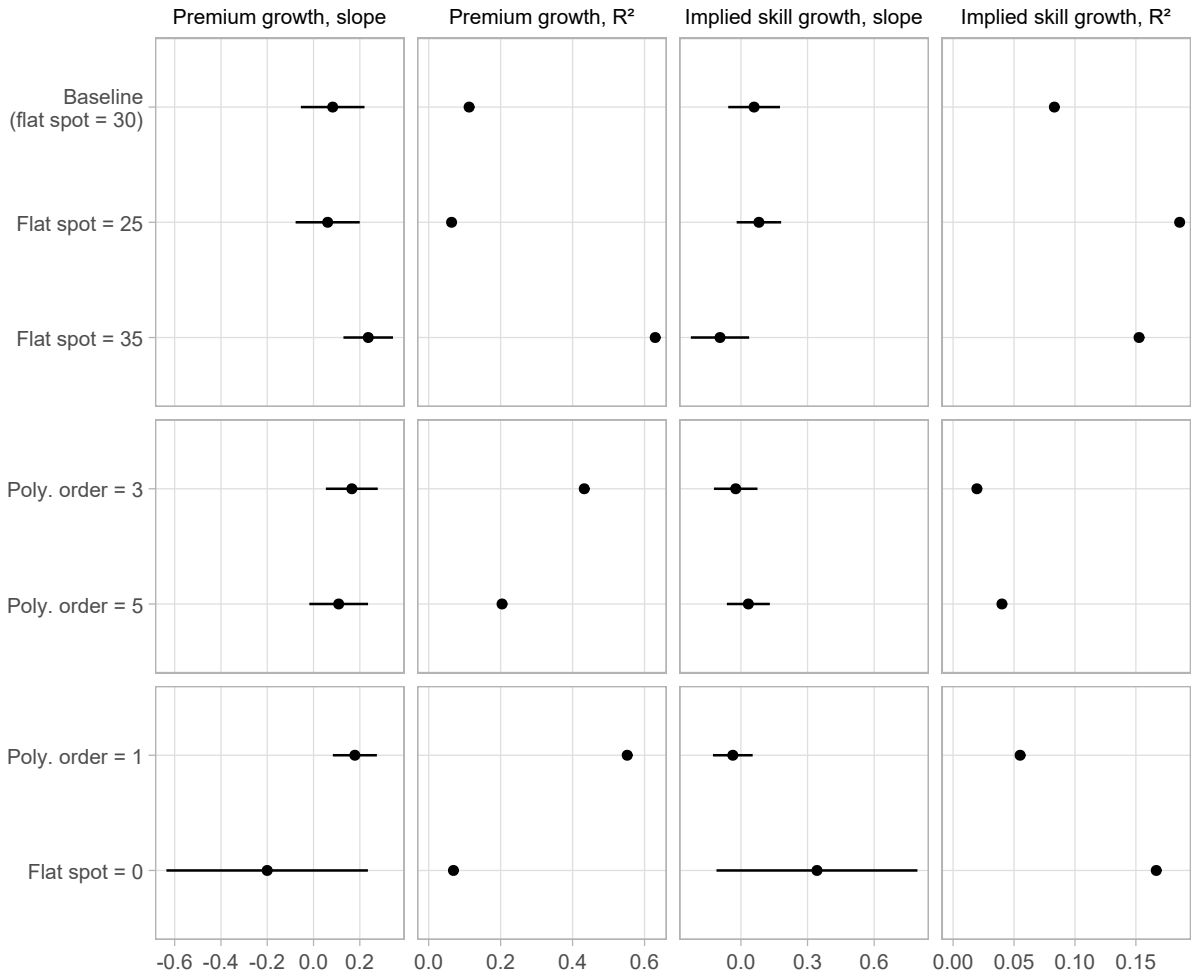


Figure D.17: Premia, skills, and employment growth—robustness checks, US

Notes: The table reports the coefficients, standard errors, and coefficients of determination from separately regressing cumulative estimated wage premia and the implied change in mean skills (growth in average wage minus premium growth) against the change in log employment at the occupation level for different sets of premia estimates. See the text for descriptions of how these estimates are produced. The weight assigned to each occupation is determined by the employment share in the first year. We use original survey weights when calculating occupation size and mean log wage.