Do Workers Sort to Firms or to Occupations?

Luke Heath Milsom* Shihang Hou*

August 17, 2026

Abstract

On average, high-wage workers work in high-wage firms. We ask whether this reflects workers matching to high-paying employers or holding high-paying occupations that those employers house. Using French and German matched employer–employee data, a flexible worker–job fixed-effects model, and correcting for limited mobility bias, we find that over 80% of worker–job sorting is between occupations rather than to firms within occupations. This is despite considerable within-occupation, across-firm dispersion in premia. Decomposing further, we find that in Germany, two-thirds of between-occupation sorting reflects differences across occupations in the average firm premium within that occupation. Since 1999, the documented rise in worker–firm sorting is consistent with a growing separation of occupations across firms rather than improved worker–firm matching.

*Heath Milsom: KU Leuven and the Institute for Fiscal Studies, luke.heathmilsom@kuleuven.be. Hou: Institute for Employment Research (IAB), shihang.hou@iab.de. With thanks to Abi Adams, Jesper Bagger, Stephane Bonhomme, Wolfgang Dauth, Binta Zahra Diop, Lucas Finamor, Martin Friedrich, Simon Janssen, Vatsal Khandelwal, Adrian Lerche, Sanghamitra Mukherjee, Barbara Petrongolo, Simon Quinn, Malte Sandner, Margaret Stevens, Martin Weidner, Verena Wiedemann, and Hannah Zillessen for helpful feedback and comments in the course of writing this paper. Access to confidential German data was made possible by the IAB. Access to confidential French data was made possible within a secure environment provided by CASD – Centre d'accès sécurisé aux données (Ref. 10.34724/CASD).

1 Introduction

Two central facts in the study of wage inequality are that some firms pay persistently more than others (Abowd et al., 1999; Mortensen, 2003; Card et al., 2013, 2016; Song et al., 2019), and that some occupations command higher pay than others (Autor et al., 2003; Goos and Manning, 2007; Acemoglu and Autor, 2011). These literatures have developed largely in parallel, with only relatively few papers considering both margins jointly. We argue that they cannot be studied separately, because the assignment of occupations to firms is itself a central feature of the wage structure. In the universe of German and French administrative data, occupations are extraordinarily segregated across firms: knowing a worker's employer resolves two-thirds of the uncertainty about their detailed occupation, a degree of segregation that survives drastic coarsening of the occupational classification and has risen over the past two decades. When occupations and firms overlap this tightly, an occupation-blind analysis cannot distinguish workers selecting high-paying employers from workers selecting high-paying lines of work.[1]

This paper disentangles the two margins. We define a "job" as a firm–occupation pair and estimate a two-way fixed effects model in the tradition of Abowd et al. (1999) (AKM), recovering worker and job effects from matched employer–employee data covering workers in Germany and in France. We then ask two questions. First, along which margin do workers sort? We use the law of total covariance to split worker–job covariance into two components. First, there is a within-occupation component, the covariance between worker and job wage components holding occupation fixed, often emphasised in the two-way fixed effects literature (Abowd et al., 1999; Card et al., 2013; Song et al., 2019). Second, there is a between-occupation component: the covariance at the occupational level between mean within-occupation worker quality and pay. This object is central to Roy-style models of occupational selection (Heckman and Sedlacek, 1985; Keane and Wolpin, 1997). Throughout we correct for limited-mobility bias using the leave-out estimator of Kline et al. (2020) (Andrews et al., 2008; Bonhomme et al., 2023), which we extend to conditional moments and to the channels of our projection decomposition, and probe exogeneity using event-study designs following Card et al. (2013). Second, we ask: what are the returns to the sorting margin made of? By projecting our estimated job fixed effects into occupation and firm components, we then further decompose the between-occupation component into an occupation-premium channel and a firm-premium channel.

First, we find that the answer to the first question is that sorting to occupations accounts

[1] Consider computer scientists at Google. An occupation-blind model reads their presence there as high-ability workers matching with a high-premium employer. But if Google is one of the few firms housing elite software engineering roles, the operative choice may have been the occupation, with the employer arriving as part of the bundle. The wage data alone cannot separate the two readings without occupational information.

for more of total wage variance than sorting to firms within occupations. In both Germany and France, sorting to occupations accounts for over 80% of total worker–job sorting. We find small shares of variance due to within-occupation sorting even despite large within-occupation heterogeneity in firm premia. High-wage workers do work at high-wage firms — but mainly because they work in high-wage occupations, far less because, within their occupation, they sort into high-premium employers, despite significant available gains on this dimension.

Second, in response to the second research question, we find that the returns to the occupational sorting margin mostly reflect firm premia. Two-thirds of between-occupation sorting arises because the occupations that attract high-ability workers have systematically higher average firm premia. The answers to the two questions are seemingly contradictory, but can be reconciled by noting that occupations are not randomly assigned across firms. Occupations in which high-wage workers work are disproportionately housed at high-premium firms. We find an occupation-premium-to-firm-premium employment-weighted slope of 0.43. Workers therefore do collect firm premia, but they do so through occupational membership rather than through competition within occupational labour markets. Choosing an occupation is, de facto, choosing a distribution of employers. An occupation-blind model reads the resulting covariance as worker–firm matching. This is partly because its firm effect mechanically absorbs each firm's mix of occupation premia, and partly because true firm premia covary with worker ability through the occupations firms choose to house.

Finally, we extend the decomposition over four sub-periods spanning 1999–2022, motivated by the documented rising role of firms in explaining wage inequality (Card et al., 2013; Song et al., 2019). We find that the rise in measured sorting is concentrated almost entirely in the firm-premium channel. Job-premium dispersion itself, by contrast, follows a hump shape, returning by 2017–2022 close to its 1999–2004 level, while occupational segregation across firms rises throughout. This combination, rising covariance through the firm channel with flat premium dispersion, is the pattern one expects if the assignment of occupations to firms became more aligned with the premium structure, rather than premia widening or matching improving. It is consistent with accounts of outsourcing and the fissuring of the workplace (Weil, 2014; Goldschmidt and Schmieder, 2017; Handwerker, 2023; Cortes et al., 2024).

The moments we document are relevant for theories of the labour market in two main ways. First, we find that sorting is predominantly a worker–occupation phenomenon as in Roy-type frameworks (Lee, 2005; Lee and Wolpin, 2006; Yamaguchi, 2012). Models of worker–firm assignment (Lise et al., 2016; Bagger and Lentz, 2019; Lentz et al., 2023; Lamadon et al., 2024) must rationalise the lack of within-occupation sorting despite considerable firm-premium variance, perhaps through search frictions, rationing, or compensating amenities. Second, large and within-occupation dispersed firm premia that are uncorrelated with worker quality are most naturally interpreted as firm wage policies: rents shared under internal-equity constraints (Card et al., 2016, 2018). As a result, the key endogenous object becomes the firm's occupational boundary, as the only way to escape paying a high premium to everyone is to push low-value occupations outside the firm. This is consistent with our finding that the firm premium channel within between-occupation sorting has increased over time.

Our paper contributes to the large literature that has used two-way fixed effect models to study worker–firm pay structure and its implications for wage inequality. Much of this literature has used two-way fixed effects to decompose wage inequality into components relating to workers, firms, and worker–firm covariance (Abowd et al., 1999, 2002; Card et al., 2013; Song et al., 2019; Babet et al., 2025; Barth et al., 2016; Bonhomme et al., 2023; Kline, 2024; Engbom and Moser, 2022; Bassier, 2023; Palladino et al., 2025; Haltiwanger et al., 2024). In the case of sorting, the main message from this line of work has been that high-wage workers work in high-wage firms. We contribute three empirical findings which refine and elaborate upon this result. First, we find that high-wage workers work in high-wage firms mainly in high-wage occupations, with the occupation-sorting channel accounting for over 80% of total worker–job sorting. Second, despite low levels of within-occupation sorting, we find that there are quantitatively significant amounts of wage dispersion within occupations, suggesting that the low levels of within-occupation firm sorting are not driven by the lack of firm premia. Third, we find that the colocation of high-wage occupations and high-wage firms contributes substantially to total wage inequality, and that the recent increase in worker–firm sorting may have been driven by increased occupation fissuring between firms (Handwerker, 2023; Cortes et al., 2024; Håkanson et al., 2021).

Our paper also contributes to the literature that jointly explores the role of occupations and firms in wage determination using an AKM approach. Prior work that incorporates occupations typically relies on more restrictive specifications, such as log-additive separability between firm and occupation effects (Torres et al., 2018) or interactions defined over coarse occupation categories (Goldschmidt and Schmieder, 2017). In contrast, our job fixed-effects specification allows wage premia to vary flexibly at the firm–occupation level, without imposing additive structure across firm or occupation main effects. This flexibility is important because additive restrictions mechanically limit within-occupation variation in firm wage premia and can compress estimated wages for workers, particularly in high-paying occupations employed at lower-paying firms. By allowing unrestricted firm–occupation heterogeneity, our approach captures richer patterns of wage dispersion and sorting that are obscured by more parsimonious specifications. Empirically, we find that match effects are important; firm–occupation match effects account for 15% of the total variance of job fixed effects after correcting for limited-mobility bias.

Finally, our work provides reduced-form moments which we believe to be informative for the large and growing literature studying labour market matching. A substantial body of work studies worker–firm sorting using models of assignment and matching, emphasising complementarities between worker ability and firm productivity as a source of wage dispersion and inequality (Lise et al., 2016; Bagger and Lentz, 2019; Lamadon et al., 2024). Parallel work highlights the role of worker–occupation sorting, in which comparative advantage and task-specific skills shape occupational choice and earnings profiles (Lee, 2005; Lee and Wolpin, 2006; Yamaguchi, 2012). Our results suggest that worker–occupation sorting, rather than worker–firm sorting, is the dominant margin underlying observed wage covariance patterns. First, we document substantial dispersion in firm wage premia within occupations, but little covariance between worker and firm wage components once occupation is held fixed. Second, the colocation of firms and occupations has largely been modelled in sorting models as a primitive, or exogenously given. Our decomposition suggests that this is an important determinant of the wage structure.

The paper proceeds as follows. Section 2 presents the econometric approach. Section 3 describes the French and German administrative data used in our decomposition. Section 4 presents the main decomposition results, robustness, and validation exercises. Section 5 relates our results to the broader labour economics literature on wage determination and wage inequality. Finally, Section 6 concludes.

2 Separately identifying worker–firm and worker–occupation wage sorting

We consider a framework that allows us to separately identify the sorting of high-wage workers to high-wage firms and to high-wage occupations, two margins that the standard wage-decomposition literature bundles together. To do this, we build directly on the AKM framework (Abowd et al., 1999; Card et al., 2013; Song et al., 2019), which splits log wages into additive worker and workplace components. Our single innovation is to replace the model's firm fixed effect with a *job* fixed effect, where a job j is a firm–occupation pair, $j \in \mathcal{J} = \mathcal{F} \times \mathcal{O}$. Letting $J(i,t)$ denote the job held by worker i in period t, the specification is given in equation 1. We apply the model to log wages residualised on a set of time-varying worker covariates X_{it}.[2]

$$\ln(w_{it}) = \alpha_i + \lambda_{J(i,t)} + \varepsilon_{it} \tag{1}$$

[2] For computational reasons we follow the recommendation in Kline et al. (2020) and apply our models to $\ln(w_{it}) - X_{it}\hat{\beta}$ throughout, where X_{it} comprises a cubic age profile and year fixed effects.

Here α_i is a worker effect, $\lambda_{J(i,t)}$ a job effect, and ε_{it} a residual. The job effect extends the worker–firm two-way fixed effects model in two ways. First, it lets a firm pay different premia to different occupations: a software developer can earn a higher premium than an accountant at the same firm. Second, it lets the same occupation be paid differently across firms, even firms with the same overall pay level. The standard worker–firm model is the special case in which the job effect is constant across occupations within a firm, so equation 1 nests it. The model admits the classic AKM decomposition of the variance of log wages,

$$\mathbb{V}[\ln(w_{it})] = \underbrace{\mathbb{V}(\alpha_i)}_{\substack{\text{Variance due to}\\ \text{individual}\\ \text{heterogeneity}}} + \underbrace{\mathbb{V}(\lambda_{J(i,t)})}_{\substack{\text{Variance due to job}\\ \text{heterogeneity}}} + \underbrace{2\cdot\mathrm{Cov}(\alpha_i, \lambda_{J(i,t)})}_{\substack{\text{Variance due to}\\ \text{workers sorting into}\\ \text{jobs}}} + \mathbb{V}(\varepsilon_{it}). \quad (2)$$

2.1 Decomposing worker–job sorting

The job effect bundles two margins that we want to distinguish between. A job is a firm–occupation pair, so the premium λ_j attached to it can reflect the firm a worker is employed at, the occupation they hold, or the particular pairing of the two. The aggregate worker–job covariance $\mathrm{Cov}(\alpha_i, \lambda_{J(i,t)})$ in equation 2 inherits this ambiguity: it tells us that high-ability workers hold high-premium jobs, but not whether they do so by sorting to high-premium firms, to high-premium occupations, or to firm–occupation combinations that happen to pay well together.[3] Consider Hospital A, a specialised surgical centre, and Hospital B, a general nursing clinic. Two forces may push up average pay at Hospital A in ways that an occupation-blind model cannot distinguish. First, surgeons command a market-wide premium over nurses regardless of where they work; if both hospitals pay the standard rate for each role, this is a pure *occupation* premium and involves no firm-level sorting at all. Second, suppose Hospital A is simply a more generous employer, paying every role — surgeon and orderly alike — above the market rate. Because able surgeons are concentrated there, this generosity raises their pay and so contributes to the covariance between worker ability and firm pay: it is a genuine channel of measured worker–firm sorting, running through the joint distribution of occupations and firms. But it is not the within-occupation matching of able workers to high-paying firms that the covariance is usually taken to represent — the premium accrues to whoever works at Hospital A, not to its surgeons for being able. A worker–firm model has no occupation axis on which to record either force. Observing

[3] In our paper, we interpret the covariance as measuring worker–firm sorting in terms of pay premia in line with studies like Song et al. (2019); Card et al. (2013). Recent research has shown that there could be a less straightforward relationship between worker fixed effects and firm productivity, both theoretically (Eeckhout and Kircher, 2011; Lopes de Melo, 2018), and empirically (Lochner and Schulz, 2024). We discuss our results in light of this work in Section 5.

only that the average worker at Hospital A out-earns the average worker at Hospital B, it attributes the whole gap to a firm-wide premium, and — because able workers (surgeons) are concentrated at Hospital A — records a large positive covariance between worker ability and the firm effect.

The first approach we adopt is to condition sorting on workers' occupations; if firms are bundles of occupations, then the question of whether able workers match to high-premium *firms* has a clean answer only once we hold the occupation fixed. To disentangle these channels, we apply the law of total covariance, partitioning the joint distribution of worker and job effects by occupation (o):

$$\mathrm{Cov}(\alpha_i, \lambda_{J(i,t)}) = \underbrace{\mathbb{E}\left[\mathrm{Cov}(\alpha_i, \lambda_{J(i,t)} \mid o)\right]}_{\text{Within-occupation sorting}} + \underbrace{\mathrm{Cov}\left(\mathbb{E}[\alpha_i \mid o], \mathbb{E}[\lambda_{J(i,t)} \mid o]\right)}_{\text{Between-occupation sorting}} \tag{3}$$

The first term, $\mathbb{E}\left[\mathrm{Cov}(\alpha_i, \lambda_{J(i,t)} \mid o)\right]$, captures within-occupation sorting: the extent to which higher-ability workers are systematically matched to higher-paying firms within the same occupation. The second term, $\mathrm{Cov}\left(\mathbb{E}[\alpha_i \mid o], \mathbb{E}[\lambda_{J(i,t)} \mid o]\right)$, captures between-occupation sorting: the covariance between average worker ability and average job pay across occupations. We can also decompose the variance of job fixed effects, $\mathbb{V}(\lambda_{J(i,t)})$, using the law of total variance, partitioning the distribution of job fixed effects by occupation into two components: the variance of job fixed effects within occupations between firms $\mathbb{E}\left[\mathbb{V}\left(\lambda_j \mid o\right)\right]$, and the variance of job fixed effects between occupations $\mathbb{V}(\mathbb{E}[\lambda_j \mid o])$.

While this partition isolates the genuinely within-occupation margin of sorting, it does not clarify the between-occupation component. Between-occupation sorting can arise because able workers select into occupations that are paid well everywhere—a worker–occupation sorting story, with no role for firms—or because those occupations are concentrated in firms that pay a premium to all their staff, the segregation channel. To separate the firm and occupation components within the estimated job fixed effect, we write the job effect as its employment-weighted linear projection onto occupation and firm indicators,

$$\lambda_j = \gamma_o + \psi_f + \Omega_{of}, \tag{4}$$

where γ_o is the occupation component, ψ_f the firm component, and Ω_{of} a residual match component that is orthogonal to the occupation and firm spaces by construction.[4] Substituting equation 4 into equation 3, and using the orthogonality of the projection (γ_o is constant within an occupation, and $\mathbb{E}[\Omega_{of} \mid o] = 0$ for every o), we can split worker–job sorting into

[4] Here ψ_f is the firm component of the job-premia surface; it is distinct from — and need not equal — the firm effect of a standard worker–firm AKM model, which we denote ϕ_f^{firm}.

six interpretable channels, two of which are identically zero:

$$\begin{array}{l|ccc} & \gamma_o \text{ (occupation)} & \psi_f \text{ (firm)} & \Omega_{of} \text{ (match)} \\ \hline \text{Within-occ} & 0 & \mathbb{E}[\mathrm{Cov}(\alpha_i, \psi_f \mid o)] & \mathbb{E}[\mathrm{Cov}(\alpha_i, \Omega_{of} \mid o)] \\ \text{Between-occ} & \mathrm{Cov}(\mathbb{E}[\alpha_i \mid o], \gamma_o) & \mathrm{Cov}(\mathbb{E}[\alpha_i \mid o], \mathbb{E}[\psi_f \mid o]) & 0 \end{array} \tag{5}$$

The rows of equation 5 sum to the within- and between-occupation terms of equation 3; the columns attribute total worker–job sorting to occupation premia, firm premia, and match effects. This decomposition thus splits worker–job sorting into four interpretable components: (i) sorting between workers and occupations ($\mathrm{Cov}(\mathbb{E}[\alpha_i \mid o], \gamma_o)$), (ii) sorting between workers and firms within occupations ($\mathbb{E}[\mathrm{Cov}(\alpha_i, \psi_f \mid o)]$), (iii) the covariance across occupations between average worker ability $\mathbb{E}[\alpha_i \mid o]$ and the average premium of their employers $\mathbb{E}[\psi_f \mid o]$, capturing the sorting of high-ability occupations into high-premium firms ($\mathrm{Cov}(\mathbb{E}[\alpha_i \mid o], \mathbb{E}[\psi_f \mid o])$), and (iv) a within-occupation match channel: able workers sorting to pairings that pay above their firm and occupation components ($\mathbb{E}[\mathrm{Cov}(\alpha_i, \Omega_{of} \mid o)]$).

2.2 Identification and estimation

Our main objective is to identify the decompositions in equation 5. To do this, we first estimate model 1 and the decomposition in equation 2 as the basis for our subsequent decompositions. The job effects are identified only up to a normalisation, and only on the largest connected set of jobs; we estimate $\{\widehat{\alpha}_i, \widehat{\lambda}_j\}$ from workers' moves between jobs. There are at least two main econometric concerns. First, the fixed effects themselves could be biased, and the subsequent decompositions could inherit this bias. Second, as the fixed effects are only identified from movers across jobs, limited-mobility bias will cause bias in quadratic transformations of the estimated fixed effects.

First, the estimates of the fixed effects are unbiased under the assumption of exogenous moves conditional on worker and job fixed effects. The validity of this assumption for the classic worker–firm framework is discussed in detail in Card et al. (2013), and in our paper, we focus on what changes in our context relative to the standard AKM model. As identification comes from movers across jobs, a sufficient condition is: $\forall j, \Pr(J(i,t) = j \mid \varepsilon) = \Pr(J(i,t) = j)$; this is therefore what we focus our attention on. There are two main ways this assumption could be violated. First, there might be match effects that are not captured by worker and job fixed effects, i.e., if workers sort to jobs on the basis of a worker–job match-specific characteristic not captured by α_i and λ_j. Second, temporary variation in wages may be correlated with the job that workers perform. A concern of this type in Card et al. (2013) is that the statistical model is incompatible with models of the labour market where workers move to jobs due to high transitory wage offers not related to the firm fixed effect. In our model, an additional problem of this kind is if temporary occupation-specific productivity shocks lead to substantial movement between occupations in the period. We probe these assumptions using event studies and by focusing only on cross-firm moves following Card et al. (2013) in Section 4.3 and find that these concerns do not appear to be substantiated in our setting.

Second, Andrews et al. (2008) and Andrews et al. (2012) pointed out that incidental parameter bias arising from small numbers of moves of workers between jobs might bias second-order moments of fixed effects from two-way fixed effect regressions, terming this "limited mobility bias". Bonhomme et al. (2023) show that limited-mobility bias is empirically large in practice. This means that standard plug-in estimates of the standard decomposition in equation 2 as well as the decomposition due to the law of total covariance in equation 3 may be significantly biased. We address this problem by applying the corrections proposed by Kline et al. (2020) (hereafter KSS). Intuitively, Kline et al. (2020)'s approach is to unbiasedly estimate the noise in the fixed effects and use this to debias estimates of the variance components using these noisy fixed effects. This approach has the disadvantage that we can now only estimate fixed effects in the largest leave-one-out connected set. This reduces the effective sample size and restricts estimation to a particular sub-sample of the data containing better-connected workers and firms.

We first apply the KSS correction to obtain unbiased estimates of the second moments appearing in the worker–job variance decomposition, including the variance of worker fixed effects, the variance of job fixed effects, and the covariance between worker and job fixed effects in equation 2. To obtain KSS-corrected estimates of the terms in equations 3 and 5, as well as the analogous between–within decomposition of job-effect variance, we note that each target can be expressed in the quadratic form required by Kline et al. (2020). In particular, let

$$\widehat{\theta} = \begin{pmatrix} \widehat{\alpha} \\ \widehat{\lambda} \end{pmatrix}$$

collect the worker and job fixed effects from the original worker–job model. Each target moment q can then be written as

$$\widehat{m}_q = \widehat{\theta}' Q_q \widehat{\theta},$$

where Q_q is a known weight matrix determined by the employment weights, occupation assignments, and the particular component being estimated. This representation follows because occupation-level averages of the fixed effects, as well as the occupation, firm, and match components γ_o, ψ_f, and Ω_{of}, are linear functions of the original worker and job fixed effects. The KSS correction therefore applies directly to each target moment, with only the weight matrix Q_q changing across components.[5] Since this is a direct application of the Kline et al. (2020) correction, it does not involve any additional sample restrictions beyond what is already required to correct the variance components in equation 2.

For our main specification, we construct these weight matrices directly for both the parent between–within moments and the channels of equation 5. For exercises requiring repeated estimation across periods and occupation codings, we retain an alternative implementation of the parent between–within decomposition. Specifically, we estimate the KSS-corrected conditional covariance and variance within each occupation and aggregate them using occupation employment shares:

$$\widehat{C}^{\text{KSS}}_{\text{within}} = \sum_{o} \widehat{s}_o \,\widehat{\text{Cov}}^{\text{KSS}} \left(\alpha_i, \lambda_{J(i,t)} \mid o\right),$$

and

$$\widehat{V}^{\text{KSS}}_{\text{within}} = \sum_{o} \widehat{s}_o \,\widehat{\mathbb{V}}^{\text{KSS}} \left(\lambda_{J(i,t)} \mid o\right),$$

where $\widehat{s}_o$ is occupation o's employment share. The corresponding between-occupation components are obtained from the adding-up identities

$$\widehat{C}^{\text{KSS}}_{\text{between}} = \widehat{C}^{\text{KSS}}_{\text{total}} - \widehat{C}^{\text{KSS}}_{\text{within}}, \qquad \widehat{V}^{\text{KSS}}_{\text{between}} = \widehat{V}^{\text{KSS}}_{\text{total}} - \widehat{V}^{\text{KSS}}_{\text{within}}.$$

This conditional-moment implementation targets the same parent between–within decomposition but does not produce the more detailed occupation, firm, and match channels. They are used for the over-time and granularity exercises (Section 4.2, Appendix F). As an additional diagnostic, we also calculate plug-in between-occupation moments directly from employment-weighted occupation averages of the estimated fixed effects. Appendix D reports plug-in and corrected estimates side by side for every channel, quantifying limited-mobility bias channel by channel. The alternative implementations give the same substantive conclusions, providing reassurance that these substantive conclusions are not driven by how the decompositions are implemented.

One might instead recover γ_o and ψ_f by estimating an additive three-way worker–firm–occupation model directly (Torres et al., 2018), rather than as a projection of the job effect. This is akin to substituting equation 4 into equation 1, and estimating the three-way fixed effects model by OLS. We do not, for two reasons. First, estimating this model by OLS will force the match effects term Ω_{of} into the residual, implicitly assuming it away: $\mathbb{E}[\Omega_{of} + \varepsilon_{it} \mid o, f] = 0$. In practice, we find that match effects account for an empirically non-trivial share

[5] In the case of the KSS-corrections of the terms in equation 5, the trace terms are computed with a Johnson–Lindenstrauss random-projection approximation as in Kline et al. (2020), with draws shared across channels.

of total wage variance. Second, while the identification of AKM-style two-way fixed effect models with mover designs is well understood, the identification of such three-way fixed effects models is not transparent; for example, it is not clear what the three-way equivalent of the largest connected set necessary for identification would be. Defining $(\gamma_o, \psi_f, \Omega_{of})$ as a projection of λ_j imposes no restriction on match effects and makes the components exact linear functions of the job effect, so that they inherit the identification of the job model itself.

3 Data

We establish our results using administrative matched employer–employee data from two of the largest labour markets in Europe. Our main setting is Germany, where we observe a full-count panel covering 1999–2022, which gives us both a large cross-section and a long time series. We complement it with France as an external-validity setting, whose data have two features that make them a useful cross-check. First, French wages are uncensored, whereas German wages are top-coded at the social-security ceiling. Second, French occupation codes are cross-checked by INSEE against raw job titles provided by employers, thus reducing the likelihood of miscoding. We therefore confirm that our German findings replicate in France, where occupation is measured in a quite different way; this reassures us that the results reflect the labour market rather than the particular way occupation is recorded in any one dataset. Our main analysis uses 2017–2022 for Germany and 2015–2019 for France. In both countries, we check robustness to occupation measurement by using coarser occupation codes, arguing that while occupation miscoding may be possible in the case of very specific occupations with hundreds of categories, it is far less likely at the 1-digit or 2-digit level.

3.1 German sample

Our German data come from the Employee Histories (BEH), which contain employment spells for all workers outside the civil service. These spell data are converted into a worker-year panel, largely following the conventional data cleaning process suggested by Stüber et al. (2023). The primary limitation is the top-coding of wages above the social-security contribution ceiling, which we address by imputing censored wages following Card et al. (2013); this may lead us to understate sorting in higher-wage occupations. The key variables are real annual earnings, the establishment identifier,[6] and occupation codes based on the 5-digit KldB 2010 classification (1,286 categories in the full classification). Following Bonhomme et al. (2023), our sample is full-time male workers aged 25–60 in West Germany, and we residualise log annual wages on a cubic in age and year fixed effects.

[6] An establishment is roughly a firm–location unit and is distinct from both the legal firm and the workplace; see Card et al. (2013).

In the German data, workers' occupations are self-reported. Because occupation is not central to social security administration, employers may not always update a worker's occupation promptly, so within-firm occupation moves are likely under-reported. This measurement concern — which would, if anything, attenuate the role we find for occupations — is one motivation for the French external-validity exercise. In 2011, Germany also switched from the older KldB 1988 classification frame to the KldB 2010 classification frame, leading to a structural break in measurement in that period.

3.2 French sample

The main dataset underlying our French data is the "Base tous salariés" (BTS) data (Insee, 2024), a series of cross-sectional matched employer–employee datasets covering the universe of French workers (we remove those in government employment). We follow Babet et al. (2025) in chaining these repeated cross-sections together into a quasi-panel tracking individuals over time, by matching their data in time t in the year t data to their information for time $t-1$ in the year $t+1$ data. This methodology allows for over 95% of individuals in each cross-section to be matched. One limitation of this approach is that those who are out of the labour force for more than one calendar year cannot be matched to their previous employment history and are instead given a new individual identifier.

To make our results comparable to the cross-setting study of AKM results in Bonhomme et al. (2023), our sample is restricted to full-time male workers aged 25–60 in metropolitan France, satisfying standard hours and earnings thresholds. Our decomposition is performed on log annual wages, residualised on a cubic in age and year fixed effects. The key variables that we use are real annual earnings, the firm identifier (SIREN), and the 4-digit occupational classification (PCS), which provides 430 distinct categories. A worker's occupation is collected from compulsory monthly employer surveys, where reported occupation titles are cross-checked against reported occupation codes using specialist INSEE software. In the 10% or so of cases where the codes do not agree, additional correction processes are used.[7] This process of asking employers rather than employees, using occupation descriptions, and cross-checking the occupation coding, reduces the probability of incorrect occupation categorisation. Summary statistics for the French samples, across the five estimation samples and both periods, are reported in Appendix Table 6.

[7] More information can be found in an INSEE "Statistical Mail" available at https://www.insee.fr/fr/information/3647029?sommaire=3647035.

3.3 Summary statistics

Table 1 reports summary statistics for the German sample across the three estimation samples: the full sample, the largest connected set, and the largest leave-one-out (LOO) set of the worker–job model, on which the KSS correction is feasible. The full sample contains 65.5 million observations, 15.0 million workers, and 1.4 million firms. The worker–job LOO set retains two-thirds of observations (66%) but a smaller share of workers (61%), firms (34%), and jobs (23%).

While the KSS correction is important in correcting for limited mobility bias, its substantial cost is that it is feasible only on a selected group of firms which, on average, have more workers but fewer occupations per firm. As shown in the lower panel of Table 1, the median worker's firm size rises from 65 to 95 from the full to the LOO sample, the employment share in firms with fewer than 50 workers falls from 45% to 38%, the mean number of occupations per firm falls from 26 to 17, and the employment share in single-occupation firms rises from 8% to 15%. Despite this, the mean and residualised variance of wages are nearly unchanged across the samples, as is the composition of worker moves. Appendix Table 7 gives the full firm-size and occupation distributions for the included and excluded samples: the excluded observations are disproportionately in small, occupationally specialised firms, though the occupation mix and the wage structure by occupation are largely unchanged. Appendix Tables 3–5 report the same statistics for the three earlier German periods, where the selection has the same character.

Table 1: Summary statistics, Germany (2017–2022), across estimation samples

	Full	Connected	Jobs-LOO
Observations (m)	65.55	51.59	43.48
Workers (m)	14.97	10.54	9.15
Firms (m)	1.44	0.93	0.49
Jobs (m)	5.36	2.90	1.23
Mean log daily wage	4.77	4.79	4.80
Var. resid. log wage	0.248	0.240	0.239
Moves (% of obs)	11.3	13.2	11.7
across firms (% of moves)	85.0	85.6	84.5
across occ. (% of moves)	66.0	66.4	63.0
both (% of moves)	51.0	52.0	47.5
Selection characteristics			
Median worker's firm size	65	79	95
Empl. share, firms <50 (%)	45.0	41.4	37.9
Mean occ. per firm (empl.-wtd.)	25.85	20.24	16.60
Empl. share, 1-occ. firms (%)	8.2	10.9	14.7

Notes: Counts in millions. "Full" is the sample after the restrictions of Section 3.1; "Connected" and "Jobs-LOO" are the largest connected and leave-one-out connected sets of the worker–job model. Move rows are shares. Median worker's firm size is the employment-weighted median; occupations per firm counts distinct 5-digit occupations per firm-year, employment-weighted.

As a benchmark, Table 2 reports the standard worker–firm AKM decomposition in Germany, which most of the previous literature estimates, by plug-in across the relevant estimation samples. The worker–job leave-one-out restriction keeps a particular sub-population, and a natural concern is that this selection, rather than the labour market, drives our results: if the connectivity requirement dropped precisely the firms in which high-ability workers match to high-premium firms *within* occupations, our later finding that worker–firm sorting operates mainly *between* occupations would be an artefact of the sample rather than a feature of the data. Table 2 characterises the direction of any such selection. Moving from the standard firm leave-one-out set to the worker–job leave-one-out set used in our main analysis does not lower measured worker–firm sorting; if anything it *raises* it, from 12% to 15% of variance, with the worker and firm variances little changed.[8] The other measured variance components, i.e., the variances of the worker and firm fixed effects, are also remarkably similar across the two LOO samples. We therefore conclude that the worker–job LOO set that underlies our main results should be seen as a more connected sample than the full

[8]The same ranking holds under the Kline et al. (2020) correction, with $\mathrm{Cov}(\alpha_i, \phi_f^{\mathrm{firm}})$ of 0.018 on the firm leave-one-out set against 0.022 on the worker–job set.

dataset, and that, if anything, measured sorting should be lower in the full dataset.

Table 2: Worker–firm AKM decomposition across estimation samples (Germany, 2017–2022)

	Firms conn.	Jobs conn.	Firms LOO	Jobs LOO
Total residualised variance, Var(y)	0.24 (100%)	0.24 (100%)	0.23 (100%)	0.24 (100%)
Var. worker effect	0.15 (65.2%)	0.15 (62.3%)	0.15 (64.3%)	0.15 (61.5%)
Var. firm effect	0.039 (16.5%)	0.036 (14.9%)	0.032 (13.8%)	0.033 (13.8%)
Sorting, 2 Cov	0.021 (8.8%)	0.031 (13.1%)	0.028 (12.1%)	0.036 (15.1%)
Residual	0.022 (9.5%)	0.023 (9.7%)	0.023 (9.8%)	0.023 (9.6%)

Notes: Plug-in (uncorrected) decomposition of the standard worker–firm AKM model. Each cell reports the raw variance component (two significant figures, in residualised log-wage units) with its share of the column's total residualised variance, $\mathbb{V}(y)$, in parentheses; the top row reports $\mathbb{V}(y)$ itself, the 100% reference. Columns are the firm-model connected set, the worker–job connected set, the firm-model leave-one-out (LOO) set, and the worker–job LOO set used in the main analysis, each normalised by its own $\mathbb{V}(y)$. Sorting is $2\,\mathrm{Cov}(\alpha, \phi^{\mathrm{firm}})$.

3.4 Segregation of occupations across firms

A premise of our analysis is that occupations and firms must be considered jointly because occupations are segregated across firms. We quantify this with the Theil-H index of segregation (Theil and Finizza, 1971; Theil, 1972), which measures how much of the uncertainty about a worker's occupation is resolved by knowing their firm: it is zero if every firm reflects the economy-wide occupation mix and one if firms are perfectly occupationally specialised. We normalise the index to $[0, 1]$ and compute it on firms with more than ten workers to prevent our results from being driven entirely by very small firms.

Occupations are highly segregated across firms in both countries. Figure 1 plots the index over the studied periods. Knowing a worker's firm resolves about two-thirds of the uncertainty about their occupation; even at the 3-digit level it removes about 63% of the uncertainty in Germany and France alike. Segregation is moreover robust to drastic coarsening of the classification. The figure also plots the index at the 1-digit level — a handful of broad occupational groups — where knowing a worker's firm in Germany still resolves about 57% of the uncertainty about their occupation.[9] That so much segregation survives even at the broadest occupational grouping makes it implausible that the pattern is an artefact of fine, possibly noisily coded, occupation cells. The index has also risen over time in both countries. This rising segregation leaves open the possibility that recently documented increases in the contribution of firms to wage inequality (Song et al., 2019) partly reflect rising occupational segregation across firms — a question we take up in Section 4.2.

[9] The French index falls further at the 1-digit level, to around 40%, because the top level of the PCS classification collapses to far fewer categories than the German KldB; the two are therefore not directly comparable at this coarse level.

Figure 1: Segregation of occupations across firms in studied periods

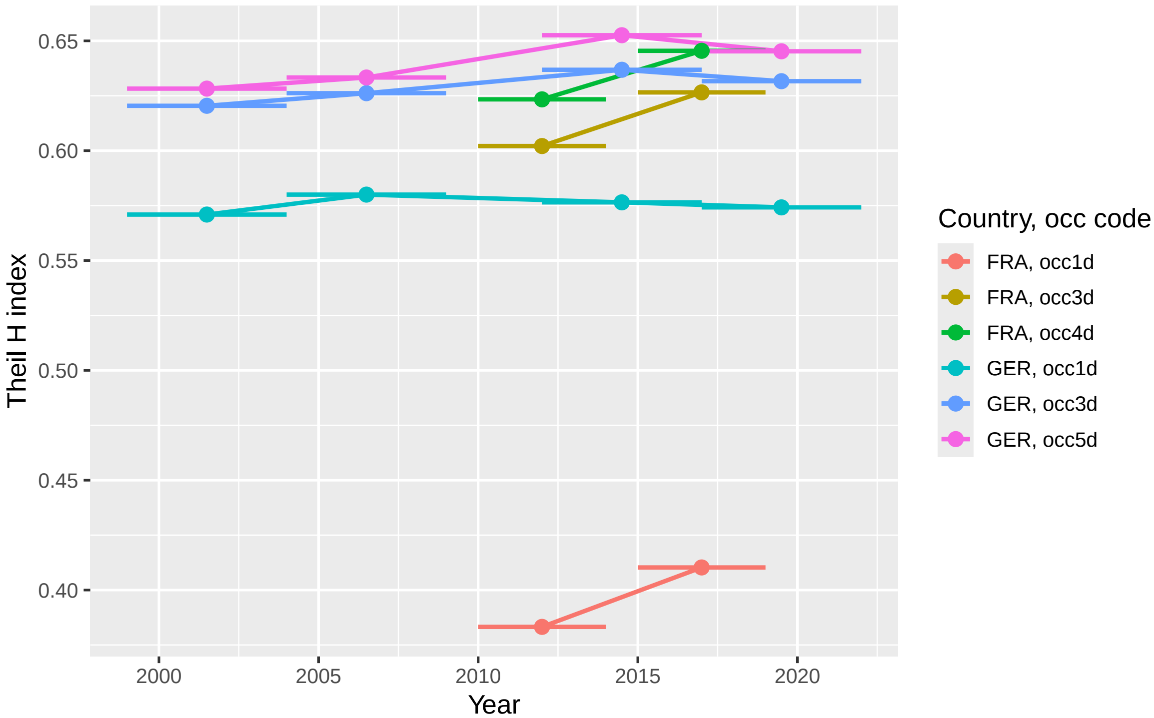

Notes: The figure reports the Theil-H index of occupational segregation across firms in Germany (1999–2022) and France (2010–2019). The index measures the employment-weighted reduction in the entropy of the occupation distribution from conditioning on firm membership, normalised by overall entropy, so that $H \in [0,1]$: it is zero if every firm reflects the economy-wide occupation mix and one if firms are perfectly occupationally specialised. In each country the index is computed under three occupation codings: a 1-digit grouping, a coarser 3-digit classification, and the most detailed classification used in the analysis (5-digit for Germany based on KldB 2010, 4-digit for France based on PCS). The 1-digit level is not directly comparable across countries, as the top level of the French PCS collapses to far fewer categories than the German KldB. The sample is restricted to firms with more than ten workers in the relevant period; worker and variable definitions follow Section 3.

4 Decomposition results

We now bring the framework of Section 2 to the data. We estimate the model on the main period (2017–2022), defining jobs at the finest, 5-digit occupation coding. Throughout we correct for limited-mobility bias using the leave-out estimator of Kline et al. (2020); the uncorrected plug-in estimates are in Appendix C, alongside both shares and absolute variance components, as recommended by Kline (2024).

Figure 2 summarises the main results. For both France and Germany, and for both the worker–firm and worker–job models, the left-hand set of grey bars in Figure 2 reports the estimated contribution of worker heterogeneity, $\widehat{\mathbb{V}}(\hat{\alpha}_i)$, to the overall log-wage variance. The middle set of bars reports the contribution of firm- or job-level heterogeneity: in grey, firm-level heterogeneity $\widehat{\mathbb{V}}(\hat{\phi}_f^{\text{firm}})$ from the worker–firm model; in blue, job-level heterogeneity from the worker–job model, where the two stacked bars sum to $\widehat{\mathbb{V}}(\hat{\lambda}_j)$. We split this job-level variance by the law of total variance into a within-occupation component (dark blue), $\widehat{\mathbb{E}}\big[\widehat{\mathbb{V}}(\hat{\lambda}_j \mid o)\big]$, and a between-occupation component (light blue), $\widehat{\mathbb{V}}\big(\widehat{\mathbb{E}}[\hat{\lambda}_j \mid o]\big)$. The right-hand set of bars reports the analogous covariance decomposition. In grey we show worker–firm sorting, $2\,\widehat{\mathrm{Cov}}(\hat{\alpha}_i, \hat{\phi}_f^{\text{firm}})$; in blue, the analogous worker–job sorting, $2\,\widehat{\mathrm{Cov}}(\hat{\alpha}_i, \hat{\lambda}_j)$, which we split by the law of total covariance into within-occupation sorting (dark blue), $2\,\widehat{\mathbb{E}}\big[\widehat{\mathrm{Cov}}(\hat{\alpha}_i, \hat{\lambda}_j \mid o)\big]$, and between-occupation sorting (light blue), $2\,\widehat{\mathrm{Cov}}\big(\widehat{\mathbb{E}}[\hat{\alpha}_i \mid o], \widehat{\mathbb{E}}[\hat{\lambda}_j \mid o]\big)$.

Figure 2: Decomposition results

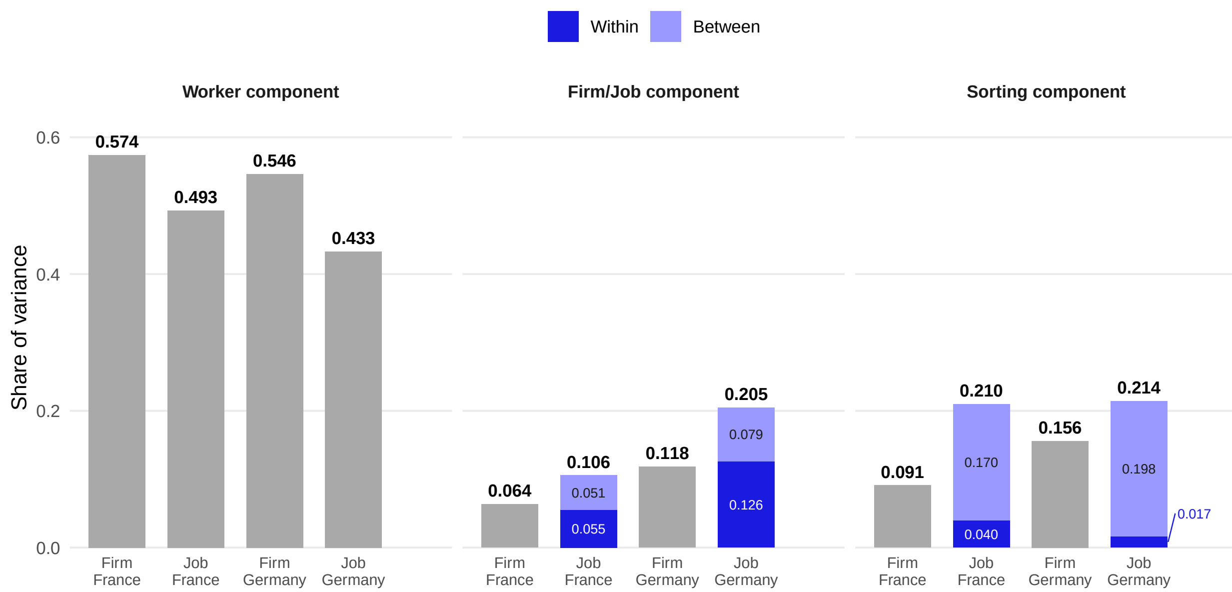

Notes: Leave-out (KSS) corrected variance components as shares of the total residualised log-wage variance, for the worker–firm AKM model ("Firm", grey) and the worker–job model ("Job") in France and Germany. The worker–job model's job-effect variance and sorting components are split into within-occupation (dark) and between-occupation (light) parts by the law of total covariance; the worker component and the worker–firm bars are not occupation-split. Sorting is $2\,\mathrm{Cov}$. Germany is the main 2017–2022 sample; France is the 2015–2019 external-validity sample. Bar labels are rounded independently and may not sum exactly to the printed totals.

We begin by comparing the two models at the level of the standard decomposition, before turning to the occupational split that is our main contribution. Moving from the worker–firm AKM model to the worker–job model widens the room for both workplace heterogeneity and sorting. In Germany the variance attributed to the workplace rises from 11.8% to 20.5% of log-wage variance, and total worker–workplace sorting rises from 15.6% to 21.4%; France shows the same direction of change. The reason is that the job effect, unlike the firm effect, can vary by occupation: it absorbs occupation pay premia that the occupation-blind firm effect either leaves in the residual or loads imperfectly onto the firm.

To understand what this larger sorting is made of, we apply the between/within-occupation decomposition of equation 3. Strikingly, we find that, holding occupation fixed, the covariance of worker and job effects contributes just 1.7% of variance in Germany. The overwhelming majority of the measured sorting, 19.8 of 21.4 percentage points, is between occupations.

France tells the same story, with 17.0 of 21.0 points between occupations. Set against the worker–firm benchmark the contrast is sharp: standard AKM attributes 15.6% of wage variance to the sorting of high-wage workers to high-wage firms, but once occupation is held fixed the within-occupation sorting component shrinks to 1.7% of variance. What AKM accounts as the matching of able workers to high-paying firms operates overwhelmingly through which occupations those workers hold, not through a within-occupation firm-matching premium.

Applying the same between-/within-occupation decomposition to the variance of job fixed effects, we find that the results are not driven by there being no residual variation of firm pay premia within occupations. In Germany the job-effect variance splits into 7.9% between occupations and 12.6% within; in France, 5.1% and 5.5%. The within-occupation component is the larger of the two in both countries: different firms pay materially different premia for workers in the same occupation, so the firm margin offers real wage differences even after occupation is netted out.

4.1 Firm-premium versus occupation-premium channels of between-occupation sorting

Having established that worker–job sorting is overwhelmingly a between-occupation phenomenon, a natural question is what between-occupation sorting itself consists of. As our theoretical discussion illustrates, between-occupation sorting could be high, even if there were no occupation premia, when high-wage occupations are overwhelmingly segregated into high-wage firms. To distinguish between worker–occupation sorting and occupation–firm segregation, we first project the estimated job fixed effect linearly into occupation, firm, and match components from Section 2, $\lambda_j = \gamma_o + \psi_f + \Omega_{of}$ (equation 4). This splits both the sorting covariance $\mathrm{Cov}(\alpha_i, \lambda_j)$ and the job-premium variance $\mathbb{V}(\lambda_j)$ into occupation, firm, and match channels, each in turn split between and within occupation, according to equation 5.

We report these sub-components with the leave-out correction of Section 2.2 applied to every channel; Appendix D reports the corresponding plug-in estimates side by side, quantifying limited-mobility bias channel by channel. To keep the comparison meaningful we estimate everything, including the standard worker–firm AKM benchmark, on the same worker–job leave-one-out sample and correct all components, so the AKM firm sorting $\mathrm{Cov}(\alpha_i, \hat{\phi}_f^{\text{firm}})$ and variance $\mathbb{V}(\hat{\phi}_f^{\text{firm}})$ furnish a like-for-like baseline against which to read the decomposition. The results are presented in Figure 3.

Our main finding is that the direct occupation-premium channel, $\mathrm{Cov}\big(\mathbb{E}[\alpha_i \mid o], \gamma_o\big)$, accounts for about one-third of between-occupation sorting (0.0083 of 0.0237), while the remaining two-thirds runs through the firm-premium channel, $\mathrm{Cov}\big(\mathbb{E}[\alpha_i \mid o], \mathbb{E}[\psi_f \mid o]\big) =$

0.0154: the occupations into which high-ability workers sort are high-paying largely because they are concentrated at firms paying high premia to their entire payroll. Sorting of workers to firms within occupations is, by contrast, far smaller — a corrected covariance of 0.0055, about a third of the size of its between-occupation counterpart, and partly offset by a negative within-occupation match channel (−0.0036).

How should we interpret the between-occupation firm-premium term? One interpretation is that it is best read not as a separate margin of sorting but as a consequence of two other features of the data: (i) workers sort across occupations on ability, accounting for about a third of the total sorting term, and (ii) the occupations that attract high-ability workers are disproportionately housed at high-premium firms — Figure 5 documents the analogous relationship for occupation premia. Given this firm–occupation structure, a positive covariance between worker ability and firm premium arises even though, within occupations, the tendency of high-ability workers to be at high-premium firms is comparatively weak.

Figure 3: Worker–job sorting $\mathrm{Cov}(\alpha_i, \lambda_j)$ by channel, against the AKM firm-sorting baseline (Germany, 2017–2022)

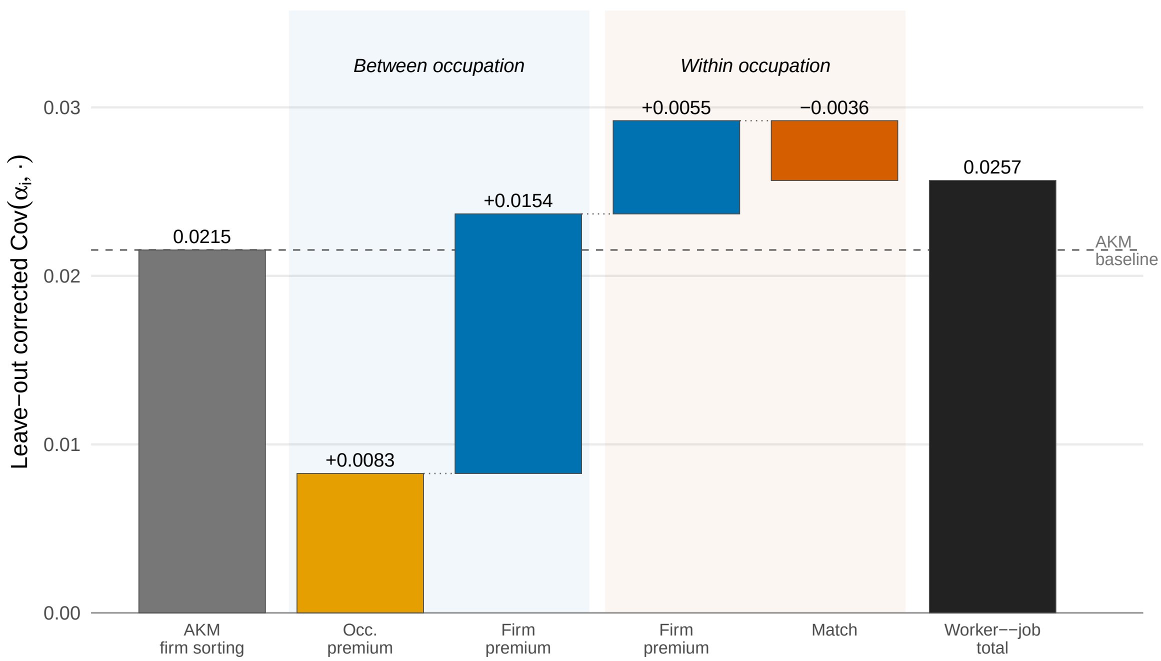

Notes: Waterfall building leave-out corrected worker–job sorting $\mathrm{Cov}(\alpha_i, \lambda_j)$ from its cross-decomposition cells, grouped into between- and within-occupation blocks (occupation premium γ_o, firm premium ψ_f, match Ω_{of}). The dashed line and leftmost bar give the standard worker–firm AKM sorting $\mathrm{Cov}(\alpha_i, \hat{\phi}_f^{\text{firm}})$. All components are leave-out (KSS) corrected — the channels using the extension described in Section 2.2, the AKM baseline using the standard leave-out estimator — and estimated on the same worker–job leave-one-out sample, so the channels sum exactly to the corrected total underlying Figure 2. Plug-in estimates for every cell are reported in Appendix D.

Figure 4 applies the same decomposition to the variance of the job premia, $\mathbb{V}(\lambda_j)$. Here the within-occupation firm-premium component is the largest single term: firms pay materially different premia for the same occupation, so there is substantial within-occupation firm wage dispersion. This dispersion exists, but high-wage workers only weakly sort to those high-premium firms (conditional on occupation), so the corresponding within-occupation firm-premium covariance is far smaller than this dispersion would allow. The match component, which an additive worker–firm–occupation model discards entirely, accounts for 15% of job-premium variance after correction.[10]

Figure 4: Worker–job premium variance $\mathbb{V}(\lambda_j)$ by channel, against the AKM firm-variance baseline (Germany, 2017–2022)

Notes: As Figure 3, for the variance of the job premia $\mathbb{V}(\lambda_j)$. Variance of a sum carries an occupation × firm interaction term, $2\,\mathrm{Cov}(\gamma_o, \mathbb{E}[\psi_f \mid o])$, shown as a separate (between-occupation) step. All components are leave-out (KSS) corrected; the correction is largest for the match component, whose plug-in value is inflated by limited-mobility bias (Appendix D).

Figure 18 in the Appendix presents an overview of all the decompositions performed above.

4.1.1 Occupations where wages are high cluster at high-premium firms

The firm-premium channel of between-occupation sorting rests on a specific claim: that occupations where wages are high are themselves concentrated at firms that pay high premia. The cross term $2\,\mathrm{Cov}(\gamma_o, \mathbb{E}[\psi_f \mid o])$ that appears in the variance decomposition (Figure 4) captures whether high-premium occupations cohere with high-premium firms — a sufficient condition for the claim that high-wage occupations cluster at high-premium firms. To study

[10] Limited-mobility bias inflates plug-in variance components, most severely for the match component: the correction reduces the match variance from 0.0132 to 0.0073 and the within-occupation firm variance from 0.0279 to 0.0229 (Appendix D).

this object, Figure 5 plots, for each occupation, the mean firm premium of the firms that employ it, $\mathbb{E}[\psi_f \mid o]$, against the occupation's own premium γ_o.

The relationship is positive: occupations that command a wage premium are also employed disproportionately at high-premium firms (an employment-weighted slope of 0.43 after leave-out correction, $\text{Cov}(\gamma_o, \mathbb{E}[\psi_f \mid o]) = 0.0021$). In a standard worker–firm AKM model, the firm effect absorbs this correlation: a firm that employs high-premium occupations is mechanically assigned a higher firm effect, even if it pays no more than average within each occupation. High-ability workers' sorting into those occupations then shows up as worker–firm sorting, when it is more accurately worker–occupation sorting at firms specialised in high-premium occupations.

Figure 5: High-premium occupations are employed at high-premium firms (Germany, 2017–2022)

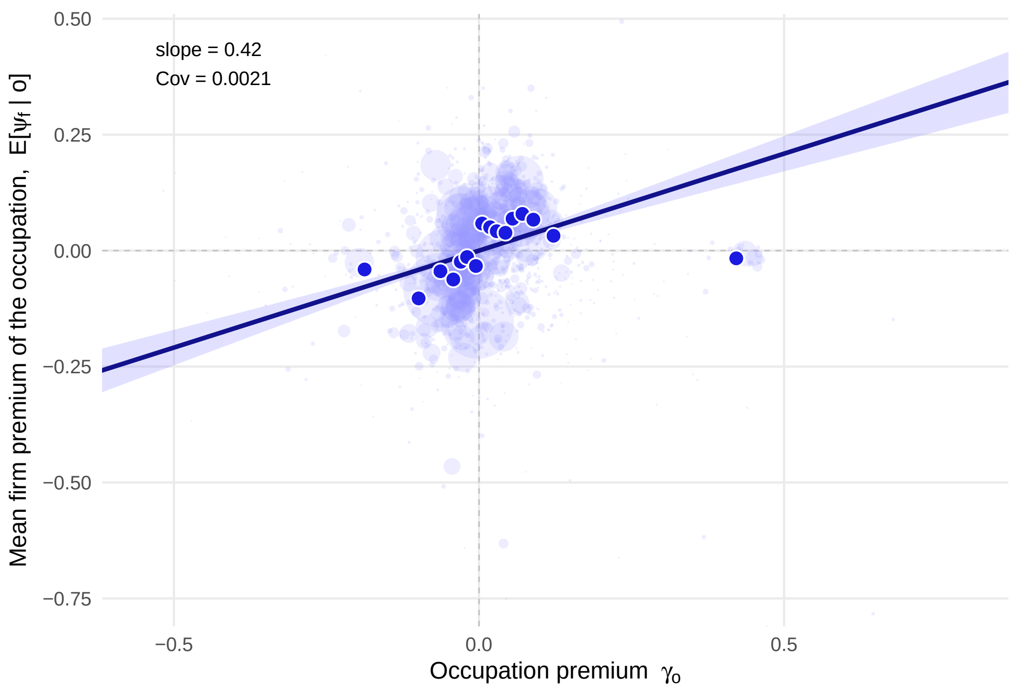

Notes: Each faint bubble is an occupation, with area proportional to employment; the horizontal axis is the occupation premium γ_o and the vertical axis the employment-weighted mean firm premium of the firms employing that occupation, $\mathbb{E}[\psi_f \mid o]$, both from the linear projection $\lambda_j = \gamma_o + \psi_f + \Omega_{of}$. Solid points are an employment-weighted binscatter (16 bins of equal occupation count); the line is the employment-weighted regression with a 95% confidence band; dashed lines mark the employment-weighted means. The plotted components are plug-in; the leave-out corrected slope reported in the text (0.43) is nearly identical to the plug-in slope shown (0.42). A handful of small occupations with extreme premia fall outside the plotting window. Germany, worker–job model, 2017–2022.

4.1.2 Firm–occupation match effects are substantial

We find that in contrast to the three-way fixed effect regression proposed by Torres et al. (2018), the firm–occupation match effect is not negligible in determining the overall wage structure. We showed earlier that the match effect accounts for 15% of the total variance in job fixed effects after leave-out correction (22% by plug-in). In this section, we analyse where in the firm–occupation space these match effects concentrate, both to understand the structure of pay and because the additive worker–firm–occupation model discussed in Section 2 would set them to zero by construction. Figure 6 plots the mean match effect across cells defined by firm and occupation mean-wage quantiles.

The additive approximation is accurate over most of the surface: across the bulk of cells (green) the job premium lies within 0.025 log points of the additive prediction $\gamma_o + \psi_f$. We find that an additive model would severely under-predict wages for high-wage occupations in low-wage firms; in these cells, the additive model under-predicts actual wages by more than 0.10 log points in the most extreme cases. Negative match effects, where the additive model over-predicts, are rarer and scattered.

This localisation matters for the decomposition. Because match effects are confined to the tails rather than spread across the surface, they contribute only modestly to the aggregate match channel in Figures 3 and 4; but they are economically meaningful where they occur, and an additive specification — which forces $\Omega_{of} = 0$ — would miss them entirely.

Figure 6: Match effects Ω_{of} across firm and occupation mean-wage quantiles (Germany, 2017–2022)

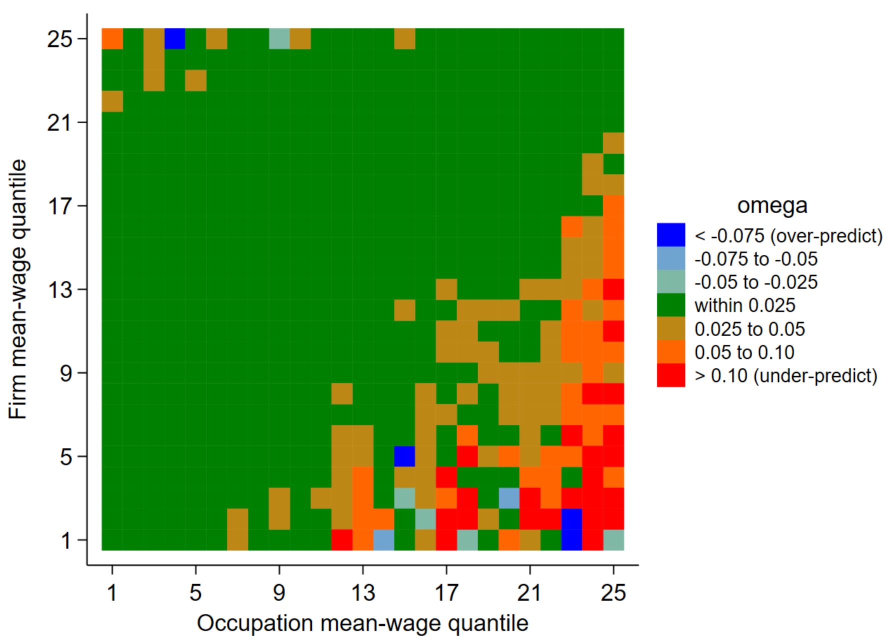

Notes: Mean match effect $\Omega_{of} = \hat{\lambda}_j - (\hat{\gamma}_o + \hat{\psi}_f)$ — the gap between the estimated worker–job premium and the additive sum of its occupation and firm components — in cells defined by 25 quantiles of firm mean wage (vertical) and occupation mean wage (horizontal). Green cells lie within ±0.025 log points of the additive prediction; warm colours mark cells where the additive model under-predicts (positive match effects), cool colours where it over-predicts. Germany, worker–job model, 2017–2022.

4.2 Time trends

The literature has recently been interested in how the role of firms in explaining wage inequality has changed over time (Card et al., 2013; Song et al., 2019; Engbom and Moser, 2022; Lachowska et al., 2023). We extend our analysis to four overlapping periods from 1999.[11] Figure 7 shows how job-level heterogeneity (top row) and worker–job sorting (bottom row) have evolved, each decomposed into between- and within-occupation components. Because the same 2011 reclassification refined the detailed occupation codes, leading the number of 5-digit occupations to roughly quintuple, we report the decomposition at both the 5-digit and the coarser 3-digit coding. The 3-digit series, whose code count barely changes at the break, is more directly comparable over time, while the 5-digit series serves as an upper bound on the between-occupation contribution.[12]

The top row confirms that, in our worker–job model, we replicate the finding of Card et al. (2013) that the dispersion of workplace wage premia rose through the 2000s, with both the between- and the within-occupation components contributing. Extending the exercise to 2022, however, job-level heterogeneity flattens and then falls, so that the final period (2017–2022) returns close to the level of the first (1999–2004); the most recent period also includes the COVID-19 pandemic, which may affect the estimates in hard-to-predict ways. This pattern is visible at both codings.

In terms of the development of sorting over time, our main qualitative conclusion is that the increase in sorting first documented in Card et al. (2013) is primarily driven by an increase in between-occupation covariance. The between-occupation component rises markedly and accounts for the bulk of both the level and the increase at either coding. However, the size of within-occupation covariance over time has also increased, albeit to a smaller extent than between-occupation covariance. At the 5-digit coding this within-occupation component is small, while at the coarser 3-digit coding it is larger and itself rising — but it remains far smaller than the between-occupation component throughout.

[11] In 2010–11 a new occupational classification was introduced in Germany, inducing an abnormally large number of apparent occupation moves; we exclude those years because occupation changes then need not reflect genuine moves.

[12] The over-time exhibits apply estimators that are available uniformly across all period–coding cells: Figure 7 uses the conditional-moment leave-out estimator described in Section 2.2, and Figure 8 reports plug-in channel estimates. The projection-based corrected channel estimator is computationally demanding and could not be completed for every period–coding cell. Where both are available, the plug-in and corrected estimates of the between-occupation components — which carry the trends of interest — are nearly identical (shown for the main specification in Appendix D), and because the limited-mobility correction is of similar magnitude across periods, changes over time difference out most of the remaining bias, so the trends remain interpretable.

Figure 7: Changes in variance of job fixed effects and sorting, 1999–2022 (Germany)

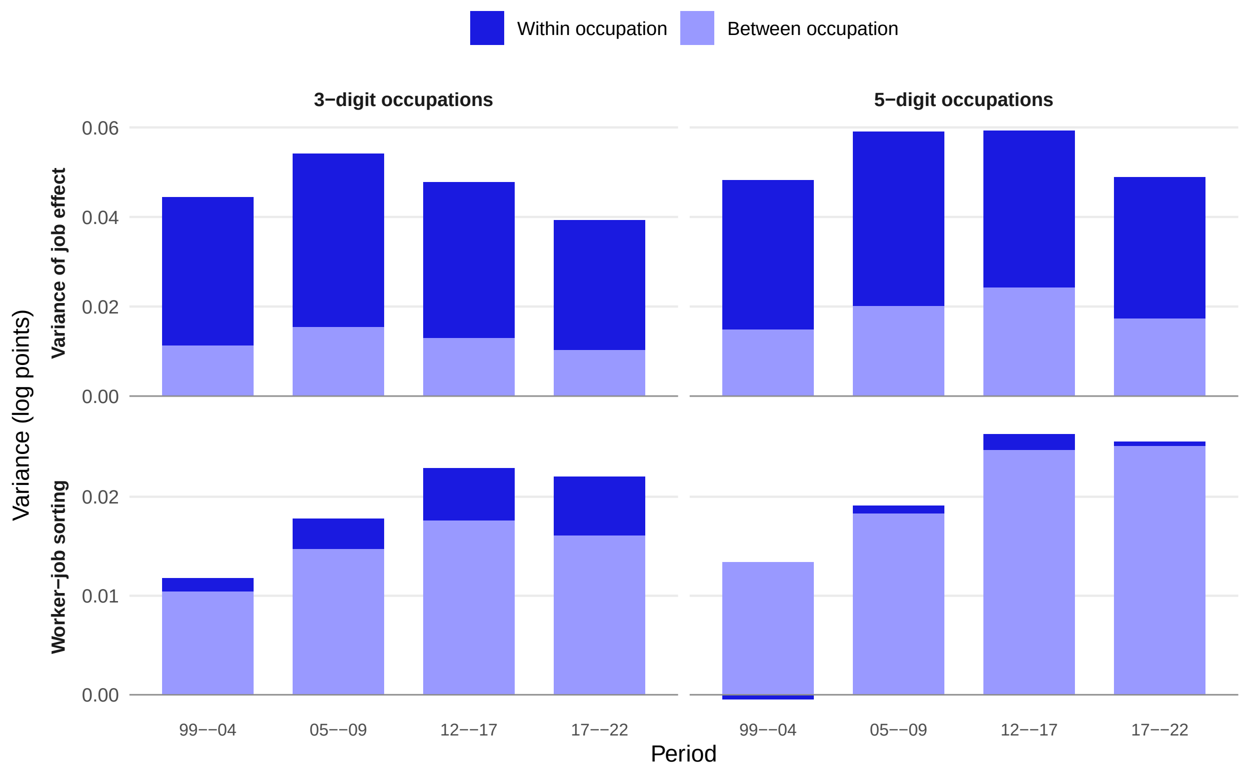

Notes: Leave-out (KSS) corrected decomposition of the variance of the job effect $\mathbb{V}(\lambda_j)$ (top row) and of worker–job sorting $\mathrm{Cov}(\alpha_i, \lambda_j)$ (bottom row) into between-occupation (light) and within-occupation (dark) components, in variance levels, for four German periods. Columns report the 3-digit and the 5-digit occupation codings. The years 2010–11 are excluded because of the change in occupation classification; at that break the number of detailed 5-digit codes rose sharply (247 to 1,233), which mechanically inflates the between-occupation component, whereas the 3-digit count changed little (115 to 144).

Figure 8 further decomposes the between-occupation variance of job fixed effects and covariance of worker and job fixed effects according to equation 5, based on the linear projection of job fixed effects into firm and occupation fixed effects. The growth in between-occupation worker–job sorting (bottom row) is concentrated in the firm-premium channel, cohering with papers which document rising firm–occupation fissuring over time (Bilal and Lhuillier, 2021; Handwerker, 2023). The variance of job premia (top row) is dominated throughout by the firm-premium component and traces the same hump-shaped path as the aggregate in Figure 7. These are plug-in estimates, and as before the 5-digit series overstates the firm-premium and between-occupation component around the 2011 reclassification.

Figure 8: Channel decomposition of job-premium variance and worker–job sorting over time (Germany, 1999–2022)

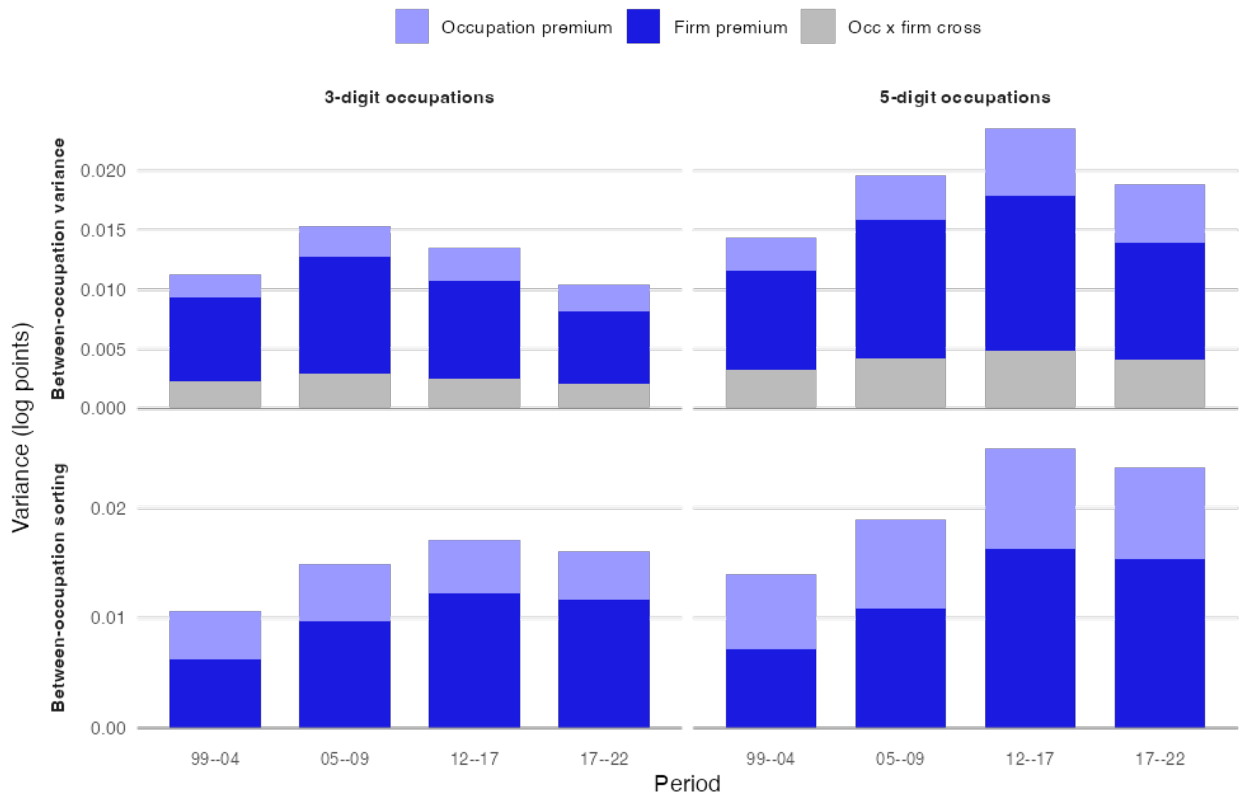

Notes: The figure decomposes the *between-occupation* components of Figure 7 into the channels of the linear projection $\lambda_j = \gamma_o + \psi_f + \Omega_{of}$ (Section 4.1). The top row is the between-occupation variance of the job premium, $\mathbb{V}(\mathbb{E}[\lambda_j \mid o]) = \mathbb{V}(\gamma_o) + \mathbb{V}(\mathbb{E}[\psi_f \mid o]) + 2\,\mathrm{Cov}(\gamma_o, \mathbb{E}[\psi_f \mid o])$, split into an occupation-premium component (γ_o), a firm-premium component ($\mathbb{E}[\psi_f \mid o]$), and the occupation $\times$ firm cross term. The bottom row is between-occupation worker–job sorting, $\mathrm{Cov}(\mathbb{E}[\alpha_i \mid o], \mathbb{E}[\lambda_j \mid o]) = \mathrm{Cov}(\mathbb{E}[\alpha_i \mid o], \gamma_o) + \mathrm{Cov}(\mathbb{E}[\alpha_i \mid o], \mathbb{E}[\psi_f \mid o])$; being a covariance it is linear in the job premium and so carries no cross term. The between-occupation match channel (Ω_{of}) is identically zero by construction ($\mathbb{E}[\Omega_{of} \mid o] = 0$) and is omitted; the negative match contribution to sorting shown in Figure 3 is entirely within-occupation. Columns report the decomposition at the 3-digit and 5-digit occupation codings. Estimates are plug-in (uncorrected); the 5-digit series overstates the firm-premium and cross components around the 2010–11 reclassification (Section 4.2). Germany, 1999–2022.

4.3 Diagnostics and robustness checks

4.3.1 Exogeneity of moves

To interpret the job fixed effects as pay premia, we require that, conditional on the estimated effects, workers' moves between jobs are exogenous. We assess this with the event-study logic of Card et al. (2013). If the additive model is correct, a worker who moves from a high-premium to a low-premium job should experience a step change in pay that is roughly the mirror image of the change for the reverse move; asymmetric or drifting wage paths would instead signal match effects or selection on transitory shocks that the fixed effects do not capture. We implement the test by sorting jobs into quartiles of the leave-out mean job fixed effect and plotting average wages in the two years on either side of a move between quartiles.

Figure 9 reports the test for Germany. The diagonal panels — moves between jobs in the same quartile — are flat, and the off-diagonal panels show clean step changes at the move that are symmetric across the diagonal: the wage gain for the most extreme upward move (first to fourth quartile) closely mirrors the loss for the reverse move, with stable wages in the years before and after. This is the pattern the additive model predicts and matches the worker–firm literature (Card et al., 2013, 2018; Bonhomme et al., 2023).

Figure 9: Event study around job moves, clustering by leave-out mean job fixed effect (Germany, 2017–2022)

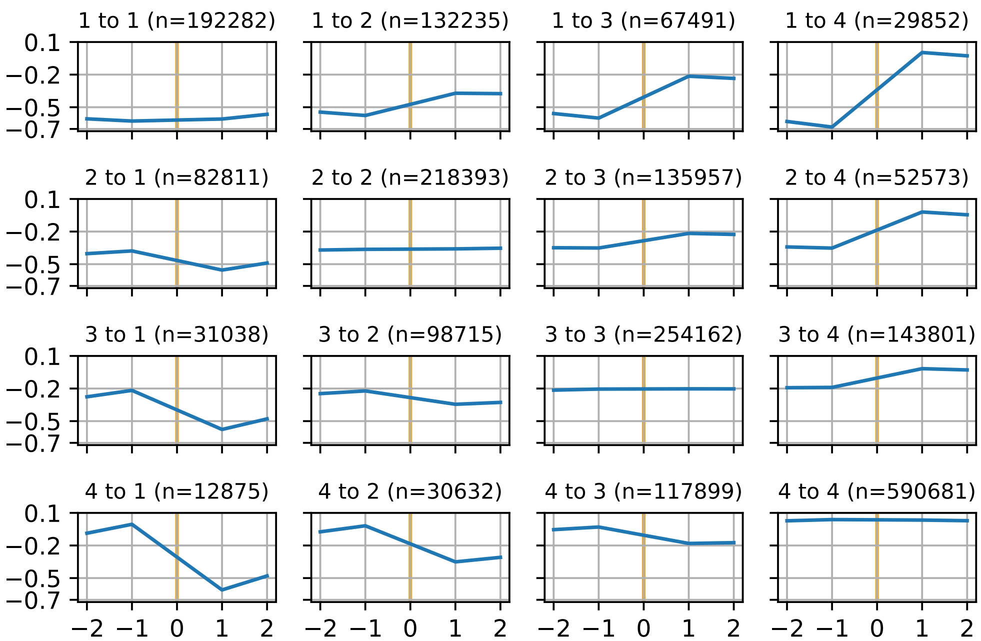

Notes: Each panel plots average residualised log wages in event time around a move between jobs, for moves from one quartile of the leave-out job fixed effect to another (origin quartile in rows, destination in columns; cell counts in titles). Jobs are sorted by the leave-out mean job fixed effect. Workers must hold the origin job for two years before and the destination job for two years after the move. Germany, worker–job model, 2017–2022.

A natural concern is that pooling all moves could mask differences in exogeneity across move types. The worker–firm literature treats moves *across firms* as conditionally exogenous, but our job effect is also identified from moves *across occupations within a firm*, and these may be more suspect: a within-firm occupation change is often a promotion, which firms may grant on the basis of private information about a worker's productivity, generating exactly the kind of match effect or selection that would invalidate the test. Figure 10 splits the event study into within-occupation across-firm moves (panel a, the classic AKM firm move) and across-occupation within-firm moves (panel b, the occupation change). Both reproduce the symmetric step pattern of the pooled figure. The firm moves behave as the AKM literature finds; more notably, the within-firm occupation moves — though rarer, so the extreme cells are thinner — show the same symmetry, with no systematic pre-move drift or asymmetry between gains and losses. We find no evidence that within-firm occupation changes are differentially endogenous, which supports identifying the occupation component of the job premium from these moves. The same patterns hold when jobs are clustered by their leave-out mean wage rather than the estimated fixed effect, and when the test is run on the independently coded French data; we report both sets of figures in Appendix E.

Figure 10: Event study by move type (Germany, 2017–2022)

(a) Within-occupation, across-firm (b) Across-occupation, within-firm

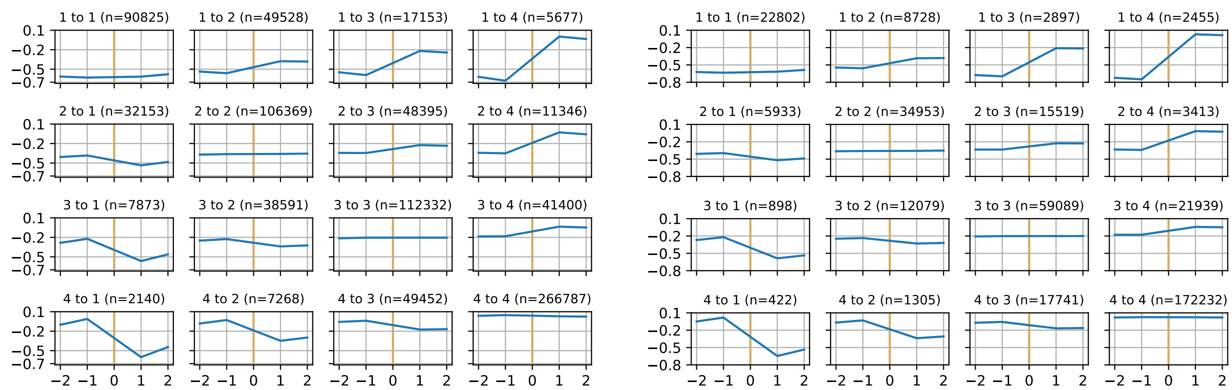

Notes: As Figure 9, restricting to moves that change the firm but not the occupation (panel a) and moves that change the occupation but not the firm (panel b). Germany, worker–job model, 2017–2022.

4.3.2 Robustness to coarser occupation categories

A natural worry is that our results are driven by fine, possibly mismeasured, occupation codes: if firms fill in detailed codes inconsistently, apparent across-occupation structure could be an artefact of reporting. We address this by re-estimating the decomposition with progressively coarser occupation codes, from the finest classification down to a single digit, given that coarser occupation codes are less likely to be misclassified. Coarsening the code can only move variation from the between- to the within-occupation component, so the between-occupation components shrink as we coarsen by construction. The test is therefore not whether the between-occupation share is invariant (it cannot be), but whether the two substantive conclusions survive aggregation.

Appendix Figures 15 and 16 report the full variance and covariance decompositions from 1- through 5-digit occupation codes (1 through 4 for France) in both countries. We find that our main conclusions are robust to using coarser occupational codes. First, the within-occupation dispersion of job premia is substantial at every granularity: the within-occupation variance of the job effect is large and essentially flat across codings (around 0.028–0.032 in Germany and 0.011–0.014 in France), so the finding that firms pay materially different premia for the same broad occupation is not an artefact of fine codes. Second, a large share of worker–job sorting remains between-occupation even at the coarsest codings. France is the more demanding test here precisely because its occupations are independently verified: even at the 1-digit level — where occupations collapse to a handful of broad groups — 67% of the wage inequality attributable to sorting between workers and jobs is still due to sorting across occupations, and 19% of the variance attributable to job heterogeneity remains between-occupation. That so much between-occupation sorting survives aggregation all the way to one digit makes occupation mismeasurement an implausible account of our findings. The same robustness holds for the channel cross-decomposition of Section 4.1: Appendix Figure 17 re-computes the occupation-premium, firm-premium, and match channels at 3- versus 5-digit codes and shows that the firm-premium channel — both its variance and its covariance with worker ability — is essentially unchanged ($\mathrm{Cov}(\alpha, \psi_f)$ of 0.021 at the 5-digit versus 0.020 at the 3-digit coding, leave-out corrected), while only the occupation-premium channel scales with granularity and the match channel stays small. The firm-dispersion and firm-sorting results are thus not artefacts of how finely occupations are coded.

5 Discussion

The central finding of the traditional AKM literature is often summarised by the statement that high-wage workers work at high-wage firms (Song et al., 2019; Kline, 2024). Our paper elaborates upon this conclusion in three ways. First, high-wage workers do work at high-wage firms, principally by holding high-wage occupations. In Germany, 19.8 of the 21.4 percentage points of residual log-wage variance attributed to worker–job sorting lie between occupations. Although we find that within occupations high-wage workers only weakly match to high-wage firms, this is not because of a lack of pay premia dispersion within occupations. Second, two-thirds of this between-occupation sorting reflects differences in average firm premia across occupations. These two facts are reconciled by the assignment of occupations to firms.

Our results qualify rather than overturn the conventional AKM decomposition. The standard worker–firm covariance remains a valid description of the association between permanent worker wage components and firm wage components in an occupation-blind model; rather, our work points out that the AKM sorting term corresponds not only to the kind of worker–firm sorting studied widely in the literature, but also to the occupational allocation of workers and the non-random allocation of those occupations across firms. Our findings have numerous implications for the wider literature on wage determination. In this discussion, we focus on two central conclusions. First, worker–firm sorting has been the central axis of theoretical and structural work on sorting in recent years; our work shows the importance of understanding whether the complementarities identified therein are driven by worker–occupation complementarities like those identified within the occupational sorting literature. Second, contrary to the findings of large worker–firm sorting in the worker–firm AKM model, we find substantial wage dispersion within occupations but little evidence of wage sorting within occupations.

Sorting along worker, firm, and occupation dimensions. In part motivated by the proliferation of AKM estimates of worker–firm sorting, a large literature has developed measuring the extent of worker–firm complementarity and sorting in the economy (Eeckhout and Kircher, 2011; Lise et al., 2016; Hagedorn et al., 2017; Bagger and Lentz, 2019; Lentz et al., 2023; Bonhomme et al., 2019; Borovičková and Shimer, 2024). Our findings complicate this picture by emphasising the role of worker–occupation sorting as a more quantitatively important driver of wage inequality. Models of worker–occupation sorting are of course not new and predate AKM: the Roy model generated a large assignment literature (Roy, 1951; Sattinger, 1993), in which workers select into the task where their comparative advantage lies and pay differences across occupations reflect the market value of the work performed. The task-based framework (Autor et al., 2003; Acemoglu and Autor, 2011; Goos and Manning, 2007) supplies an account of why those relative values shift over time, and dynamic models of occupational choice and occupation-specific human capital (Keane and Wolpin, 1997; Yamaguchi, 2012) an account of why an assignment persists once made.

A first-order open question posed by our results is: are previously measured worker–firm complementarities driven implicitly by worker–occupation complementarity? A large literature shows that AKM covariances do not generally identify complementarities in underlying primitives, such as worker ability and firm productivity (Eeckhout and Kircher, 2011; Hagedorn et al., 2017; Lopes de Melo, 2018; Borovičková and Shimer, 2024). The same caveat applies to each part of our decomposition. For example, the larger between-occupation covariance does not establish that worker–occupation complementarities are stronger than worker–firm complementarities. This calls for models of sorting which jointly feature both firm and occupational components, and our decomposition supplies reduced-form moments for that exercise: such a model should reproduce large firm-premium dispersion within occupational markets, a weak covariance between that dispersion and the portable worker component, and worker–occupation covariance channelled predominantly through the firm premia of employing firms. On the firm side, recent theoretical and structural work relaxes the assumption that workers and jobs can be ordered by scalar types. Workers instead possess bundles of skills and jobs require different bundles of tasks, so that sorting is dimension-specific and governed by comparative advantage and worker–job complementarities (Lindenlaub, 2017; Lise and Postel-Vinay, 2020; Lindenlaub and Postel-Vinay, 2023). Indeed, Lise and Postel-Vinay (2020) measure job requirements using occupation-level task information. Viewed through this framework, a job is naturally indexed by both its employer and its occupation. Firm-based sorting statistics that suppress the occupational dimension may therefore combine the matching of worker skills to occupational task bundles with residual sorting across firms within occupations.

A further implication of our analysis is that an important part of measured wage inequality is generated by the alignment of the occupation-to-firm assignment with the ability composition of occupations and, secondarily, with occupation premia. This coheres with two literatures that ask which occupations belong inside the firm. The first, on outsourcing and fissuring (Abraham and Taylor, 1996; Goldschmidt and Schmieder, 2017; Bilal and Lhuillier, 2021; Handwerker, 2023; Cortes et al., 2024), studies which occupations firms shed to outside contractors, and finds the motive is frequently to avoid paying firm-level premia to occupations peripheral to the firm's core activity. The second, on knowledge-based hierarchies (Garicano, 2000; Garicano and Rossi-Hansberg, 2006; Caliendo et al., 2015), studies how firms allocate knowledge across layers of workers, so that a firm's occupational composition is an outcome of its size, productivity and organisational choices rather than a technological given: larger and more productive firms support more layers and are correspondingly top-heavy in high-knowledge occupations. The two accounts generate the alignment we measure from opposite directions: the first by subtraction, as high-premium firms divest themselves of their low-premium occupations, and the second by addition, as productive firms build the upper layers that low-premium occupations do not occupy.

One implication is that previous measurements of the increasing importance of firm sorting (Song et al., 2019) may reflect increasing occupational segregation over time, rather than improved matching efficiency. The evolution of the decomposition reinforces this organisational interpretation. The rise in worker–job sorting since the early 2000s occurs primarily between occupations, and the dispersion of job premia is roughly flat over the same period. This pattern is consistent with a realignment of the occupation-to-firm assignment rather than a widening premia or improved matching. The result is consistent with a growing separation of occupations across firms, including the forms of fissuring documented in the outsourcing literature.

Wage dispersion without sorting within occupations. A second central finding is the contrast between dispersion and sorting within occupations. Substantial job-premium dispersion remains among workers recorded in the same detailed occupation: it accounts for 12.6% of residual log-wage variance in Germany and 5.5% in France, and constitutes the larger component of job-premium variance in both countries. Moreover, the absolute within-occupation variance is relatively stable as the occupation classification is coarsened. AKM firm effects have been interpreted as capturing rent-sharing, an efficiency-wage premium, or strategic wage posting (Card et al., 2013); as the outcome of bargaining (Card et al., 2016), of rent-sharing over quasi-rents (Kline et al., 2019), and of imperfect competition in the labour market (Lamadon et al., 2022); and as the pay component of a broader compensation package including amenities (Sorkin, 2018; Card et al., 2018; Sockin, 2022; Humlum et al., 2025). Our within-occupation estimates reinforce the core finding of this literature: firms differ persistently and substantially in the premia they pay, and this dispersion is not an artefact of occupational composition.

At the same time, the German estimates show little tendency for workers with higher permanent wage components to hold the higher-premium jobs within their occupation, while the French estimates show that within-occupation sorting accounts for only 20% of total worker–job sorting. This stands in contrast to the worker–firm sorting emphasised in much of the structural literature. One possibility is that the weak covariance partly reflects measurement. German wages are censored at the social-security ceiling and are imputed above it following Card et al. (2013). Residual attenuation could be especially consequential in highly paid occupations, where within-occupation sorting may be strongest. The French data do not have the same censoring problem and nevertheless assign most sorting to the between-occupation component, although the remaining 4.0 percentage points of within-occupation covariance are economically non-negligible. Limited-mobility bias is another natural concern: early findings of weak worker–firm sorting were substantially affected by noise in estimated fixed effects (Andrews et al., 2008, 2012). Our leave-out correction changes the within-occupation covariance substantially relative to its plug-in value but still produces only a small positive estimate in Germany. Measurement concerns therefore qualify the magnitude rather than the finding itself: the French data are uncensored and independently coded and deliver the same ranking, and the leave-out correction moves the German estimate from negative (-6.7% of variance) to a small positive value ($+1.7\%$) rather than uncovering substantial sorting that noise had masked.

The null also admits economic interpretations. Search frictions, limited information about employer wage policies, geographic mobility costs, or occupation-specific labour markets may prevent workers from reaching the highest-paying employer within their occupation. Collective bargaining, rent sharing, or other institutional wage-setting arrangements may generate broad firm premia without those premia being finely allocated according to workers' permanent wage components. Alternatively, the scalar worker effect may omit occupation-specific or multidimensional skills on which firms and workers do sort. In that case, economically meaningful matching could occur within occupations without producing a large covariance between the job premium and the particular worker component estimated here.

Sorting over non-wage job attributes provides another possible explanation. AKM firm effects have been interpreted many ways (Card et al., 2013, 2016; Kline, 2024; Lamadon et al., 2022), but they may also form part of a broader compensation package containing non-wage amenities (Sorkin, 2018; Card et al., 2018; Sockin, 2022; Humlum et al., 2025). If workers select firms on total utility rather than wages alone, wage-premium sorting need not reveal utility sorting. Recent evidence however has tended to find that valuable amenities and wages often covary positively across firms, so accounting for amenities can widen rather than narrow dispersion (Sockin, 2022; Maestas et al., 2023). These aggregate relationships do not rule out type-specific compensating differentials within occupations, but they imply that amenities cannot automatically be invoked to convert weak wage sorting into strong compensating utility sorting.

Finally, it is also possible that there is less worker–firm complementarity than previously believed, and that part of what existing structural estimates attribute to worker–firm complementarity reflects worker–occupation complementarity operating through the joint distribution of occupations and firms. As discussed in the previous section, two-way fixed effect regressions are not well-placed to uncover such underlying complementarity, and adjudicating it requires a structural approach that combines sorting across occupations and across firms; the moments documented here provide reduced-form targets for that exercise.

6 Conclusion

The two-way fixed effects literature summarises wage sorting with the statement that high-wage workers work at high-wage firms. In this paper we show that at least in the German and French settings, this happens primarily because workers work in high-wage occupations at these firms. Estimating a worker–job fixed-effects model where a job is a firm–occupation pair on the universe of German and French matched employer–employee data, we find that over 80% of worker–job sorting lies between occupations. Within occupations, the allocation of workers to firms is only weakly related to ability, even though within-occupation dispersion in firm premia is the largest single component of job-premium variance. High-wage workers do end up at high-wage firms, but they get there by holding high-wage occupations, not by sorting into positions at high-premium employers within their occupation.

The returns to the occupational margin are nonetheless mostly made of firm premia: two-thirds of between-occupation sorting arises because the occupations that attract high-ability workers are concentrated in firms that pay high premia to their entire payroll. These two findings are reconciled by the assignment of occupations to firms. Occupations are severely segregated across firms (knowing a worker's employer resolves two-thirds of the uncertainty about their detailed occupation). Choosing an occupation is therefore, de facto, choosing a distribution of employers, and the firm premia workers collect accrue through occupational membership rather than through competition within occupational labour markets. Over time, the rise in measured sorting since the early 2000s is concentrated in this firm-premium channel while job-premium dispersion itself is flat. This pattern is consistent with a labour market reorganising which occupations firms house (fissuring of the workplace (Weil, 2014; Goldschmidt and Schmieder, 2017)), rather than one where matching workers to firms becomes more efficient.

Methodologically, our single innovation on the standard AKM model is to replace the firm fixed effect with a job fixed effect, so that the model nests the worker–firm specification while letting premia vary freely at the firm–occupation level. An auxiliary projection then splits the job premium into occupation, firm, and match components without imposing additive separability. This flexibility matters because match effects account for 15% of job-premium variance even after a leave-out correction. Event studies around job moves (including within-firm occupation changes) display the symmetric step changes the model implies. We account for limited-mobility bias using the leave-out estimator of Kline et al. (2020), extended to the conditional moments and projection channels our decompositions require, and show that the core results survive coarsening the occupation classification to a single digit and replicate in the independently coded, uncensored French data.

These results have implications for both theory and policy. Models of worker–firm assignment must rationalise large firm-premium dispersion within occupations alongside a near-zero ability gradient across it, and should treat the firm's occupational boundary as an endogenous object central to the wage structure rather than a primitive. For policy, the ability gradient in pay runs through occupational attachment, so the levers on that margin are the institutions governing access to occupations: training, licensing, and educational tracking. The inequality remaining within occupations is premium dispersion allocated largely orthogonally to ability; if these premia are rents, the instruments that reach it are wage-setting institutions (sectoral bargaining, pay transparency) rather than policies helping workers reach better firms. And if firms escape payroll-wide rent-sharing by pushing low-premium occupations outside their boundaries, rent-sharing becomes increasingly regressive and measured firm-level inequality can rise without any change in any firm's wage policy. This implies that policies affecting boundary status, such as equal treatment for outsourced workers and joint-employer liability, are the natural response.

References

Abowd, J. M., R. H. Creecy, and F. Kramarz (2002, March). Computing Person and Firm Effects Using Linked Longitudinal Employer-Employee Data. Longitudinal Employer-Household Dynamics Technical Papers 2002-06, Center for Economic Studies, U.S. Census Bureau.

Abowd, J. M., F. Kramarz, and D. N. Margolis (1999). High wage workers and high wage firms. *Econometrica 67*(2), 251–333.

Abraham, K. G. and S. K. Taylor (1996). Firms' use of outside contractors: Theory and evidence. *Journal of Labor Economics*, 394–424.

Acemoglu, D. and D. Autor (2011). Skills, tasks and technologies: Implications for employment and earnings. In *Handbook of labor economics*, Volume 4, pp. 1043–1171. Elsevier.

Andrews, M. J., L. Gill, T. Schank, and R. Upward (2008, 03). High Wage Workers and Low Wage Firms: Negative Assortative Matching or Limited Mobility Bias? *Journal of the Royal Statistical Society Series A: Statistics in Society 171*(3), 673–697.

Andrews, M. J., L. Gill, T. Schank, and R. Upward (2012). High wage workers match with high wage firms: Clear evidence of the effects of limited mobility bias. *Economics Letters 117*(3), 824–827.

Autor, D. H., F. Levy, and R. J. Murnane (2003). The skill content of recent technological change: An empirical exploration. *The Quarterly journal of economics 118*(4), 1279–1333.

Babet, D., O. Godechot, and M. G. Palladino (2025). In the land of akm: Explaining the dynamics of wage inequality in france.

Bagger, J. and R. Lentz (2019). An empirical model of wage dispersion with sorting. *The Review of Economic Studies 86*(1), 153–190.

Barth, E., A. Bryson, J. C. Davis, and R. Freeman (2016). It's where you work: Increases in the dispersion of earnings across establishments and individuals in the united states. *Journal of Labor Economics 34*(S2), S67–S97.

Bassier, I. (2023). Firms and inequality when unemployment is high. *Journal of Development Economics 161*, 103029.

Bilal, A. and H. Lhuillier (2021). Outsourcing, inequality and aggregate output. Technical report, National Bureau of Economic Research.

Bonhomme, S., K. Holzheu, T. Lamadon, E. Manresa, M. Mogstad, and B. Setzler (2023). How much should we trust estimates of firm effects and worker sorting? *Journal of Labor Economics 41*(2), 291–322.

Bonhomme, S., T. Lamadon, and E. Manresa (2019). A distributional framework for matched employer employee data. *Econometrica 87*(3), 699–739.

Borovičková, K. and R. Shimer (2024). Assortative matching and wages: The role of selection. Technical report, National Bureau of Economic Research.

Caliendo, L., F. Monte, and E. Rossi-Hansberg (2015). The anatomy of french production hierarchies. *Journal of Political Economy 123*(4), 809–852.

Card, D., A. R. Cardoso, J. Heining, and P. Kline (2018). Firms and labor market inequality: Evidence and some theory. *Journal of Labor Economics 36*(S1), S13–S70.

Card, D., A. R. Cardoso, and P. Kline (2016). Bargaining, sorting, and the gender wage gap: Quantifying the impact of firms on the relative pay of women. *The Quarterly Journal of Economics 131*(2), 633–686.

Card, D., J. Heining, and P. Kline (2013). Workplace heterogeneity and the rise of west german wage inequality. *The Quarterly journal of economics 128*(3), 967–1015.

Cortes, G. M., D. Dabed, A. Oliveira, and A. Salomons (2024). Fissured firms and worker outcomes. CLEF, Canadian Labour Economics Forum.

Eeckhout, J. and P. Kircher (2011). Identifying sorting—in theory. *The Review of Economic Studies 78*(3), 872–906.

Engbom, N. and C. Moser (2022). Earnings inequality and the minimum wage: Evidence from brazil. *American Economic Review 112*(12), 3803–3847.

Garicano, L. (2000). Hierarchies and the organization of knowledge in production. *Journal of political economy 108*(5), 874–904.

Garicano, L. and E. Rossi-Hansberg (2006). Organization and inequality in a knowledge economy. *The Quarterly journal of economics 121*(4), 1383–1435.

Goldschmidt, D. and J. F. Schmieder (2017). The rise of domestic outsourcing and the evolution of the german wage structure. *The Quarterly Journal of Economics 132*(3), 1165–1217.

Goos, M. and A. Manning (2007). Lousy and lovely jobs: The rising polarization of work in britain. *The review of economics and statistics 89*(1), 118–133.

Hagedorn, M., T. H. Law, and I. Manovskii (2017). Identifying equilibrium models of labor market sorting. *Econometrica 85*(1), 29–65.

Håkanson, C., E. Lindqvist, and J. Vlachos (2021). Firms and skills: the evolution of worker sorting. *Journal of Human Resources 56*(2), 512–538.

Haltiwanger, J., H. R. Hyatt, and J. R. Spletzer (2024). Rising top, falling bottom: Industries and rising wage inequality. *American Economic Review 114*(10), 3250–3283.

Handwerker, E. W. (2023, None). Outsourcing, occupationally homogeneous employers, and wage inequality in the united states. *Journal of Labor Economics 41*(S1), 173–203.

Heckman, J. J. and G. Sedlacek (1985). Heterogeneity, aggregation, and market wage functions: an empirical model of self-selection in the labor market. *Journal of political Economy 93*(6), 1077–1125.

Humlum, A., M. Rasmussen, and E. K. Rose (2025). Firm premia and match effects in pay vs. amenities. Technical report, National Bureau of Economic Research.

Insee (2024). Base Tous Salariés : fichiers Salariés - 2022.

Keane, M. P. and K. I. Wolpin (1997). The career decisions of young men. *Journal of political Economy 105*(3), 473–522.

Kline, P., N. Petkova, H. Williams, and O. Zidar (2019). Who profits from patents? rent-sharing at innovative firms. *The Quarterly Journal of Economics 134*(3), 1343–1404.

Kline, P., R. Saggio, and M. Sølvsten (2020). Leave-out estimation of variance components. *Econometrica 88*(5), 1859–1898.

Kline, P. M. (2024). Firm wage effects: A review of literature and methods. *Working Paper*.

Lachowska, M., A. Mas, R. Saggio, and S. A. Woodbury (2023). Do firm effects drift? evidence from washington administrative data. *Journal of Econometrics 233*(2), 375–395.

Lamadon, T., J. Lise, C. Meghir, and J.-M. Robin (2024). Labor market matching, wages, and amenities. Technical report, National Bureau of Economic Research.

Lamadon, T., M. Mogstad, and B. Setzler (2022). Imperfect competition, compensating differentials, and rent sharing in the us labor market. *American Economic Review 112*(1), 169–212.

Lee, D. (2005). An estimable dynamic general equilibrium model of work, schooling, and occupational choice. *International Economic Review 46*(1), 1–34.

Lee, D. and K. I. Wolpin (2006). Intersectoral labor mobility and the growth of the service sector. *Econometrica 74*(1), 1–46.

Lentz, R., S. Piyapromdee, and J.-M. Robin (2023). The anatomy of sorting—evidence from danish data. *Econometrica 91*(6), 2409–2455.

Lindenlaub, I. (2017). Sorting multidimensional types: Theory and application. *The Review of Economic Studies 84*(2), 718–789.

Lindenlaub, I. and F. Postel-Vinay (2023). Multidimensional sorting under random search. *Journal of political Economy 131*(12), 3497–3539.

Lise, J., C. Meghir, and J.-M. Robin (2016). Matching, sorting and wages. *Review of Economic Dynamics 19*, 63–87.

Lise, J. and F. Postel-Vinay (2020). Multidimensional skills, sorting, and human capital accumulation. *American Economic Review 110*(8), 2328–76.

Lochner, B. and B. Schulz (2024). Firm productivity, wages, and sorting. *Journal of Labor Economics 42*(1), 85–119.

Lopes de Melo, R. (2018). Firm wage differentials and labor market sorting: Reconciling theory and evidence. *Journal of Political Economy 126*(1), 313–346.

Maestas, N., K. J. Mullen, D. Powell, T. von Wachter, and J. B. Wenger (2023, July). The value of working conditions in the united states and the implications for the structure of wages. *American Economic Review 113*(7), 2007–47.

Mortensen, D. (2003). *Wage dispersion: why are similar workers paid differently?* MIT press.

Palladino, M. G., A. Roulet, and M. Stabile (2025). Narrowing industry wage premiums and the decline in the gender wage gap. *Labour Economics 94*, 102693.

Roy, A. D. (1951). Some thoughts on the distribution of earnings. *Oxford economic papers 3*(2), 135–146.

Sattinger, M. (1993). Assignment models of the distribution of earnings. *Journal of economic literature 31*(2), 831–880.

Sockin, J. (2022). Show me the amenity: Are higher-paying firms better all around?

Song, J., D. J. Price, F. Guvenen, N. Bloom, and T. Von Wachter (2019). Firming up inequality. *The Quarterly journal of economics 134*(1), 1–50.

Sorkin, I. (2018). Ranking firms using revealed preference. *The Quarterly Journal of Economics 133*(3), 1331–1393.

Stüber, H., W. Dauth, and J. Eppelsheimer (2023). A guide to preparing the sample of integrated labour market biographies (siab, version 7519 v1) for scientific analysis. *Journal for Labour Market Research 57*(1), 7.

Theil, H. (1972). Statistical decomposition analysis: With applications in the social and administrative sciences. *(No Title)*.

Theil, H. and A. J. Finizza (1971). A note on the measurement of racial integration of schools by means of informational concepts.

Torres, S., P. Portugal, J. T. Addison, and P. Guimaraes (2018). The sources of wage variation and the direction of assortative matching: Evidence from a three-way high-dimensional fixed effects regression model. *Labour Economics 54*, 47–60.

Weil, D. (2014). *The Fissured Workplace: Why Work Became So Bad for So Many and What Can Be Done to Improve It.* Cambridge, MA: Harvard University Press.

Yamaguchi, S. (2012). Tasks and heterogeneous human capital. *Journal of Labor Economics 30*(1), 1–53.

A Summary statistics for other periods

A.1 Germany

Tables 3–5 report the same summary statistics as the main-text Table 1 for the three earlier German periods. Sample sizes and wage moments are comparable across periods, and the leave-one-out selection has the same character throughout — a shift towards larger, less occupationally diverse firms.[13]

Table 3: Summary statistics, Germany (1999–2004), across estimation samples

	Full	Connected	Jobs-LOO
Observations (m)	60.34	49.91	43.49
Workers (m)	13.62	10.12	8.97
Firms (m)	1.56	1.00	0.56
Jobs (m)	4.03	2.31	1.10
Mean log daily wage	4.77	4.80	4.81
Var. resid. log wage	0.227	0.210	0.204
Moves (% of obs)	10.8	12.2	11.1
across firms (% of moves)	88.2	88.6	88.6
across occ. (% of moves)	49.2	49.5	45.7
both (% of moves)	37.3	38.0	34.3
Selection characteristics			
Median worker's firm size	64	78	94
Empl. share, firms <50 (%)	45.5	42.1	38.7
Mean occ. per firm (empl.-wtd.)	18.38	14.48	12.37
Empl. share, 1-occ. firms (%)	12.1	14.8	18.5

Notes: As in Table 1, for 1999–2004.

[13] The increase in occupations per firm between the 2004–2009 and 2012–2017 panels reflects the 2011 change in the KldB occupation classification, which roughly quintupled the number of detailed occupations, rather than a change in firm structure; we exclude 2010–11 from the time-series analysis for the same reason (Section 4.2).

Table 4: Summary statistics, Germany (2004–2009), across estimation samples

	Full	Connected	Jobs-LOO
Observations (m)	57.78	45.63	38.93
Workers (m)	12.81	9.13	7.93
Firms (m)	1.49	0.88	0.47
Jobs (m)	3.78	1.95	0.89
Mean log daily wage	4.74	4.78	4.80
Var. resid. log wage	0.258	0.240	0.236
Moves (% of obs)	9.2	10.8	9.7
across firms (% of moves)	89.1	89.6	89.2
across occ. (% of moves)	48.3	48.7	44.4
both (% of moves)	37.4	38.2	33.7
Selection characteristics			
Median worker's firm size	62	78	96
Empl. share, firms <50 (%)	45.9	41.9	38.1
Mean occ. per firm (empl.-wtd.)	17.26	12.64	10.65
Empl. share, 1-occ. firms (%)	12.2	15.9	20.1

Notes: As in Table 1, for 2004–2009.

Table 5: Summary statistics, Germany (2012–2017), across estimation samples

	Full	Connected	Jobs-LOO
Observations (m)	62.47	47.97	40.11
Workers (m)	14.03	9.70	8.32
Firms (m)	1.40	0.86	0.45
Jobs (m)	4.79	2.45	1.03
Mean log daily wage	4.75	4.77	4.79
Var. resid. log wage	0.261	0.252	0.252
Moves (% of obs)	9.9	11.8	10.4
across firms (% of moves)	87.2	87.2	85.8
across occ. (% of moves)	62.4	63.2	59.7
both (% of moves)	49.6	50.4	45.5
Selection characteristics			
Median worker's firm size	65	80	97
Empl. share, firms <50 (%)	45.1	41.3	37.6
Mean occ. per firm (empl.-wtd.)	24.98	19.36	16.34
Empl. share, 1-occ. firms (%)	9.0	12.3	16.4

Notes: As in Table 1, for 2012–2017.

A.2 France

Table 6 reports summary statistics for the French samples in the two periods we study (Panel A, 2015–2019; Panel B, 2010–2014), across the five estimation samples used to form the worker–firm and worker–job connected and leave-one-out sets. Because the French data are accessed under the constraints noted in Section 3.2, the table reproduces the available exported statistics — counts and wage moments — and so carries fewer rows than its German counterparts.

Table 6: Summary statistics, France, across estimation samples

	Full	Firms conn.	Firms LOO	Jobs conn.	Jobs LOO
Panel A. 2015–2019					
Observations (m)	48.01	40.78	38.75	35.93	31.53
Workers (m)	14.25	10.28	9.78	9.06	8.09
Firms (m)	1.25	0.65	0.38	0.57	0.27
Jobs (m)	4.48	3.17	2.60	2.04	0.82
Mean log annual wage	10.39	10.42	10.43	10.42	10.43
Var. log annual wage	0.25	0.24	0.23	0.23	0.23
Mean log hourly wage	2.96	2.97	2.98	2.97	2.98
Var. log hourly wage	0.21	0.20	0.20	0.20	0.19
Var. resid. log annual wage	0.24	0.22	0.22	0.22	0.21
Panel B. 2010–2014					
Observations (m)	46.10	39.47	37.29	33.15	28.78
Workers (m)	12.66	9.40	8.89	7.93	7.03
Firms (m)	1.15	0.58	0.33	0.50	0.22
Jobs (m)	4.99	3.56	2.92	2.21	0.82
Mean log annual wage	10.34	10.37	10.38	10.36	10.38
Var. log annual wage	0.24	0.23	0.23	0.23	0.23
Mean log hourly wage	2.89	2.91	2.92	2.91	2.92
Var. log hourly wage	0.19	0.19	0.19	0.19	0.18
Var. resid. log annual wage	0.23	0.22	0.22	0.21	0.21

Notes: Counts in millions. Columns are the full sample and the connected and leave-one-out sets of the worker–firm and worker–job models. Panel A covers 2015–2019, Panel B covers 2010–2014.

B Occupation and firm-size distributions across samples

Table 7 gives a more detailed breakdown of the firm-size and occupation distributions of the German sample than the scalar summaries in Table 1, comparing the full sample, the worker–job leave-one-out (LOO) sample used in estimation, and the excluded observations (those in the full sample but not the LOO set). The excluded observations are concentrated in small firms — 31% are in firms with fewer than ten workers, against 13% in the LOO sample — consistent with the leave-one-out restriction dropping poorly connected, peripheral employers. Reassuringly for our analysis, however, the occupation *mix* and, in particular, the *wage structure across occupations* are similar in the included and excluded samples: the ordering and approximate levels of mean wages by occupation group are preserved. The selection thus operates on which firms are observed, not on the occupational wage structure that our decomposition exploits.

Table 7: Balance of included and excluded samples (Germany, 2017–2022)

	Full	Jobs-LOO	Excluded
Firm size (employment shares, %)			
1–9	19.1	12.9	31.4
10–49	25.9	25.0	27.8
50–249	28.2	29.7	25.2
250+	26.8	32.5	15.6
Occupation group (KldB 1-digit, employment %)			
1 Agriculture, forestry, farming	1.7	1.1	3.0
2 Production & manufacturing	35.0	36.9	31.3
3 Construction & building technology	10.1	8.5	13.0
4 Science, geography, informatics	6.7	6.9	6.3
5 Traffic, logistics, safety	16.2	19.1	10.5
6 Commercial services, sales, hospitality	8.0	6.9	10.3
7 Business, accounting, law, admin.	14.6	13.5	16.8
8 Health, social work, education	5.3	4.9	5.9
9 Humanities, media, arts, culture	2.4	2.1	2.9
Mean log wage by occupation group			
1 Agriculture, forestry, farming	4.38	4.41	4.36
2 Production & manufacturing	4.77	4.81	4.69
3 Construction & building technology	4.62	4.63	4.60
4 Science, geography, informatics	5.07	5.12	4.95
5 Traffic, logistics, safety	4.46	4.47	4.46
6 Commercial services, sales, hospitality	4.80	4.87	4.71
7 Business, accounting, law, admin.	5.09	5.16	4.97
8 Health, social work, education	4.81	4.90	4.67
9 Humanities, media, arts, culture	4.96	5.01	4.89

Notes: Excluded := observations in the full sample that are not in the worker–job LOO set, computed as the employment-weighted difference of the full and LOO columns. Occupation groups are 1-digit KldB 2010 (the ∼0% military group is omitted). Firm size is measured within each sample, so the excluded firm-size column is an approximation; occupation shares and occupation wages are worker-level and exact.

C Variance decomposition tables

Tables 8 and 9 report the decomposition behind Figure 2 in tabular form, for Germany and France. Each gives the worker–firm AKM model and the worker–job model under both the plug-in and the leave-out (KSS) estimator, as shares of total residualised log-wage variance, and splits the worker–job model's job-effect variance and sorting covariance between and within occupations. The KSS worker–job columns reproduce the figure; the plug-in columns are the uncorrected estimates referred to in Section 4.

Table 8: Variance decomposition: Germany (2017–2022)

	Worker–firm		Worker–job	
	Plug-in	KSS	Plug-in	KSS
Total residualised variance, Var(y)	0.233	0.233	0.239	0.239
Shares of Var(y) (%)				
Var. worker effect	64.3	54.6	53.2	43.3
Var. firm/job effect	13.8	11.8	25.1	20.5
between-occupation	—	—	7.9	7.9
within-occupation	—	—	17.2	12.6
Sorting (2 Cov)	12.1	15.6	13.0	21.4
between-occupation	—	—	19.7	19.8
within-occupation	—	—	-6.7	1.7
Var. residual	9.8	17.9	8.8	14.8

Notes: Decomposition of residualised log-wage variance for the worker–firm AKM model (job = firm) and the worker–job model (job = firm $\times$ 5-digit occupation), Germany 2017–2022, under the plug-in and the leave-out (KSS) estimator of Kline et al. (2020). Each model is estimated on its own leave-one-out connected set, so $\mathbb{V}(y)$ differs slightly across models; component rows are shares of that model's $\mathbb{V}(y)$. The worker–job model's job-effect variance is split between/within occupations by the law of total variance and the sorting by the law of total covariance; the worker–firm model is occupation-blind. The worker–job KSS variance, sorting, and occupation-split cells use the projection-based corrected estimator of Section 2.2, so they are exactly consistent with the corrected channel decomposition of Section 4.1; the worker row uses the standard leave-out estimate, and the worker–job KSS residual is computed as the remainder of $\mathbb{V}(y)$. Sorting is $2\,\mathrm{Cov}(\alpha, \cdot)$. Components may not sum exactly due to rounding.

Table 9: Variance decomposition: France (2015–2019)

	Worker–firm		Worker–job	
	Plug-in	KSS	Plug-in	KSS
Total residualised variance, Var(y)	0.205	0.205	0.198	0.198
Shares of Var(y) (%)				
Var. worker effect	77.9	57.4	63.9	49.3
Var. firm/job effect	10.1	6.4	16.5	10.6
between-occupation	—	—	7.9	5.1
within-occupation	—	—	8.6	5.5
Sorting (2 Cov)	2.2	9.1	10.4	21.0
between-occupation	—	—	8.4	17.0
within-occupation	—	—	2.0	4.0
Var. residual	9.9	27.1	9.3	19.2

Notes: As Table 8, for France (job = firm × 4-digit occupation).

D Plug-in versus leave-out corrected channel decomposition

Table 10 reports every channel of the projection decomposition of equation 5, together with its parent between/within components, under both the plug-in and the leave-out corrected estimator of Section 2.2, for the main German specification (2017–2022, 5-digit occupations, worker–job leave-one-out sample). Reading the two panels of columns side by side quantifies limited-mobility bias channel by channel. The between-occupation channels, which average estimation noise over large occupation cells, are essentially unaffected by the correction: between-occupation sorting is 19.7% of variance by plug-in and 19.8% corrected, with each of its two channels virtually unchanged. The correction is instead concentrated within occupations, where cells are small and leverages large: on the variance side, the within-occupation firm-premium component falls from 0.0279 to 0.0229 and the match component from 0.0132 to 0.0073; on the sorting side, the within-occupation firm-premium covariance rises from 0.0016 to 0.0111 (as 2 Cov) and the match covariance shrinks from −0.0177 to −0.0071, so that total within-occupation sorting moves from −6.7% to +1.7% of variance. The corrected channels sum exactly to the corrected parent components by construction (Section 2.2).

Table 10: Plug-in versus leave-out corrected channel decomposition (Germany, 2017–2022)

	Plug-in		Leave-out (KSS)	
	Level	Share (%)	Level	Share (%)
Panel A: worker–job sorting, $2\,Cov(\alpha_i,\cdot)$				
Between-occ.: occupation premium, γ_o	0.0164	6.9	0.0165	6.9
Between-occ.: firm premium, $\mathbb{E}[\psi_f \mid o]$	0.0308	12.9	0.0308	12.9
Between-occupation sorting	*0.0473*	*19.7*	*0.0474*	*19.8*
Within-occ.: firm premium, ψ_f	0.0016	0.7	0.0111	4.6
Within-occ.: match, Ω_{of}	-0.0177	-7.4	-0.0071	-3.0
Within-occupation sorting	*-0.0161*	*-6.7*	*0.0040*	*1.7*
Total worker–job sorting	0.0312	13.0	0.0513	21.4
Panel B: variance of the job effect, $\mathbb{V}(\lambda_j)$				
Between-occ.: occupation premium, $\mathbb{V}(\gamma_o)$	0.0049	2.0	0.0049	2.0
Between-occ.: firm premium, $\mathbb{V}(\mathbb{E}[\psi_f \mid o])$	0.0099	4.1	0.0098	4.1
Between-occ.: occ. × firm cross, $2\,\mathrm{Cov}(\gamma_o, \mathbb{E}[\psi_f \mid o])$	0.0041	1.7	0.0041	1.7
Between-occupation job variance	*0.0189*	*7.9*	*0.0188*	*7.9*
Within-occ.: firm premium, $\mathbb{E}[\mathbb{V}(\psi_f \mid o)]$	0.0279	11.7	0.0229	9.6
Within-occ.: match, $\mathbb{V}(\Omega_{of})$	0.0132	5.5	0.0073	3.1
Within-occupation job variance	*0.0411*	*17.2*	*0.0302*	*12.6*
Total job-effect variance	0.0600	25.1	0.0491	20.5

Notes: Plug-in and leave-out (KSS) corrected estimates of each channel of the projection $\lambda_j = \gamma_o + \psi_f + \Omega_{of}$ (equation 5), Germany 2017–2022, worker–job leave-one-out sample, 5-digit occupations. Levels are variance contributions in residualised log-wage units (2 Cov for the sorting rows of Panel A); shares are percentages of the sample's total residualised variance $\mathbb{V}(y) = 0.239$. Italicised rows are the parent between/within-occupation components of the laws of total covariance and variance (equation 3); corrected channels sum exactly to corrected parents by construction (Section 2.2), although displayed values may not sum exactly due to rounding. Channels that are identically zero by construction are omitted.

E Event studies

This appendix collects two robustness exercises for the exogeneity test of Section 4.3.1. First, Figures 11 and 12 repeat the German event studies clustering jobs by their leave-out mean *wage* rather than by the estimated job fixed effect, to confirm that the symmetric step pattern is not an artefact of binning on noisily estimated effects. Second, Figures 13 and 14 reproduce the test in the French data, where occupations are coded independently of employers, as an external check.

Figure 11: Event study around job moves, clustering by leave-out mean job wage (Germany, 2017–2022)

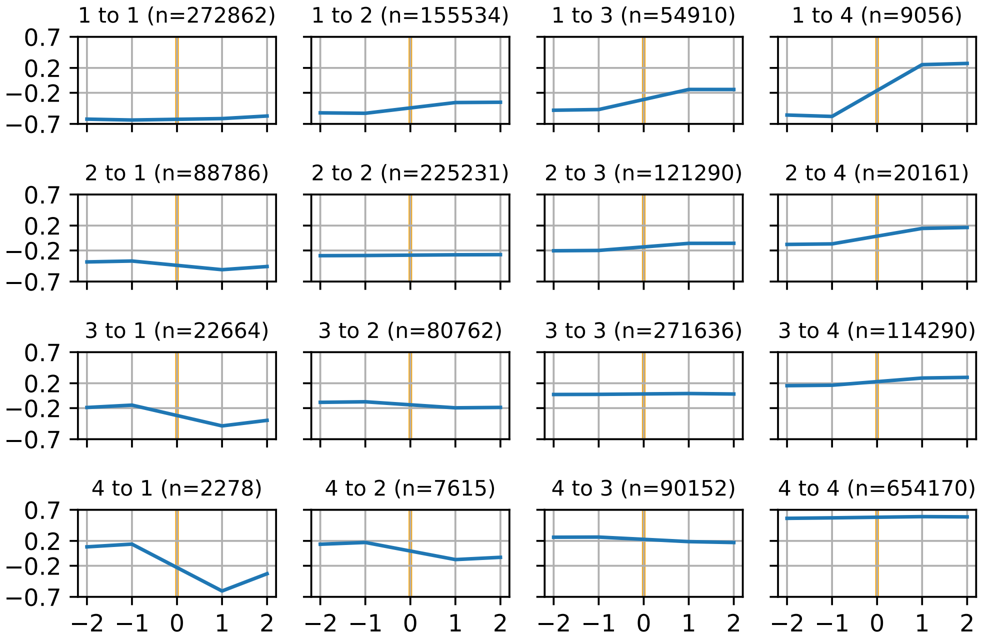

Notes: As Figure 9, but jobs are sorted into quartiles of the leave-out mean *wage* of the job rather than the estimated job fixed effect. Each panel plots average residualised log wages in event time around a move between quartiles (origin in rows, destination in columns; cell counts in titles). Germany, worker–job model, 2017–2022.

Figure 12: Event study by move type, clustering by leave-out mean job wage (Germany, 2017–2022)

(a) Within-occupation, across-firm (b) Across-occupation, within-firm

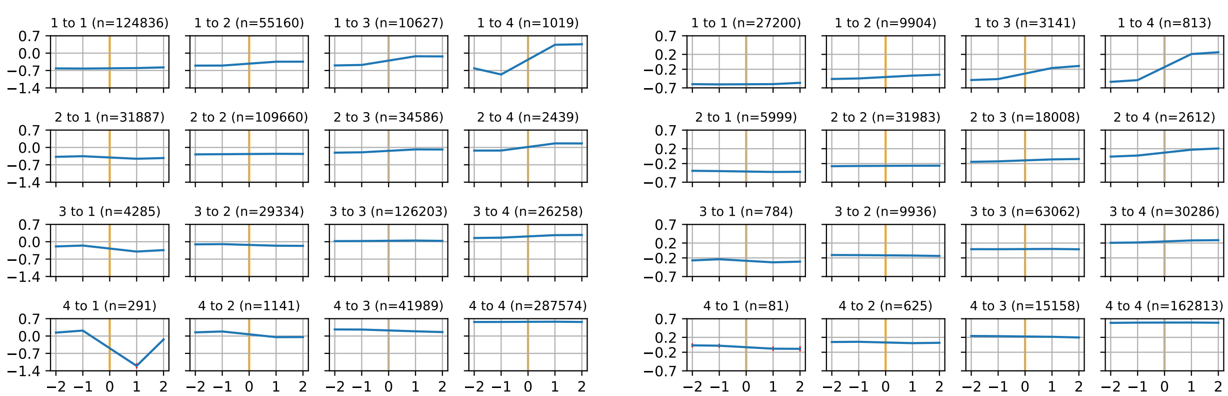

Notes: As Figure 10, but clustering by the leave-out mean wage of the job rather than the estimated job fixed effect; restricting to moves that change the firm but not the occupation (panel a) and the occupation but not the firm (panel b). Germany, 2017–2022.

Figure 13: Event study around job moves (France)

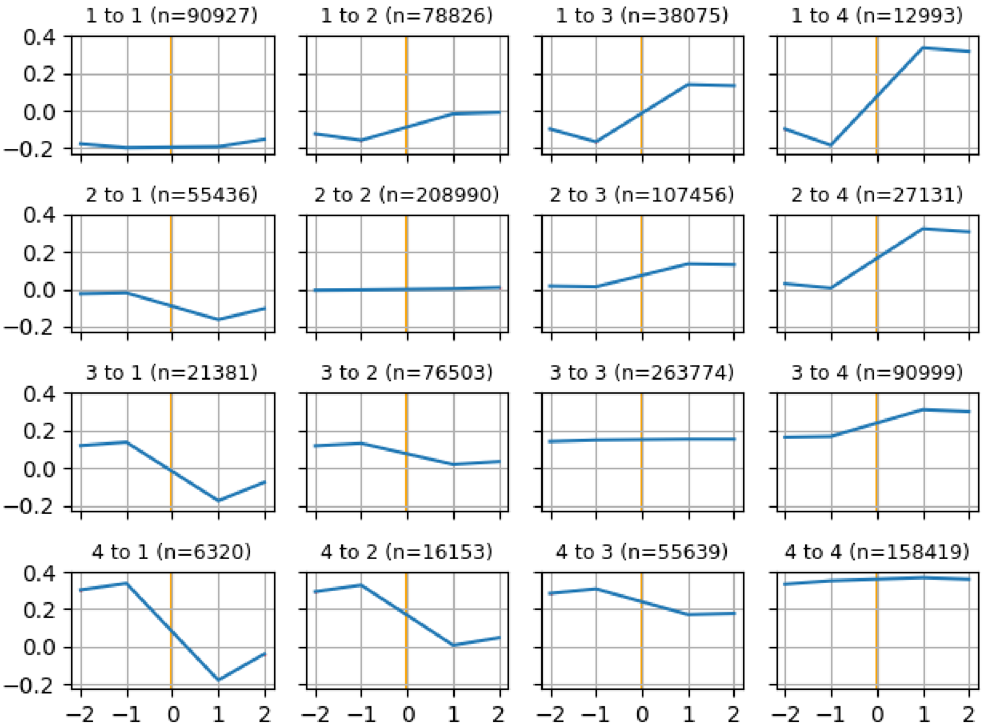

Notes: The exogeneity test of Section 4.3.1 estimated on the French data. Jobs are sorted into quartiles of the leave-out job fixed effect; each panel plots average residualised log wages in event time around a move between quartiles (origin in rows, destination in columns; cell counts in titles).

Figure 14: Event study by move type (France)

(a) Across-firm moves (b) Across-occupation moves

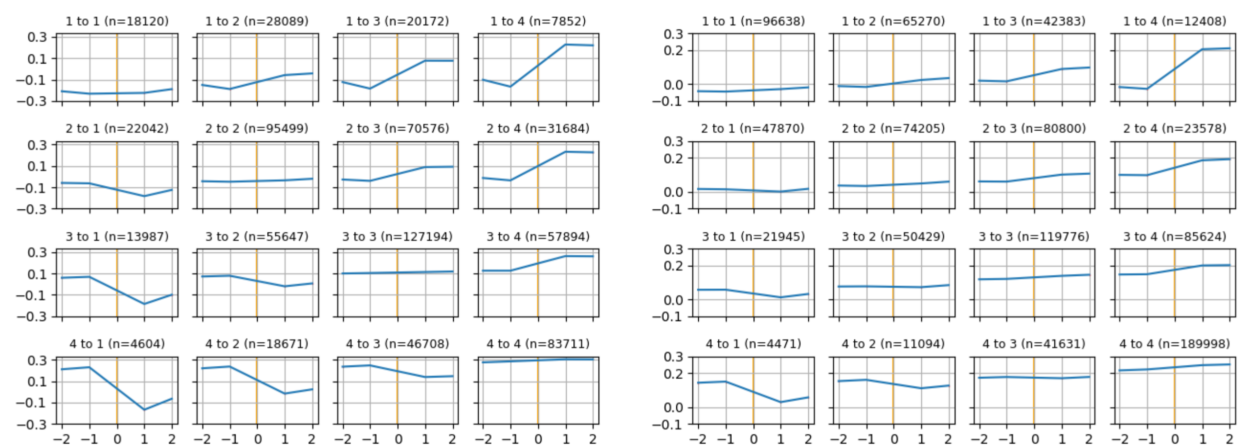

Notes: As Figure 13, restricting to moves that change the firm (panel a) and moves that change the occupation (panel b), in the French data.

F Granularity of occupation coding

These figures re-estimate the worker–job decomposition at progressively coarser occupation codes and report the between- versus within-occupation split of the job-effect variance (Figure 15) and of worker–job sorting (Figure 16). Germany is shown at 1- through 5-digit KldB 2010 codes and France at 1- through 4-digit PCS codes; all components are KSS-corrected. These sweeps use the conditional-moment leave-out estimator, which is computationally feasible at every coding; the more demanding projection-based channel correction of Section 2.2 is applied at the 3- and 5-digit codings in Figure 17.

Figure 15: Job-variance decomposition by occupation granularity

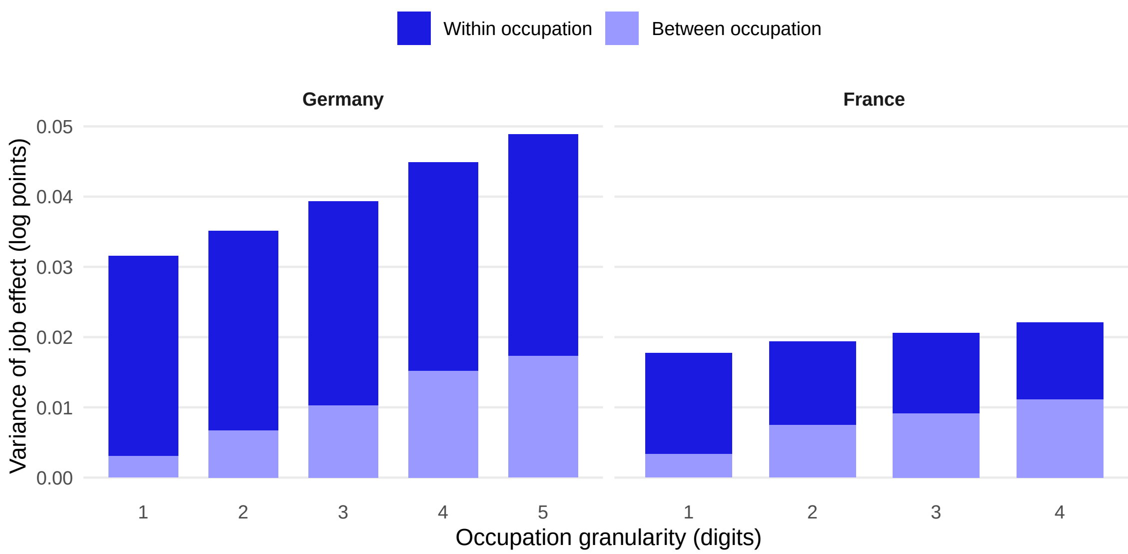

Notes: Each bar decomposes the variance of the estimated job fixed effect, $\mathbb{V}(\lambda)$, into between- and within-occupation components by the law of total variance, at the occupation granularity given on the horizontal axis (digits of the classification). Germany (2017–2022, KldB 2010, 1–5-digit) and France (PCS, 1–4-digit); KSS-corrected. The within-occupation component is the employment-weighted mean of the occupation-conditional job-effect variances; the between-occupation component is the variance of occupation-level mean job effects.

Figure 16: Worker-job covariance decomposition by occupation granularity

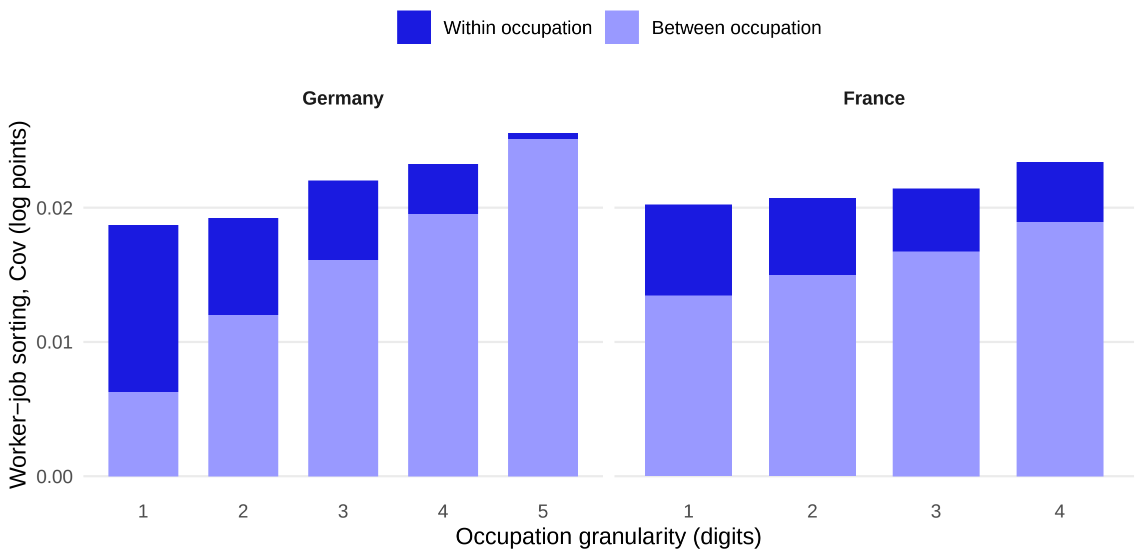

Notes: As Figure 15, decomposing worker–job sorting $\mathrm{Cov}(\alpha, \lambda)$ into between- and within-occupation components by the law of total covariance. As the occupation code coarsens (right to left within each panel), between-occupation differences are mechanically reassigned to the within-occupation component, so the between-occupation share falls by construction; a substantial between-occupation component nonetheless survives even at one digit, most clearly in the independently coded French data.

Figure 17: Channel cross-decomposition by occupation granularity (Germany)

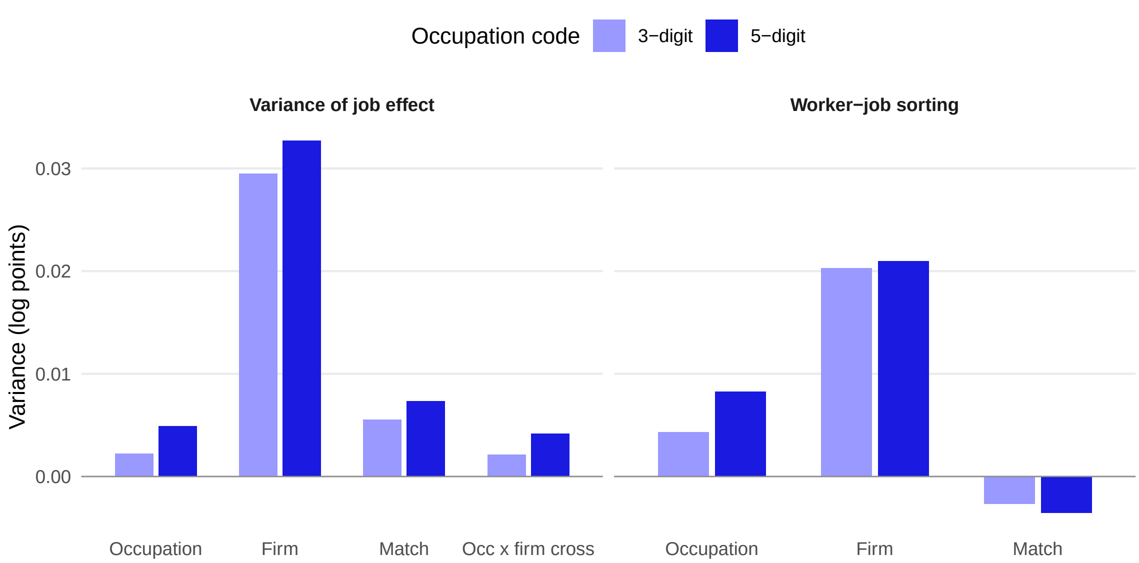

Notes: Leave-out corrected decomposition of the linear projection $\lambda_j = \gamma_o + \psi_f + \Omega_{of}$ into occupation-premium (γ), firm-premium (ψ), and match (Ω) channels, at genuine 3-digit (144 occupations) versus 5-digit (1,233 occupations) KldB codes. Left panel: the variance of the job effect $\mathbb{V}(\lambda)$, where the "occ × firm cross" bar is $2\,\mathrm{Cov}(\gamma_o, \mathbb{E}[\psi_f \mid o])$. Right panel: worker–job sorting $\mathrm{Cov}(\alpha, \lambda)$, which carries no cross term. Germany, 2017–2022; leave-out corrected (Section 2.2), comparable to the waterfalls of Section 4.1. The firm-premium channel is near-identical across codings; the occupation-premium channel scales with granularity; the match channel is small throughout.

G Decomposition overview figure

Figure 18: Log wage decomposition tree (Germany, 2017–2022)

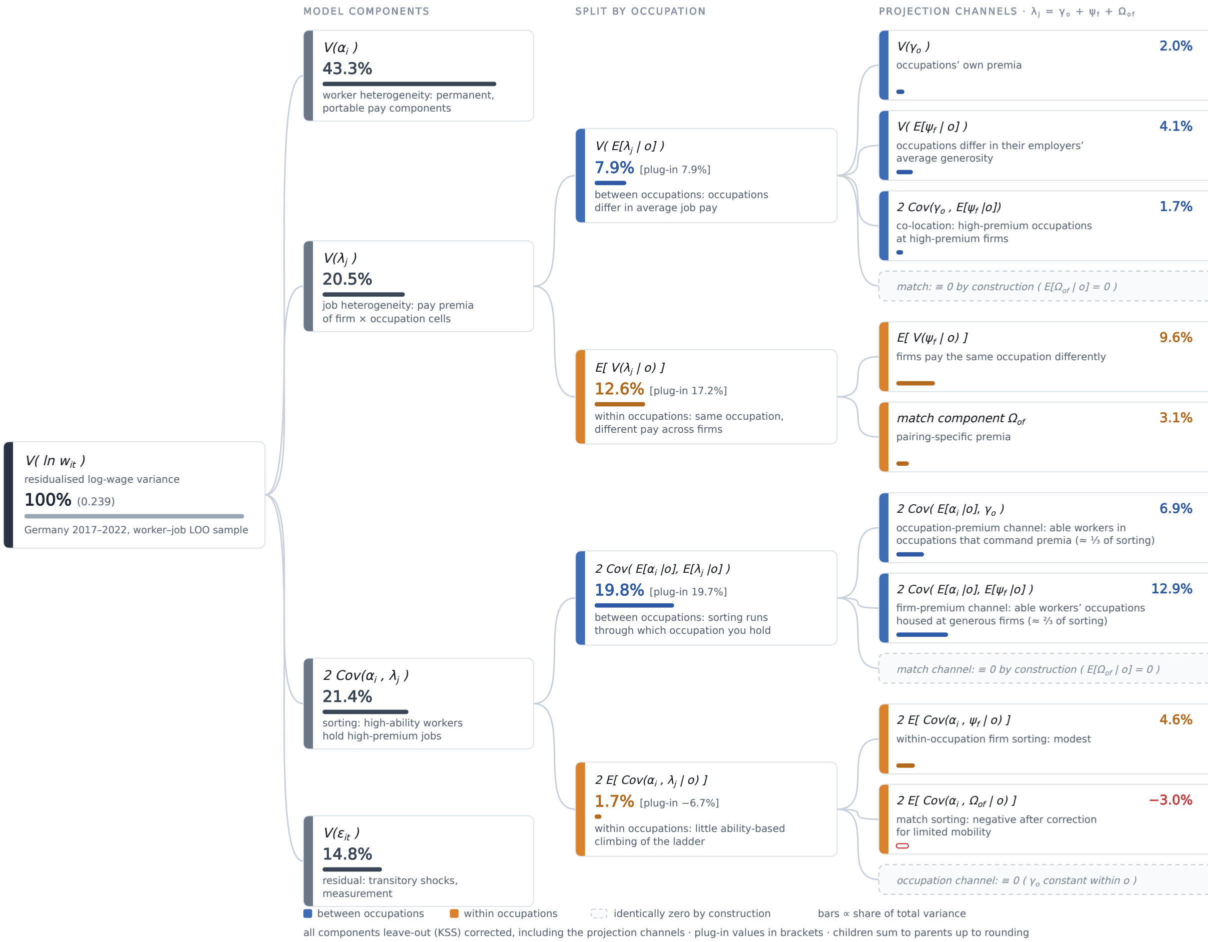

Notes: The figure reports the complete decomposition of residualised log-wage variance for the worker–job model (Germany, 2017–2022, worker–job leave-one-out sample), read left to right: the first column decomposes $\mathbb{V}(y)$ into worker, job, sorting, and residual components (equation 1); the second splits the job-variance and sorting components between and within occupations by the laws of total variance and covariance (equation 3); the third splits each of these into occupation-premium, firm-premium, and match channels via the employment-weighted linear projection $\lambda_j = \gamma_o + \psi_f + \Omega_{of}$ (equations 4–5). Every component is expressed as a share of total residualised variance (0.239); sorting components are reported as 2 Cov throughout; horizontal bars are proportional to shares, with negative values drawn as outlines. All components, including the projection channels, are leave-out corrected (Kline et al., 2020) using the extension described in Section 2.2, so children sum to their parent components up to rounding; plug-in values are shown in square brackets at the occupation split, and plug-in estimates for every channel are reported in Appendix D. The occupation × firm cell in the variance branch is the doubled cross term $2\,\mathrm{Cov}(\gamma_o, \mathbb{E}[\psi_f \mid o])$ that arises in the variance of a sum; the sorting branch, being linear in the job premium, carries no cross term. Dashed cells mark channels that are identically zero by construction of the projection: $\mathbb{E}[\Omega_{of} \mid o] = 0$ for every o, and γ_o is constant within occupations (equation 5). The grouping of the within-occupation variance channels follows Figure 4. First- and second-column values correspond to Table 8 and Figure 2; third-column values are the components of Figures 3 and 4 expressed as shares of total variance. Components may not sum exactly due to rounding.