Do Workers Sort to Firms or to Occupations?

Luke Heath Milsom* Shihang Hou*

February 19, 2026

Abstract

In this paper, we find that high-wage workers sort mainly to high-wage occupations and not to high-wage firms, and that at least half of the previously documented sorting to firms can be attributed to the segregation of occupations across firms. To reach these conclusions, we leverage the universe of matched employee-employer data from France and Germany and estimate a flexible two-way worker-job fixed effects model of log wages. We then isolate worker sorting to firms by studying the within-occupation across-firm covariance between worker and job fixed effects.

*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 some confidential data, on which this work is based, 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 some occupations command higher pay than others (Autor et al., 2003; Goos and Manning, 2007; Acemoglu and Autor, 2011). In this paper, we argue that it is crucial for work on wage inequality to consider both of these margins in conjunction because high-paying occupations are unevenly distributed across firms. This leads to the risk that focusing solely on firms conflates returns from sorting into high-paying firms with the returns from sorting into high-paying occupations.[1] The potential for confounding is empirically important; our analysis of German and French administrative data reveals that occupations are highly segregated across firms—so much so that identifying an individual's employer resolves two-thirds of the uncertainty regarding their four-digit occupation.

To disentangle firm- and occupation-level sources of wage variation and sorting patterns, we build upon the two-way fixed effects framework of Abowd et al. (1999) (AKM). We begin by defining a "job" as a firm-occupation pair, and estimate a two-way model recovering worker and job fixed effects. We then use the Law of Total Covariance to decompose the covariance of worker-job fixed effects into two separate components: within-occupation covariance between worker and firm wage components, and the cross-occupation covariance between worker and occupation wage components. This decomposition separates worker-job covariance into two elements that are conceptually relevant to different theoretical literatures. The first component captures worker–firm covariance conditional on occupation, corresponding to the familiar "high-wage workers work in high-wage firms" pattern emphasized in the two-way fixed effects literature (e.g. Abowd et al., 1999; Card et al., 2013; Song et al., 2019). The second component captures worker–occupation covariance, which is central to Roy-style models of occupational selection (Heckman and Sedlacek, 1985; Keane and Wolpin, 1997).[2]

We implement this decomposition using the universe of French and German matched employee-employer administrative data. We apply the leave-one-out method of Kline et al. (2020) to correct for biases from limited worker mobility (Andrews et al., 2008; Bonhomme et al., 2023), and extend this methodology to correct for biases in conditional moments. We check for exogeneity using the event-study approach of Card et al. (2013), and show

[1] For example, computer scientists may join Google not because it offers a generous firm-wide pay premium to all its workers, but because it is one of the few employers that offers elite software engineering roles. Even if Google paid industry-average wages in these high-value occupations, a researcher using the firm-premium approach would estimate a large Google premium, simply due to its concentration of high-value occupations.

[2] Throughout, we use sorting to refer to reduced-form covariance in wage components rather than to structural complementarities, consistent with its use in e.g. Card et al. (2013) and Song et al. (2019). In the literature, authors have also used the term "sorting" to reflect the sorting of high ability workers to more productive firms due, e.g., to complementarities in production. We discuss the relation of our results to that literature in section 4.3.

robustness to using coarser occupation classifications, and different time periods.

In both the French and German settings, we find that sorting to occupations is quantitatively over four times as important as sorting to firms, and that over half of the sorting to firms found in a standard AKM decomposition can be explained by the clustering of occupations across firms. Drawing on our French data, we find that while the canonical model attributes 9.1% of wage variance[3] to worker-firm sorting, our decomposition finds that sorting across firms within occupations accounts for only 4.0%. In contrast, the sorting of workers to occupations accounts for nearly 17.0% of the variance.

Our model also reveals that the total importance of workplace premia is larger than previously understood. We find that within-occupation firm heterogeneity accounts for 5.0% of wage variance, a figure comparable to the 6.4% estimate for the entire firm effect in the standard model. However, we find that differences in pay premia between occupations account for a further 6.0% of total log wage variance, bringing the total share of log wage variance explained by job fixed effects to 11.0%. Our findings are quantitatively similar in the German context and are robust to using coarser occupational classifications. Finally, following a recent literature that has documented the increasing role of firms in explaining wage inequality (Song et al., 2019; Card et al., 2013), we find that these increases are due both to increases in between occupation heterogeneity, and increases in within-occupation cross-firm heterogeneity.

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 to 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). These decompositions are often interpreted as reflecting firm-specific pay policies or non-wage amenities (Card et al., 2018; Alvarez et al., 2018). We contribute to this literature by showing that such interpretations require accounting explicitly for occupations. Incorporating occupations reveals substantial wage dispersion across firms, even conditional on occupation, consistent with earlier findings, but also substantially lower covariance between worker and firm wage components within occupation groups. This suggests that occupational allocation is an important margin through which wage inequality and sorting patterns arise, and that controlling for occupation is a necessary step for interpreting firm wage premia as reflecting firm-specific amenities or pay policies.

Our paper also contributes to the large literature that studies labour market matching

[3] We present all results relative to residualised log wage variance, where we control for a cubic age profile and year fixed effects.

(Sattinger, 1993; Postel-Vinay and Robin, 2002; Shimer, 2005). A substantial body of work studies worker–firm sorting within models of assignment and matching, emphasizing complementarities between worker ability and firm productivity as a source of wage dispersion and inequality (Lise et al., 2016; Bagger and Lentz, 2019; Lentz et al., 2023; 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 is the dominant margin underlying observed wage inequality and wage covariance patterns. In particular, while we find substantial dispersion in firm wage premia even conditional on occupation, we document relatively low covariance between worker and firm wage components within occupation groups, indicating limited sorting of high-wage workers into high-wage firms once occupation is held fixed.

A related literature is that studying the identification of underlying worker-firm complementarities from matched employer-employee data (Eeckhout and Kircher, 2011; Hagedorn et al., 2017; Borovičková and Shimer, 2017; Lopes de Melo, 2018; Sorkin, 2018; Bonhomme et al., 2019; Borovičková and Shimer, 2024). This literature suggests that the covariance of worker and firm fixed effects might not identify sorting primitives because, among other reasons, (i) wages are not always monotonic in worker ability, and (ii) selection leads to AKM estimating an average treatment effect on the treated. Our results suggest an additional perspective: a significant share of productive complementarities operates through worker–occupation matching rather than worker–firm matching, implying that firm-level models may misattribute the locus of sorting. An important direction for future research is therefore to assess whether worker–job complementarities primarily arise at the occupational level or within firms, conditional on occupation.

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, 2017a; Lamadon et al., 2022). In contrast, our job–fixed-effect 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.

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. Finally, section 5 concludes.

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

We consider a framework that allows us to identify sorting between workers and firms, flexibly accounting for occupations. To do this, we augment the classic AKM framework (Abowd et al., 1999; Card et al., 2013; Song et al., 2019) to decompose log-wages into components due to individual, i, and job, j effects, where jobs are defined by firm-occupation pairs $j \in \mathcal{J} = \mathcal{F} \times \mathcal{O}$. Define the assignment function $J(i,t) = j$. This specification is given in equation 1. Throughout, we also condition on a set of time-varying worker covariates X_{it}, such as age and year fixed effects. These have been omitted in this discussion for notational brevity.[4]

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

This framework extends the worker-firm two-way fixed effects model in two dimensions. First, it allows firms to pay varying wage premia across occupations. For example, software developers can receive a higher premium than accountants within the same firm. Second, it allows firm-occupation-specific pay premia: software developers at Google can be paid a higher premium than software developers at other firms, even those with the same firm-pay premia.[5] The specification in equation 1 allows a simple decomposition of the observed variance of log wages as given in equation 2 below.

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

The term $\mathrm{Cov}(\alpha_i, \lambda_{J(i,t)})$ aggregates two distinct economic mechanisms: the sorting of workers into specific occupations and the sorting of workers into specific firms within those

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

[5] In particular, this framework is more flexible than a three-way fixed-effects specification with occupation, γ_o, and firm ψ_f fixed effects. To see this, note that we can write $\lambda_j = \gamma_o + \psi_f + \Omega_{of}$. Therefore, while the three-way fixed effects model captures the covariance between workers and occupations ($\mathrm{Cov}(\alpha_i, \gamma_o)$), and workers and firms ($\mathrm{Cov}(\alpha_i, \psi_f)$), it does not capture the covariance between workers and firm-occupation match effects ($\mathrm{Cov}(\alpha_i, \Omega_{of})$. This missing component can be intuitively understood as the sorting of high-wage workers to high-wage firm-occupation pairs.

occupations.[6] To disentangle these channels, we apply the Law of Total Covariance, partitioning the joint distribution of worker and job effects by occupation (o):

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

The first term, $\mathbb{E}\left[\text{Cov}(\alpha_i, \lambda_J(i,t)|o)\right]$, 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, $\text{Cov}\left(\mathbb{E}[\alpha_i|o], \mathbb{E}[\lambda_{J(i,t)}|o]\right)$, captures between-occupation sorting: the covariance between average worker ability and average job pay across occupations. We 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 $\mathbb{E}\left[\text{Var}\left(\lambda_j|o\right)\right]$, and the variance of job fixed effects between occupations $\text{Var}\left(\mathbb{E}\left[\lambda_j|o\right]\right)$.

The relevance of this distinction can be illustrated by comparing hypothetical firms, each providing healthcare. Consider Hospital A, a specialized surgical center, and Hospital B, a general nursing clinic. Suppose that, due to market-wide occupational premia, surgeons earn significantly more than nurses regardless of where they work. If both hospitals pay the standard market rate for each role, there is zero within-occupation sorting: a surgeon (or nurse) earns the same premium at either facility. However, an occupation-blind model—which ignores the o in $j \in \mathcal{F} \times \mathcal{O}$—would observe that the average worker at Hospital A earns more than the average worker at Hospital B. Because the model cannot 'see' the occupational composition, it incorrectly attributes this difference to a generic firm-wide pay premium (ψ_f). Consequently, the model would find a large, positive covariance between worker ability and firm effects, simply because high-ability workers (surgeons) are concentrated at Hospital A. Our decomposition would instead suggest this to be a statistical artifact of occupational specialization rather than true firm-level sorting. By partitioning the covariance, we correctly assign this effect to the between-occupation component, preventing the inflation of the firm's role in wage determination.

We next turn to the practical problem of recovering our object of interest, $S_{JFE} = \mathbb{E}[\text{Cov}(\alpha_i, \lambda_j|o)]$, from the data. Using the universe of matched employee-employer data, we can recover $\{\widehat{\alpha}_i, \widehat{\lambda}_j\}$ using equation 2. Note that fixed effects can only be recovered in a relative sense, and therefore, we can only estimate them for the largest connected set of jobs. S_{JFE} could be biased for two reasons. First, the fixed effects themselves could be

[6] 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 on this work in section 4.3.

biased, and S_{JFE} 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, such as S_{JFE}.

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: $\forall j, Pr(J(i,t) = j|\varepsilon) = Pr(J(i,t) = j)$; 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 sub-section 4.2 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 variable 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) shows 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 estimates results for a particular sub-sample of the data containing more well-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 recover the within-occupation components of the law of total covariance, we construct bias-corrected estimates of the conditional covariance $\mathrm{Cov}(\alpha_i, \lambda_{J(i,t)} \mid o)$ by aggregating the relevant KSS-corrected quadratic forms across worker-job observations within each occupation.[7] The within-occupation sorting component is then obtained by taking the employment-weighted mean of these occupation-specific covariances across occupations. An analogous procedure is used to estimate the within-occupation variance of job fixed effects, $\mathrm{Var}(\lambda_j \mid o)$.As a robustness check, we also compute plug-in estimates of the between-occupation components implied by the law of total covariance, i.e.: $\mathrm{Cov}\big(\mathbb{E}[\alpha_i \mid o], \mathbb{E}[\lambda_{J(i,t)} \mid o]\big)$ and $\mathrm{Var}(\mathbb{E}[\lambda_j \mid o])$, using occupation-level averages of the estimated fixed effects. These alternative computations yield quantitatively similar results, providing reassurance that our implementation of the decomposition is not sensitive to the specific way in which the moments are constructed.

3 Data

We establish our empirical conclusions using administrative data from two of the largest labour markets in Europe: France and Germany. These datasets offer complementary strengths: the French data provides complete, uncensored wage information, but has to be constructed into a panel by stitching together individual cross-sections — a process with a roughly 95% success rate. The German data is available as a full-count panel, but the wage variables available are top-coded. Establishing our core results in both contexts, despite their different institutional settings and data strengths and weaknesses, underscores the robustness and generality of our findings. We compute results for the periods 2015-2019 for France, and 2017-2022 for Germany.

Our empirical strategy leverages information about a worker's occupation; we thus describe how the occupation variable was collected separately for each country. In both cases, we check for the robustness of our results in two ways: (1) we use more aggregated occupational categories as there is less data error in the coding of more aggregate occupations, and (2) we estimate fixed effects based only on changes in occupation due to cross-firm moves, which are also less likely to be coded with error.

3.1 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 except those in government employment. We follow Babet et al. (2025) in

[7] See Azkarate-Askasua and Zerecero (2024) for an alternative, computationally faster approach when computing multiple corrections, as we do here.

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, residualized 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.[8] This process of asking employers rather than employees, using occupation descriptions, and checking coding, reduces the probability of incorrect occupation categorisation.

3.2 German Sample

Our German data comes from the Employee Histories (BEH), which contains employment spells for all workers outside the civil service. Its primary limitation is the top-coding of wages above the social security contribution ceiling, which we address by imputing censored wages using the established methodology of Card et al. (2013). This imputation may lead to the underestimation of sorting in higher-wage occupations. Key variables are real annual earnings, the establishment identifier[9], and the highly granular occupation codes based on the 5-digit KldB 2010 occupational classification (with 1286 categories in the full classification). Our sample selection mirrors the French data: full-time male workers aged 25-60 in West Germany. As with the French data, we use log annual wages residualized on an age cubic and year fixed effects.

The occupational variable is collected from filings by employers from social security records; since one's occupation is not important for this purpose, there is a perception that employers may not always update their workers' occupations in a timely manner. As a result, occupational moves within firms are likely to be under-reported but do not suffer

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

[9] An establishment can be thought of as roughly a firm-industry group and is distinct from both firm and branch. See Card et al. (2013) for a discussion of this issue.

from potential issues surrounding independent interviewing (Carrillo-Tudela and Visschers, 2023).

3.3 Summary Statistics

Table 1 shows some summary statistics from the two samples used in the main analysis, that is, 2015-19 for France and 2017-22 for Germany. After implementing the restrictions described, our main sample covers 48.0 million observations consisting of 14.3m workers, 1.3m firms, and 4.5m jobs in France, and 54.6 million observations consisting of 14.2m workers, 1.4m firms, and 4.9m jobs in Germany.

One disadvantage of using the Kline et al. (2020) correction of second-order moments of estimated fixed effects is that it is only feasible for the leave-one-out connect set (LOO) for the worker-jobs model. We present summary statistics for the connected sub-sample, for which identification of the model is feasible, in column 2 and for the leave-one-out connected subsample, for which the KSS correction is feasible in column 3.

Once we consider the leave-one-out connected set for the worker-jobs model, we have significantly fewer observations (around 66% of the overall total for the French data and 60% for the German data). The LOO connected set for the jobs model, which underlies our main specification contains 57% of workers, 21% of firms, and 18% of jobs in the full data in France (56% of workers, 29% of establishments, and 19% of jobs in Germany). On the other hand, we find that the mean and variance of earnings is similar across the samples. The mean in the leave-one-out sample is about 0.04 log points higher in terms of annualised wages in France and Germany, and the residualised log-wage variance is somewhat smaller, by 0.03 and 0.01 respectively.

We also provide summary statistics on the number of moves we observe in the data. We report the total number of times workers are observed to change either their firm or occupation, as well as the number of moves that are between firms, the number of moves that are between occupations, and the number of moves that are between both firms and occupations. In the full data, we find that 14.5% of observations involve moves in France and 10.6% of observations involve moves in Germany. Of these moves, 51.6% involve moves between firms, 79.0% involve moves between occupations, and so 30.5% involve changes in both firms and occupations (84.2%, 65.1%, and 49.4% for Germany).

We find that the ratio of moves to observations in the LOO connected set is comparable to that in the full data. For the LOO connected set, 14.1% of all observations involve moves, of which 57.1% are moves between firms, 74.5% are moves between occupations, and 31.8% are moves between both firms and occupations in France (10.9%, 84.1%, 61.7%, and 45.8% in Germany). Thus, despite only containing 2/3 of the full sample, the leave-one-out connected sample nevertheless resembles the full sample in terms of the composition of moves across firms and occupations.

Table 1: Summary statistics for the French and German samples

	Full Sample	Connected Sample	Leave-one-out Connected Sample
France (2015-2019)			
Total observations (m)	48.01	35.93	31.53
Total workers (m)	14.25	9.06	8.09
Total firms (m)	1.25	0.57	0.27
Total jobs (m)	4.48	2.04	0.82
Mean log annual wage	10.39	10.42	10.43
Var log annual wage	0.25	0.23	0.23
Mean log hourly wage	2.96	2.97	2.98
Var log hourly wage	0.21	0.20	0.19
Var residualised log hourly wage	0.24	0.22	0.21
N moves (m)	6.97	5.69	4.43
N firm moves (m)	3.59	3.25	2.53
N occ moves (m)	5.50	4.38	3.30
N firm + occ moves (m)	2.13	1.94	1.41
Germany (2017-2022)			
Total observations (m)	54.56	40.16	32.96
Total workers (m)	14.22	9.44	7.95
Total estabs (m)	1.35	0.78	0.39
Total jobs (m)	4.91	2.30	0.92
Mean log daily wage	4.78	4.80	4.82
Var log daily wage	0.27	0.26	0.26
Var residualised log daily wage	0.25	0.24	0.24
N moves (m)	5.79	5.00	3.58
N estab moves (m)	4.88	4.27	3.01
N occ moves (m)	3.77	3.29	2.21
N estab + occ moves (m)	2.86	2.56	1.64

Notes: This table presents summary statistics for the two most recent periods studied in both Germany and France. Statistics on the number of observations and number of moves are presented in millions. Column 1 presents statistics for the full sample after the restrictions described in section 3, while columns 2 and 3 describe statistics for the largest connected set and the largest leave-one-out connected set, respectively.

3.4 Segregation of Occupations Across Firms

In this paper, we argue that when decomposing wage inequality, it is important to account for occupations and firms jointly because occupations are segregated across firms. To quantify the extent of this segregation, we turn to information theory and compute the Theil-H index across firms (Theil and Finizza, 1971; Theil, 1972). The Theil-H index is a measure of segregation between multiple groups and can be interpreted as how much uncertainty about a worker's occupation would be resolved if one knew which firm they worked at. If occupations were perfectly segregated across firms, then knowing a worker's firm would be equivalent to knowing their occupation. If, on the other hand, occupations in each firm reflected the overall distribution of occupations in the economy, knowing the worker's firm would not improve one's guesses about their occupation. We present the index normalised with respect to the overall amount of uncertainty in workers' occupations, so that the index lies between 0 and 1. To compute the index, we use firms with more than 10 workers, so as to ensure that the findings are not driven solely by very small firms.

We find that occupations are highly segregated across firms in both Germany and France. Figure 1 plots the Theil-H index in Germany and France over the studied time periods. Knowledge of which firm a worker is currently employed at removes about two-thirds of the uncertainty about their occupation in both Germany and France. While reducing the fineness of occupational classification lowers the measured segregation, the effect is not substantial; for example, at the 3-digit level, knowing one's firm still removes about 63.2% of uncertainty in Germany and 62.7% in France. This high level of occupational segregation across firms indicates that considering firms or occupations separately will conflate the two margins.

We also find a consistent upward trend in occupational segregation across firms in both Germany and France — workers are increasingly sorted into firms based on their specific occupations. While Germany shows a steady long-term rise before seeing a slight moderation in 2017-22, France has experienced a notably sharper increase in recent years, with both countries converging at a point where knowing a worker's firm resolves nearly two-thirds of the uncertainty regarding their occupation. This rising segregation (Goldschmidt and Schmieder, 2017b; Handwerker, 2023) leaves open the possibility that recently documented increases in the contribution of firm heterogeneity to inequality (Song et al., 2019) may simply be due to rising occupation segregation across firms. We investigate this in section 4.1.

Figure 1: Segregation of occupations across firms in studied periods

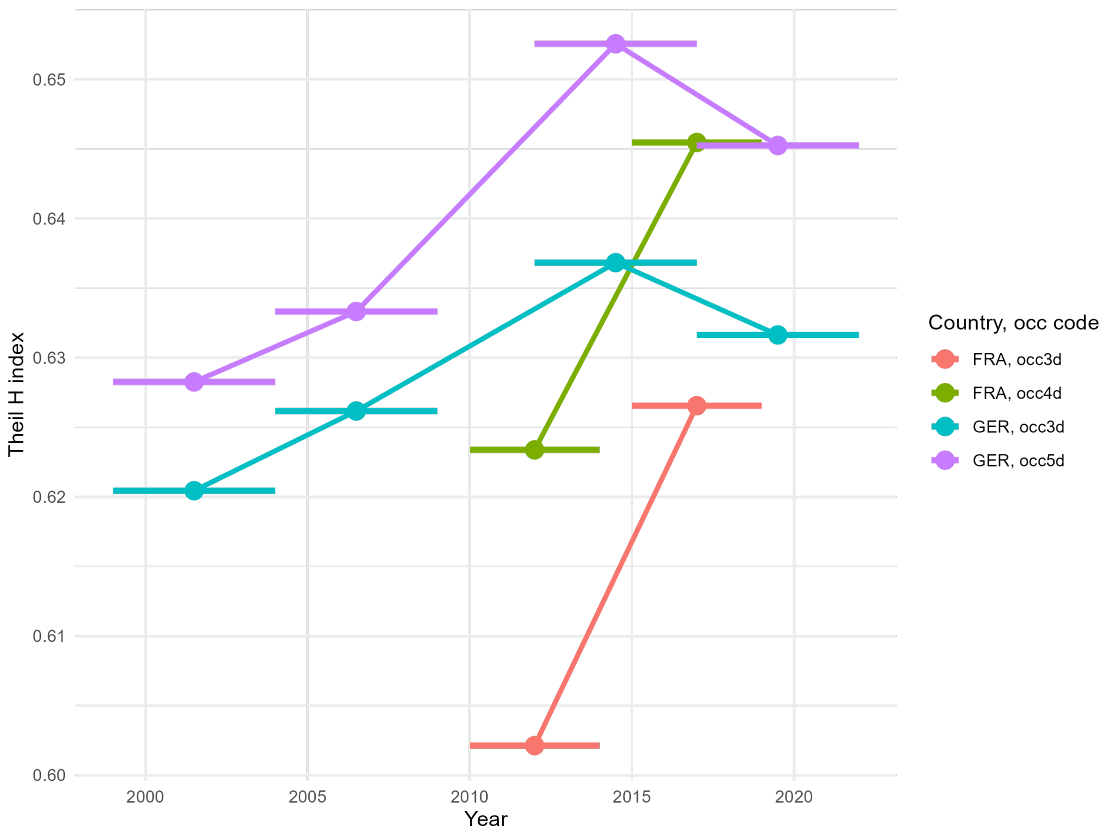

Notes: This figure reports the Theil H-index of occupational segregation across firms in Germany and France. The index is computed using the joint distribution of employment across firms f and occupations o, where n_{fo} denotes the number of worker–year observations in occupation o at firm f, $n_f = \sum_o n_{fo}$ denotes firm size, and $N = \sum_f n_f$ denotes total employment in the relevant sample–period cell. Let $p_o = \sum_f n_{fo}/N$ denote the economy-wide occupation share and $p_{o|f} = n_{fo}/n_f$ denote the within-firm occupation share. Define the entropy of the overall occupation distribution as $E = -\sum_o p_o \log p_o$, and the within-firm entropy as $E_f = -\sum_o p_{o|f} \log p_{o|f}$. The Theil H-index is then computed as the employment-weighted reduction in entropy from conditioning on firm membership, normalised by overall entropy, $H = \frac{E - \sum_f (n_f/N)\, E_f}{E}$, so that $H \in [0, 1]$. In all computations, the sample is restricted to firms with more than 10 workers in the relevant period, and the underlying worker sample and variable definitions follow the restrictions described in Section 3. The figure reports H under two occupation codings in each country: a coarser 3-digit classification and the most detailed classification used in the analysis (5-digit for Germany based on KldB 2010 and 4-digit for France based on PCS).

4 Decomposition results

In this section, we present our main decomposition results for both the French and German data side-by-side. We decompose log-wage variance using the fixed effects estimated from equation 1, correcting for limited mobility bias and using the law of total variance and covariance to separate between-occupation (e.g. $\text{Cov}\left(\mathbb{E}[\alpha_i|o], \mathbb{E}[\lambda_{J(i,t)}|o]\right)$) and within-occupation across-firm components (e.g. $\mathbb{E}\left[\text{Cov}(\alpha_i, \lambda_J(i,t)|o)\right]$). We also display the results from employing the same procedure on the standard AKM worker-firm model for comparison.

Figure 2 summarises the main results. In table 4 and 5 in the appendix, we present the same results in table form and report the results in shares and absolute variance components as suggested in Kline (2024). For both France and Germany, for both the worker-firm and worker-job models figure 2 shows the estimated contribution of worker heterogeneity, $\widehat{\mathbb{V}}(\hat{\alpha}_i)$, to overall log-wage variance in the left-hand-side set of gray bars. In the middle set of bars we show the estimated contribution of firm or job level heterogeneity. In gray we show firm-level heterogeneity, $\widehat{\mathbb{V}}(\hat{\phi}_f)$, from the worker-firm model. The sum of the two blue bars corresponds to the overall contribution of job-level heterogeneity, $\widehat{\mathbb{V}}(\hat{\lambda}_j)$. We further decompose this into the estimated component due to within-occupation heterogeneity, $\widehat{\mathbb{V}}(\widehat{\mathbb{E}}[\hat{\alpha}_i|o], \widehat{\mathbb{E}}[\hat{\lambda}_j|o])$, in darker blue, and the estimated component due to between-occupation heterogeneity, $\widehat{\mathbb{E}}[\widehat{\mathbb{V}}(\hat{\alpha}_i, \hat{\lambda}_j|o)]$, in lighter blue. Finally, on the right-hand side we show the analogous covariance decomposition. In gray bars we show the degree of sorting between workers and firms in the firm level model, measured as $2\widehat{\text{Cov}}(\hat{\alpha}_i, \hat{\phi}_f)$. Then, in blue, we show the analogous quantity in the worker-job model, $2\widehat{\text{Cov}}(\hat{\alpha}_i, \hat{\lambda}_j)$. Finally, we decompose this quantity into between-occupation sorting in dark blue, $2\widehat{\text{Cov}}(\widehat{\mathbb{E}}[\hat{\alpha}_i|o], \widehat{\mathbb{E}}[\hat{\lambda}_j|o])$, and within-occupation sorting in light blue, $2\widehat{\mathbb{E}}[\widehat{\text{Cov}}(\hat{\alpha}_i, \hat{\lambda}_j|o)]$.

Figure 2: Decomposition Results

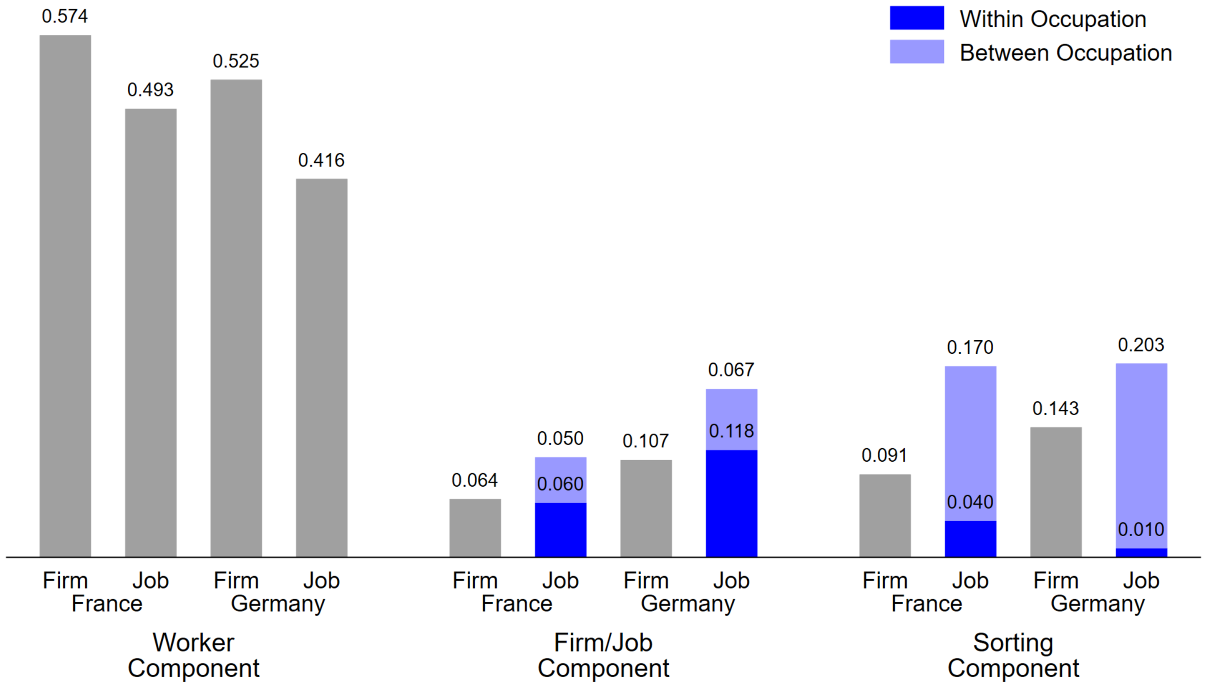

Notes: This figure reports variance and covariance decompositions of residualised log wages using fixed effects from the worker-firm and worker-job models. Bars are expressed as shares of total residualised log-wage variance. Fixed-effect second moments are computed using bias-corrected estimators that adjust for limited worker mobility in connected sets. The left group of bars shows the component due to worker heterogeneity. The middle group shows the contribution of workplace/job heterogeneity. In gray we report the firm-effect variance from the standard AKM worker–firm model. In blue we report the job-effect variance from the worker–job model decomposed using the law of total variance with respect to occupation o into a between-occupation component (dark blue) given by the variance of occupation-specific means, and a within-occupation component (light blue) given by within-occupation across-firm heterogeneity. The right group of bars shows sorting, measured as twice the worker–firm (or worker–job) covariance. Gray bars report this quantity from the worker–firm model. Blue bars report sorting in the worker–job model, further decomposed by occupation using the law of total covariance into the *between-occupation* sorting component (dark blue) and the *within-occupation* sorting component (light blue). Results are shown separately for France (2015-19) and Germany (2017-22).

We find that the sorting of workers to firms within occupations accounts for a small proportion of log-wage variance (4% in the French sample, and 1% in the German sample), significantly less than that recovered by a worker-firm AKM model. Sorting to firms is of a quantitatively unimportant magnitude. On the other hand, sorting across occupations accounts for a relatively large proportion of log-wage variance (17% in the French sample and 20% in the German sample). These results suggest that worker-occupation sorting may underlie much of what was previously attributed to worker-firm sorting in decompositions based on the standard AKM model. The sorting of workers to firms within occupations is only 44.0% as large as the estimated sorting of workers to firms in the AKM model in France and only 7.0% as large in Germany.

Figure 2 also shows that firm heterogeneity explains a similar proportion of log wage variance in the worker-job and worker-firm models in both France and Germany (in France around 6%, whereas in Germany around 11%). In both countries, we also find a significant role for occupation heterogeneity (5% in the French sample, and 6.7% in the German sample). Our results thus support the findings in prior work that there is substantial wage premia between firms even within finely defined occupation groups. However, there is little sorting of workers to firms within occupations, even given this variance. Finally, in both France and Germany, we find that worker-heterogeneity explains significantly less log wage variance in the worker-job as opposed to the worker-firm model, reducing from 57% to 49% in France and from 53% to 42% in Germany.

These results show that sorting is mainly to occupations, and not firms, in line with models of occupation-choice rather than those of firm sorting. A full discussion of the implications of these findings is presented in section 4.3.

4.1 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). In the French setting, due to data restrictions, we cannot leverage a sufficiently long time series. However, using the German data, we can perform our decomposition every five years since 1999.[10] In figure 3 we present how the contribution of job-level heterogeneity and worker-job sorting has changed over time, decomposing each component into that due to within and between occupation components.

Figure 3 panel (a) confirms that we replicate, in our worker-job model, the findings of Card et al. (2013) and Song et al. (2019) that the dispersion of firm wage premia increased over the 2000's. Decomposing this into that due to within and across occupation changes, we find that both are quantitatively important. However, extending this exercise until 2022 we find a flattening and then a decrease in the explanatory power of job level heterogeneity. Again, decreases in both across-occupation heterogeneity and within-occupation, across-firm heterogeneity are qualitatively important. In the final period 2017-2022, we find a similar decomposition to the first period 1999-2004. It should, however, be noted that the most recent period includes the Covid-19 pandemic, which may skew results in hard-to-predict ways.

Figure 3 panel (b) shows that the degree of sorting between workers and jobs has almost doubled over time, and that all of this increase can be attributable to the greater sorting of workers to occupations. While we confirm that sorting between workers and firms within occupations has also increased over this period (as previously documented in Card et al. (2013)), it is of a quantitatively unimportant magnitude. This increase in occupation-level

[10] In 2010-11, a new occupational classification was introduced in Germany, inducing an abnormally large number of apparent occupational moves. We exclude those years from the analysis because occupation changes in these years will not always reflect actual occupational moves.

sorting, coupled with the increase in the segregation of occupations across firms documented in section 3.4 explain the increases in sorting. This provides evidence in line with the recent literature that has attributed increased between-firm inequality to changes in the boundaries of the firm (Bilal and Lhuillier, 2021; Handwerker, 2023; Cortes et al., 2024).

Figure 3: Changes in variance of job fixed effects and sorting from 1999-2022 in Germany

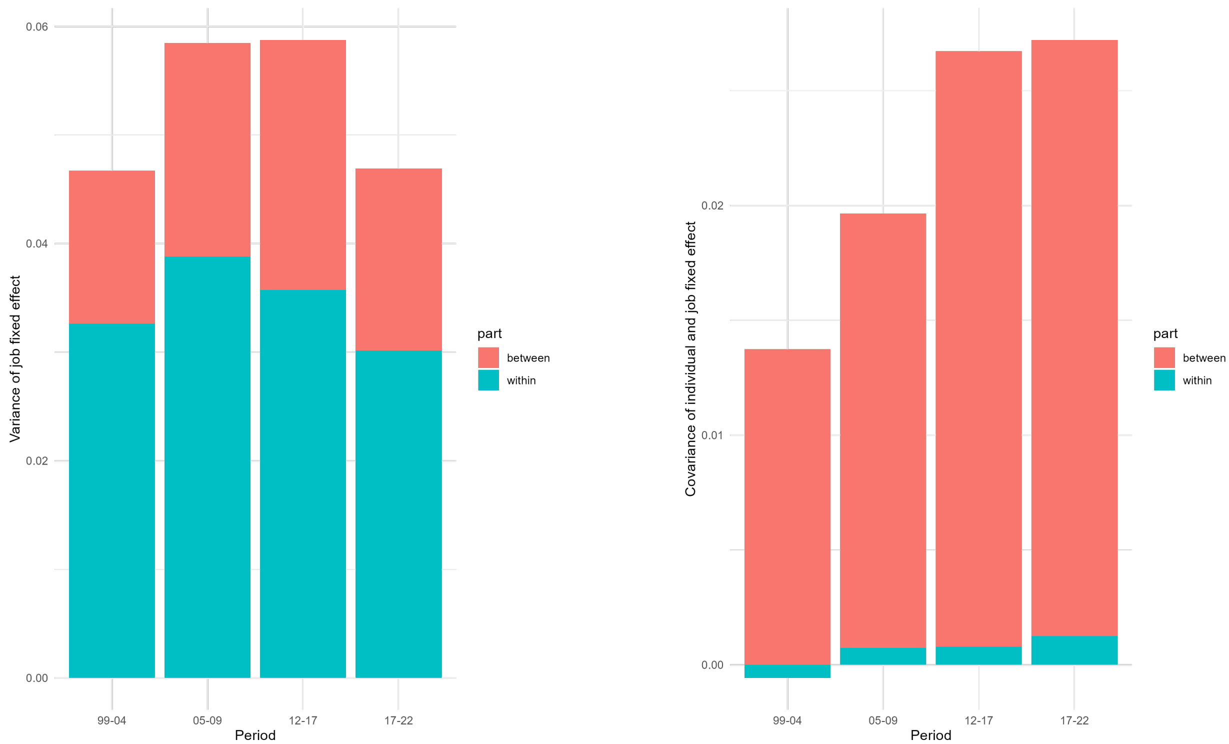

(a) Variance of job fixed effects

(b) Covariance of individual and job fixed effects

Notes: These figures plot the variance of job fixed effects (in panel 3a) and the covariance of individual and job fixed effects (in panel 3b. We decompose these terms into between-occupation (orange) and within-occupation (blue) terms as described in the text. We plot these statistics for four time periods in Germany, 1999-2004, 2005-2009, 2012-2017, 2017-2022. The years 2010 and 2011 are excluded because of an occupational change which led to abnormally large numbers of occupational moves.

4.2 Diagnostics and robustness checks

In this section, we discuss diagnostic tests of the econometric assumptions underlying our interpretations in section 4 and probe the robustness of this analysis to various specification changes. In addition to the analysis described in this sub-section, we find that the qualitative takeaways of our main decomposition result are also robust to using log hourly wages instead of annual wages, as well as using a sample of only women. These results are available on request.

4.2.1 Exogeneity of moves

As described in section 2, to credibly interpret the job fixed effects as pay premia, we require exogenous moves of workers to jobs conditional on the estimated fixed effects. We follow the logic of Card et al. (2013) by considering wage changes around job changes. If the econometric model specified is correct, then when a worker moves from a job with a high wage premium to a job with a low wage premium, we should see a step-wise relative change in their earnings that is roughly equal to the negative of an analogous move from the low-wage premium job to the high wage premium job. On the other hand, if there were match effects or temporary wage effects not captured by the fixed effects, then we should expect such moves between jobs not to lead to a symmetric effect on earnings. Card et al. (2013) produces a diagnostic for the identification assumption based on this idea as follows: they first cluster firms into four clusters by their wage premia, and then study wage changes when workers move between firms in these clusters. They argue that if the identification assumption holds, moves between firm clusters should produce relatively symmetric wage changes.

We implement these event-study "tests" by categorising the jobs that individuals move into or out of into four earnings quartiles using the leave-out job-specific mean fixed effect. Figure 4 then plots how average wages move around the event of a job switch between each of the four categories. Focusing on the most extreme moves from the first to the fourth quartile and from the fourth to the first, there is a clear symmetry in the impact. The other cells, although less extreme, show a similar symmetry. We produce an analogous diagram categorising jobs using average job wages in figure 6 in the appendix, which shows similar results. Symmetric wage changes when moving between different quality jobs, stable wages when moving between similar quality jobs, and stability in the years around a move all give credence to the key identifying assumptions and follow a pattern of similar results in the worker-firm literature Bonhomme et al. (2023); Card et al. (2013, 2018).

Finally, one might be worried that the averages presented in the event study might obfuscate different degrees of exogeneity between different types of moves; e.g. moves across firms may be conditionally exogenous as is accepted in the AKM literature, while moves across occupations within a firm may not be, particularly if firms have more private information about workers which allow them to achieve significant match effects. We therefore plot a version of the event study using only moves within firms across occupations in figure 7 in the appendix. Even for these moves, we still observe the symmetric wage gains and losses implied by the two-way fixed effect model when workers move across occupations in different fixed effect quartiles.

Figure 4: Event study around job moves, clustering by leave-out job mean fixed effect

(a) France

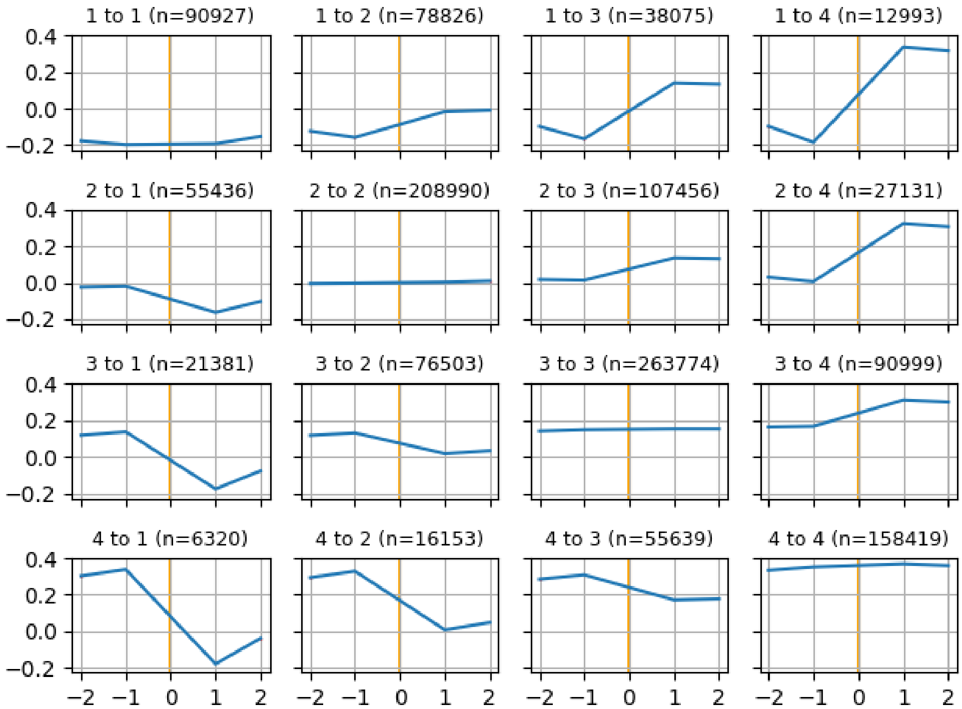

(b) Germany

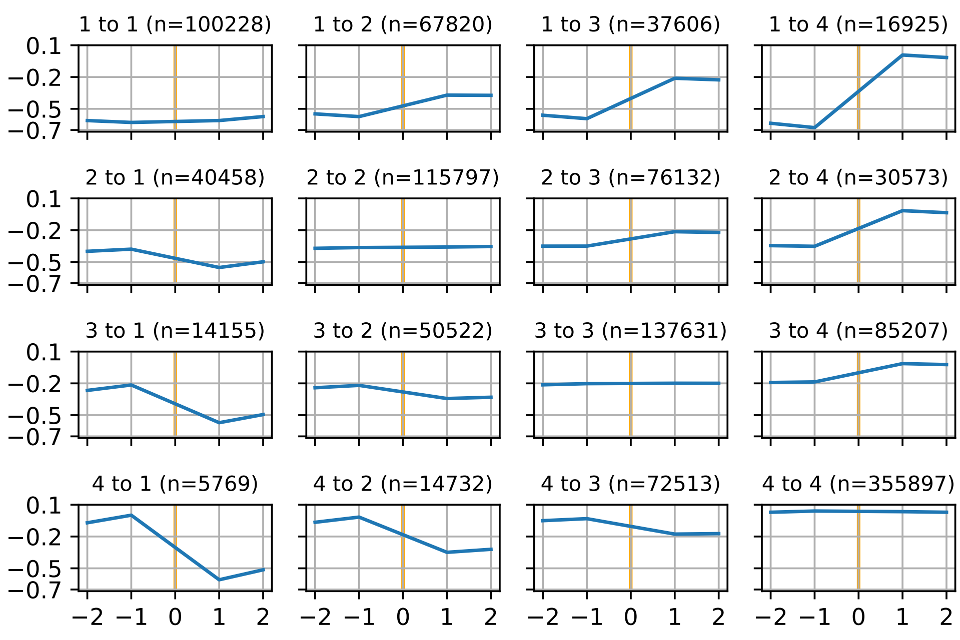

Notes: These figures show the impact on average wages around the event of job movement. Each cell shows the average wage change associated with a movement event from one quartile to another quartile of the average job fixed effect distribution. Following Card et al. (2013), we cluster jobs into quartiles by computing the mean leave-out job fixed effect within the job, excluding the own firm. Only those who remain in their old job for two years before and their new jobs for two years after the move event are included. The number of switchers in each cell is given in the cell title. Panel (a) shows the results for France, and panel (b) shows the results for Germany.

4.2.2 Linearity of worker and job effects

A possible criticism of the approach we take is that the two-way fixed effect regression imposes an inappropriate log-additive functional form on worker and job fixed effects. The log-additive structure we impose on the data could mask important heterogeneity, mechanisms, or disallow potentially relevant theoretical channels. To check for this, we again follow the approach in Card et al. (2013) and show in figure 10 in the appendix that the mean of the residuals of the two-way FE regression are near zero on average, and by job and worker cells. We find that the specification performs well in this test.

4.2.3 Performance relative to additive firm and occupation fixed effects

Another method used in the literature to incorporate occupations into an AKM framework is to add occupation fixed effects additively to a model with worker and firm fixed effects (Torres et al., 2018). Relative to our preferred specification, this approach extrapolates the pay premium for a given job as the simple sum of its firm and occupation components, thereby ignoring potential match-specific idiosyncratic returns. While the additive specification is more parsimonious and significantly less data-intensive—requiring only the estimation of $|\mathcal{F}| + |\mathcal{O}|$ parameters rather than the up to $|\mathcal{F} \times \mathcal{O}|$ coefficients in our flexible job-fixed-effect model—this efficiency comes at a structural cost. Figure 5 uses French data and plots the difference between estimated job fixed effects (λ_j) from our preferred worker-job model and the sum of estimated firm (ψ_f) and occupation (γ_o) fixed effects from a log-additive worker-firm-occupation model.

Figure 5: Contribution of worker-job match effects by firm and occupation mean pay quantiles

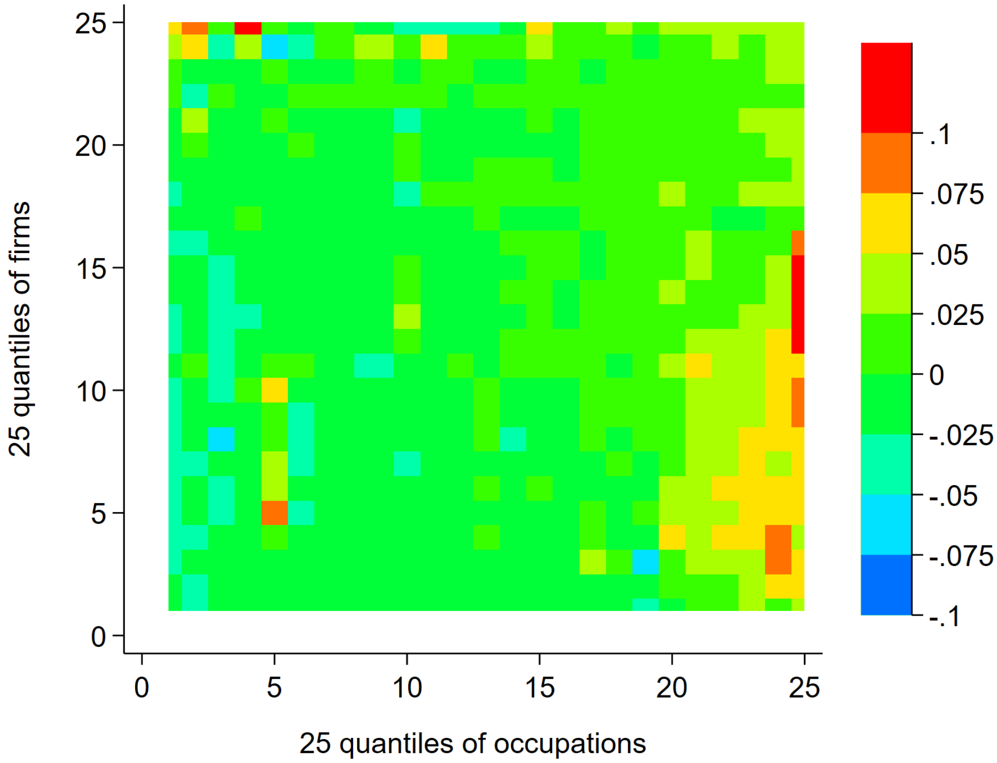

Notes: The figure shows a heatmap of the difference between the estimated job fixed effect in our main specification, and the sum of a firm fixed effect and an occupation fixed effect in an auxiliary specification, which we interpret as match effects between occupation and firms. Green squares imply that if the job fixed effect were replaced by the additive firm and occ fixed effects, the resulting predicted log wage would be within 0.05 log points of the true predicted log wage with job FE. Red implies that using the additive specification would underpredict the true wage by over 0.1 log points and blue implies that overprediction by over 0.075 log points. The x-axis indexes 25 quantiles of occupations, sorted by the mean wage within the occupation, and the y-axis indexes 25 quantiles of firms, sorted by the mean wage within the firm.

The heatmap demonstrates that while the additive model serves as a strong approximation for the majority of firm-occupation pairs (indicated by the predominant green cells within ± 0.05 log points), it systematically fails to capture significant match effects at the extremes of the distribution. These deviations, which represent cases where specific firms pay a premium or discount to specific occupations above their baseline rates, would be misattributed in a more parsimonious model. In particular, wages for high-pay occupations at low-pay firms seem to be systematically higher than an additive firm-occupation would suggest. By utilizing a flexible job fixed effect, we ensure that these match-specific components do not confound our estimates of worker-firm and worker-occupation sorting.

4.2.4 Robustness to coarser occupation categories

Another concern is that occupations could be measured with error in our data. First, if we use very fine-grained occupational codes, it could be difficult for firms to consistently fill in workers' occupations accurately. Apparent segregation across firms could reflect differences in reporting standards and practices across firms and not actual differences in the work done. To check for this, we reproduce our main results using coarser occupation definitions; whereas at the 4- or 5-digit level mistakes in occupation classification may be made, this seems more improbable the more aggregated occupation groups become.

Details of results and decompositions for one-, two-, three-, and four-digit occupation classifications are given in the online appendix. Figure 8 compares the decomposition of the variance of job fixed effects into between- and within-occupation components while figure 9 compares the similar decomposition of the covariance of worker and job fixed effects. Naturally, as the occupation classification becomes coarser, across-occupation differences will mechanically explain less variation, even in the absence of measurement error. However, we find that even at the one-digit occupation level in the French setting, 63% of the wage inequality that can be attributed to sorting between workers and jobs is due to sorting across occupations, and 14% of the wage inequality that can be attributed to job-heterogeneity is due to across-occupation differences.

4.3 Discussion

The implications of our analysis depend on the intended interpretation of the two-way fixed effect framework. Under a descriptive interpretation, the AKM firm fixed effect answers "which firms pay more on average," while the worker-firm covariance identifies whether high-wage workers are concentrated at those high-wage firms. Our paper contributes to this view by identifying a critical confounding margin: the worker's occupation. We demonstrate that high-wage firms are frequently those that disproportionately employ high-wage occupations. Consequently, documented worker-firm sorting largely reflects the sorting of high-ability workers into high-paying occupations rather than a firm-specific matching premium within those occupations. Furthermore, we provide evidence that the observed historical rise in worker-firm covariance is driven primarily by increased worker-occupation sorting and the intensifying segregation of occupations across firm boundaries.

Beyond a descriptive interpretation, firm fixed effects have been interpreted as "[potentially representing] rent-sharing, an efficiency wage premium, or strategic wage posting behaviour" (Card et al., 2013). For example, Card et al. (2018) interprets AKM fixed effects as capturing preference for firm-level amenities. We contribute to this literature by verifying that even within fine occupation categories, there is still substantial dispersion of firm fixed effects, suggesting that variation in firm-level pay premia are not solely driven by differences in occupational composition. However, our analysis also points to the need for papers to take into account differences in occupational composition before making conclusions about the specific economic explanation for wage dispersion being studied.

Finally, our findings provide a new perspective on the literature that uses AKM covariances to identify primitives of labor market sorting. We show that the locus of sorting is more critical than previously understood. By applying the Law of Total Covariance, we demonstrate that sorting to occupations is quantitatively over four times as important as sorting to firms. In the French context, for instance, sorting across occupations accounts for 17.0% of wage variance, whereas sorting to firms within those occupations accounts for only 4.0%. This suggests that the "sorting" documented in standard AKM models is largely a result of Roy-style self-selection into tasks and occupations rather than the firm-level matching typically emphasized in search-and-matching models.[11]

A growing literature has sought to understand the relation between AKM covariances and primitives in matching models of the labour market (Eeckhout and Kircher, 2011; Hagedorn et al., 2017; Borovičková and Shimer, 2017; Lopes de Melo, 2018; Borovičková and Shimer, 2024). These papers have largely concluded that covariances from AKM models do not identify primitives from a model of sorting between workers and firms. Thus, it is inappropriate to view our results as providing definitive evidence that complementarities between workers and firms are driven by the worker's occupation. We view our paper as identifying an important locus of sorting that future attempts to recover these primitives must account for. Specifically, if worker-job complementarities identified in structural approaches arise primarily from worker-occupation matching rather than worker-firm matching, then standard firm-level models may misattribute the source of productive efficiency. By demonstrating that sorting to occupations is the dominant driver of observed wage covariance, we provide a descriptive foundation for future research to explore whether significant complementarities exist within occupations or if they are primarily a feature of occupational selection.

5 Conclusion

In this paper, we argue that the degree to which high-wage workers sort to high-wage firms has been overestimated in AKM-style decompositions. This overestimation is due to high-wage occupations clustering in high-wage firms. Not explicitly accounting for a worker's occupation, therefore, leads to sorting to high-wage occupations being mistaken for sorting

[11] It should be noted that sorting to occupations in our main empirical analysis captures both sorting across sets of broad skills (programmer vs butcher vs mechanic, for example) and vertical occupation classifications within skill sets. In robustness analysis, we consider coarser occupation classifications that capture more of the former notion and find our qualitative conclusions are unchanged.

to high-wage firms.

We extend the standard AKM model by estimating worker-job two-way fixed effects instead of worker-firm two-way fixed effects, where jobs are occupation-firm pairs. We show using event studies around job changes that wages experience step-changes consistent with the fixed effects model when they move from high-wage jobs to low-wage jobs. We account for limited-mobility bias using the leave-one-out variance estimator due to Kline et al. (2020), and demonstrate robustness of our core results to using coarser occupation codes, considering different time periods, and considering different data definitions.

We show that quantitatively, sorting of workers to occupations accounts for far more of the total log wage variance than sorting of workers to firms within occupations. Estimates of worker-firm sorting from standard AKM models are substantially higher than estimates of sorting of workers to firms within occupations in our model, suggesting that much of what was previously considered worker-firm sorting may have been worker-occupation sorting instead. Second, we show that even after accounting for occupations, there is considerable variation of firm wage premia within occupations.

Some authors have interpreted the firm premia found in the previous literature has evidence of firm market power arising from workers' heterogeneous preferences for firms (Card et al., 2018). In this sense, authors have argued that pay policies like minimum wages could reduce overall wage inequality (Alvarez et al., 2018). Our results suggest that another piece of the puzzle is occupational sorting. A compelling direction for policy may be to understand to what extent pay premia within occupations reflect differences in the market value for tasks (Acemoglu and Autor, 2011) or rent-seeking, e.g. through occupational licensing. Furthermore, it would be interesting to understand more fully whether the identified worker-occupation covariances are driven by productivity (i.e. more productive workers sorting to more demanding and highly compensated jobs) or rent-seeking (i.e. more capable people being better away to compete for highly compensated occupations protected by regulatory barriers).

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.

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.

Alvarez, J., F. Benguria, N. Engbom, and C. Moser (2018). Firms and the decline in earnings inequality in brazil. *American Economic Journal: Macroeconomics 10*(1), 149–189.

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.

Azkarate-Askasua, M. and M. Zerecero (2024). No more limited mobility bias: Exploring the heterogeneity of labor markets. *Working Paper*.

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 (2017). High wage workers work for high wage firms. Technical report, National Bureau of Economic Research.

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

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.

Carrillo-Tudela, C. and L. Visschers (2023). Unemployment and endogenous reallocation over the business cycle. *Econometrica 91*(3), 1119–1153.

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.

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

Goldschmidt, D. and J. F. Schmieder (2017b, 04). 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.

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.

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., 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.

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.

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

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.

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.

Postel-Vinay, F. and J.-M. Robin (2002). Equilibrium wage dispersion with worker and employer heterogeneity. *Econometrica 70*(6), 2295–2350.

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

Shimer, R. (2005). The assignment of workers to jobs in an economy with coordination frictions. *Journal of political Economy 113*(5), 996–1025.

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.

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.

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

A Summary Statistics for Other Periods

A.1 France

Table 2 reports descriptive statistics for the French matched employer–employee data over the 2010–2014 period, constructed using the same data-processing pipeline, sample restrictions, and variable definitions as in the main French analysis. The table is presented across five alternative estimation samples that differ only by the connectedness requirements implied by the fixed-effect structure. Specifically, the first column summarizes the restricted full sample. The second and third columns summarize, respectively, the largest connected set and the largest leave-one-out connected set (LOO) under the worker–firm specification. The fourth and fifth columns summarize, respectively, the largest connected set and the largest leave-one-out connected set under the worker–job specification, where jobs are defined at the firm–occupation level, consistently with the main text.

Within each sample column, Table 2 reports the number of observations, the number of unique workers, firms, and jobs, and a set of wage moments. Wage moments include the mean and variance of log annual wages and log hourly wages, together with the variance of residualised log hourly wages, where residualisation follows the control structure described in the main text. The table additionally reports mobility counts, including the number of observed job-to-job moves and their decomposition into firm changes, occupation changes, and moves involving simultaneous changes in both firm and occupation. The table notes indicate that individuals are linked across yearly cross-sections using the same chaining procedure as in the main French panel construction.

Table 2: Summary statistics, 2010-14

	Full data	Firms connected set	Firms LOO set	Jobs connected set	Jobs LOO set
2010-14					
N obs	46,103,347	39,467,309	37,286,700	33,154,058	28,775,013
N workers	12,664,742	9,400,612	8,890,944	7,934,805	7,034,652
N firms	1,146,870	580,256	329,961	502,751	222,969
N jobs	4,989,214	3,561,234	2,921,078	2,209,007	821,744
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 residualised log hourly wage	0.19	0.19	0.19	0.19	0.18
Var log hourly wage	0.23	0.22	0.22	0.21	0.21
N moves	8,503,206	2,954,239	2,694,455	6,003,822	4,562,592
N firm moves	3,131,395	2,954,239	2,694,455	2,760,518	2,116,072
N occ moves	7,199,647	2,660,426	2,395,688	4,855,056	3,569,830
N firm + occ moves	1,827,836	1,735,055	1,550,929	1,611,752	1,123,310

Notes: This table shows the summary statistics for data from 2010 to 2014, cleaned in the same way as our main sample, which we use as a robustness check. The underlying data is from yearly BTS data files, French administrative matched employee-employer data. Individuals are mapped over time using the procedure and kindly provided programs in Babet et al. (2025). Sample construction and restrictions are discussed in the main text.

A.2 Germany

Table 3 presents descriptive statistics for the West German sample across four multi-year windows—1999–2004, 2004–2009, 2012–2017, and 2017–2022—using the same unit definitions and sample restrictions as in the main German analysis. For each period, the table reports statistics for three nested estimation samples: the restricted full sample, the largest connected set under the main worker–job structure, and the largest leave-one-out observation connected set (LOO) under the same worker–job structure. The LOO sample is the one required for the leave-one-out variance-component corrections implemented in the main decomposition analysis.

For each period-by-sample cell, Table 3 reports the number of observations and the number of unique workers, firms/establishments, and jobs. It then reports wage moments for daily wages (including the average and standard deviation), as well as the standard deviation of residualised daily wages, where residualisation is performed using year fixed effects and a cubic age profile, as described in the paper. Finally, the table summarizes mobility by reporting the total number of job-to-job moves and the decomposition of moves into firm moves, occupation moves, and moves involving both firm and occupation changes.

Table 3: Descriptive statistics for entire West German sample, in four periods from 1999-2022

	1999-2004	2004-09	2012-17	2017-22
Total observations (m)	60.34	57.78	62.47	54.56
Total workers (m)	13.62	12.81	14.03	14.22
Total firms (m)	1.56	1.49	1.40	1.35
Total jobs (m)	4.03	3.78	4.79	4.91
Average daily wages	4.77	4.74	4.75	4.78
SD daily wages	0.49	0.53	0.53	0.52
SD resid. daily wages	0.48	0.51	0.51	0.50
N moves	6.51	5.34	6.16	5.79
N firm moves	5.74	4.76	5.37	4.88
N occ moves	3.20	2.58	3.85	3.77
N firm + occ moves	2.43	2.00	3.06	2.86
Largest connected set				
Total observations (m)	48.69	44.81	47.24	40.16
Total workers (m)	9.89	8.98	9.56	9.44
Total firms (m)	0.99	0.87	0.85	0.78
Total jobs (m)	2.24	1.91	2.41	2.30
Average daily wages	4.80	4.78	4.77	4.80
SD daily wages	0.47	0.51	0.52	0.51
SD resid. daily wages	0.46	0.49	0.50	0.49
N moves	5.65	4.62	5.37	5.00
N firm moves	5.08	4.18	4.71	4.27
N occ moves	2.79	2.26	3.42	3.29
N firm + occ moves	2.23	1.82	2.76	2.56
Largest leave-out observation connected set				
Total observations (m)	42.22	38.06	39.32	32.96
Total workers (m)	8.73	7.76	8.17	7.95
Total firms (m)	0.54	0.45	0.44	0.39
Total jobs (m)	1.05	0.86	1.00	0.92
Average daily wages	4.81	4.81	4.79	4.82
SD daily wages	0.46	0.50	0.52	0.51
SD resid. daily wages	0.45	0.49	0.50	0.49
N moves	4.43	3.53	3.94	3.58
N firm moves	3.98	3.18	3.40	3.01
N occ moves	2.02	1.58	2.37	2.21
N firm + occ moves	1.57	1.22	1.82	1.64

Notes: This table presents summary statistics for all four periods for three samples - the entire population, the largest connected set of the main worker-job specification, and the largest leave-out observation set of the main worker-job specification. We present the total number of observations, the total number of unique workers, the total number of unique firms and the total number of unique jobs in millions, the total number of moves, as well as the average and standard deviation of the daily (imputed) wage as well as the standard deviation of the daily wage residualised on year fixed effects and a cubic age profile.

B Variance decomposition table

This appendix section reports the full numerical variance decompositions corresponding to the graphical summaries in the main text. In each country, the table presents the decomposition of residualised log wage dispersion under both a worker–firm specification and the preferred worker–job specification, and it reports the corresponding variance and covariance objects in both level and share form.

Table 4 decomposes the variance of residualised log wages in the French sample under two alternative two-way fixed-effect models. Under the firm model, residualised log wages are decomposed into worker and firm fixed-effect components and their covariance, along with the residual variance, following the standard AKM structure. Under the job model, the firm effect is replaced by a job fixed effect defined at the firm–occupation level, and the decomposition is reported in the same format. For each component, the table reports the variance contribution in levels and the corresponding share of total residualised log-wage variance.

In addition, Table 4 implements the within- versus between-occupation decompositions used in the main text. In the job model, the variance of job fixed effects is decomposed by the law of total variance into a between-occupation component and a within-occupation component. Similarly, the worker–job covariance ("sorting") term is decomposed by the law of total covariance into between-occupation and within-occupation contributions, using the same conditioning and notation as in the main text.

Table 4: Variance decomposition: France

Component		Firm model			Job model		
		Variance	Proportion of total variance	Proportion of component variance	Variance	Proportion of total variance	Proportion of component variance
Worker	Total	0.118	0.574	1	0.098	0.493	1
Job/ Firm	Total	0.013	0.064	1	0.021	0.106	1
	Within Occupation				0.011	0.06	0.52
	Between Occupation				0.010	0.05	0.48
Sorting	Total	0.019	0.091	1	0.042	0.210	1
	Within Occupation				0.008	0.04	0.19
	Between Occupation				0.034	0.17	0.81
Error & controls	Total	0.056	0.271	1	0.038	0.192	1

Notes: This table decomposes log wage variance in the French sample using the worker-job model described in the main text. It also shows the results from applying the law of total variance and law of total covariance to further decompose job-level components into occupation and firm effects.

Table 5 reports the analogous set of variance and covariance decompositions for the German sample. As in Table 4, results are reported for both the firm model and the job model, with components displayed in levels and as shares of total residualised log-wage variance. For the job model, Table 5 further reports the decomposition of job fixed-effect variance into within-occupation and between-occupation components via the law of total variance, and it reports the decomposition of the worker–job covariance term into within-occupation and between-occupation components via the law of total covariance, mirroring the structure and definitions used in the main text.

Table 5: Variance decomposition: Germany

Component		Firm model			Job model		
		Variance	Proportion of total variance	Proportion of component variance	Variance	Proportion of total variance	Proportion of component variance
Worker	Total	0.130	0.525	1	0.106	0.416	1
Job/ Firm	Total	0.027	0.107	1	0.047	0.184	1
	Within Occupation				0.030	0.12	0.64
	Between Occupation				0.017	0.07	0.36
Sorting	Total	0.035	0.143	1	0.054	0.213	1
	Within Occupation				0.002	0.05	0.01
	Between Occupation				0.051	0.20	0.95
Error & controls	Total	0.056	0.226	1	0.048	0.187	1

Notes: This table decomposes log wage variance in the German sample using the worker-job model described in the main text. It also shows the results from applying the law of total variance and law of total covariance to further decompose job-level components into occupation and firm effects.

C Event studies

Figure 6 presents event-study profiles of wage dynamics around job-to-job moves, organized by the wage level of the origin and destination jobs. Jobs are assigned to quartiles based on the mean leave-out job wage measure, constructed analogously to the leave-out statistics used in standard two-way fixed-effect diagnostics.

The figure is arranged as a four-by-four grid of panels, where each cell corresponds to an origin quartile and a destination quartile of this leave-out job mean wage distribution. Within each cell, the plotted series reports the mean wage outcome by event time relative to the move, over the event-time window indicated in the figure. The underlying sample is restricted to workers who remain in the origin job for two years prior to the move and remain in the destination job for two years following the move, so that pre- and post-move paths are traced over a balanced window around the transition. Each cell title reports the origin and destination quartiles and the number of switchers contributing to the cell. The figure is shown separately for France (panel (a)) and Germany (panel (b)), with an identical cell structure across countries.

Figure 6: Event study around job moves, clustering by leave-out job mean wage

(a) France

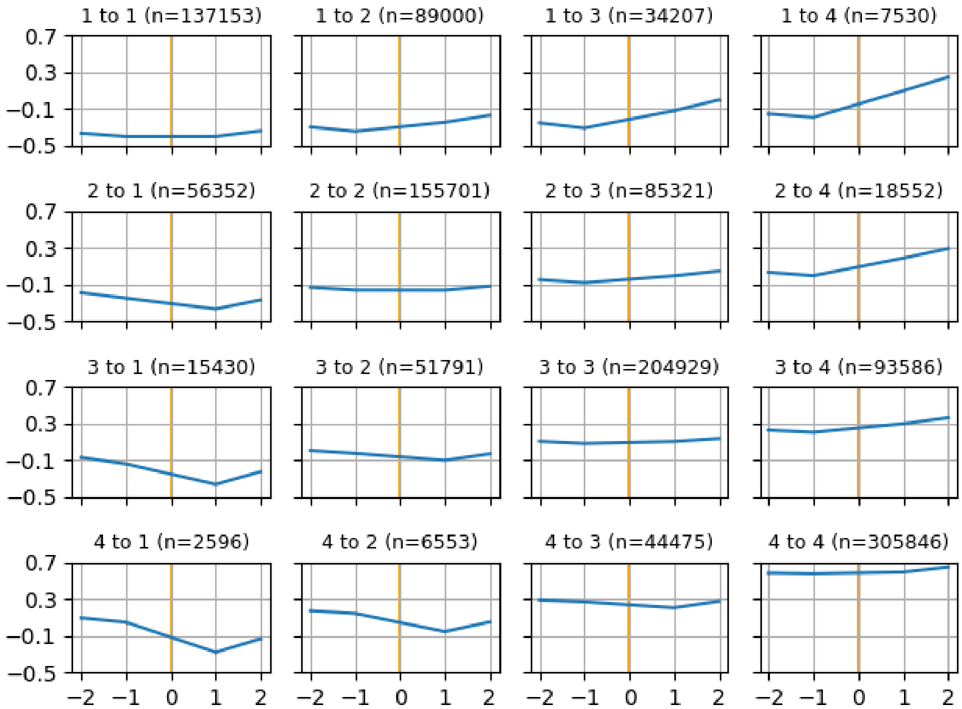

(b) Germany

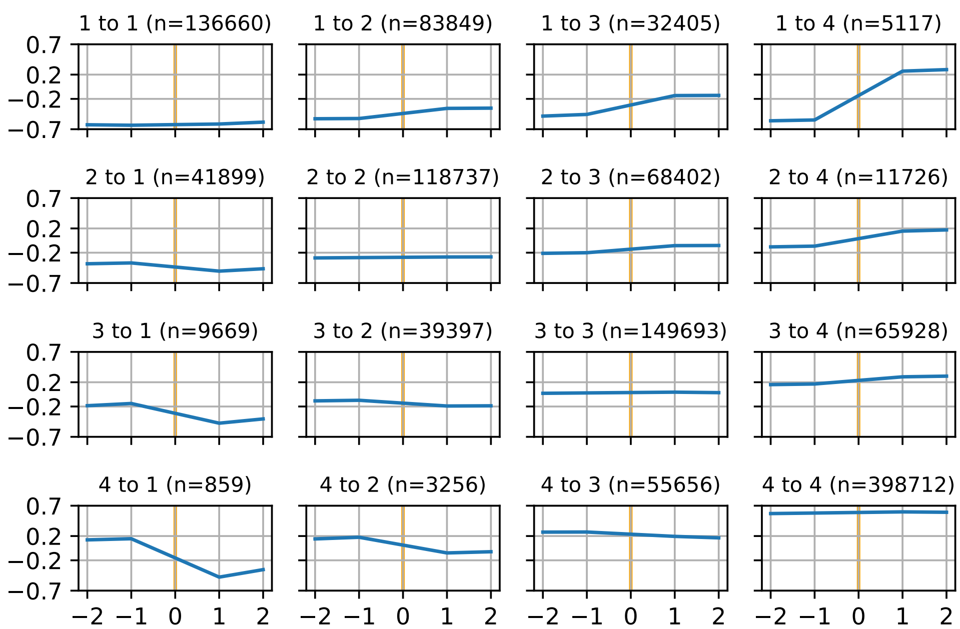

Notes: These figures show the impact on average wages around the event of job movement. Each cell shows the average wage change associated with a movement event from one quartile to another quartile of the average job wage distribution. Following Card et al. (2013), we cluster jobs into quartiles by computing the mean leave-out job wage. Only those who remain in their old job for two years before and their new jobs for two years after the move event are included. The number of switchers in each cell is given in the cell title. Panel (a) shows the results for France, and panel (b) shows the results for Germany.

Figure 7 reproduces the event-study design while restricting attention to moves that occur within firms and involve an occupation change, i.e., transitions across firm–occupation jobs holding the firm identifier fixed. In this figure, quartile assignments are based on the estimated job fixed effect from the preferred worker–job two-way fixed-effect model, rather than on the leave-out job mean wage measure. The figure again uses a four-by-four grid, where each cell corresponds to an origin and destination quartile of the estimated job fixed effect distribution. For each cell, the series plots the mean wage outcome by event time around the move using the event-time window indicated in the figure, and each cell title reports the number of switchers contributing to that origin–destination cell. The layout is reported separately for France (panel (a)) and Germany (panel (b))

Figure 7: Event study around job moves, including only moves across occupations within firms, clustering by estimated job fixed effect

(a) France

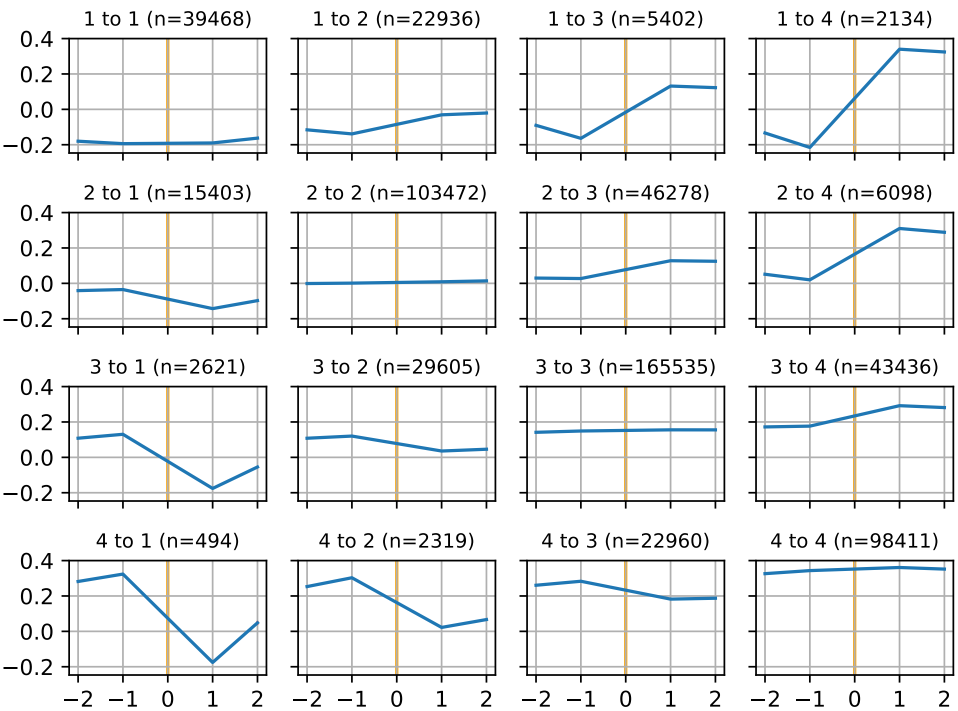

(b) Germany

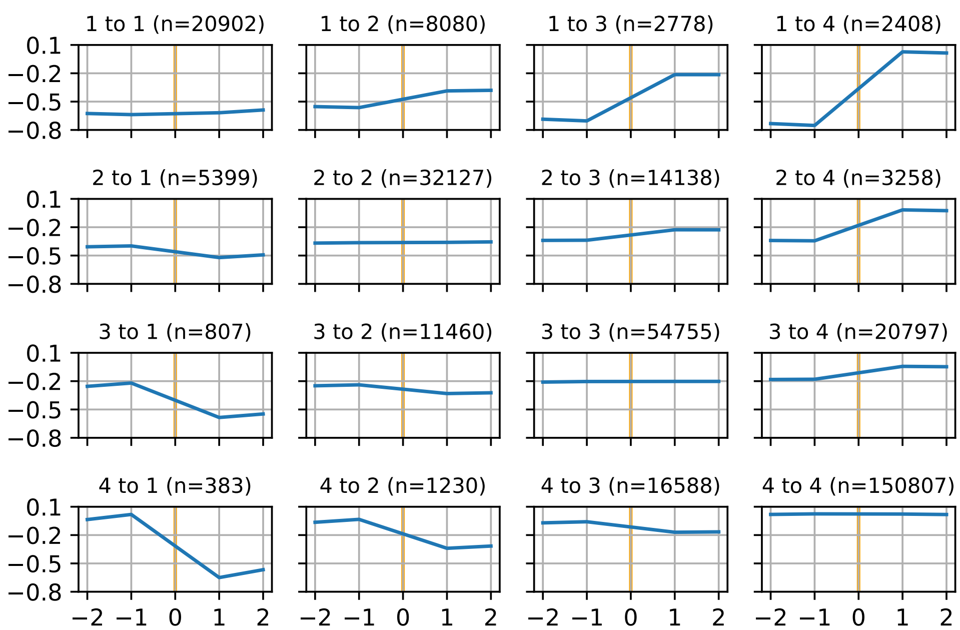

Notes: These figures show the impact on average wages around the event of job movement across occupations within firms. Each cell shows the average wage change associated with a movement event from one quartile to another quartile of the average job fixed effect distribution. Following Card et al. (2013), we cluster jobs into quartiles by computing the mean leave-out job fixed effect within the job excluding. Only those who remain in their old job for two years before and their new jobs for two years after the move event are included. The number of switchers in each cell is given in the cell title. Panel (a) shows the results for France, and panel (b) shows the results for Germany.

D Granularity of occupation coding

Figure 8 reports how the variance of estimated job fixed effects is partitioned into within-occupation and between-occupation components under alternative occupation-code granularities. For each occupation definition considered, the figure displays a bar corresponding to the total variance of the estimated job fixed effect and decomposes that variance using the law of total variance into a between-occupation component and a within-occupation component, consistent with the decomposition in the main text. The occupation-code granularity associated with each decomposition is indicated in the labels placed adjacent to the corresponding bar, allowing the figure to be read as a sequence of decompositions across increasingly coarse occupation classifications.

Figure 8: Decomposing job variance into within and between occupation components by different levels of occupation granularity

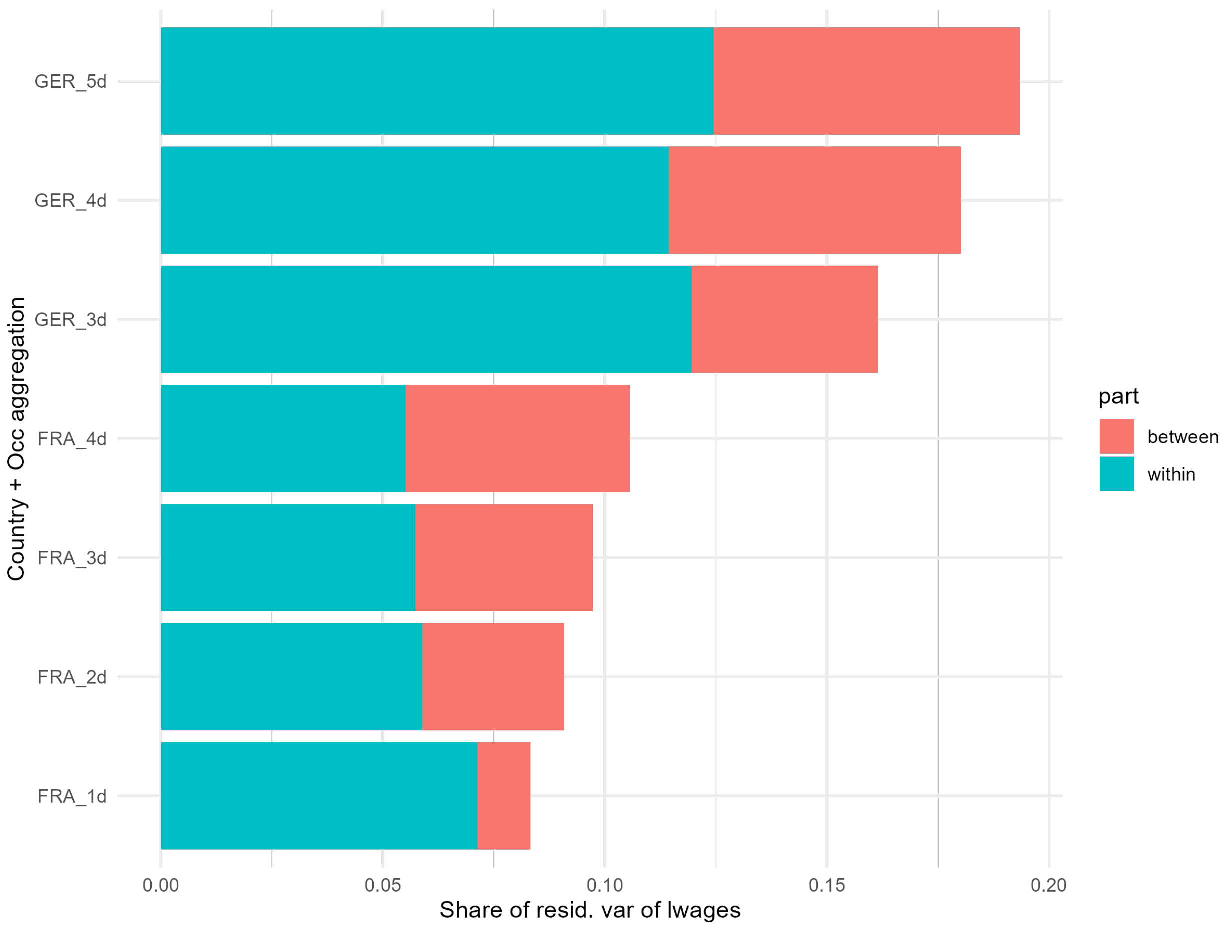

Notes: This figure shows the decomposition of job variance into between and within occupation components for various definitions of occupation at different granularities. The dimension of granularity is indicated on the left-hand side of each bar.

Figure 9 provides the analogous set of decompositions for the covariance between worker and job fixed effects, i.e., the worker–job sorting term. For each occupation-code granularity, the figure decomposes the worker–job covariance using the law of total covariance into a between-occupation component—capturing covariance in occupation-level conditional means—and a within-occupation component—capturing the expected covariance conditional on occupation. The granularity used for each decomposition is reported in the labels aligned with the corresponding bar, so that the figure can be read as documenting the within- versus between-occupation partition of the sorting covariance across increasingly coarse occupation classifications.

Figure 9: Decomposing worker-job covariance into within and between occupation components by different levels of occupation granularity

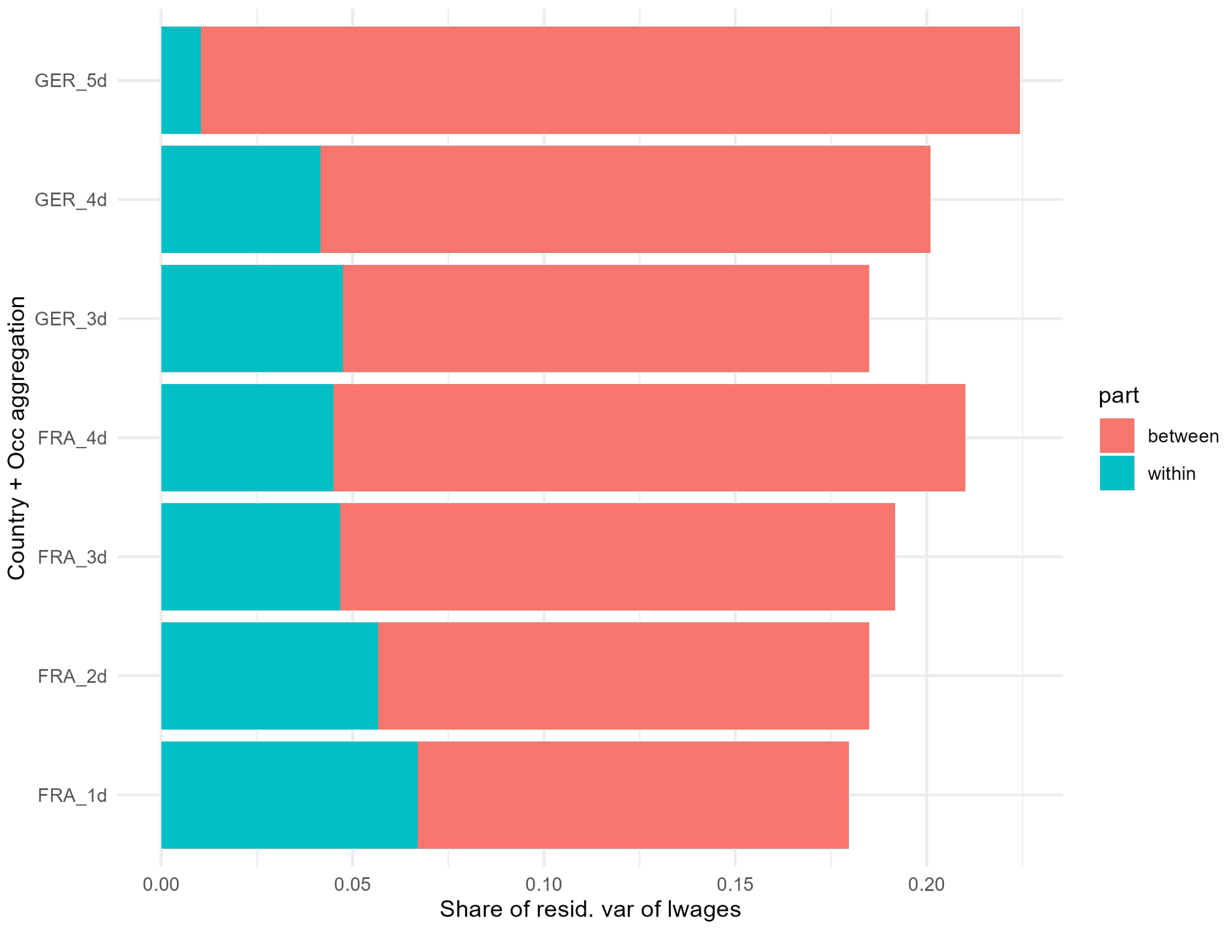

Notes: This figure shows the decomposition of worker-job covariance into between and within occupation components for various definitions of occupation at different granularities. The dimension of granularity is indicated on the left-hand side of each bar.

E Linearity of the worker-job two-way fixed effect model

Figure 10 shows a heatmap of the estimated residual of the two-way fixed effect model averaged within each of 625 cells defined by 25 worker fixed effect quantiles and 25 job fixed effect quantiles. Following the diagnostic test proposed by Card et al. (2013), this plot assesses the validity of the log-additive separability assumption in the AKM framework. Under this assumption, the match effect is zero, meaning the expected value of the residual term should be zero, conditional on the worker and job fixed effects. Therefore, if the additive model is a good approximation, the average residuals within each cell should be close to zero and exhibit no systematic pattern (e.g., higher values for high-worker/high-job combinations). The predominance of green across the heatmap indicates that the mean residuals are indeed small in magnitude and lack a distinct pattern across the distribution of worker and job effects, providing support for the additive specification in this setting.

Figure 10: Residuals from the worker-job model by worker and job fixed effect quantiles

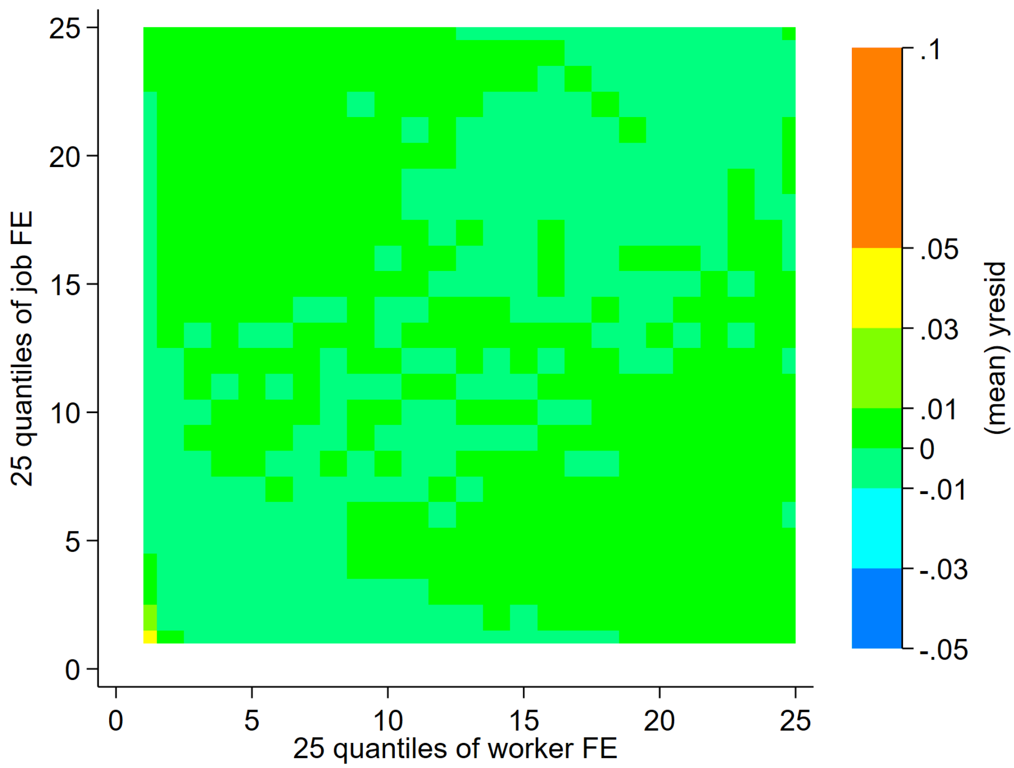

Notes: The heatmap displays the mean residuals from the two-way worker-job fixed effect model (Equation 1), aggregated into a 25 × 25 grid of worker and job fixed effect quantiles.