

Discussion Paper Series

IZA DP No. 18502

March 2026

Unpacking the Wage Sorting Trend

Jesper Bagger

University of Edinburgh, Aarhus University,
IZA@LISER and RFBerlin

Malthe Elholm

Aarhus University

Jonas Maibom

Aarhus University, IZA@LISER and RFBerlin

Rune Majlund Vejlin

Aarhus University, IZA@LISER and RFBerlin

The IZA Discussion Paper Series (ISSN: 2365-9793) ("Series") is the primary platform for disseminating research produced within the framework of the IZA@LISER Network, an unincorporated international network of labour economists coordinated by the Luxembourg Institute of Socio-Economic Research (LISER). The Series is operated by LISER, a Luxembourg public establishment (établissement public) registered with the Luxembourg Business Registers under number J57, with its registered office at 11, Porte des Sciences, 4366 Esch-sur-Alzette, Grand Duchy of Luxembourg.

Any opinions expressed in this Series are solely those of the author(s). LISER accepts no responsibility or liability for the content of the contributions published herein. LISER adheres to the European Code of Conduct for Research Integrity. Contributions published in this Series present preliminary work intended to foster academic debate. They may be revised, are not definitive, and should be cited accordingly. Copyright remains with the author(s) unless otherwise indicated.



Unpacking the Wage Sorting Trend*

Abstract

Using 1980–2019 Danish matched employer–employee data, we unpack the rise in wage sorting — the correlation between worker and firm wage fixed effects (Abowd et al., 1999) — from 0.06 to 0.18. The rise is driven entirely by reallocation of employment from persistently low-sorting to persistently high-sorting firms, with the average sorting contribution of any given firm remaining stable over time. A decomposition shows that 60% reflects reallocation among surviving firms and 40% firm turnover through entry and exit. Regression analysis identifies firm entry and exit and industry reallocation as the dominant firm-side drivers, and rising educational attainment as the key workerside factor — reflecting concentration of educated workers in high-sorting firms rather than a systematic tendency of educated workers to form high-sorting matches across all employers. Event studies establish direct job-to-job moves as the primary mechanism through which reallocation is implemented at the worker-level

JEL classification

E24, J21, J31

Keywords

wage inequality, wage sorting, firm dynamics, employment reallocation, job-to-job mobility, matched employer–employee data

Corresponding author

Rune Majlund Vejlin

rvejlin@econ.au.dk

* We thank Michele Battisti, Rasmus Lentz, Espen Moen, Mikkel Sølvsten, and seminar, workshop and conference participants for useful discussions and comments. Our data access is through the ECONAU project database (Department of Economics and Business Economics, Aarhus University). The paper was supported by the Danish Research Council grant 9038-00004B on “Sorting in the Labour Market.”

1 Introduction

Matched employer-employee data consistently show that much of the rise in wage inequality reflects growing between-firm wage dispersion (Hoffmann et al., 2020; Card et al., 2013; Song et al., 2019). A central driver of the increase in between-firm dispersion is wage sorting: the tendency of high-wage workers to cluster in high-wage firms has strengthened markedly over the past four decades (Card et al., 2013; Song et al., 2019; Lentz et al., 2023). Yet the margins through which wage sorting has strengthened remain poorly understood. Has the distribution of workers changed? Has the distribution of firms shifted? Or has employment reallocated toward certain types of firms through differential growth, entry and exit? Through which labor-market transitions—direct job-to-job moves or unemployment spells—has this reallocation occurred? Identifying these margins is essential for understanding the sources of rising inequality and for disciplining models of firm heterogeneity and firm dynamics.

Using matched employer-employee data for the Danish business sector, we document a marked increase in wage sorting from 1980 to 2019: the correlation between worker and firm wage fixed effects, estimated using a Abowd et al. (1999) two-way fixed effect regression, rises from 0.06 to 0.18. The rise is driven entirely by reallocation of employment from persistently low-sorting to persistently high-sorting firms—60 percent through reallocation among surviving firms and 40 percent through firm turnover. Firm entry and exit, industry reallocation, and rising educational attainment account for most of the trend. Direct job-to-job moves are the key facilitating mechanism through which this reallocation is implemented at the worker level.

Establishing these results requires a framework that connects aggregate wage sorting—a cross-sectional covariance—to worker-firm match-level wage sorting contributions that can be studied across firms, workers, demographic groups, and labor market transitions. We exploit the fact that the cross-sectional wage sorting covariance decomposes additively into match-level wage sorting contributions, defined as the product of worker and firm wage types centered at their cross-sectional means. The additive structure allows us to apply three complementary methods to analyze the wage sorting trend. First, decomposition tools from the productivity growth literature separate the contributions of firm entry, firm exit, and reallocation among surviving firms. Second, we project the match-level wage sorting contribution onto firm and worker characteristics in a regression framework—controlling simultaneously for observed characteristics such as industry and educational attainment, but also unobserved permanent worker and firm heterogeneity through a two-way fixed

effects specification—thereby isolating the independent contribution of each factor to the sorting trend. Third, event-study designs exploit the match-level wage sorting contributions to address a complementary question: which labor market transitions—direct job-to-job moves or job changes with an intervening unemployment spell—are actually implementing the reallocation at the worker level.¹

The rise in wage sorting could in principle reflect two distinct forces: changes in the types of workers employed by a given firm over time, or reallocation of employment across firms with different sorting levels. We provide the first decomposition of the wage sorting trend into these channels, which we refer to as composition (of employment within a firm) and reallocation (of employment across firms), further separating the reallocation channel into reallocation among surviving firms and firm turnover via entry and exit. We find that the composition channel is negligible—the average match-level sorting contribution of a firm remains stable over its lifetime. The 12 correlation point 1980–2019 rise in wage sorting is instead driven entirely by the reallocation channel: employment has shifted from firms that persistently form low-sorting matches toward those that persistently form high-sorting matches. Roughly 60 percent of the total increase in wage sorting reflects reallocation among surviving firms—low-sorting firms contracting and high-sorting firms expanding. The remaining 40 percent reflects firm turnover—the exit of persistently low-sorting firms and the entry of persistently high-sorting firms. These findings identify firm dynamics as the central margin through which wage sorting has increased.

Turning to the match-level wage sorting regression analysis, we examine which types of firms and workers drive the reallocation behind the rise in wage sorting. The dominant force is firm-side: compositional shifts in firm characteristics account for about 55 percent of the 12 correlation point rise in wage sorting between 1980 and 2019. Within this firm-side contribution, the most important factor is firm turnover—entering cohorts of firms form systematically higher-sorting matches than the exiting cohorts they replace, accounting for nearly 80 percent of the firm-side effect. Industry reallocation, particularly the secular decline of lower-sorting manufacturing, accounts for the remaining 20 percent. On the worker side, rising educational attainment explains most of the remaining 45 percent of the 1980–2019 rise in wage sorting. However, controlling for unobserved permanent worker and firm heterogeneity shows that this reflects not an inherent tendency of highly educated workers to sort positively, but rather their concentration in persistently high-sorting firms.

¹The worker and firm fixed effects from the [Abowd et al. \(1999\)](#) two-way fixed effect regression are well-defined statistical objects, but their structural interpretation is not straightforward and requires further assumptions ([Eckhout and Kircher, 2011](#); [Borovičková and Shimer, 2024](#)). We discuss the relationship between our approach and structural approaches in more detail below.

The event-study analysis examines the worker-level transitions through which reallocation is implemented. Direct job-to-job moves generate large and persistent increases in individual sorting contributions, and account for the bulk of the reallocation driving the aggregate trend. This effect is especially pronounced when workers join entering firms or separate from exiting ones, directly connecting firm dynamics to worker mobility. Transitions through unemployment generate smaller and less persistent sorting gains and play a more limited role.

In summary, we confirm that rising wage sorting is an important source of rising between-firm wage inequality and establish three novel empirical facts about the rise in wage sorting: wage sorting remains largely stable over time for a given firm; firm dynamics, entry, and exit are therefore the central margins facilitating the rise in wage sorting; and job-to-job mobility is an important reallocation mechanism at the worker level. These findings are important for understanding the sources of rising inequality and theories seeking to account for the rise in between-firm wage inequality should address these facts.

Related literature. A large and well-established literature, with roots in [Davis and Haltiwanger \(1991\)](#) and [Dunne et al. \(2004\)](#), documents the central role of firms in rising wage inequality ([Faggio et al., 2010](#); [Bagger et al., 2013](#); [Barth et al., 2016](#); [Mueller et al., 2017](#); [Alvarez et al., 2018](#); [Song et al., 2019](#); [Håkanson et al., 2021](#); [Morin, 2023](#)). These studies consistently show that much of the rise in wage and earnings inequality reflects growing between-firm wage variance, a finding that holds across countries and time periods. While this literature extensively documents between-firm variance and often highlights the role of sorting, it generally does not isolate the firm-level margins through which sorting has strengthened. In a highly influential study, [Card et al. \(2013\)](#), building on the worker-firm log wage variance decomposition of [Abowd et al. \(1999\)](#), establishes that rising sorting between workers and firms is a central driver of rising wage inequality.

[Card et al. \(2013\)](#) show that rising wage inequality in West Germany reflects increases in worker heterogeneity, establishment heterogeneity, and the tendency of high-wage workers to sort into high-wage establishments, each contributing in roughly equal measure. They note that newer establishments entering after the mid-1990s show increasing heterogeneity in wage premiums, but stop short of a formal decomposition of the sorting trend—a gap we address in two ways. First, we develop a match-level decomposition that connects the aggregate sorting trend to its firm-entry, firm-exit, and—among surviving firms—employment composition changes within firms and employment reallocation across firms. Second, we use a wage sorting regression analysis to identify which types of firms and workers are driving the trend.

These extensions enable a direct analysis of the wage sorting trend and tightly connects firm dynamics with increasing wage sorting.

While our decomposition framework operates at the firm level, a related set of papers documents the role of industry reallocation. At the industry level, [Haltiwanger et al. \(2024\)](#) and [Briskar et al. \(2025\)](#) document that rising between-industry earnings dispersion accounts for the bulk of the rise in overall earnings inequality in the United States and Italy respectively. [Mouton \(2024\)](#) studies the wage sorting trend in West Germany and finds that industry reallocation, particularly the decline of manufacturing, is an important driver. None of these papers, however, isolates the contributions of firm entry, exit, and survivor reallocation to the sorting trend. Our wage sorting regression analysis dissects the sorting trend into contributions from compositional changes across industries, firm age classes, firm size classes as well as entry exit cohorts, and still finds that industry reallocation, particularly the decline of manufacturing, accounts for a meaningful share of the Danish wage sorting trend—consistent with their findings.

Our analysis also sheds light on the findings of [Sorkin and Wallskog \(2021\)](#), who shows that successive cohorts of U.S. firms enter more dispersed in pay and remain so throughout their lives, with cohort effects accounting for 50–100 percent of the rise in between-firm earnings inequality, and connect this to declining business dynamism ([Decker et al., 2014, 2020](#)). Using Danish data, we show that wage sorting is stable over time for a given firm, implying that employment reallocation and firm entry and exit are likely mechanisms driving the rise in wage sorting and between-firm wage dispersion; hence, cohort effects should loom large in between-firm wage dispersion. We further quantify the role of other compositional changes across years and show that job-to-job transitions have played an important role in mediating the reallocation of employment from low- to high-sorting firms.

Structural approaches study sorting on productivity types and derive its implications for wages (e.g., [Eeckhout and Kircher, 2011](#); [Hagedorn et al., 2017](#); [Lopes de Melo, 2018](#); [Bagger and Lentz, 2019](#); [Borovičková and Shimer, 2024](#)), but a key finding of this literature is that worker and firm fixed effects from two-way fixed effects regressions do not cleanly reflect underlying productivity types; moreover, firm effects in wages may also reflect amenities ([Sorkin, 2018](#); [Taber and Vejlin, 2020](#); [Lentz et al., 2023](#); [Gola and Zhao, 2024](#)). Given these complications, and without direct data on productivity or amenities, a natural alternative is to remain agnostic about the underlying primitives and instead identify worker and firm types directly from wages and mobility patterns. [Bonhomme et al. \(2019\)](#) and [Lentz et al. \(2023\)](#) take this approach, classifying workers and firms into a small number of discrete types; this sidesteps the incidental parameter problem by reducing the parameter space to a small

number of types rather than estimating a fixed effect for each individual worker and firm. However, firm entry and exit are not a natural fit for their frameworks: [Bonhomme et al. \(2019\)](#) explicitly exclude firm entry and exit in their empirical application, noting that these events do not map naturally to their model; and [Lentz et al. \(2023\)](#) note that non-stationarity in firm group sizes remains an issue to be addressed in future work. Our paper puts firm entry and exit at the center of the analysis, exploiting the [Abowd et al. \(1999\)](#) two-way fixed effects regression to develop a match-level decomposition of the wage sorting trend that identifies the firm- and worker-level margins driving it—job-to-job moves, unemployment spells, and firm entry and exit—at the level of individual worker-firm matches.

2 Data

This section describes our matched employer-employee data (MEE) for Denmark and describes trends in wage inequality and firm and worker demographics in the business sector over the 40-year observation period from 1980 to 2019.

2.1 MEE data for the Danish business sector

The data sources. Our data are sourced from the IDA database (Integreret Database for Arbejdsmarkedsforskning). IDA is a matched employer-employee dataset that covers all legal residents in Denmark aged 15-74 during 1980–2019. The unit of observation is a person-year, and a worker’s labor market status and their employer assignment reflect their circumstances in the last week of November. Persons are identified by their (anonymized) social security number. Firms are identified by their (anonymized) business registry ID. IDA provides annual information on the average hourly wage in each last-week-of-November-job, and firm and worker characteristics. Further details are provided in [Appendix A](#).

Firm IDs. IDA identifies firms by anonymized registration numbers in Danish business registries, which introduces two potential sources of spurious firm entry and exit. First, a firm may change its registration number for organizational or administrative reasons without any material change in its economic activity. Second, IDA relies on the SE-registry (Stamregistret for Erhvervsdrivende) from 1980 to 1998 and the CV-registry (Centrale Virksomheds Register) from 1999 to 2019, with the 1999 data break causing spurious firm entry and exit events. To mitigate the impact of these issues on our analysis, we develop a stable firm ID based on the person IDs of the firms’ workforces. The procedure only counts changes

in firm IDs when accompanied by substantial changes in the underlying workforce. Further details are provided in Appendix A.

Selecting the analysis data. We first discard observations in IDA (i.e., person-years) where an individual is not employed. Second, we drop individuals when they are above the age of 60. Third, we define labor market entry to take place at age 19 or at the age an individual graduates from their highest completed education, whichever occurs last, and discard pre-entry observations. Fourth, we retain only observations on jobs in the NACE 2.0 business sector.² Fifth, we delete observations with missing information on relevant variables and trim the bottom and top percentiles of the distribution of hourly wages in each year.³ Lastly, we delete firms represented by fewer than 10 person-year observations and, for identification purposes, include only the largest connected set of workers and firms.⁴

The resulting analysis data contain 36,690,944 observations with 3,071,943 workers and 187,941 firms. Defining a “mover” as a worker who is observed to work with at least two different firms during sample period, Table 1 shows that the number of observations, firms, workers, and movers in our analysis data is stable across the four decades covered by our data. We discuss the role of job movers in Section 3.4.

We inflate wages to 2010 prices by the CPI.

2.2 Descriptive statistics

We next describe trends in wage inequality and firm and worker characteristics in the Danish business sector from 1980 to 2019.

2.2.1 Wage inequality trends

Let w_{it} be the log wage for worker i in year t , and let $\hat{r}_{it} := w_{it} - \mathbf{x}'_{it}\hat{\boldsymbol{\beta}}$ be residual log wages residualized with respect to year effects and age controls interacted with education

²The NACE 2.0 business sector contains the following industries: Mining and Quarrying; Manufacturing; Electricity, Gas and Steam; Water supply etc.; Construction; Wholesale; Transportation; Accommodation and Food Service; Information and Communication; Finance and Insurance; Real Estate; Professional and Scientific work, and Administration and Support Services.

³We drop pre-2008 wage observations that Statistics Denmark deems unreliable prior to trimming.

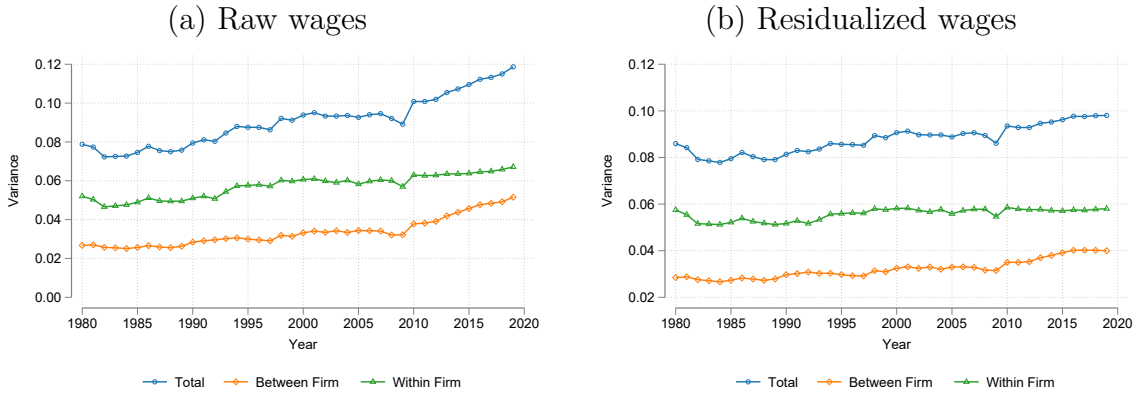
⁴Citing Abowd et al. (2002, p. 3): “When a group of persons and firms is connected, the group contains all the workers who ever worked for any of the firms in the group and all the firms at which any of the workers were ever employed. In contrast, when a group of persons and firms is not connected to a second group, no firm in the first group has ever employed a person in the second group, nor has any person in the first group ever been employed by a firm in the second group.”

Table 1: Number of observations, firms, workers, and movers in the analysis data

	Full sample	1980–1989	1990–1999	2000–2009	2010–2019
Observations	36,690,944	8,435,120	9,051,112	9,516,359	9,688,353
Firms	187,941	86,482	101,781	112,757	99,497
Workers	3,071,943	1,463,821	1,498,520	1,599,039	1,711,298
Movers	1,816,509	915,755	1,148,752	1,278,041	1,147,686

Notes: A “mover” is a worker who is observed to work with at least two different firms within the 40-year sample period. The mover definition is based on the full 1980–2019 period throughout, including for the decadal statistics. For example, a worker employed at one firm in 1980 and another in 1990 is classified as a mover and counted as such in both the 1980–1989 and 1990–1999 subpanels.

Figure 1: Total, between- and within-firm log wage variance in Denmark, 1980–2019



Notes: Residual log wage is $\hat{r}_{it} := w_{it} - \mathbf{x}'_{it}\hat{\beta}$ where \mathbf{x} includes year dummies and a cubic polynomial in worker-age interacted with education dummies, see (3) for details.

dummies.⁵ Figure 1, panel (a) shows that $\text{Var}(w_{it}|t)$ increased from 0.079 in 1980 to 0.119 in 2019, a 51 percent increase, averaging 1 percent per year.⁶ Panel (b) shows that $\text{Var}(\hat{r}_{it}|t)$ increased by 14 percent from 0.086 to 0.099 between 1980 and 2019—and, in fact, by 25 percent (from 0.078 to 0.098) between 1984 and 2016. Wage inequality in Denmark is low, but has been rising quickly.

Wages vary both within and between firms. Denote by $J(i, t)$ the firm employing worker

⁵Education indicators capture the highest completed level: primary education, upper secondary education, vocational education, short higher education, bachelor’s degree, and master’s or PhD degrees. The wage regression used for residualization also includes worker and firm fixed effects; for this descriptive analysis, however, these components are absorbed into \hat{r}_{it} . See (3) and the accompanying discussion for details.

⁶Song et al. (2019, p. 16, Figure II, panel A) reports that the variance of US log real earnings (20-60 year olds earning at least the minimum wage for one-quarter full time, employed in firms with 20+ employees) increased from 0.6 to 0.8 during 1978-2013, a 32 percent increase, averaging to 0.8 percent per year.

i in year t , and let \mathcal{I}_t be the set of workers employed in year t . The set of workers at firm j in year t is then $\mathcal{I}_{jt} = \{i \in \mathcal{I}_t : J(i, t) = j\}$. We use a tilde (\sim) to indicate sample averages of estimated quantities, with subscripts specifying the relevant group. For instance, the average log wage residual at firm j in year t is

$$\tilde{r}_{jt} = \frac{1}{|\mathcal{I}_{jt}|} \sum_{i \in \mathcal{I}_{jt}} \hat{r}_{it}, \quad (1)$$

where $|\mathcal{I}_{jt}|$ is the number of workers in \mathcal{I}_{jt} .

Using this notation, the cross-sectional variance of log wages (or residuals) can be decomposed into within- and between-firm components. Specifically, for residualized log wages:⁷

$$\mathbb{V}ar(\hat{r}_{it} | t) = \underbrace{\mathbb{V}ar(\hat{r}_{it} - \tilde{r}_{J(i,t)t} | t)}_{\text{Within-firm}} + \underbrace{\mathbb{V}ar(\tilde{r}_{J(i,t)t} | t)}_{\text{Between-firm}}. \quad (2)$$

Figure 1 shows the within- and between-firm variance of raw and residual log wages. While within-firm variance is larger overall, between-firm variance grew faster from 1980 to 2019: raw log wages saw a 0.015 increase within firms versus 0.025 between firms, and for residual log wages, within-firm variance rose by only 0.001 compared to a 0.012 increase between firms. Consequently, rising between-firm variance accounts for 62% of the increase in raw log wage variance and 90% of the increase in residual variance over 1980–2019 (68% for 1984–2016).⁸ These patterns highlight the important role of between-firm differences in Denmark’s rising wage inequality.

2.2.2 Firm and worker demographic trends

Figure 2 shows firm demographics trends in panels (a)–(d), and trends in worker demographics in panels (e)–(f) for 1980–2019.

Looking first at the firm demographic trends in panels (a)–(d), the most salient feature is the decline of the Manufacturing employment share from 38 percent in 1980 to 21 percent in 2019 that is mirrored by an expansion of service sector employment, in particular in Professional Services (from 4 percent 1980 to 9 percent in 2019), Administrative and Support Service Activities (from 1 to 7 percent), and Information and Communication (from 2 to 7 percent). Average firm size exhibits a negative trend over much of the observation period. On

⁷For raw log wages, using the standard bar-notation to denote sample averages of observed variables, the within- and between-firm decomposition reads $\mathbb{V}ar(w_{it} | t) = \mathbb{V}ar(w_{it} - \bar{w}_{J(i,t)t} | t) + \mathbb{V}ar(\bar{w}_{J(i,t)t} | t)$.

⁸Song et al. (2019) similarly find that two-thirds of the 1978–2013 rise in U.S. earnings inequality occurs between firms.

average, 1.4 percent of the business sector employment is in “entry firms” (firms in their first year of existence), and 1.8 percent is in “exit firms” (firms in their last year of existence), with both entry and exit rates exhibiting inverse U-shaped trends. The distribution of employment across firm location regions is stable during our observation period.

Turning to worker demographics, panels (e) and (f) document changes in the composition of the business sector workforce. The workforce has become substantially more educated: the proportion of workers with primary school education dropped from 42% to 15%, while those with a Bachelor’s degree increased from 5% to 13%, and those with a Master’s degree grew from 2% to 15%. Women’s share of business sector employment rose from 29% in 1980 to 36% in 2019. The average age of workers rose from 39 in 1980 to 43 in 2019.

In sum, the Danish business sector has undergone a significant structural transformation over the past four decades: employment has shifted from manufacturing to services, firms have become smaller, firm entry and exit has reallocated employment across firms, and the workforce has become older, more female, and substantially more educated.

3 Wage sorting

In this section, we define high- and low-wage workers and wage sorting, document that rising wage sorting is a substantial source of increasing between-firm wage inequality, and introduce the match-level wage sorting contributions that underpin our analysis.

3.1 High- and low wage workers and firms and wage sorting

Let \mathbf{x}_{it} be vector of regressors pertaining to worker i in year t that includes a constant and a set of time-varying observable regressors specified below, and let \mathbf{d}_{it} and \mathbf{f}_{it} denote the vectors of worker and firm indicators.⁹ We consider the [Abowd et al. \(1999\)](#) two-way fixed-effects log wage regression

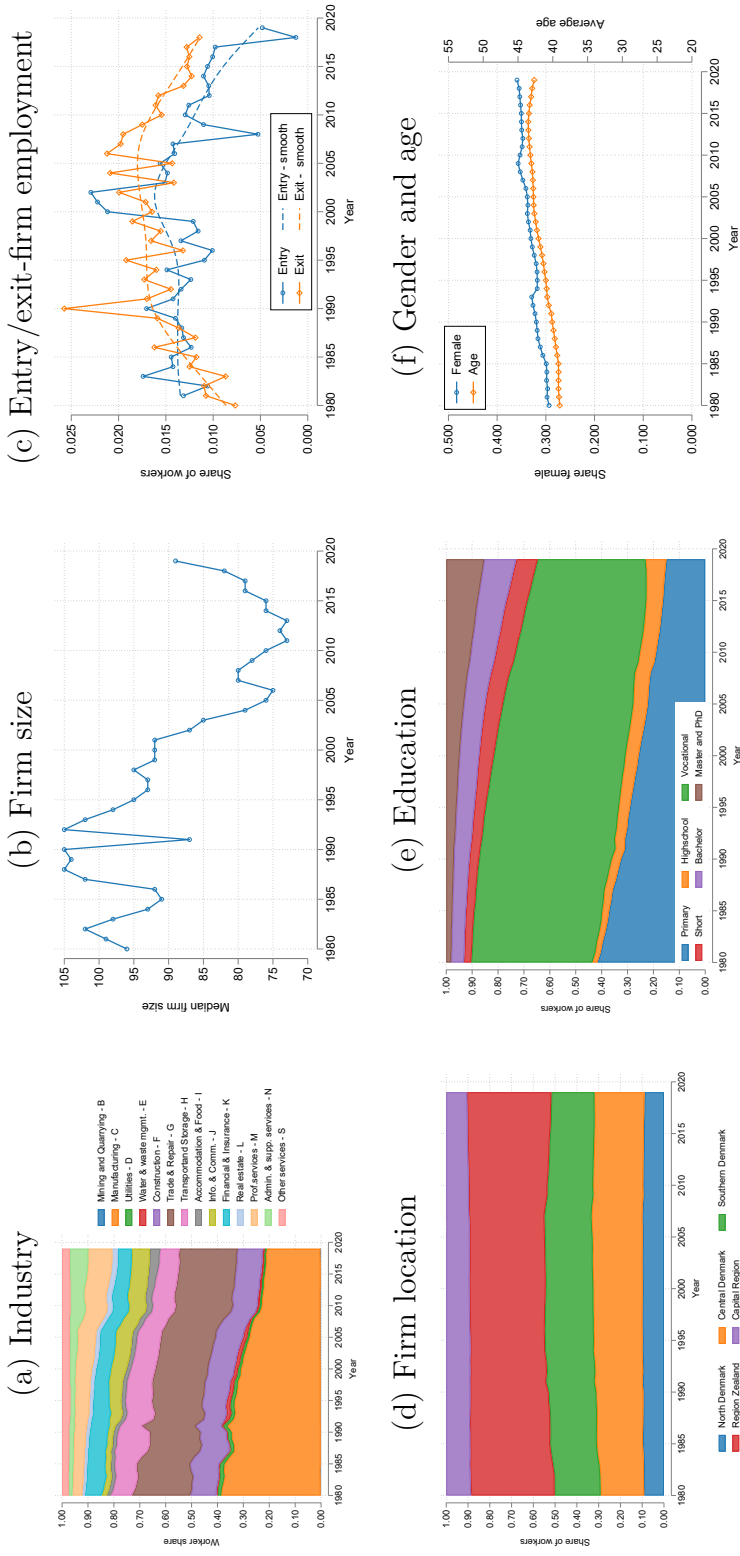
$$w_{it} = \alpha + \mathbf{x}'_{it}\boldsymbol{\beta} + \mathbf{d}'_{it}\boldsymbol{\theta} + \mathbf{f}'_{it}\boldsymbol{\psi} + \varepsilon_{it}, \quad \varepsilon_{it} \perp [\mathbf{1}, \mathbf{X}, \mathbf{D}, \mathbf{F}], \quad (3)$$

where α is a constant term, $\boldsymbol{\beta}$ is conformable with \mathbf{x} , which includes year indicators and a cubic polynomial in worker age both interacted with education dummies,¹⁰ $\boldsymbol{\theta} = [\theta_i]$ and

⁹That is, $\mathbf{f}_{it} = [f_{it}^j]$ where $f_{it}^j = 1$ if $J(i, t) = j$ and $f_{it}^j = 0$ otherwise. Similarly, $\mathbf{d}_{it} = [d_{it}^{i'}]$ where $d_{it}^{i'} = 1$ if $i = i'$ and $d_{it}^{i'} = 0$ otherwise.

¹⁰The age polynomial in (3) is re-centered at age 45 (see [Card et al., 2018](#)), and the coefficient on the linear age term is normalized to zero because the linear age term, year dummies, and the full set of worker

Figure 2: Firm and worker demographic trends (1980–2019)



Notes: The figure show demographic trends in both firm and (employed) worker characteristics. All firm characteristics are employment weighted. Panel (a) show the employment share across NACE 2.0 industry codes; panel (b) show the trend in the median firm size; panel (c) plots the share of workers employed in a firm entering (blue circle) and a firm exiting (orange diamond) the economy. Entry and exit are defined as the first (last) year a firm is observed with non-zero employment. The dotted line plots the smoothed employment rates after applying the Hodrick-Prescott filter with a filtering parameter of 100. Panel (d) plots the distribution of employment across firm location regions. Panel (e) shows the distribution of workers at different educational levels; and panel (f) plots the share of females (blue circles) and the average age (orange diamonds) of workers.

$\boldsymbol{\psi} = [\psi_j]$ are vectors of worker and firm fixed effects, ε_{it} the error term, and $\mathbf{1}$ is a vector of 1s, $\mathbf{X} = [\mathbf{x}'_{it}]$, $\mathbf{D} = [\mathbf{d}'_{it}]$ and $\mathbf{F} = [\mathbf{f}'_{it}]$ are the stacked matrices of observable regressors and worker and firm indicators.

The orthogonality condition $\varepsilon_{it} \perp [\mathbf{1}, \mathbf{X}, \mathbf{D}, \mathbf{F}]$ defines $\boldsymbol{\beta}$, $\boldsymbol{\theta}$, and $\boldsymbol{\psi}$ as linear projection coefficients, implying that the residual is uncorrelated with workers' past, present, and future firm assignments—an assumption known as exogenous mobility (Abowd et al., 1999; Card et al., 2013). Identification of worker and firm fixed effects relies on workers who switch firms. We estimate (3) over the full 40-year period, imposing time-invariant worker and firm effects. Although least squares estimation is straightforward, limited worker mobility can bias the estimated effects; however, as we discuss further below our findings on rising wage sorting are robust to the standard leave-out correction of Kline et al. (2020).

Let $\hat{\boldsymbol{\beta}}$ denote the estimate of $\boldsymbol{\beta}$ in (3), which we referenced above when we defined the residualized log wages as $\hat{r}_{it} := w_{it} - \mathbf{x}'_{it}\hat{\boldsymbol{\beta}}$.¹¹ Furthermore, let $\hat{\theta}_i$ and $\hat{\psi}_j$ denote the estimates of the worker- i fixed effect and the firm- j fixed effect, respectively, which we refer to as worker and firm wage types. A high-wage worker earns systematically more than predicted by \mathbf{x} and the wage type of their employer, while a high-wage firm systematically pays wages above those predicted by \mathbf{x} and worker wage types.

Wage sorting is typically measured as the covariance (or correlation) between worker and firm wage types across matches, i.e. $\text{Cov}(\hat{\theta}_i, \hat{\psi}_{J(i,t)})$. It captures the extent to which the labor market assortatively matches workers and firms by wage type: sorting is positive when high-wage workers are employed by high-wage firms, and negative when high-wage workers are employed by low-wage firms.

3.2 Between-firm wage inequality and wage sorting trends

Our object of interest is the change in relevant aggregate economic variables—such as the wage sorting correlation—between a base year s and a later year $t > s$, where we typically set $s = 1980$. For a generic time series y_t , we denote the s -to- t change by

$$\Delta_s y_t := y_t - y_s.$$

With this notation, and using (3), the s -to- t change in the between-firm variance in

indicators are collinear.

¹¹The residualized log wages \hat{r} net out aggregate productivity and life-cycle trends from w and is, according to (3), the sum of the worker wage type $\hat{\theta}$, the firm wage type $\hat{\psi}$, and the residual $\hat{\varepsilon}$: $\hat{r}_{jt} = \hat{\theta}_i + \hat{\psi}_{J(i,t)} + \hat{\varepsilon}_{it}$. None of our key findings depend critically on whether or how we residualize wages.

Table 2: The between-firm variance trend: segregation, firm type and wage sorting

	$t = 1990$	$t = 2000$	$t = 2010$	$t = 2019$
Between-firm \tilde{r} var. trend, $\Delta_{1980}\text{Var}(\tilde{r}_{J(i,t)t} t)$	0.001	0.004	0.007	0.012
Worker segregation trend, $\Delta_{1980}\text{Var}(\tilde{\theta}_{J(i,t)t} t)$	-0.002	-0.001	-0.000	-0.000
Firm wage type trend, $\Delta_{1980}\text{Var}(\hat{\psi}_{J(i,t)t} t)$	-0.001	-0.001	0.000	0.002
Wage sorting trend, $2\Delta_{1980}\text{Cov}(\hat{\theta}_i, \hat{\psi}_{J(i,t)t} t)$	0.002	0.004	0.005	0.005
Small-firm-year-cell noise, $\Delta_{1980}A(t)$	0.002	0.002	0.002	0.004

Notes: As described in the main text, if w_{it} be the log wage for worker i in year t , then $\hat{r}_{it} := w_{it} - \mathbf{x}'_{it}\hat{\beta}$ are the residual log wages residualized with respect to year effects interacted with education dummies, and a cubic polynomial in worker-age interacted with education dummies, where $\hat{\beta}$ is estimated on the basis of (3); finally, see (1) for definition of $\tilde{r}_{J(i,t)t}$.

residual log wages \hat{r}_{it} can be written as

$$\Delta_s \text{Var}(\tilde{r}_{J(i,t)t}|t) = \underbrace{\Delta_s \text{Var}(\tilde{\theta}_{J(i,t)t}|t)}_{\text{Segregation}} + \underbrace{\Delta_s \text{Var}(\hat{\psi}_{J(i,t)t}|t)}_{\text{Firm types}} + \underbrace{2\Delta_s \text{Cov}(\hat{\theta}_i, \hat{\psi}_{J(i,t)t}|t)}_{\text{Wage sorting}} + \underbrace{\Delta_s A(t)}_{\text{Noise}}, \quad (4)$$

where $A(t)$ is the between-firm log wage variance that stems from finite firm-year cell noise in $\tilde{\varepsilon}_{jt}$, the average residual in firm j in year t .¹²

The first term in (4) captures changes in worker wage type segregation, i.e., how strongly high- and low-wage-type workers cluster together. The second term captures changes in the dispersion of firm wage types. The third term reflects the *wage sorting trend*, which measures how the strength of assortative matching of worker and firm wage types evolves over time.¹³ The fourth term is residual variation due to finite firm-year cell sizes (cf. footnote 12).

Table 2 shows the between-firm residualized log wage variance trend decomposition (4) for $s = 1980$ and $t \in \{1990, 2000, 2010, 2019\}$. The between-firm residual log wage variance increased by 1.2 log points during 1980–2019. The contribution of 1980–2019 changes in worker segregation to the increase in between-firm log wage variance is negligible (in fact, comprising only -2.7 percent of $\Delta_{1980}\text{Var}(\tilde{r}_{J(i,t)t}|t = 2019)$), while the 1980–2019 changes

¹² Specifically, $A(t) := \text{Var}(\tilde{\varepsilon}_{J(i,t)t}|t) + 2\text{Cov}(\tilde{\varepsilon}_{J(i,t)t}, \tilde{\theta}_{J(i,t)t}|t) + 2\text{Cov}(\tilde{\varepsilon}_{J(i,t)t}, \hat{\psi}_{J(i,t)t}|t)$. Since (3) controls for year, worker, and firm dummies, least squares estimation ensures that $\tilde{\varepsilon} = \tilde{\varepsilon}_t = \tilde{\varepsilon}_i = \tilde{\varepsilon}_j = 0$ for all t , i , and j . However, as firm effects are time-invariant, least squares does not impose $\tilde{\varepsilon}_{jt} = 0$ in the sample even if it is implied by the population-analogue of the orthogonality condition $\varepsilon_{it} \perp [\mathbf{1}, \mathbf{X}, \mathbf{D}, \mathbf{F}]$. Indeed, in practice, due to finite firm-year cell-sizes, $\tilde{\varepsilon}_{jt} \neq 0$ and varies with j for given t , giving rise to $A(t)$.

¹³To align with common notation, we express the wage sorting component as $\text{Cov}(\hat{\theta}_i, \hat{\psi}_{J(i,t)t}|t)$, and not as $\text{Cov}(\tilde{\theta}_{J(i,t)t}, \hat{\psi}_{J(i,t)t}|t)$. However, note that $\text{Cov}(\hat{\theta}_i, \hat{\psi}_{J(i,t)t}|t) = \mathbb{E}(\hat{\theta}_i \hat{\psi}_{J(i,t)t}|t) - \mathbb{E}(\hat{\theta}_i|t)\mathbb{E}(\hat{\psi}_{J(i,t)t}|t) = \mathbb{E}(\mathbb{E}(\hat{\theta}_i|J(i,t), t)\hat{\psi}_{J(i,t)t}|t) - \mathbb{E}(\mathbb{E}(\hat{\theta}_i|J(i,t), t)|t)\mathbb{E}(\hat{\psi}_{J(i,t)t}|t) = \text{Cov}(\hat{\theta}_{J(i,t)t}, \hat{\psi}_{J(i,t)t}|t)$.

in firm wage type variance comprise 0.2 log points (or 19.4 percent of $\Delta_{1980} \text{Var}(\tilde{r}_{J(i,t)}|t = 2019)$). Increasing wage sorting is the largest contributor to the trend in between-firm wage variance at every time horizon. In fact, increasing wage sorting accounts for 126, 94, 76, and 46 percent of the increase in between-firm wage variance between 1980 and 1990, 2000, 2010, and 2019, respectively.

3.3 Match-level wage sorting contributions

Having established that rising wage sorting is the central driver of rising between-firm wage inequality, we now develop the analytical framework that underpins the remainder of the paper. The key ingredient is a measure of wage sorting that connects the aggregate sorting trend—a cross-sectional covariance—to the level of individual worker-firm matches, where it can be studied across firms, workers, demographic groups, and labor market transitions.

Wage sorting enters the between-firm log wage variance through the covariance term $\text{Cov}(\hat{\theta}_i, \hat{\psi}_{J(i,t)}|t)$, which is the natural measure from the preceding analysis. However, as a covariance, it is expressed in squared log-wage units, which makes it difficult to interpret. We therefore normalize by the product of the grand standard deviations of worker and firm wage types, computed over all years rather than year-by-year, to obtain the unit-less wage sorting pseudo-correlation

$$\hat{\rho}_t := \frac{1}{|\mathcal{I}_t|} \sum_{i \in \mathcal{I}_t} \frac{(\hat{\theta}_i - \tilde{\theta}_t)(\hat{\psi}_{J(i,t)} - \tilde{\psi}_t)}{\hat{\sigma}_\theta \hat{\sigma}_\psi}, \quad (5)$$

where the grand standard deviations are

$$\hat{\sigma}_\theta^2 = \frac{1}{\sum_{t=1980}^{2019} |\mathcal{I}_t|} \sum_{t=1980}^{2019} \sum_{i \in \mathcal{I}_t} (\hat{\theta}_i - \tilde{\theta})^2; \quad \hat{\sigma}_\psi^2 = \frac{1}{\sum_{t=1980}^{2019} |\mathcal{I}_t|} \sum_{t=1980}^{2019} \sum_{i \in \mathcal{I}_t} (\hat{\psi}_{J(i,t)} - \tilde{\psi})^2.$$

The use of grand rather than year-specific standard deviations ensures that $\hat{\rho}_t$ directly tracks the evolution of $\text{Cov}(\hat{\theta}_i, \hat{\psi}_{J(i,t)}|t)$ over time. This is not a restrictive choice: as Figure 3 shows, the standard deviations of worker and firm wage types are stable over the sample period, so $\hat{\rho}_t$ practically coincides with the conventional year-by-year Pearson correlation.

The definition of $\hat{\rho}_t$ in equation (5) immediately reveals the key analytical observation underlying the remainder of the paper. Since $\hat{\rho}_t$ is a mean of individual terms, we can define

the wage sorting contribution of the worker- i -firm- $J(i, t)$ match at time t as

$$\hat{\eta}_{it} := \frac{(\hat{\theta}_i - \tilde{\theta}_t)(\hat{\psi}_{J(i,t)} - \tilde{\psi}_t)}{\hat{\sigma}_\theta \hat{\sigma}_\psi}, \quad (6)$$

so that, by construction,

$$\hat{\rho}_t = \frac{1}{|\mathcal{I}_t|} \sum_{i \in \mathcal{I}_t} \hat{\eta}_{it}. \quad (7)$$

Since individual wage sorting contributions are fixed within a match, rising wage sorting must be tied to worker mobility: it is only when workers change jobs that the distribution of matches changes and aggregate sorting can rise.¹⁴ This observation has a direct interpretive implication: the rise in wage sorting reflects a labor market in which the reallocation of workers across firms has systematically moved workers toward higher-wage-sorted matches. Furthermore, the fact that aggregate wage sorting is the mean of individual match-level contributions provides a key methodological vehicle for our empirical analysis. The additive structure allows us to decompose the sorting trend using tools from the productivity growth literature, to project match-level contributions onto worker and firm characteristics in a regression framework, and to trace how sorting contributions change around individual job transitions in an event-study design—thereby pinpointing the margins through which wage sorting has strengthened.


Figure 3 plots the 1980–2019 time series of the wage sorting pseudo-correlation $\hat{\rho}_t$ (red squares, every fifth year labeled), the conventional year-by-year Pearson correlation (green triangles), and the standard deviations of worker and firm wage types (blue circles and orange diamonds). Wage sorting increased by 11.9 correlation points, from 0.061 in 1980 to 0.180 in 2019. The standard deviations of worker and firm wage types are stable over time, which is why the pseudo-correlation and the Pearson correlation practically coincide throughout the sample period.

3.4 Robustness analysis: Limited mobility bias

High-dimensional fixed effect regressions are subject to the incidental parameter problem if the number of fixed effects grows with the sample size, so each fixed effect is estimated from a number of observations that does not grow fast enough to eliminate estimation noise. The two-way fixed effect structure of the [Abowd et al. \(1999\)](#) regression makes the incidental

¹⁴Inspecting (6), a match-level sorting contribution can change over time only if the cross-sectional means of the worker and firm fixed effects, $\hat{\theta}_i$ and $\hat{\psi}_t$, change—but these means are stable over time and their contribution to the overall wage sorting trend is negligible.

Figure 3: The wage sorting correlation trend



Notes: Red squares plot the wage sorting pseudo-correlation, $\hat{\rho}_t$. Green triangles plot the Pearson correlation of worker and firm wage types. Blue circles and orange diamonds plot the standard deviations of worker and firm wage types, $\hat{\theta}$ and $\hat{\psi}$.

parameter problem particularly acute—separating permanent worker and firm heterogeneity requires worker mobility, so precision depends not on total observations but on the number of job switchers and the structure of the mobility network—the bipartite graph of workers and firms connected by job switching. When mobility is limited, estimation noise contaminates any variance or covariance computed from the fixed effects: variances are biased upward and the covariance between worker and firm effects—the basis of the wage sorting correlation—is biased downward. These biases are referred to as limited mobility bias (Andrews et al., 2008; Kline et al., 2020; Bonhomme et al., 2023).¹⁵

¹⁵The present robustness analysis addresses an estimation issue: whether limited mobility bias endangers the premise of our analysis, namely a rising wage sorting trend. A structural interpretation of the Abowd et al. (1999) fixed effects raises separate concerns—exogenous mobility, monotonicity of the structural wage function, and permanence—which we do not address directly. We note that Morin (2023) cannot reject exogenous mobility on Danish data using the Card et al. (2013) test, though Borovičková and Shimer (2024) show this test has low power; and that the rising sorting trend is robust to estimating (3) on consecutive

In light of this, we perform two robustness checks of the wage sorting trend in Figure 3. First, we document worker mobility and firm connectivity in our data. Second, we apply the bias correction of [Kline et al. \(2020\)](#) to the wage sorting trend. Both exercises lead to the same conclusion: the rise in wage sorting is not an artifact of limited mobility bias.

3.4.1 Dense and stable mobility networks

The severity of the limited mobility bias diminishes with the degree of worker mobility: when each worker moves many times across many firms, no single worker-firm match has much leverage over the estimated worker and firm fixed effects, estimation noise averages out across many identifying observations, and the bias in all variance components—including the downward bias in the wage sorting covariance—is reduced. Separately, [Bonhomme et al. \(2023\)](#) find that while the level of the bias can be large, it appears stable over time in U.S. data. If the same holds in our Danish data, the bias will affect the level of $\hat{\rho}_t$ but not its trend, and the premise of our analysis—that wage sorting has been rising over 1980–2019—is not endangered by limited mobility bias. Whether this stability assumption is warranted in our setting depends on whether firm connectivity—the density of the mobility network linking workers and firms—is itself stable over time. Table B.1 in Appendix B confirms that connectivity in our data is indeed high and stable.

3.4.2 Limited mobility bias correction

[Kline et al. \(2020\)](#) develop a bias correction for variance and covariance components in high-dimensional fixed effect regressions like (3). The correction computes each match’s contribution to the covariance between worker and firm fixed effects using leave-out estimators: both the worker and firm fixed effects entering each match’s contribution are estimated excluding all observations from that match, so that idiosyncratic wage shocks in that match cannot contaminate either estimated effect, breaking the mechanical negative covariance that drives the bias. We apply the correction over a series of multi-year observation windows—either five non-overlapping 8-year subpanels or 33 rolling 8-year windows—and construct a time series of bias-corrected wage sorting correlations to recover the adjusted wage sorting trend.

The robustness analysis is reported in Appendix B. Two findings emerge. First, limited mobility bias is sizable in short panels: in the 8-year subpanels, the bias-corrected correlation is roughly 10 points higher than the uncorrected estimate. Second, and most importantly for our purposes, the bias-corrected trends from both specifications closely track the full 40-year

8-year subpanels or a rolling panel, see Figure B.1 in Appendix B.

baseline, confirming that the rise in wage sorting is not an artifact of limited mobility bias—a finding consistent with [Bonhomme et al. \(2023\)](#).¹⁶

Taken together, the robustness analysis lead to a clear conclusion: the rise in wage sorting documented in [Figure 3](#)—the premise of our analysis—is not an artifact of limited mobility bias. Worker mobility and firm connectivity in our data are high and stable, limiting the scope for bias in the first place, and formal bias corrections confirm that the aggregate wage sorting trend is robust to limited mobility bias in levels. The [Kline et al. \(2020\)](#) limited mobility bias correction operates on the aggregate wage sorting correlation $\hat{\rho}_t$, but does not correct each of the individual match-level contributions $\hat{\eta}_{it}$ that are central to our analysis: the correction targets aggregate variances and covariances, not the individual match-level contributions, and correcting these aggregate statistics does not recover unbiased estimates of the match-level contributions, even when the aggregate correction is valid. There is no off-the-shelf correction for $\hat{\eta}_{it}$, and while a leave-out correction could in principle be constructed, its feasibility is in question: it would need to be applied to each of the millions of individual matches in our data, and each one of the millions of corrections is a function of the entire data set. We therefore conduct our analysis using the non-corrected wage sorting correlation $\hat{\rho}_t$ and its match-level contributions $\hat{\eta}_{it}$, defined in [\(7\)](#) and [\(6\)](#). The high and stable rates of worker mobility and firm connectivity in our data, and the robustness of the aggregate wage sorting trend to the [Kline et al. \(2020\)](#) correction established above, are our best available assurance that limited mobility bias does not distort our conclusions.

4 Wage sorting and firm dynamics

The rise in wage sorting is well-documented and robust ([Bagger et al., 2013](#); [Card et al., 2013](#); [Song et al., 2019](#)), but the economic forces behind it remain poorly understood. We address this gap by applying decomposition methods from the productivity growth literature ([Baily et al., 1992](#); [Olley and Pakes, 1996](#); [Melitz and Polanec, 2015](#)) to the firm-level representation of aggregate wage sorting that follows directly from the additive structure of $\hat{\rho}_t$ introduced in [Section 3](#).

Since $\hat{\rho}_t$ is the mean of individual match-level contributions $\hat{\eta}_{it}$, it can equivalently be

¹⁶The window structure also delivers a robustness check distinct from the bias correction itself. By estimating worker and firm fixed effects within each window rather than over the full 40-year panel, we implicitly allow for time-varying worker and firm wage types, relaxing the assumption of permanent fixed effects in [\(3\)](#). The close alignment between the uncorrected window-based and full-panel trends suggests that this assumption is not materially affecting our conclusions, consistent with existing evidence that firm effects are highly persistent over time ([Lachowska et al., 2023](#)).

written as an employment-weighted average of firm-level average sorting contributions. Let \mathcal{J}_t denote the set of firms with positive employment in year t , and let $\omega_{jt} \equiv |\mathcal{I}_{jt}|/|\mathcal{I}_t|$ denote firm j 's employment share in year t . Then:

$$\hat{\rho}_t = \sum_{j \in \mathcal{J}_t} \omega_{jt} \tilde{\eta}_{jt}, \quad (8)$$

where $\tilde{\eta}_{jt}$ is the average wage sorting contribution of matches involving firm j in year t :

$$\tilde{\eta}_{jt} := \frac{1}{|\mathcal{I}_{jt}|} \sum_{i \in \mathcal{I}_{jt}} \hat{\eta}_{it}. \quad (9)$$

Using the wage sorting correlation representation (8) and (9), we ask how changes in firm-level average sorting contributions $\tilde{\eta}_{jt}$, employment shares ω_{jt} , and firm turnover each contribute to the year- s -to-year- t wage sorting trend $\Delta_s \rho_t$.

4.1 Firm dynamics, entry and exit, and the wage sorting trend

The year- s -to-year- t wage sorting trend $\Delta_s \rho_t$ now decomposes into four channels. The first arises when the set of active firms and the employment distribution are held fixed—that is, with $\mathcal{J}_t = \mathcal{J}_s$ and $\omega_{jt} = \omega_{js}$ for all j —but the composition of matches within firms changes, so that the average wage sorting contribution $\tilde{\eta}_{jt} \neq \tilde{\eta}_{js}$ for some firm j . We refer to this as the *composition channel*.

The second channel arises when the set of active firms and the average sorting contributions of their matches are held fixed—that is, with $\mathcal{J}_t = \mathcal{J}_s$ and $\tilde{\eta}_{jt} = \tilde{\eta}_{js}$ for all j —but workers are reallocated in a way that alters the employment distribution, so that $\omega_{jt} \neq \omega_{js}$ for some firm j . We refer to this as the *reallocation channel*.

The third and fourth channels arise when the set of active firms changes, so that $\mathcal{J}_t \neq \mathcal{J}_s$. We refer to these as the *firm entry* and *firm exit* channels. These channels operate through selection on wage sorting contributions: wage sorting increases between year- s and year- t if entering firms tend to form matches with higher wage sorting contributions, or exiting firms tend to form matches with lower contributions, than the firms that survive from year- s to year- t .

Formally, define the following sets of firms:

- *Surviving firms*: $\mathcal{S}_{s,t} := \mathcal{J}_t \cap \mathcal{J}_s$, i.e., firms with employment in both year- s and year- t .
- *Exiting firms*: $\mathcal{X}_{s,t} := \mathcal{J}_s \setminus \mathcal{J}_t$, i.e., firms with employment in year- s but not in year- t .

- *Entering firms:* $\mathcal{E}_{s,t} := \mathcal{J}_t \setminus \mathcal{J}_s$, i.e., firms with employment in year- t but not in year- s .

Furthermore, let $\omega_t^{\mathcal{G}} := \sum_{j \in \mathcal{G}} \omega_{jt}$ be the year- t employment share of a generic set of firms \mathcal{G} , and, provided $\omega_t^{\mathcal{G}} > 0$, define the year- t wage sorting in matches involving firms in \mathcal{G} as:

$$\hat{\rho}_t^{\mathcal{G}} := \sum_{j \in \mathcal{G}} \frac{\omega_{jt}}{\omega_t^{\mathcal{G}}} \tilde{\eta}_{jt}. \quad (10)$$

A firm that is in operation in year- s either survives to year- t or exits, and thus: $\omega_s^{\mathcal{S}_{s,t}} + \omega_s^{\mathcal{X}_{s,t}} = 1$; likewise, a firm that is in operation in year- t either existed in year- s or entered since then, so $\omega_t^{\mathcal{S}_{s,t}} + \omega_t^{\mathcal{E}_{s,t}} = 1$.¹⁷ Hence, wage sorting in year- s is $\hat{\rho}_s = \hat{\rho}_s^{\mathcal{S}_{s,t}} + \omega_s^{\mathcal{X}_{s,t}} (\hat{\rho}_s^{\mathcal{X}_{s,t}} - \hat{\rho}_s^{\mathcal{S}_{s,t}})$ and wage sorting in year- t is $\hat{\rho}_t = \hat{\rho}_t^{\mathcal{S}_{s,t}} + \omega_t^{\mathcal{E}_{s,t}} (\hat{\rho}_t^{\mathcal{E}_{s,t}} - \hat{\rho}_t^{\mathcal{S}_{s,t}})$, and the year- s -to- t wage sorting trend can be expressed as:

$$\Delta_s \hat{\rho}_t = \hat{\rho}_t^{\mathcal{S}_{s,t}} - \hat{\rho}_s^{\mathcal{S}_{s,t}} + \omega_t^{\mathcal{E}_{s,t}} (\hat{\rho}_t^{\mathcal{E}_{s,t}} - \hat{\rho}_t^{\mathcal{S}_{s,t}}) + \omega_s^{\mathcal{X}_{s,t}} (\hat{\rho}_s^{\mathcal{S}_{s,t}} - \hat{\rho}_s^{\mathcal{X}_{s,t}}). \quad (11)$$

The term $\hat{\rho}_t^{\mathcal{S}_{s,t}} - \hat{\rho}_s^{\mathcal{S}_{s,t}}$ on the right-hand-side of (11) measures the change in wage sorting among firms that survive from year- s to year- t and may be further decomposed into composition- and reallocation channels. Indeed, adding and subtracting $\sum_{j \in \mathcal{S}_{s,t}} (\omega_{js}/\omega_s^{\mathcal{S}_{s,t}}) \tilde{\eta}_{jt}$ to (11), using the definition (10), and rearranging, yields


$$\begin{aligned} \Delta_s \rho_t = & \underbrace{\sum_{j \in \mathcal{S}_{s,t}} \frac{\omega_{js}}{\omega_s^{\mathcal{S}_{s,t}}} (\tilde{\eta}_{jt} - \tilde{\eta}_{js})}_{\text{Composition}} + \underbrace{\sum_{j \in \mathcal{S}_{s,t}} \left(\frac{\omega_{jt}}{\omega_t^{\mathcal{S}_{s,t}}} - \frac{\omega_{js}}{\omega_s^{\mathcal{S}_{s,t}}} \right) \tilde{\eta}_{jt}}_{\text{Reallocation}} \\ & + \underbrace{\omega_t^{\mathcal{E}_{s,t}} (\hat{\rho}_t^{\mathcal{E}_{s,t}} - \hat{\rho}_t^{\mathcal{S}_{s,t}})}_{\text{Firm entry}} + \underbrace{\omega_s^{\mathcal{X}_{s,t}} (\hat{\rho}_s^{\mathcal{S}_{s,t}} - \hat{\rho}_s^{\mathcal{X}_{s,t}})}_{\text{Firm exit}}. \quad (12) \end{aligned}$$

Equation (12) decomposes the wage sorting trend into four economically distinct channels—composition, reallocation, firm entry, and firm exit—each capturing a different margin through which the labor market can generate rising wage sorting. Two methodological choices are worth noting. First, following Melitz and Polanec (2015), entering and exiting firms are compared to contemporaneous survivors at the point of entry or exit, which avoids dependence on the base period and eases interpretation of the entry and exit contributions even when there is an increasing aggregate trend in the data. Second, we decompose the survivor component into composition and reallocation channels in the spirit of Baily et al. (1992), rather than using the covariance term of Olley and Pakes (1996). While the latter

¹⁷Formally, $\mathcal{S}_{s,t} \cap \mathcal{X}_{s,t} = \emptyset$ and $\mathcal{S}_{s,t} \cup \mathcal{X}_{s,t} = \mathcal{J}_s$ and $\mathcal{S}_{s,t} \cap \mathcal{E}_{s,t} = \emptyset$ and $\mathcal{S}_{s,t} \cup \mathcal{E}_{s,t} = \mathcal{J}_t$.

has a precise allocative efficiency interpretation in productivity decompositions, this interpretation does not naturally extend to wage sorting, making the [Baily et al. \(1992\)](#)-style decomposition more descriptively meaningful in our setting.

Figure 4: The 1980-to- t wage sorting trend decomposition



Notes: The figure plots decomposition (12) with $s = 1980$ and $t = 1981, 1982, \dots, 2019$, thus decomposing $\Delta_{1980}\hat{\rho}_t = \hat{\rho}_t - \hat{\rho}_{1980}$ for each year t .

We implement decomposition (12) with $s = 1980$ and $t = 1981, 1982, \dots, 2019$, thus decomposing $\Delta_{1980}\hat{\rho}_t = \hat{\rho}_t - \hat{\rho}_{1980}$ for each year t . Figure 4 plots this decomposition and delivers a clear message: the composition channel is negligible throughout, meaning that the average wage sorting contribution of a match within a given firm remains largely stable over time. Instead, the rise in wage sorting reflects employment shifting toward firms that persistently form high-sorting matches—both through reallocation among surviving firms and through the exit of low-sorting firms and entry of high-sorting firms. Early in the sample, firm exit is the dominant force; the reallocation channel grows rapidly from around 2000, becoming the dominant contributor thereafter. Firm entry contributes positively throughout, though more modestly than exit.

Table 3: The 1980–to– t wage sorting trend decomposition

	$t = 1990$	$t = 2000$	$t = 2010$	$t = 2019$
The 1980–to– t wage sorting trend, $\Delta_{1980}\hat{\rho}_t$	0.034	0.086	0.112	0.119
Composition channel (survivors)	0.008	0.010	−0.014	−0.003
Reallocation channel (survivors)	0.003	0.014	0.065	0.071
Entry channel	0.008	0.029	0.027	0.011
Exit channel	0.015	0.033	0.035	0.040

Notes: The tables shows decomposition (12) with $s = 1980$ and $t \in (1990, 2000, 2010, 2019)$, thus decomposing $\Delta_{1980}\hat{\rho}_t = \hat{\rho}_t - \hat{\rho}_{1980}$.

Table 3 quantifies the decomposition at four horizons. From 1980 to 2019, the wage sorting correlation increased by 11.9 points, virtually all of which reflects the reallocation, firm entry, and firm exit channels: the reallocation channel among surviving firms contributes 60 percent, the firm exit channel 34 percent, and the firm entry channel 9 percent. The composition channel is negligible throughout, contributing just -0.3 points (-3%) over the full period and turning negative from 2010 onward.

A potential concern with any long-horizon decomposition like our implementation of (12) with $s = 1980$ and, say, $t = 2019$ is that the employment shares of entering and exiting firms mechanically increase as the horizon grows, since firms are not infinitely lived. This is not driving our results: firm exit is already quantitatively important at short horizons (e.g., 1980–1990), and at longer horizons the combined entry and exit contribution declines relative to the reallocation channel, which grows in absolute value. Appendix C confirms this using a decadal implementation with a fixed 10-year horizon. Appendix C also shows that our findings are robust to alternative decomposition approaches following Melitz and Polanec (2015), Griliches and Regev (1995), and Foster et al. (2001).

In sum, our decomposition (12) paints a clear and robust picture: the rise in wage sorting is not driven by changes in the composition of matches within firms, but by the reallocation of employment toward firms that persistently form high-sorting matches—through the expansion of surviving high-sorting firms, the entry of new high-sorting firms, and the exit of low-sorting firms. Firm dynamics are the central margin through which wage sorting—and thus wage inequality—has increased.

5 Demographic changes and the wage sorting trend

The decomposition in Section 4 established that the rise in wage sorting reflects employment reallocation toward firms that on average engage in high-sorting matches, with firm entry and exit playing a central role. One might ask, however, whether this finding reflects genuine firm-side heterogeneity or whether it is driven by high-sorting workers increasingly concentrating in certain firms—which would mechanically raise those firms’ average sorting contributions and show up as firm-side reallocation in decomposition (12), even if the underlying force is worker-side. The negligible composition channel in Table 3 already argues against this interpretation: surviving firms’ average sorting contributions remain stable over time even as their workforces turn over. The two complementary approaches in this section build on this by identifying which types of firms and workers underpin the reallocation, and by conditioning on permanent worker and firm heterogeneity directly to further rule out that the firm-side reallocation documented in Section 4 is a statistical artifact of worker-side forces.

5.1 Observed demographic changes and the wage sorting trend

We start by projecting match-level wage sorting contributions $\hat{\eta}_{it}$ onto vectors of observable firm and worker characteristics, \mathbf{z} and \mathbf{w} , and a full set of year indicators. The firm characteristics $\mathbf{z}' = (\mathbf{z}'_1, \mathbf{z}'_2, \dots, \mathbf{z}'_5)$ comprise NACE section industry indicators (\mathbf{z}_1), region of location indicators (\mathbf{z}_2), firm-size indicators (\mathbf{z}_3), and entry- and exit-year indicators (\mathbf{z}_4 and \mathbf{z}_5). A firm’s entry (exit) year is defined as its first (last) year observed in the data, so that entry year 1980 denotes firms observed in 1980 or earlier and exit year 2019 denotes firms observed in 2019 or later. The worker characteristics $\mathbf{w}' = (\mathbf{w}'_1, \mathbf{w}'_2, \mathbf{w}'_3)$ comprise education-group fixed effects (\mathbf{w}_1), a gender indicator (\mathbf{w}_2), and age indicators (\mathbf{w}_3). Formally,

$$\hat{\eta}_{it} = \alpha + \sum_{k=1}^5 \mathbf{z}'_{k,J(i,t)t} \gamma_k + \sum_{\ell=1}^3 \mathbf{w}'_{\ell,it} \zeta_{\ell} + \mathbf{g}'_{it} \boldsymbol{\tau} + \epsilon_{it}; \quad \epsilon_{it} \perp [\mathbf{1}, \mathbf{Z}, \mathbf{W}, \mathbf{G}] = \mathbf{0}, \quad (13)$$

where α is a constant, γ_k and ζ_{ℓ} are coefficient vectors conformable with \mathbf{z}'_k and \mathbf{w}'_{ℓ} , $\boldsymbol{\tau}$ is a vector of year fixed effects, ϵ_{it} is the error term, and $\mathbf{1}$ is a vector of 1s, $\mathbf{Z} = [\mathbf{z}'_{it}]$, $\mathbf{W} = [\mathbf{w}'_{it}]$ and $\mathbf{G} = [\mathbf{g}'_{it}]$. The coefficients in (13) capture the partial linear associations between wage sorting contributions and each included characteristic, and are consistently estimated by ordinary least squares.

Both the regression (13) and the decomposition of the wage sorting trend that follows rest on the match-level wage sorting contributions $\hat{\eta}_{it}$ introduced in Section 3: because aggregate

wage sorting is the mean of these contributions and the residuals from (13) average to zero by construction, aggregating the fitted regression immediately decomposes the 1980–to– t wage sorting trend $\Delta_{1980}\hat{\rho}_t$ into three components:

$$\Delta_{1980}\hat{\rho}_t = \underbrace{\sum_{k=1}^5 (\bar{z}'_{k,t} - \bar{z}'_{k,1980}) \hat{\gamma}_k}_{\text{Firm demographics}} + \underbrace{\sum_{\ell=1}^3 (\bar{w}'_{\ell,t} - \bar{w}'_{\ell,1980}) \hat{\zeta}_\ell}_{\text{Worker demographics}} + \underbrace{\hat{\tau}_t - \hat{\tau}_{1980}}_{\text{Residual trend}}, \quad (14)$$

where $\bar{z}'_{k,t}$ and $\bar{w}'_{\ell,t}$ denote the employment-weighted average firm and worker characteristics across matched worker-firm pairs in year t . The firm and worker demographic components capture how changes over time in the composition of employment across firm and worker types translate into changes in aggregate wage sorting, in proportion to the estimated coefficients $\hat{\gamma}_k$ and $\hat{\zeta}_\ell$.

The residual trend $\hat{\tau}_t - \hat{\tau}_{1980}$ captures changes in wage sorting not accounted for by observed characteristics—reflecting shifts in unobserved firm or worker traits that correlate with sorting contributions, or broader changes in matching patterns. Note that if none of the included characteristics associate with wage sorting contributions, then $\hat{\gamma} \approx \mathbf{0}$ and $\hat{\zeta} \approx \mathbf{0}$, and the residual trend simply reproduces the raw wage sorting trend.

Table 4: Firm and worker demographic changes and the wage sorting trend

	$t = 1990$	$t = 2000$	$t = 2010$	$t = 2019$
Wage sorting trend, $\Delta_{1980}\hat{\rho}_t$	0.034	0.086	0.112	0.119
Firm demographics	0.016	0.046	0.050	0.067
Industry	−0.001	0.005	0.011	0.013
Region of location	−0.002	−0.003	−0.002	0.000
Firm size	−0.002	−0.002	−0.006	0.002
Entry cohort	0.008	0.021	0.023	0.028
Exit cohort	0.012	0.025	0.024	0.025
Worker demographics	0.004	0.013	0.032	0.055
Educational attainment	0.002	0.011	0.029	0.053
Female	0.000	0.000	0.000	0.000
Age	0.001	0.002	0.004	0.002
Residual wage sorting trend	0.015	0.027	0.030	−0.003

Notes: Bold entries are aggregate contributions; the rows below each aggregate report the contributions of the individual sub-components.

Table 4 presents the decomposition (14) of the wage sorting trend based on the observed firm and worker characteristics in (13), for $t = 1990, 2000, 2010,$ and 2019 . Changes in firm

demographics account for 6.7 correlation points, or around 56 percent, of the 11.9-point rise in wage sorting over 1980–2019, and account for roughly half of the trend at shorter horizons. Changes in worker demographics contribute 5.5 points, or approximately 45 percent, over the full period, with their role growing steadily over time. The residual trend accounts for around 45 percent of the 1980–1990 increase but declines to essentially zero by 2019, indicating that observed firm and worker characteristics account for nearly the entire rise in sorting over the full period. We discuss the sub-components of the firm and worker demographic contributions in the following two sub-sections.

5.1.1 Observable firm demographic changes

Turning to the five sub-components of the firm demographics trend—industry, location, firm size, and entry and exit cohort effects—the central role of firm turnover again emerges, now shown conditional on explicit controls for industry, location, size, and worker demographics. The entry and exit cohort effects account for essentially the entire firm demographics component of the wage sorting trend in 1980–1990 and 1980–2000, more than 90 percent in 1980–2010, and 5.3 of the 6.7 correlation-point increase (nearly 80 percent) over the full 1980–2019 period. This confirms that the role of firm turnover in driving the rise in wage sorting is not an artifact of particular types of firms entering or exiting in terms of industry, region, or size. Industry reallocation contributes modestly but with a growing role: it is small and negative in 1980–1990, accounts for around 20 percent of the firm demographics trend in 1980–2000, and reaches 1.3 correlation points—nearly 20 percent of the 6.7-point firm demographics contribution—over the full period. Changes in employment composition across firm size categories and regions contribute essentially nothing.


Figure 5 plots the estimated entry and exit cohort effects and industry fixed effects from (13), providing further insight into the mechanisms behind these findings.¹⁸ Panel (a) shows that post-1980 entrants form matches with wage sorting contributions roughly 2-10 correlation points higher than pre-1980 entrants, the baseline group. Firms exiting between 1980 and 1995 are associated with contributions roughly 10 points lower than later-exiting firms, with exit effects near zero after 1995.¹⁹ Together, these patterns confirm that firm turnover raises wage sorting by reallocating employment away from low-sorting exiting firms toward high-sorting entering firms. Panel (b) reveals pronounced heterogeneity in wage sorting contributions across industries: with employment shifting from low-sorting Manufacturing toward high-sorting advanced services such as Information and Communication (see Figure

¹⁸Appendix D reports the remaining coefficient estimates from (13).

¹⁹The baseline group are firms observed in 2019 or later.

2), industry reallocation exerts steady upward pressure on aggregate wage sorting.

Figure 5: The wage sorting contribution projection: Entry cohort, exit cohort and industry



Notes: The plotted estimated coefficients on observable firm characteristics are obtained from (13).

5.1.2 Observable worker demographic changes

Turning to the three sub-components of the worker demographics trend—educational attainment, age, and gender—educational attainment is by far the dominant force. Changes in educational composition contribute little to the 1980–1990 trend but account for 1.1 correlation points in 1980–2000 and 2.9 points in 1980–2010, representing the entire worker demographics contribution at that horizon. Over the full 1980–2019 period, changes in educational attainment adds 5.3 correlation points to the wage sorting trend (Table 4). The estimated education effects from (13) (see Appendix D for the estimates) show that workers with a bachelor’s degree or above form matches with wage sorting contributions substantially higher than less educated workers (an education premium of around 20 correlation points). Therefore, as the share of highly educated workers rose from around 7 percent in 1980 to more than 25 percent in 2019 (Figure 2), this compositional shift exerted strong upward pressure on aggregate wage sorting.

The finding that rising educational attainment accounts for a substantial share of the wage sorting trend sits in apparent tension with the negligible composition channel in Section 4. If more educated workers have higher sorting contributions, why does the average sorting contribution within surviving firms remain stable as their workforces become more educated over time? The resolution is that rising educational attainment raises aggregate sorting not by changing the composition of workers within firms, but by reallocating workers across firms—educated workers concentrate in already high-sorting firms. Rising education

therefore shows up in the reallocation channel of Section 4, not the composition channel, and the analysis above identifies education as a key worker characteristic involved in that reallocation. The next sub-section examines this more closely by conditioning on permanent worker and firm heterogeneity directly, effectively asking whether the education gradient in wage sorting reflects an inherent tendency of educated workers to form high-sorting matches, or whether it reflects their concentration in firms that persistently engage in high-sorting matches regardless of which workers they employ.

5.2 Unobserved heterogeneity and the wage sorting trend

We address this question by projecting match-level sorting contributions onto worker and firm fixed effects, summarizing the persistent component of sorting contributions on each side of the market. Since the firm fixed effects absorb all time-invariant firm characteristics, including industry, and the worker fixed effects absorb all time-invariant worker characteristics, including educational attainment, this approach controls for observed and unobserved permanent heterogeneity on both sides of the market. The two-way structure therefore allows us to ask, holding firm-side heterogeneity fixed, whether workers with persistently high sorting contributions have become more numerous over time, and symmetrically, holding worker-side heterogeneity fixed, whether firms with persistently high sorting contributions have gained employment share. In particular, the education gradient identified in the previous sub-section may reflect unobserved firm characteristics that are correlated with education, and the fixed effects approach controls for such confounders directly.

To proceed, we project match-level wage sorting contributions onto worker, firm, and year indicators:

$$\hat{\eta}_{it} = \alpha + \mathbf{d}'_{it}\boldsymbol{\kappa} + \mathbf{f}'_{it}\boldsymbol{\phi} + \mathbf{g}'_{it}\boldsymbol{\tau} + \xi_{it}, \quad \xi_{it} \perp [\mathbf{1}, \mathbf{D}, \mathbf{F}, \mathbf{G}], \quad (15)$$

where α is a constant, $\boldsymbol{\kappa} = [\kappa_i]$ is the vector of worker fixed effects, $\boldsymbol{\phi} = [\phi_j]$ the vector of firm fixed effects, $\boldsymbol{\tau} = [\tau_t]$ the vector of year fixed effects, ξ_{it} the residual, $\mathbf{1}$ a stacked vector of ones, and $\mathbf{D} = [\mathbf{d}'_{it}]$, $\mathbf{F} = [\mathbf{f}'_{it}]$, and $\mathbf{G} = [\mathbf{g}'_{it}]$ are the stacked indicator matrices. Identification of κ_i and $\phi_{J(i,t)}$ hinges on worker mobility and is ensured under the same conditions that identify the worker and firm wage types in (3); estimation is by least squares.²⁰

We refer to κ_i and $\phi_{J(i,t)}$ as the *permanent worker and firm sorting components* and

²⁰By the projection property, ξ_{it} is orthogonal to all regressors across all periods, so the worker and firm fixed effects are defined so that the residual is uncorrelated with past, present, and future firm indicators for each worker.

Table 5: The wage sorting trend and worker and firm wage sorting components

	$t = 1990$	$t = 2000$	$t = 2010$	$t = 2019$
Wage sorting trend, $\Delta_{1980}\hat{\rho}_t$	0.034	0.086	0.112	0.119
Firm wage sorting component	0.019	0.063	0.091	0.122
Worker wage sorting component	0.012	0.017	0.008	-0.031
Residual wage sorting trend	0.003	0.006	0.014	0.027

Notes: The tables shows decomposition (16) with $t \in (1990, 2000, 2010, 2019)$.

denote their estimates $\hat{\kappa}_i$ and $\hat{\phi}_{J(i,t)}$. They summarize the persistent component of sorting contributions on each side of the market—the between-worker and between-firm variation in average sorting contributions that is stable across matches—while the year fixed effects and the residual capture within-worker and within-firm variation in sorting contributions.

As with the observables regression, the additive structure of $\hat{\rho}_t$ means that aggregating (15) immediately yields a decomposition of the 1980-to- t wage sorting trend into contributions from changes in the average firm sorting component, changes in the average worker sorting component, and a residual trend:

$$\Delta_{1980}\hat{\rho}_t = \underbrace{\tilde{\phi}_t - \tilde{\phi}_{1980}}_{\text{Firm sorting}} + \underbrace{\tilde{\kappa}_t - \tilde{\kappa}_{1980}}_{\text{Worker sorting}} + \underbrace{\hat{\tau}_t - \hat{\tau}_{1980}}_{\text{Residual trend}}. \quad (16)$$

Table 5 decomposes the 1980-to- t wage sorting trend using (16). Changes in the average firm sorting component dominate throughout: $\tilde{\phi}_t - \tilde{\phi}_{1980}$ accounts for roughly 55, 75, and 80 percent of the 1980–1990, 1980–2000, and 1980–2010 trends, respectively, and contributes 12 correlation points over 1980–2019, thus accounting for the entire long-run rise in wage sorting. Changes in the average worker sorting component $\tilde{\kappa}_t - \tilde{\kappa}_{1980}$ contribute positively early on—35 percent of the 1980–1990 trend—but shrink steadily, turning negative after 2010 and subtracting 3.1 correlation points over 1980–2019. The residual trend contribution $\hat{\tau}_t - \hat{\tau}_{1980}$, which captures changes in wage sorting not accounted for by shifts in the composition of worker and firm sorting components, is modest throughout, accounting for 2.7 correlation points over the full period.

This stands in contrast to the observables regression (13), which attributed roughly 45 percent of the trend to worker demographics, primarily rising educational attainment (Table 4). The contrast reflects the different sources of identifying variation in the two regressions: the observables regression (13) identifies education effects from cross-sectional variation, while the two-way fixed effects regression (15) identifies worker sorting components from

within-worker variation across jobs, holding permanent firm heterogeneity fixed. The latter shows that highly educated workers do not have systematically higher permanent sorting components than less educated workers. Instead, they tend to concentrate in firms with persistently high sorting components—thus resolving of the apparent tension between the education result in Section 5.1 and the negligible composition channel in Section 4.

6 Job changes and wage sorting

The decomposition and regression analyzes in Sections 4 and 5 identified the firm-level margins driving the rise in wage sorting—reallocation among surviving firms and firm turnover via entry and exit—and the firm and worker characteristics that underpin them. We now turn to a complementary question: through which labor market transitions is this reallocation being implemented at the worker level? By construction, aggregate wage sorting $\hat{\rho}_t$ (5) increases if workers on average experience gains in wage sorting contributions when they change jobs. Whether these gains differ systematically across types of transitions—direct job-to-job moves versus transitions through unemployment, and moves involving entering or exiting firms—is an empirical question that the event study is designed to answer.

The results show that direct job-to-job moves are associated with the largest and most persistent gains in wage sorting contributions, particularly moves to entering firms and away from exiting firms, while transitions through unemployment play a more limited role.

6.1 The event study data

The event of interest is a change of employer. For each observed job change, we refer to the dissolved match as the origin-job and the newly formed match as the destination-job, and set event time $s = 0$ in the first year a worker is observed in the destination-job. A job change is included in the event study data if the worker is observed at $s \in \{-3, -2, \dots, 2\}$ and the origin-job was ongoing at $s = -1$ and $s = -2$.²¹

If the worker involved in a job change is receives either social assistance or unemployment benefits between $s = -1$ and $s = 0$, we record a job-to-unemployment-to-job (JUU) change; otherwise, we record a job-to-job (JJ) change. We additionally record whether the origin-job is with a firm that exits at $s = 0$ and whether the destination-job is with a firm that entered at $s = 0$. Each decade contains 257,000–443,000 job changes by 247,000–394,000 workers,

²¹Overall, a little more than 50 percent of all job changes in our data have a stable origin-job satisfying this requirement.

of which 72–83 percent are JJ-changes. Around 30,000–67,000 job changes involve workers leaving an exiting firm, and 17,000–47,000 involve workers moving to an entering firm. Full summary statistics are reported in Table E.1 in the Appendix E.

6.2 Event study evidence on job changes and wage sorting

As with the decomposition and regression analyzes in the previous sections, the event study exploits the match-level wage sorting contributions $\hat{\eta}_{it}$ introduced in Section 3: because each $\hat{\eta}_{it}$ is an individual observation, it can serve directly as the dependent variable in an event study regression that traces how sorting contributions evolve around job changes.

Formally, let D_{it}^s be the indicator that worker i is subject to a job change event between year $t-s-1$ and year $t-s$. We estimate the event time profile of wage sorting contributions by regressing $\hat{\eta}_{it}$ onto the vector of event time indicators ($D_{it}^{-3}, D_{it}^{-2}, D_{it}^0, D_{it}^1, D_{it}^2$):

$$\hat{\eta}_{it} = \alpha + \sum_{\substack{s=-3 \\ s \neq -1}}^2 \delta_s D_{it}^s + \epsilon_{it}; \quad \mathbb{E}(\epsilon_{it} | D_{it}^{-3}, D_{it}^{-2}, D_{it}^0, D_{it}^1, D_{it}^2) = 0, \quad (17)$$

where $(\delta_{-3}, \delta_{-2}, \delta_0, \delta_1, \delta_2)$ is the event time profile normalized relative to $s = -1$: δ_s measures the average wage sorting contribution s years relative to the job change. We estimate (17) separately for mutually exclusive event types—JJ vs. JUJ changes, moves to entering vs. non-entering firms, and moves from exiting vs. non-exiting firms—using the event study data for 1980–1989, 1990–1999, 2000–2009, and 2010–2019.

Table 6 presents the immediate impact δ_0 and the two-year impact δ_2 for each event type and decade. Full event time profiles are reported in Figure E.2 and Table E.2 in the Appendix E.

Wage sorting contributions: JJ vs. JUJ transitions. Table 6, top panel shows that during the 1980s and 1990s, when aggregate wage sorting was rising rapidly, JJ-changes generated large and persistent gains in wage sorting contributions: immediate gains of 6.0 and 6.3 correlation points in the 1980s and 1990s, respectively, declining modestly to 3.7 and 3.5 points two years later. JUJ-changes, by contrast, yielded much smaller gains— just 1.0 immediately in the 1980s and 1.7 correlation points immediately in the 1990s, compared to the 6.0 and 6.3 points for JJ-changes in the same decades. In the 2000s and 2010s, as the aggregate wage sorting trend slowed, the gains from job changes were smaller and less persistent overall. In the 2000s, JUJ transitions generated somewhat larger immediate

Table 6: The impact of a job change on the match-level wage sorting contribution

	JJ-changes				JUU-changes			
	1980s	1990s	2000s	2010s	1980s	1990s	2000s	2010s
$s = 0$	0.060*** (0.002)	0.063*** (0.002)	0.017*** (0.002)	0.018*** (0.002)	0.010** (0.003)	0.017*** (0.003)	0.036*** (0.003)	0.018** (0.005)
$s = 2$	0.037*** (0.003)	0.035*** (0.002)	-0.002 (0.002)	0.014*** (0.002)	-0.004 (0.004)	0.018*** (0.004)	0.025*** (0.004)	-0.010 (0.005)
	To entry firm				To non-entry firm			
	1980s	1990s	2000s	2010s	1980s	1990s	2000s	2010s
$s = 0$	0.059*** (0.008)	0.095*** (0.007)	0.028*** (0.006)	0.040*** (0.011)	0.044*** (0.002)	0.047*** (0.002)	0.019*** (0.002)	0.017*** (0.002)
$s = 2$	0.013 (0.008)	0.052*** (0.007)	0.000 (0.006)	-0.026** (0.011)	0.027*** (0.002)	0.029*** (0.002)	0.003 (0.002)	0.013*** (0.002)
	From exit firm				From non-exit firm			
	1980s	1990s	2000s	2010s	1980s	1990s	2000s	2010s
$s = 0$	0.099*** (0.006)	0.060*** (0.004)	-0.002 (0.003)	0.029*** (0.006)	0.038*** (0.002)	0.051*** (0.002)	0.024*** (0.002)	0.017*** (0.002)
$s = 2$	0.112*** (0.007)	0.075*** (0.005)	-0.014* (0.004)	0.015* (0.007)	0.014*** (0.002)	0.023*** (0.002)	0.006** (0.002)	0.009*** (0.002)

Notes: Standard errors in parentheses, clustered at the individual level. Statistical significance is denoted by stars: * indicates significance at the 10% level, ** at the 5% level, and *** at the 1% level. s is event time normalized so $s = 0$ for the year where the job change took place, and all effects are measured relative to $s = -1$. JJ-changes refers to job-to-job changes; JUU-changes refers to job-to-unemployment-to-job changes. An entry firm comes into existence in the year of the job change. An exit firm is not in existence in the year after the job change.

gains than JJ transitions (3.6 versus 1.7 correlation points), reversing the pattern from earlier decades. In the 2010s, immediate gains were similar for JJ and JUU transitions (1.8 points each), but the effect persisted for JJ transitions (1.4 points at $s = 2$) while fading for JUU transitions (-1.0 points, insignificant). Overall, the evidence points to direct job-to-job mobility as an important mechanism through which the reallocation of workers toward higher-sorting matches was implemented, particularly during the period of fastest wage sorting growth in the 1980s and 1990s.

Wage sorting contributions: Job changes and firm entry. The middle panel of Table 6 splits job changes into those where the worker moves to an entering firm and those where they move to a non-entering firm. The pattern is broadly similar across the two groups, but moves to entering firms generated particularly large and persistent gains in the 1990s: workers transitioning to an entering firm experienced an immediate wage sorting contribution gain of 9.5 correlation points on average, compared to just 4.7 points for moves

to non-entering firms in the same decade. This suggests that the entry of high-sorting firms in the 1990s—a period of rapid aggregate wage sorting growth—was accompanied by strong sorting gains for the workers who joined them.

Wage sorting contributions: Job changes and firm exit. The bottom panel of Table 6 splits job changes according to whether the worker leaves an exiting firm or a surviving firm. We find that the wage sorting gains realized when leaving an exiting firm are higher than those from leaving a surviving firm. In the 1980s and 1990s, moves from exiting firms yielded immediate gains of 9.9 and 6.0 correlation points, compared to 3.8 and 5.1 points for moves from surviving firms. In the 2000s, the immediate impact of leaving an exiting firm is essentially zero, whereas workers leaving non-exiting see an immediate increase of 2.4 correlation points. In the 2010s there is a modest difference between leaving an exiting and non-exiting firm with the former yielding an increase of 2.9 correlation points and the latter 1.7 points on the immediate impact.

In sum, the event study evidence shows that the rise in wage sorting was implemented primarily through direct job-to-job moves, particularly during the 1980s and 1990s when sorting was rising most rapidly. Moves away from exiting firms generate persisting sorting gains in the 1980s, and to a smaller degree in the 1990s. Similarly, moves to entering firm generated large sorting gains in the 1990s and to a small degree in the 1980s and 2000s. Transitions through unemployment played a more limited role throughout our sample period.

Together, the event study findings paint a coherent picture: the rise in wage sorting reflects a labor market in which direct job-to-job mobility was the key vehicle through which workers moved from low- to high-sorting matches, with firm entry and exit systematically shaping the sorting contributions of the realized matches.

7 Conclusion

Using matched employer–employee data for Denmark from 1980 to 2019, we provide a comprehensive empirical analysis of the economic forces behind the rise in wage sorting. The central finding is that a firm’s average wage sorting contribution per match is stable over time, and that the wage sorting trend reflects reallocation of employment from firms that persistently form low-sorting matches toward those that persistently form high-sorting matches. Reallocation among surviving firms accounts for roughly 60 percent of the increase, while firm turnover—through the exit of low-sorting firms and the entry of high-sorting firms—accounts for the remaining 40 percent.

Regression analysis of match-level wage sorting contributions onto observed firm and worker characteristics shows that firm turnover, captured by entry and exit cohort effects, accounts for the bulk of the firm-side contribution, with industry reallocation—particularly the decline of lower-sorting manufacturing—contributing modestly but with a growing role over time. Rising educational attainment accounts for most of the worker-side contribution. However, a two-way fixed effects specification suggests that this reflects not an inherent tendency of highly educated workers to sort positively, but rather their concentration in firms that tend to engage in matches with high wage sorting contributions.

Event-study analysis of individual job changes shows that direct job-to-job moves generate the largest and most persistent gains in match-level sorting contributions, particularly moves to entering firms and away from exiting firms. Transitions through unemployment play a more limited role. The contribution of firm exit to the sorting trend operates primarily through firm-side selection—exiting firms are persistently low-sorting—rather than through unusually large sorting gains at the point of transition.

Taken together, our findings confirm that rising wage sorting is an important source of rising between-firm wage inequality and establish three novel facts: wage sorting remains largely stable over time for a given firm; firm dynamics, entry, and exit are the central margins facilitating the rise in wage sorting; and job-to-job mobility is an important reallocation mechanism at the worker level. Structural models seeking to account for the rise in between-firm wage inequality should address these facts.

References

- Abowd, J. M., R. H. Creecy, and F. Kramarz (2002). Computing Person and Firm Effects Using Linked Longitudinal Employer-Employee Data. Technical Paper 2002-06, 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.
- 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). High Wage Workers and Low Wage Firms: Negative Assortative Matching or Limited Mobility Bias? *Journal of the Royal Statistical Society A(2008)* 17(Part 3), 673–697.
- Bagger, J. and R. Lentz (2019). An Empirical Model of Wage Dispersion with Sorting. *Review of Economic Studies* 86(1), 153–190.
- Bagger, J., K. L. Sørensen, and R. Vejlin (2013). Wage Sorting Trends. *Economics Letters* 118(1), 63–67.
- Baily, M. N., C. Hulten, D. Campbell, T. Bresnahan, and R. E. Caves (1992). Productivity Dynamics in Manufacturing Plants. *Brookings Papers on Economic Activity. Microeconomics 1992*, 187.
- 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(2), 67–97.
- 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. NBER Working Paper No 33184.
- Briskar, J., R. L. de Melo, E. D. Porto, J. V. R. Mora, and C. Tealdi (2025). The Role of Industries in Rising Inequality. Mimeo.
- 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), 13–70.
- Card, D., J. Heining, and P. Kline (2013). Workplace Heterogeneity and the Rise of West. *The Quarterly Journal of Economics* 128(3), 967–1015.

- Davis, S. J. and J. Haltiwanger (1991). Wage Dispersion Between and Within U.S. Manufacturing Plants, 1963-1986. NBER Working Paper No 3722.
- Decker, R. A., J. Haltiwanger, R. S. Jarmin, and J. Miranda (2014). The Role of Entrepreneurship in US Job Creation and Economic Dynamism. *Journal of Economic Perspectives* 28(3), 3–24.
- Decker, R. A., J. C. Haltiwanger, R. S. Jarmin, and J. Miranda (2020). Changing Business Dynamism and Productivity: Shocks versus Responsiveness. *American Economic Review* 110(12), 3952–3990.
- Dunne, T., L. Foster, J. Haltiwanger, and K. R. Troske (2004). Wage and Productivity Dispersion in United States Manufacturing: The Role of Computer Investment. *Journal of Labor Economics* 22(2), 397–429.
- Eeckhout, J. and P. Kircher (2011). Identifying Sorting—In Theory. *Review of Economic Studies* 78(3), 872–906.
- Faggio, G., K. G. Salvanes, and J. V. Reenen (2010). The Evolution of Inequality in Productivity and Wages: Panel Data Evidence. *Industrial and Corporate Change* 19(6), 1919—1951.
- Foster, L., J. C. Haltiwanger, and C. J. Krizan (2001). Aggregate Productivity Growth: Lessons from Microeconomic Evidence. In *New Developments in Productivity Analysis*, pp. 303–372. University of Chicago Press.
- Gola, P. and Y. Zhao (2024). A Firm Link: Overall, Between- and Within-Firm Inequality Through the Lens of a Sorting Model. Mimeo.
- Griliches, Z. and H. Regev (1995). Firm Productivity in Israeli Industry 1979–1988. *Journal of Econometrics* 65(1), 175–203.
- 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.
- Hoffmann, F., D. S. Lee, and T. Lemieux (2020). Growing Income Inequality in the United States and Other Advanced Economies. *Journal of Economic Perspectives* 34(4), 52–78.
- Kline, P., R. Saggio, and M. Sølvssten (2020). Leave-Out Estimation of Variance Components. *Econometrica* 88(5), 1859–1898.
- 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).

- Lentz, R., P. Suphanit, and J.-M. Robin (2023). The Anatomy of Sorting—Evidence from Danish Data. *Econometrica* 91(6), 2409–2455.
- Lopes de Melo, R. (2018). Firm Wage Differentials and Labor Market. *Journal of Political Economy* 126(1), 313–346.
- Melitz, M. J. and S. Polanec (2015, June). Dynamic Olley-Pakes Productivity Decomposition with Entry and Exit. *The RAND Journal of Economics* 46(2), 362–375.
- Morin, A. (2023). Workplace Heterogeneity and Wage Inequality in Denmark. *Journal of Applied Econometrics* 38(1), 123–133.
- Mouton, A. (2024, Aug). What Drives Wage Sorting? Evidence From West Germany. Working Papers 112, Wake Forest University, Economics Department.
- Mueller, H. M., P. P. Ouimet, and E. Simintzi (2017). Wage Inequality and Firm Growth. *American Economic Review* 107(5), 379–383.
- Olley, G. S. and A. Pakes (1996). The Dynamics of Productivity in the Telecommunications Equipment Industry. *Econometrica* 64(6), 1263–1297.
- 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, 8). Ranking Firms Using Revealed Preference. *The Quarterly Journal of Economics* 133(3), 1331–1393.
- Sorkin, I. and M. Wallskog (2021). The Slow Diffusion of Earnings Inequality. *Journal of Labor Economics* 41(1), 95–127.
- Taber, C. and R. Vejlín (2020). Estimation of a Roy/Search/Compensating Differential Model of the Labor Market. *Econometrica* 88(3), 1031–1069.

Appendix

A Data sources

The IDA (Integreret Database for Arbejdsmarkedsforskning) data is organized in components. We make use of three IDA components: IDA-P (persons), IDA-S (establishments), and IDA-N (jobs).²² Persons are identified by their (anonymized) social security number. Firms are identified by their (anonymized) business registry ID. A firm comprises one or more establishments (physical workplaces). Establishments are identified by an establishment ID developed by Statistics Denmark. At any given point in time, an establishment is linked to one and only one firm. Our empirical analysis identifies employers at the firm-level, and we only make brief use of the establishment ID to merge the IDA components.

IDA-P contains annual individual-level information including age, gender, and accumulated labor market experience. An individual is identified by a unique, time-invariant individual ID (variable: PNR), an anonymized social security number issued to every legal resident. The unit of observation in the IDA-P files is thus a person-year.

IDA-S contains annual information on every physical workplace in Denmark, including information on location, industry, and the identity of the firm that owns the workplace. We refer to a physical workplace as an establishment. An establishment is identified by an establishment ID (variable: LBNR) developed and maintained by Statistics Denmark, and each establishment is assigned to a firm (firm ID variable: ARBGNR). The unit of observation in the IDA-S files is thus an establishment-year.

IDA-N contains annual information on every job that is ongoing in the last week of November, including the average annual wage paid to the worker. A job in IDA-N is identified by the individual and establishment IDs, i.e. by a (PNR, LBNR)-combination, and the unit of observation in the IDA-N files is thus a job-year.

We merge the 1980–2019 IDA-N files with the 1980–2019 IDA-P files year-by-year by PNR. Of course, only IDA-P observations on employed individuals can be merged with an IDA-N observation. We then merge the 1980–2019 IDA-S files onto the matched IDA-N/IDA-P data year-by-year by LBNR. We focus on firms in the Nace 2.0 business sector.²³ The main text explains how we select our analysis data from the merged IDA-N/IDA-P/IDA-S data.

²²Additionally, we use information on employment relationships in 1964-1979 identified from mandatory pension contributions to construct a stable firm ID (data-file name: EXPYEAR), see Appendix A.1

²³Mining and Quarrying; Manufacturing; Electricity, Gas and Steam; Water supply etc.; Construction; Wholesale; Transportation; Accommodation and Food Service; Information and Communication; Finance and Insurance; Real Estate; Professional and Scientific work, and Administration and Support Services.

A.1 Stable firm IDs

This appendix describes how we construct a stable firm ID from the administrative firm ID that is available in IDA. We use j to indicate generic original administrative firm IDs. We use k to indicate a generic new stable firm ID.

The data input is the 1980–2019 universe of primary employment relationships from IDA-N and data on the 1964-1979 universe of employment relationships from records containing firms’ mandatory contributions to a supplementary pension scheme (ATP: Arbejdsmarkedets Tillægspension, data-file name: EXPYEAR²⁴). Workers are identified by a worker ID (variable PNR) and firms by an administrative firm ID (variable ARBGNR).

Let $j(i, t)$ be the function that assigns the administrative firm ID of worker i in year t . That is, $j(i, t) = j$ if worker i is associated with the administrative firm ID j in year t . If worker i is not employed in year t , $j(i, t) = 0$. Let $\mathcal{A}_{j,t} = \{i : j(i, t) = j\}$ be the worker IDs of the year- t workforce of administrative firm j .

Let $k(i, t)$ be the function that assigns the stable firm ID of worker i in year t . That is, $k(i, t) = k$ if worker i is associated with stable firm ID k in year t . If worker i is not employed in year t , $k(i, t) = 0$. Let $\mathcal{S}_{k,t} = \{i : k(i, t) = k\}$ be the worker IDs of the year- t workforce of stable firm k .

To initialize our algorithm, for $t = 2019$, set $k(i, 2019) = j(i, 2019)$ for every $i \in \cup_j \mathcal{A}_{j,2019}$. Hence, the 2019 administrative firm IDs are stable firm IDs. Consider the administrative firm ID j in year $t - 1$, which we refer to as the *origin* firm ID. The year $t - 1$ workforce of the origin firm is $\mathcal{A}_{j,t-1}$. The list of year- t stable firm IDs for this set of workers is $\{k(i, t)\}_{i \in \mathcal{A}_{j,t-1}}$. We refer to these as the *destination* firm IDs. Given the 2019 initialization, and because the algorithm operates backwards in time, destination firm IDs are stable firm IDs.

Our algorithm updates the origin firm IDs sequentially backward using pair of years $t - 1$ and t starting with $t = 2019$ and ending with $t = 1965$: consider generic years $t - 1$ and t and consider generic origin (i.e. year $t - 1$ administrative) firm ID j .

- If the most frequently occurring destination firm ID is the origin firm ID, retain the origin firm ID as the stable firm ID for year $t - 1$. That is, if $\text{mode}(\{k(i, t)\}_{i \in \mathcal{A}_{j,t-1}}) = j$, set $k(i, t - 1) = j$ for every $i \in \mathcal{A}_{j,t-1}$.
- If the most frequently occurring destination firm ID is *not* the origin firm ID, i.e. if $\text{mode}(\{k(i, t)\}_{i \in \mathcal{A}_{j,t-1}}) \neq j$, update the origin firm ID according to the following rule: if each of the conditions below are met, overwrite the origin firm ID with the most

²⁴Note that this dataset does not contain information on earnings which is why it is not used in the main analysis but only used to keep track of firms entry year and workforce flows.

frequently occurring destination firm ID, i.e. set $k(i, t - 1) = \text{mode}(\{k(i, t)\}_{i \in \mathcal{A}_{j,t-1}})$ for every $i \in \mathcal{A}_{j,t-1}$. Otherwise, retain the origin firm ID, i.e. set $k(i, t - 1) = j$ for every $i \in \mathcal{A}_{j,t-1}$. The conditions are:

1. at least 25% of the origin firm workforce, $\mathcal{A}_{j,t-1}$, is present in the most frequently occurring destination firm ID, $\text{mode}(\{k(i, t)\}_{i \in \mathcal{A}_{j,t-1}})$,
 2. at least 25% of the workforce in the most frequently occurring destination firm, $\mathcal{S}_{\text{mode}(\{k(i,t)\}_{i \in \mathcal{A}_{j,t-1}}),t}$, stem from the origin firm j ,
 3. both the origin firm and the most frequently occurring destination firm has at least 5 workers, i.e. $|\mathcal{A}_{j,t-1}| \geq 5$ and $|\mathcal{S}_{\text{mode}(\{k(i,t)\}_{i \in \mathcal{A}_{j,t-1}}),t}| \geq 5$,
- When stable firm IDs have been recorded for each firm present in year $t - 1$, move on to pair of years $t - 2$ and $t - 1$.

A.2 Firm entry and exit rates with stable firm IDs

Figure A.1 plots the 1965-2019 time series of annual entry rates and exit rates. Panels (a) and (b) show employment-weighted entry rates and exit rates series. Panels (c) and (d) show the non-weighted series. The series are computed on the entire 1965-1979 EXPYEAR and 1980-2019 IDA-N panels of primary employment relationships. Each panel shows the series obtained with original administrative firm IDs and with our stable firm ID.


Panels (a) and (b) show that entry and exit rates based on the original firm ID ("Old CVR") are much higher and exhibit extreme spikes compared to those based on the stable firm ID ("New CVR"). These spikes align with major changes to the Danish business registration system (1980, 1992) and local governance reforms in the mid-2000s, which caused spurious changes to (administrative) firm IDs. The stable firm ID eliminates these anomalies, resulting in lower and more consistent rates. Panels (c) and (d) show that the difference between non-weighted entry and exit rates based on the two firm IDs is much smaller than in the employment-weighted series.

Using the stable firm ID, panels (a) and (b) in Figure A.1 show that about 2-3% of workers experience employer exits or are employed by newly established firms each year. Panels (c) and (d) reveal that, on average, 10% of firms enter or exit annually.²⁵


²⁵Note that the entry and exit rates in Figure A.1 are based on the entire IDA-N panel, which includes many small firms, while the main analysis focuses on the business sector and excludes many small firms based on connectivity conditions.

Figure A.1: Entry and exit rates with administrative firm IDs and stable firm IDs


(a) Employment-weighted entry rates




(b) Employment-weighted exit rates



(c) Non-weighted entry rates



(d) Non-weighted exit rates



Notes: All rates are computed on the entire 1965-1979 EXPYEAR and 1980-2019 IDA-N panels of employment relationships. “Old CVR” refers to the original administrative firm ID. “New CVR” refers to the constructed stable firm ID.

B Robustness analysis: Limited mobility bias

This appendix supports the limited mobility bias robustness analysis in Section 3.4. We document worker mobility and firm connectivity in our data, report the results of applying the Kline et al. (2020) bias correction to the aggregate wage sorting trend, and provide further details on both. The latter part of the analysis also shows that the rising trend in wage sorting from 1980 to 2019 is robust to allowing for time-varying worker and firm effects similar to Card et al. (2013).

B.1 Dense and stable mobility networks

Table B.1 presents summary statistics on the density and stability of the mobility networks, focusing on two simple indicators. The first is the number of employers per worker (employment-weighted, top panel). The second is the number of movers per firm, where a *mover* is defined as a worker observed at two or more firms during 1980–2019; this is reported both firm-weighted (middle panel) and employment-weighted (bottom panel). In each panel, the five columns cover the full observation period (1980–2019) and the four decades 1980–1989, 1990–1999, 2000–2009, and 2010–2019 separately. The mover definition is based on the full 1980–2019 period throughout, including for the decadal statistics. For example, a worker employed at one firm in 1980 and another in 1990 is classified as a mover and counted as such in both the 1980–1989 and 1990–1999 subpanels.

The top panel reports the employment-weighted distribution of the number of firms per worker (i.e. one observation per employee-year). Overall, workers are observed at an average of 3.39 firms (median: 3) over the full sample period. The distribution is stable across decades: within the 1980–1989, 1990–1999, 2000–2009, and 2010–2019 decades, the employment weighted average number of firms per worker is 1.71 (median: 1), 1.75 (median: 1), 1.91 (median: 2), and 1.82 (median: 1), respectively.

The middle panel tabulates the distribution of the number of movers per firm (firm-weighted, i.e. one observation per firm). Overall, the average firm employs 35.28 movers during 1980-2019 (median: 8). Again, the distribution is stable across decades: among firms appearing in 1980–1989, 1990–1999, 2000–2009, and 2010–2019, the average number of movers per firm is 21.32 (median: 6), 21.13 (median: 6), 22.21 (median: 6), and 23.11 (median: 6), respectively.

Finally, the bottom panel tabulates the employment-weighted distribution of the number of movers per firm (i.e. one observation per employee-year in a firm). Overall, the average worker works in a firm that employs 3,783 movers during 1980-2019 (median: 390). Once

Table B.1: Worker mobility and firm connectivity

	Number of firms per worker (worker-weighted)				
	Full sample	1980–1989	1990–1999	2000–2009	2010–2019
Average	3.39	1.71	1.75	1.91	1.82
10th percentile	1	1	1	1	1
50th percentile	3	1	1	2	1
90th percentile	7	3	3	4	3
	Movers per firm (firm-weighted)				
	Full sample	1980–1989	1990–1999	2000–2009	2010–2019
Average	35.28	21.32	21.13	22.21	23.11
10th percentile	2	2	1	1	2
50th percentile	8	6	6	6	6
90th percentile	53	33	32	34	36
	Movers per firm (employment-weighted)				
	Full sample	1980–1989	1990–1999	2000–2009	2010–2019
Average	3,783.15	1,626.31	2,113.61	1,739.39	1,330.75
10th percentile	15	9	10	10	9
50th percentile	390	180	192	180	163
90th percentile	10,602	5,300	7,193	5,905	3,697

Notes: A “mover” is a worker observed at two or more firms within the sample period. The mover definition is based on the full 1980–2019 period throughout, including for the decadal statistics. For example, a worker employed at one firm in 1980 and another in 1990 is classified as a mover and counted as such in both the 1980–1989 and 1990–1999 sub-panels. To comply with disclosure rules, the reported k th percentile is the average value of the 10 workers or firms that are closest to the precise k th percentile.

again, the distribution is stable across decades: among firms appearing in 1980–1989, 1990–1999, 2000–2009, and 2010–2019, the employment-weighted average number of movers per firm is 1,626 (median: 180), 2,114 (median: 192), 1,739 (median: 180), and 1,331 (median: 163), respectively.

B.1.1 Bias-corrected wage sorting trends

Kline et al. (2020) develops a bias correction for variance and covariance estimates in high-dimensional linear regressions, using “leave-out” estimators that exclude specific clusters of observations. Applied to two-way fixed effects regressions such as (3), the correction adjusts for limited mobility bias in the joint distribution of worker and firm fixed effects. We

implement this correction to recover a bias-corrected wage sorting trend by estimating worker and firm fixed effects either over five 8-year sub-panels (1980–1987, 1988–1995, 1996–2003, 2004–2011, 2012–2019) or over 33 rolling 8-year windows (1980–1987 through 2012–2019), applying the correction in each window.

The window structure also delivers a robustness check distinct from the bias correction itself. By estimating worker and firm fixed effects within each window rather than over the full 40-year panel, we implicitly allow for time-varying worker and firm wage types, relaxing the assumption of permanent fixed effects in (3). Comparing the uncorrected wage sorting trends from the window structure with the trend in Figure 3 thus serves to check whether the wage sorting trend reported in the main text is robust to the assumption of permanent worker and firm fixed effects.

Figure B.1 plots the wage sorting (pseudo-)correlations computed according to (5). We first compare the three uncorrected wage sorting trends: “AKM - baseline” (blue diamonds), which reproduces the wage sorting trend from the main text; “AKM - 8 year subpanel” (green filled circles); and “AKM - 8 year rolling window” (red squares). Limited mobility bias is more pervasive when estimating on shorter panels, manifesting as lower measured correlations. Consequently, the 8-year sub-panel and 8-year rolling window trends lie below the baseline, with the rolling window correlations particularly depressed—yielding negative point estimates for most of the sample period. Both shorter-panel trends are nonetheless increasing over the observation period, similar to the baseline trend: the 8-year sub-panel trend rises by approximately 8 correlation points (from 0.01 in 1980 to 0.09 in 2019), by construction the same trend reveals in the 8-year rolling, but it reveals how the trend develops smoothly over the time period. This further suggests that the wage sorting trend is not an artifact of the permanent worker and firm fixed effects imposed in our main text analysis.

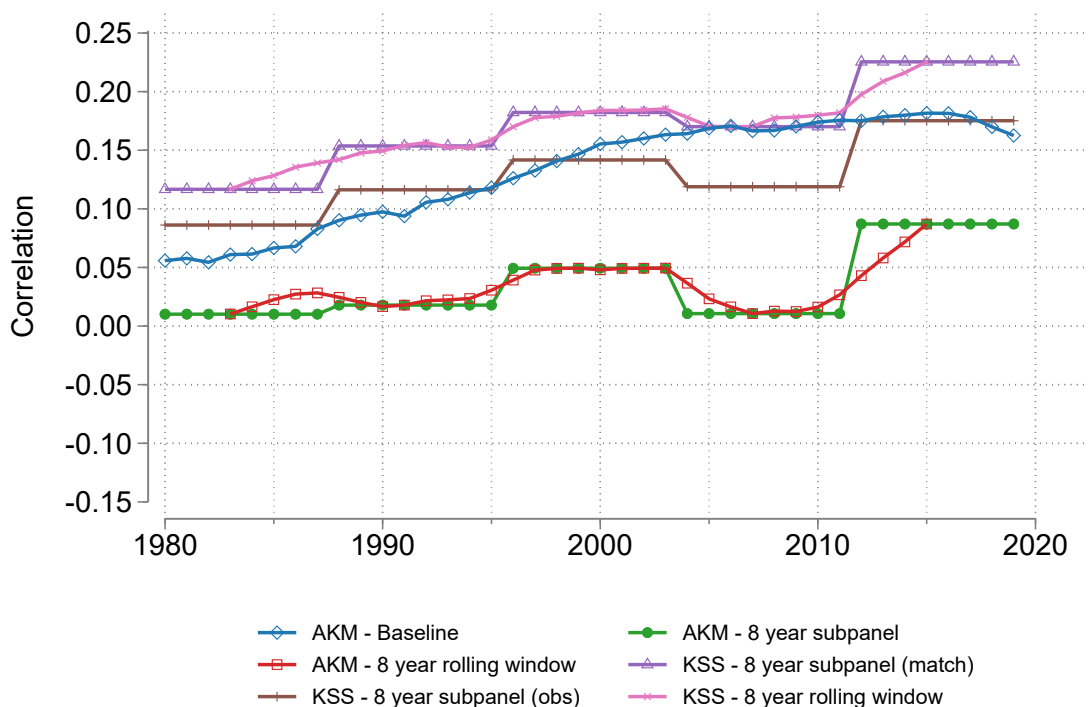
Next, we turn to the bias-corrected counterparts. The figure plots three corrected wage sorting trends: “KSS - 8-year subpanel (match)” (purple triangles) and “KSS - 8-year sub-panel (obs)” (brown plusses) show the bias-corrected versions of the “AKM - 8-year subpanel” trend, using *leave-match-out* and *leave-observation-out* estimators respectively;²⁶ “KSS - 8 year rolling window” (pink crosses) shows the bias-corrected version of the “AKM - 8 year rolling window” trend. All three corrected trends are closely aligned with the uncorrected baseline trend. The bias correction removes the downward level shift in the correlation, but has no discernible effect on the trend, which is present in both the corrected and uncorrected profiles. Hence, consistent with [Bonhomme et al. \(2023\)](#), while the level of wage sorting is

²⁶The *leave-match-out* estimators exclude all observations on a given match while the *leave-observation-out* estimators excludes only a single observation.

susceptible to limited mobility bias, the trend—that is, the change in wage sorting over time—appears robust to it.


A final concern is whether the bias correction of the variances—used to scale the covariance into a correlation—is driving the results in Figure B.1. Figure B.2 plots the covariance trends directly, leaving the scaling aside. The results are unchanged, confirming that our findings are not driven by the bias correction of the variances.

Figure B.1: Wage sorting trends (correlations) with limited mobility bias correction



Notes: “AKM - Baseline” reproduces the baseline wage sorting correlation profile from Figure 3. “AKM - 8-year subpanel” and “AKM - 8 year rolling window” show the uncorrected wage sorting trend estimated on five non-overlapping 8-year sub-panels and 33 rolling 8-year windows (1980–1987 through 2012–2019), respectively. “KSS - 8-year subpanel (match)” and “KSS - 8-year subpanel (obs)” show the [Kline et al. \(2020\)](#) bias-corrected wage sorting trends based on the five 8-year sub-panels, using leave-match-out and leave-observation-out estimators respectively. “KSS - 8 year rolling window” shows the [Kline et al. \(2020\)](#) bias-corrected wage sorting trend based on the 33 rolling 8-year windows.

Figure B.2: Wage sorting trends (covariance) with limited mobility bias correction



Notes: “AKM - Baseline” reproduces the baseline wage sorting covariance profile from the full-sample estimation. “AKM - 8 year subpanel” and “AKM - 8 year rolling window” show the uncorrected wage sorting covariance trend estimated on five non-overlapping 8-year sub-panels and 33 rolling 8 year windows (1980–1987 through 2012–2019), respectively. “KSS - 8-year subpanel (match)” and “KSS - 8-year subpanel (obs)” show the [Kline et al. \(2020\)](#) bias-corrected covariance trends based on the five 8-year sub-panels, using leave-match-out and leave-observation-out estimators respectively. “KSS - 8 year rolling window” shows the [Kline et al. \(2020\)](#) bias-corrected covariance trend based on the 33 rolling 8 year windows.

C Wage sorting and firm dynamics: Additional results

Appendix C.1 presents a decadal implementation of the main-text decomposition (12) that alleviates potential concerns about mechanical dominance of entry and exit at long horizons.

Appendices C.2 and C.3 presents three alternative decompositions of the wage sorting trend from the productivity growth literature: Melitz and Polanec (2015), Griliches and Regev (1995), and Foster et al. (2001). As described in the main text, Melitz and Polanec (2015) extends Olley and Pakes (1996) to a dynamic setting with firm entry and exit, while Griliches and Regev (1995) and Foster et al. (2001) are based on the Baily et al. (1992) decomposition framework. Across all three alternatives, the main conclusions of the paper are confirmed: the rise in wage sorting is primarily a between-firm phenomenon, with firm entry and exit playing economically meaningful roles. The appendix also

C.1 Decadal wage sorting trends

In the main-text decomposition (12) of the wage sorting trend since 1980, $\hat{\rho}_t - \hat{\rho}_{1980}$, as t moves further from 1980, the employment shares of entering and exiting firms will tend to increase. Indeed, because firms are not infinitely lived, sufficiently long-run changes in wage sorting associate exclusively to firm entry and exit.²⁷ This is a feature—not a flaw—of any long-horizon decomposition that separates surviving, entering, and exiting firms.

A concern may be whether the 40-year sample is long enough for firm entry and exit to “mechanically” dominate long-run trends (this of course also requires that $\hat{\rho}_s^{\mathcal{X}_{s,t}}$ is in fact different from survivors). Figure 4 and Table 3 show that firm exit is already quantitatively important at short horizons (e.g., 1980–1990), while firm entry becomes important from the mid-1990s onward. At longer horizons (from 2000 onward), however, the combined contribution of entry and exit declines relative to the between-firm component, which grows in absolute value, reflecting increasing wage sorting driven by labor reallocation among surviving firms.²⁸ These patterns imply that our main conclusions for the rise in wage sorting between 1980 and 2019—that it is primarily a between-firm phenomenon and that firm entry and exit play an economically meaningful role—are not artifacts of applying decomposition (12) over a long horizon.²⁹

²⁷Unless there is a firm that is infinitely lived, there is a t such that $\mathcal{S}_{1980,t} = \emptyset$.

²⁸The firm-entry component declines in absolute terms after 2001, while the firm-exit component remains approximately constant from 2000 onward.

²⁹We do not report firm-level survival statistics because doing so would require addressing left-censored spells.

This sub-section complements the main text discussion by presenting a decadal implementation of (12), in which we take $s = t - 9$ for $t = 1989, 1990, \dots, 2019$, decomposing the decadal wage sorting trend $\hat{\rho}_t - \hat{\rho}_{t-9}$ into within-firm, between-firm, firm entry, and firm exit components.

Figure C.1, panel plots the decadal change in wage sorting, $\hat{\rho}_t - \hat{\rho}_{t-9}$, for each t between 1989 and 2019. The total change (blue circles) is always positive, ranging from 1 to 5 correlation points per decade, with a downward trend towards the end of the observation period. This pattern is consistent with the 1980– t wage sorting trend shown in Figure 4. Decomposing the decadal change into composition, reallocation, firm entry, and firm exit components confirms the conclusions from Figure 4 and Table 3. In the 1990s, the increase in wage sorting is driven mainly by the exit of relatively low-wage-sorting firms. During the 2000s, the contribution of entry by relatively high-wage-sorting firms rises, becoming comparable to that of firm exit. Over the same period, the between-firm component increases steadily, becomes the dominant contributor in the 2000s, and remains so thereafter, although a drop in the between-firm component towards the end of the sample renders the firm exit component the largest in the final years. By contrast, the within-firm component is small for most of the sample and, when it becomes quantitatively relevant in the late 2000s, its contribution is negative.


C.2 The Melitz and Polanec (2015) decomposition

This sub-section derives and presents a decomposition of the wage sorting trend using Melitz and Polanec (2015), who provide a dynamic version of the Olley and Pakes (1996) decomposition that accounts for firm entry and exit. The difference between our main-text decomposition (12) and Melitz and Polanec (2015) lies in the treatment of surviving firms' contribution to the wage-sorting change, $\Delta_s \hat{\rho}_t$. In (12), we decompose incumbents' contributions into composition and reallocation components, following Baily et al. (1992). By contrast, Melitz and Polanec (2015) applies the Olley and Pakes (1996) covariance decomposition to wage sorting among surviving firms, or incumbent firms in the terminology of Olley and Pakes (1996).

To apply the Olley and Pakes (1996) decomposition to $\hat{\rho}_t^{S_{s,t}}$, define the *unweighted* cross-sectional average wage-sorting contribution in year t among surviving firms as

$$\text{Avg}_t^{S_{s,t}}(\tilde{\eta}_{jt}) := \frac{1}{|S_{s,t}|} \sum_{j \in S_{s,t}} \tilde{\eta}_{jt}, \quad (\text{C.1})$$

Figure C.1: The decadal wage sorting trend $\hat{\rho}_t - \hat{\rho}_{t-9}$ by $t = 1989, 1990, \dots, 2019$: composition, reallocation, and entry and exit channels



and the *unweighted* cross-sectional average employment share in year t across surviving firms as

$$\text{Avg}_t^{\mathcal{S}_{s,t}} \left(\frac{\omega_{jt}}{\omega_t^{\mathcal{S}_{s,t}}} \right) := \frac{1}{|\mathcal{S}_{s,t}|} \sum_{j \in \mathcal{S}_{s,t}} \frac{\omega_{jt}}{\omega_t^{\mathcal{S}_{s,t}}}. \quad (\text{C.2})$$

Since $\hat{\rho}_t^{\mathcal{S}_{s,t}} = \sum_{j \in \mathcal{S}_{s,t}} \frac{\omega_{jt}}{\omega_t^{\mathcal{S}_{s,t}}} \tilde{\eta}_{jt}$, see (10) in the main text, it follows that

$$\hat{\rho}_t^{\mathcal{S}_{s,t}} = \underbrace{\text{Avg}_t^{\mathcal{S}_{s,t}}(\tilde{\eta}_{jt}) + \sum_{j \in \mathcal{S}_{s,t}} \left[\frac{\omega_{jt}}{\omega_t^{\mathcal{S}_{s,t}}} - \text{Avg}_t^{\mathcal{S}_{s,t}} \left(\frac{\omega_{jt}}{\omega_t^{\mathcal{S}_{s,t}}} \right) \right]}_{|\mathcal{S}_{s,t}| \text{Cov} \left(\frac{\omega_{jt}}{\omega_t^{\mathcal{S}_{s,t}}}, \tilde{\eta}_{jt} \right)} \left[\tilde{\eta}_{jt} - \text{Avg}_t^{\mathcal{S}_{s,t}}(\tilde{\eta}_{jt}) \right]. \quad (\text{C.3})$$

The second term in (C.3) is the wage-sorting analogue of the [Olley and Pakes \(1996\)](#) covariance term, capturing $|\mathcal{S}_{s,t}|$ times the cross-sectional covariance in year t between employment shares and average wage-sorting contributions across surviving firms, $|\mathcal{S}_{s,t}| \text{Cov}(\omega_{jt}/\omega_t^{\mathcal{S}_{s,t}}, \tilde{\eta}_{jt})$.


Substituting (C.3) into (11) yields the [Melitz and Polanec \(2015\)](#) decomposition of the change in wage sorting from year s to year t . It attributes the wage sorting trend to shifts in the mean level and employment-weighted allocation of wage-sorting contributions among surviving firms, as well as to firm entry and exit:

$$\Delta_s \hat{\rho}_t = \underbrace{\text{Avg}_t^{\mathcal{S}_{s,t}}(\tilde{\eta}_{jt}) - \text{Avg}_s^{\mathcal{S}_{s,t}}(\tilde{\eta}_{js})}_{\text{Mean wage sorting}} + \underbrace{|\mathcal{S}_{s,t}| \left[\text{Cov} \left(\frac{\omega_{jt}}{\omega_t^{\mathcal{S}_{s,t}}}, \tilde{\eta}_{jt} \right) - \text{Cov} \left(\frac{\omega_{js}}{\omega_s^{\mathcal{S}_{s,t}}}, \tilde{\eta}_{js} \right) \right]}_{\text{Employment-sorting covariance}} + \underbrace{\omega_t^{\mathcal{E}_{s,t}} (\hat{\rho}_t^{\mathcal{E}_{s,t}} - \hat{\rho}_t^{\mathcal{S}_{s,t}})}_{\text{Firm entry}} + \underbrace{\omega_s^{\mathcal{X}_{s,t}} (\hat{\rho}_s^{\mathcal{S}_{s,t}} - \hat{\rho}_s^{\mathcal{X}_{s,t}})}_{\text{Firm exit}} \quad (\text{C.4})$$

where the wage sorting correlations among entering and exiting firms, $\hat{\rho}_t^{\mathcal{E}_{s,t}}$ and $\hat{\rho}_s^{\mathcal{X}_{s,t}}$, are as defined in (11) in the main text.

Figure C.2 presents the decomposition in (C.4) with $s = 1980$ and $t = 1981, \dots, 2019$, making it directly comparable to Figure 4. The total wage-sorting trend, $\Delta_{1980} \hat{\rho}_t$ (blue circles), and the entry and exit components coincide with those in Figure 4. The new elements are the mean wage-sorting component (orange) and the employment-wage-sorting covariance component (green), which together replace the within- and between-firm terms. The mean wage sorting component is slightly negative initially but rises steadily, becoming the main contributor after 2010. The covariance component is small for most of the period but increases after 2010; by 2019, it is the second-largest contributor to the 1980–2019 trend,

Figure C.2: The 1980-to- t wage sorting trend decomposition (Melitz and Polanec, 2015)



just below firm exit.

C.3 The Baily et al. (1992) decomposition approach

This sub-section derives the Baily et al. (1992) decomposition framework that underlies both the Griliches and Regev (1995) and Foster et al. (2001) decompositions presented below. The Baily et al. (1992) approach differs from our main-text decomposition (12) in how it treats the contribution of entering and exiting firms to the year- s -to- t wage sorting trend, $\Delta_s \hat{\rho}_t$.

The Baily et al. (1992) approach uses the same partition of active firms as in the main text and in Melitz and Polanec (2015): surviving firms $\mathcal{S}_{s,t}$, entering firms $\mathcal{E}_{s,t}$, and exiting firms $\mathcal{X}_{s,t}$. Let ω_{jt} denote firm j 's employment share in year t , and $\tilde{\eta}_{jt}$ its average wage-sorting contribution. Then, the wage sorting correlation in year $t > s$ relative to a fixed reference correlation ρ^* can be expressed as

$$\hat{\rho}_t - \rho^* = \sum_{j \in \mathcal{S}_{s,t}} \omega_{jt} (\tilde{\eta}_{jt} - \rho^*) + \sum_{j \in \mathcal{E}_{s,t}} \omega_{jt} (\tilde{\eta}_{jt} - \rho^*) \quad (\text{C.5})$$

the correlation in year $s < t$ relative to the same fixed reference correlation ρ^* reads

$$\hat{\rho}_s - \rho^* = \sum_{j \in \mathcal{S}_{s,t}} \omega_{js}(\tilde{\eta}_{js} - \rho^*) + \sum_{j \in \mathcal{X}_{s,t}} \omega_{js}(\tilde{\eta}_{js} - \rho^*), \quad (\text{C.6})$$

and, thus, the year- s -to- t wage sorting trend $\Delta_s \hat{\rho}_t$ can be expressed as:

$$\Delta_s \hat{\rho}_t = \sum_{j \in \mathcal{S}_{s,t}} \omega_{jt}(\tilde{\eta}_{jt} - \rho^*) - \sum_{j \in \mathcal{S}_{s,t}} \omega_{js}(\tilde{\eta}_{js} - \rho^*) + \sum_{j \in \mathcal{E}_{s,t}} \omega_{jt}(\tilde{\eta}_{jt} - \rho^*) + \sum_{j \in \mathcal{X}_{s,t}} \omega_{js}(\rho^* - \tilde{\eta}_{js}). \quad (\text{C.7})$$

The composition and reallocation components in our main text decomposition (12) obtain with $\rho^* = 0$ and rearranging using $\omega_t^{\mathcal{S}_{s,t}} + \omega_t^{\mathcal{E}_{s,t}} = \omega_s^{\mathcal{S}_{s,t}} + \omega_s^{\mathcal{X}_{s,t}} = 1$, in which case entry contributions are evaluated relative to wage sorting in year t and exit contributions relative to year s . Here, entering and exiting firms are compared to survivors at the point of entry or exit, which eases interpretation of the contributions from entry and exit even when there is an increasing aggregate trend in the data, see [Melitz and Polanec \(2015\)](#). We show two alternative decompositions due to [Griliches and Regev \(1995\)](#) and [Foster et al. \(2001\)](#), both implementations of the [Baily et al. \(1992\)](#) approach.


C.3.1 The [Griliches and Regev \(1995\)](#) decomposition

This sub-section derives and presents the wage sorting trend decomposition of [Griliches and Regev \(1995\)](#), which implements the [Baily et al. \(1992\)](#) approach by setting the reference correlation to $\rho^* = (\hat{\rho}_s + \hat{\rho}_t)/2$, i.e. the simple average of the wage-sorting correlations in years s and t . Adding and subtracting $\sum_{j \in \mathcal{S}_{s,t}} \left(\frac{\omega_{js} + \omega_{jt}}{2}\right) (\tilde{\eta}_{jt} - \tilde{\eta}_{js})$ and rearranging yields

$$\begin{aligned} \Delta_s \hat{\rho}_t = & \underbrace{\sum_{j \in \mathcal{S}_{s,t}} \left(\frac{\omega_{js} + \omega_{jt}}{2}\right) (\tilde{\eta}_{jt} - \tilde{\eta}_{js})}_{\text{Composition}} + \underbrace{\sum_{j \in \mathcal{S}_{s,t}} (\omega_{jt} - \omega_{js}) \left(\frac{\tilde{\eta}_{js} + \tilde{\eta}_{jt}}{2} - \frac{\hat{\rho}_s + \hat{\rho}_t}{2}\right)}_{\text{Reallocation}} \\ & + \underbrace{\sum_{j \in \mathcal{E}_{s,t}} \omega_{jt} \left(\tilde{\eta}_{jt} - \frac{\hat{\rho}_s + \hat{\rho}_t}{2}\right)}_{\text{Entry}} + \underbrace{\sum_{j \in \mathcal{X}_{s,t}} \omega_{js} \left(\frac{\hat{\rho}_s + \hat{\rho}_t}{2} - \tilde{\eta}_{js}\right)}_{\text{Exit}}. \quad (\text{C.8}) \end{aligned}$$

The first two terms in (C.8) decompose the change in wage sorting among surviving firms into composition and reallocation components. The composition term captures changes in each firm's average sorting contribution, weighted by its average employment share across years s and t . The reallocation term captures changes in employment shares across surviving

Figure C.3: The 1980–to– t wage sorting trend decomposition (Griliches and Regev, 1995)



firms with different average sorting contributions. The entry and exit terms are evaluated relative to the two-period average wage-sorting correlation $(\hat{\rho}_s + \hat{\rho}_t)/2$: entry contributes positively if the average sorting contribution of entering firms exceeds this benchmark, while exit contributes positively if the average sorting contribution of exiting firms falls below it.

Figure C.3 presents the decomposition in (C.8) with $s = 1980$ and $t = 1981, \dots, 2019$. Through the lens of this decomposition, firm entry and exit play an even more prominent role in accounting for the rise in wage sorting: the Griliches and Regev (1995) decomposition attributes 80 percent of the 1980–2019 wage sorting trend to firm entry and exit in roughly equal proportions. The composition component remains small throughout, while the reallocation component only becomes meaningful around 2000, growing to account for approximately 10 percent of the wage sorting trend by 2010.

C.4 The Foster et al. (2001) decomposition

This sub-section derives and presents the wage sorting trend decomposition of Foster et al. (2001), which implements the Baily et al. (1992) approach by setting the reference correlation


to $\rho^* = \hat{\rho}_s$, i.e. the wage-sorting correlation in the base year s . By adding and subtracting $\sum_{j \in \mathcal{S}_{s,t}} \omega_{js}(\tilde{\eta}_{jt} - \tilde{\eta}_{js}) + \sum_{j \in \mathcal{S}_{s,t}} \omega_{js}(\tilde{\eta}_{jt} - \hat{\rho}_s)$ in (C.7) and rearranging terms, the Foster et al. (2001) decomposition is:

$$\begin{aligned} \Delta_s \hat{\rho}_t = & \underbrace{\sum_{j \in \mathcal{S}_{s,t}} \omega_{js}(\tilde{\eta}_{jt} - \tilde{\eta}_{js})}_{\text{Composition}} + \underbrace{\sum_{j \in \mathcal{S}_{s,t}} (\omega_{jt} - \omega_{js})(\tilde{\eta}_{jt} - \hat{\rho}_s)}_{\text{Reallocation}} \\ & + \underbrace{\sum_{j \in \mathcal{S}_{s,t}} (\omega_{jt} - \omega_{js})(\tilde{\eta}_{jt} - \tilde{\eta}_{js})}_{\text{Emp.-share/sorting growth cross}} \\ & + \underbrace{\sum_{j \in \mathcal{E}_{s,t}} \omega_{jt}(\tilde{\eta}_{jt} - \hat{\rho}_s)}_{\text{Entry}} + \underbrace{\sum_{j \in \mathcal{X}_{s,t}} \omega_{js}(\hat{\rho}_s - \tilde{\eta}_{js})}_{\text{Exit}}. \quad (\text{C.9}) \end{aligned}$$

The first two terms in (C.9) represent the composition and reallocation components of the year- s -to-year- t wage sorting trend. As in Foster et al. (2001), the composition component captures changes in firms' average wage sorting contributions, weighted by their employment shares in year s . The reallocation component reflects changes in employment among surviving firms, evaluated relative to the year- s average wage sorting contribution: it is positive when firms with above-average wage sorting contributions in year s expand their employment shares. The third term—which is an uncentered cross-product moment that captures a notion of covariance between changes in wage sorting contributions and employment growth—is positive when firms with rising wage-sorting contributions expand while those with declining contributions contract. Finally, the contributions of firm entry and firm exit are evaluated relative to the base-period wage sorting correlation $\hat{\rho}_s$: entry contributes positively when entrants' wage sorting contributions exceed $\hat{\rho}_s$, while exit contributes positively when exiting firms' contributions fall below $\hat{\rho}_s$.

Figure C.4 presents the decomposition in (C.9) with $s = 1980$ and $t = 1981, \dots, 2019$. The overall trend (blue circles) coincides with that in Figure 4. In the Foster et al. (2001) framework, the composition and reallocation components (orange and green, respectively) are negligible at all horizons; over the full sample, they account for less than 0.01 and 0.05 correlation points, respectively, of the total 11.9-point increase in wage sorting from 1980 to 2019. The wage sorting trend is instead driven primarily by the covariance between wage sorting changes and employment growth, firm entry, and firm exit. In the 1980s, entry and exit contribute about equally while the covariance term is minor. From 1995 onward, exit remains flat whereas entry rises nearly linearly and becomes the dominant force. The

Figure C.4: The 1980-to- t wage sorting trend decomposition Foster et al. (2001)



covariance term stays modest until around 2000, then increases and surpasses exit as the second-largest contributor.

D Demographic changes and the wage sorting trend: Additional results

This appendix presents additional results from the regression analysis of match-level wage sorting contributions onto observed firm and worker characteristics described in Section 5. It reports the full set of estimated fixed effects from (13), complementing the decomposition results presented in the main text.

The reference firm operates in Manufacturing, with 1-4 employees, is located in North Denmark, entered in 1980 or earlier, and remained in operation in 2019.³⁰ The reference worker is a woman born in 1921 (59 years old in 1980), who completed primary education but with no further educational attainment, who has four years or less of tenure on the job, and four years or less of experience in the labor market.³¹

D.1 Observable firm demographic heterogeneity

Figure D.1 plots the estimated entry and exit year fixed effects γ_4 and γ_5 from (13) in panel (a), the industry fixed effects γ_1 in panel (b), the region fixed effects γ_2 in panel (c), and the firm size fixed effects γ_3 in panel (d). The estimated fixed effects represent the average differences in wage sorting contributions between matches involving a firm with a specific characteristic and the reference firm, measured in correlation points. The entry and exit year effects in panel (a) and the industry effects in panel (b) are discussed in Section 5 in the main text. Here we focus on the region and firm size effects in panels (c) and (d).

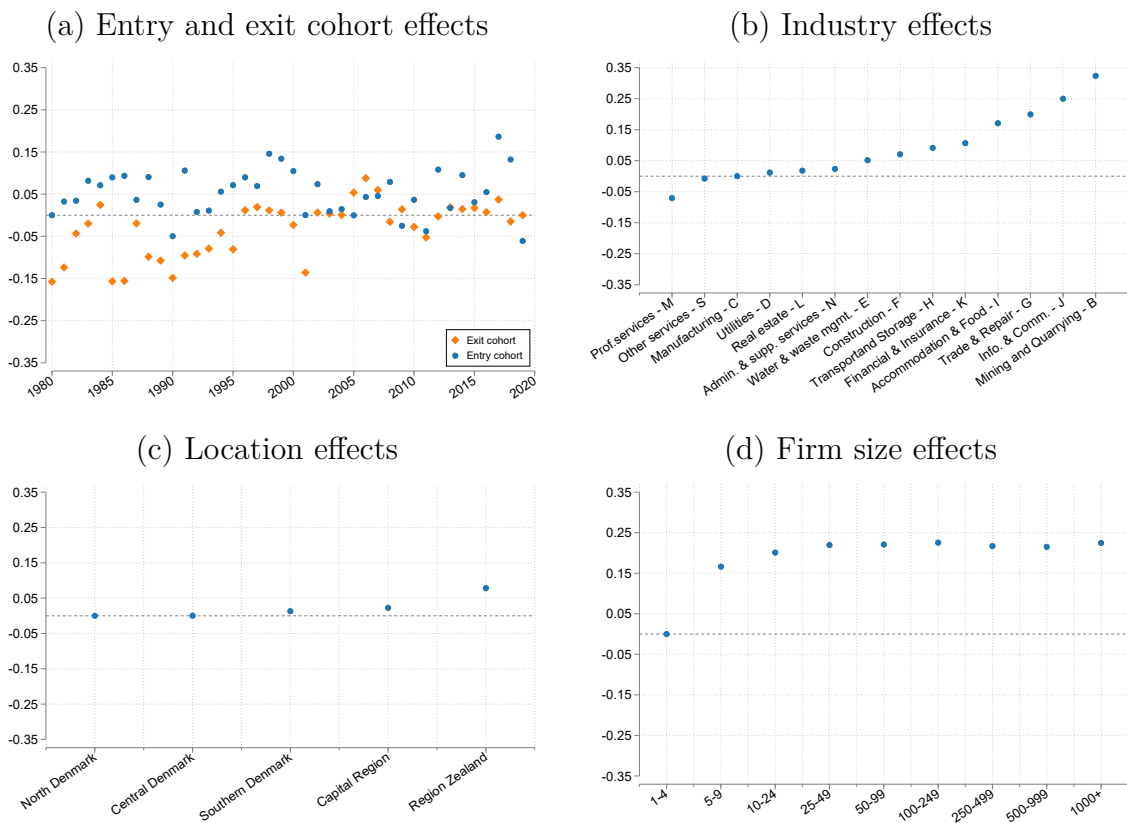
Panel (c) shows limited spatial variation in wage sorting. Firms in Region Zealand have the highest contributions, about 7 correlation points above North Denmark, the lowest, while the densely populated Capital Region only slightly exceeds North Denmark. Panel (d) shows that worker-firm matches involving very small micro firms (1-4 employees) have the lowest wage sorting contributions, though estimation error likely biases these toward zero.³² Wage sorting contributions rise sharply for firms with 5-9 employees—by about 15 correlation points—and remain stable across larger firms, including those with over 1,000 employees.

³⁰We cannot identify trend, entry cohort and age effects simultaneously and (13) normalizes firm age effects in wage sorting to zero.

³¹Worker age, time trend, and birth cohort effects cannot be identified simultaneously, and (13) normalizes worker age effects in wage sorting to zero.

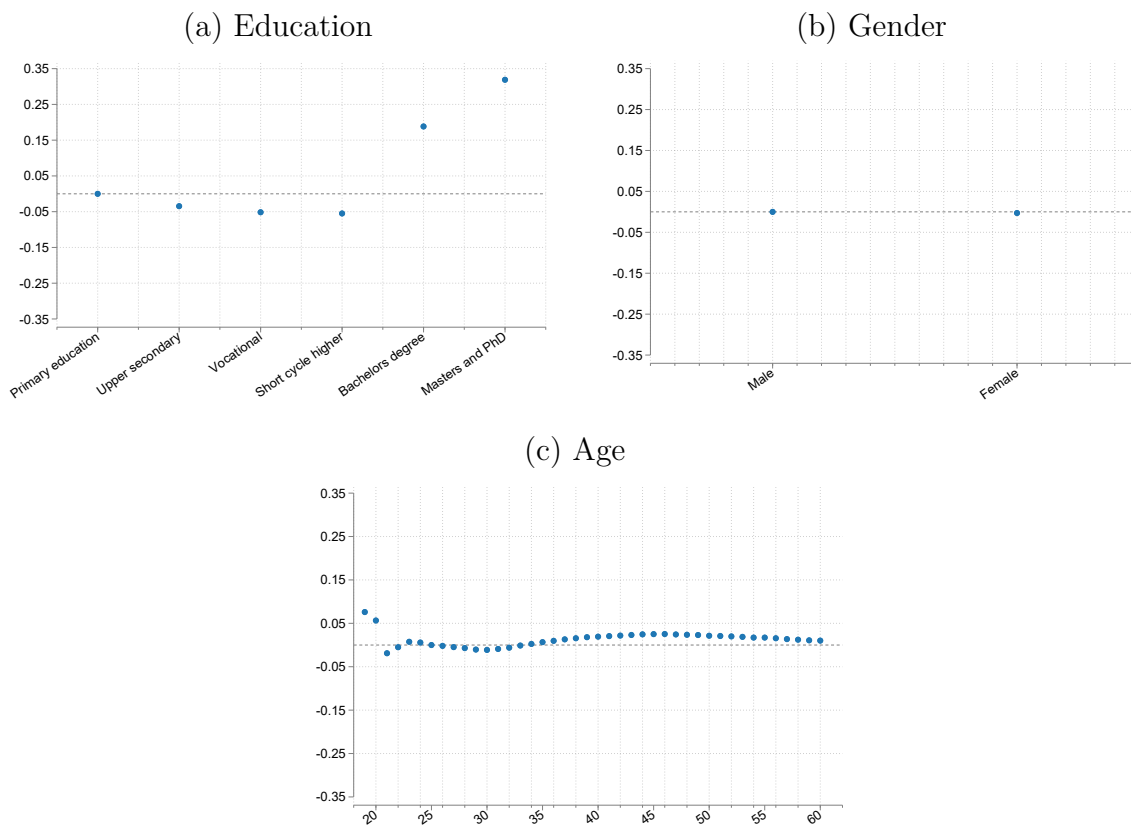
³²The firm wage fixed effects in very small micro firms (1-4 employees) are likely imprecisely estimated, thus biasing the average wage sorting contribution—the product of the firm and worker wage fixed effects—toward zero.

Figure D.1: The wage sorting contribution projection and firm characteristics: Entry cohort, exit cohort, industry, location and size



Notes: The estimated coefficients on observable firm characteristics are obtained from (13). Panels (a) and (b) are reproduced from Figure 5 in the main text.

Figure D.2: The wage sorting contribution projection and worker characteristics: Education, gender, and age



Notes: The estimated coefficients on observable worker characteristics are obtained from (13).

D.2 Observable worker demographic heterogeneity

Figure D.2 plots the estimated education effects ζ_1 from (13) in panel (a), gender effects ζ_2 in panel (b), and age effects ζ_3 in panel (c). Each coefficient represents the average difference in wage sorting contributions between matches involving a worker with a given characteristic and the reference worker, measured in correlation points. The education effects in panel (a) are discussed in Section 5. Panel (b) shows no systematic gender differences in wage sorting contributions, and panel (c) shows negligible age effects.

E Wage sorting and job changes: Additional results

This appendix provides additional results for the event study analysis of Section 6. Table E.1 reports full summary statistics for the event study data. Table E.2 reports the complete event time profiles underlying the summary in Table 6 in the main text. Figure E.1 plots the full event time profiles pooling across all years and Figure E.2 plots by decade .

Table E.1: Summary statistics for the job change event study data

	All	JJ	JUJ	Exit	Entry
Job changes, 1980–1989	269,576	194,718	74,858	33,572	30,151
Job changes, 1990–1999	342,707	267,956	74,751	56,663	38,062
Job changes, 2000–2009	443,488	366,550	76,938	67,990	47,114
Job changes, 2010–2019	257,494	214,300	43,194	30,292	16,949
Job changes, all years	1,584,298	1,263,233	321,065	235,138	175,458
Workers, 1980–1989	252,064	183,711	72,842	33,426	29,960
Workers, 1990–1999	310,923	246,045	72,501	56,083	37,796
Workers, 2000–2009	394,208	330,167	74,628	67,013	46,738
Workers, 2010–2019	247,056	206,374	42,781	30,167	16,916
Workers, all years	1,026,634	862,320	281,007	219,349	167,990


Notes: JJ refers to job-to-job changes; JUJ refers to job-to-unemployment-to-job changes. Exit refers to the origin firm exiting in the year of the job change; Entry refers to the destination firm entering in the year of the job change.

Table E.2: The impact of a job change on the wage sorting correlation contribution

	All job changes						JJ-changes						JuJ-changes										
	1980s		2000s		2010s		1980s		1990s		2000s		2010s		1980s		1990s		2000s		2010s		
Event impact at $s = -3$	-0.001 (0.001)	0.005 (0.001)	-0.006** (0.001)	-0.009*** (0.002)	-0.001 (0.001)	0.002 (0.001)	-0.009*** (0.001)	-0.013*** (0.002)	-0.002 (0.002)	0.018*** (0.003)	0.010* (0.003)	0.013* (0.004)	-0.001 (0.001)	-0.001 (0.001)	-0.001 (0.001)	0.002 (0.002)	0.018*** (0.003)	0.010* (0.003)	0.013* (0.004)	-0.002 (0.002)	0.002 (0.002)	0.018*** (0.003)	
Event impact at $s = -2$	-0.001 (0.000)	-0.001 (0.000)	-0.000 (0.000)	-0.002 (0.000)	-0.001 (0.000)	-0.001 (0.000)	-0.001 (0.000)	-0.003 (0.001)	-0.001 (0.001)	0.002 (0.001)	0.002 (0.001)	0.004 (0.001)	0.046*** (0.002)	0.053*** (0.002)	0.042*** (0.002)	0.032*** (0.002)	0.034*** (0.002)	0.026*** (0.002)	0.026*** (0.002)	0.046*** (0.002)	0.053*** (0.002)	0.042*** (0.002)	
Event impact at $s = 0$	0.046*** (0.002)	0.053*** (0.002)	0.020*** (0.001)	0.018*** (0.002)	0.060** (0.002)	0.063*** (0.002)	0.017*** (0.002)	0.018*** (0.002)	0.010** (0.002)	0.035*** (0.002)	0.037*** (0.002)	0.042*** (0.002)	0.046*** (0.002)	0.053*** (0.002)	0.020*** (0.001)	0.018*** (0.002)	0.018*** (0.002)	0.018*** (0.002)	0.018*** (0.002)	0.018*** (0.002)	0.018*** (0.002)	0.018*** (0.002)	
Event impact at $s = 1$	0.034*** (0.002)	0.042*** (0.002)	0.012*** (0.001)	0.014*** (0.002)	0.047*** (0.003)	0.049*** (0.002)	0.068*** (0.002)	0.07*** (0.002)	0.017*** (0.002)	0.018*** (0.002)	0.018*** (0.002)	0.018*** (0.002)	0.018*** (0.002)	0.018*** (0.002)	0.018*** (0.002)	0.018*** (0.002)	0.018*** (0.002)	0.018*** (0.002)	0.018*** (0.002)	0.018*** (0.002)	0.018*** (0.002)	0.018*** (0.002)	
Event impact at $s = 2$	0.026*** (0.002)	0.032*** (0.002)	0.003 (0.002)	0.010** (0.002)	0.037*** (0.003)	0.035*** (0.002)	-0.002 (0.002)	0.014*** (0.002)	0.014*** (0.002)	0.025*** (0.002)	0.025*** (0.002)	0.025*** (0.002)	0.026*** (0.002)	0.032*** (0.002)	0.003 (0.002)	0.010** (0.002)	0.018*** (0.002)	0.025*** (0.002)	0.025*** (0.002)	0.026*** (0.002)	0.032*** (0.002)	0.003 (0.002)	
	All job changes						To entry firm						To non-entry firm										
Event impact at $s = -3$	-0.001 (0.001)	0.005 (0.001)	-0.006** (0.001)	-0.009*** (0.002)	-0.001 (0.001)	0.006 (0.007)	-0.033*** (0.006)	-0.015* (0.011)	-0.001 (0.001)	0.005** (0.001)	-0.002 (0.002)	-0.008*** (0.002)	-0.001 (0.001)	-0.001 (0.001)	-0.001 (0.001)	-0.001 (0.001)	0.005** (0.001)	0.005** (0.001)	0.005** (0.001)	-0.001 (0.001)	-0.001 (0.001)	-0.002 (0.002)	-0.008*** (0.002)
Event impact at $s = -2$	-0.001 (0.000)	-0.001 (0.000)	-0.000 (0.000)	-0.002 (0.000)	-0.001 (0.000)	-0.001 (0.000)	-0.001 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)
Event impact at $s = 0$	0.046*** (0.002)	0.053*** (0.002)	0.020*** (0.001)	0.018*** (0.002)	0.059*** (0.008)	0.095*** (0.007)	0.028*** (0.006)	0.040*** (0.011)	0.044*** (0.002)	0.047*** (0.002)	0.047*** (0.002)	0.047*** (0.002)	0.047*** (0.002)	0.047*** (0.002)	0.047*** (0.002)	0.047*** (0.002)	0.047*** (0.002)	0.047*** (0.002)	0.047*** (0.002)	0.047*** (0.002)	0.047*** (0.002)	0.047*** (0.002)	0.047*** (0.002)
Event impact at $s = 1$	0.034*** (0.002)	0.042*** (0.002)	0.012*** (0.001)	0.014*** (0.002)	0.021* (0.008)	0.068*** (0.007)	0.018*** (0.006)	0.004 (0.011)	0.035*** (0.002)	0.039*** (0.002)	0.039*** (0.002)	0.039*** (0.002)	0.039*** (0.002)	0.039*** (0.002)	0.039*** (0.002)	0.039*** (0.002)	0.039*** (0.002)	0.039*** (0.002)	0.039*** (0.002)	0.039*** (0.002)	0.039*** (0.002)	0.039*** (0.002)	0.039*** (0.002)
Event impact at $s = 2$	0.026*** (0.002)	0.032*** (0.002)	0.003 (0.002)	0.010** (0.002)	0.013 (0.008)	0.052*** (0.007)	0.000 (0.006)	-0.026** (0.011)	0.027*** (0.002)	0.029*** (0.002)	0.029*** (0.002)	0.029*** (0.002)	0.029*** (0.002)	0.029*** (0.002)	0.029*** (0.002)	0.029*** (0.002)	0.029*** (0.002)	0.029*** (0.002)	0.029*** (0.002)	0.029*** (0.002)	0.029*** (0.002)	0.029*** (0.002)	0.029*** (0.002)
	All job changes						From exit firm						From non-exit firm										
Event impact at $s = -3$	-0.001 (0.001)	0.005 (0.001)	-0.006** (0.001)	-0.009*** (0.002)	0.010* (0.003)	0.025*** (0.003)	0.028*** (0.003)	0.032*** (0.005)	-0.003 (0.001)	0.002 (0.001)	-0.012*** (0.001)	-0.014*** (0.002)	-0.001 (0.001)	-0.001 (0.001)	-0.001 (0.001)	-0.001 (0.001)	0.002 (0.001)	0.002 (0.001)	0.002 (0.001)	-0.003 (0.001)	-0.001 (0.001)	-0.002 (0.001)	-0.008*** (0.002)
Event impact at $s = -2$	-0.001 (0.000)	-0.001 (0.000)	-0.000 (0.000)	-0.002 (0.000)	-0.001 (0.000)	-0.000 (0.000)	0.000 (0.000)	0.002 (0.000)	-0.001 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)
Event impact at $s = 0$	0.046*** (0.002)	0.053*** (0.002)	0.020*** (0.001)	0.018*** (0.002)	0.099*** (0.006)	0.060*** (0.004)	-0.002 (0.003)	0.029*** (0.006)	0.038*** (0.002)	0.051*** (0.002)	0.024*** (0.002)	0.017*** (0.002)	0.038*** (0.002)	0.051*** (0.002)	0.024*** (0.002)	0.024*** (0.002)	0.038*** (0.002)	0.051*** (0.002)	0.024*** (0.002)	0.038*** (0.002)	0.051*** (0.002)	0.017*** (0.002)	0.017*** (0.002)
Event impact at $s = 1$	0.034*** (0.002)	0.042*** (0.002)	0.012*** (0.001)	0.014*** (0.002)	0.105*** (0.007)	0.060*** (0.004)	-0.002 (0.004)	0.020** (0.006)	0.023*** (0.002)	0.038*** (0.002)	0.014*** (0.002)	0.013*** (0.002)	0.023*** (0.002)	0.038*** (0.002)	0.014*** (0.002)	0.014*** (0.002)	0.038*** (0.002)	0.038*** (0.002)	0.014*** (0.002)	0.038*** (0.002)	0.038*** (0.002)	0.013*** (0.002)	0.013*** (0.002)
Event impact at $s = 2$	0.026*** (0.002)	0.032*** (0.002)	0.003 (0.002)	0.010** (0.002)	0.112*** (0.007)	0.075*** (0.005)	-0.014* (0.004)	0.015* (0.007)	0.014*** (0.002)	0.023*** (0.002)	0.006** (0.002)	0.009*** (0.002)	0.014*** (0.002)	0.023*** (0.002)	0.006** (0.002)	0.006** (0.002)	0.014*** (0.002)	0.023*** (0.002)	0.014*** (0.002)	0.023*** (0.002)	0.023*** (0.002)	0.009*** (0.002)	0.009*** (0.002)


Notes: Standard errors in parentheses, clustered at the individual level. Statistical significance is denoted by stars: * indicates significance at the 10% level, ** at the 5% level, and *** at the 1% level. s is event time normalized so $s = 0$ for the year where the job change took place, and all effects are measured relative to $s = -1$. JJ-changes refers to job-to-job changes; JuJ-changes refers to job-to-unemployment-to-job changes. An entry firm comes into existence in the year of the job change. An exit firm is not in existence in the year after the job change.

Figure E.1: Match-level wage sorting contributions and job change events, all years



Notes: Each panel displays the event profiles of match-level wage sorting contributions for a job change at event time $s = 0$, and all effects are measured relative to $s = -1$. Panel (a) categorizes job changes based on whether the worker transitions are job-to-job (JJ) changes (in blue) or job-to-unemployment-to-job (JUJ) changes (in orange). Panel (b) categorizes job changes based on whether the worker transitions to an entry firm (in orange) or an established firm (in blue). Panel (c) categorizes job changes based on whether the worker transitions from an exit firm (in orange) or a surviving firm (in blue). The figures further show the 95% confidence intervals, with standard errors clustered at the individual level.

Figure E.2: Match-level wage sorting contributions and job change events by decade



Notes: Each panel displays the event profiles of match-level wage sorting contributions for a job change at event time $s = 0$, and all effects are measured relative to $s = -1$. The panels in the first row categorize job changes based on whether the worker transitions are job-to-job (JJ) changes (in blue) or job-to-unemployment-to-job (JUJ) changes (in orange). The panels in the second row categorize job changes based on whether the worker transitions to an entry firm (in orange) or an established firm (in blue). The panels in the third row categorize job changes based on whether the worker transitions from an exit firm (in orange) or a surviving firm (in blue). The figures further show the 95% confidence intervals, with standard errors clustered at the individual level.