The Scarring Effects of Firm Shutdowns on Workers' Wages: A Distributional Perspective*

Johannes Seebauer[†] Matteo Targa[‡] Johannes König[§] Maximilian Longmuir[¶]

April 8, 2024

Abstract

To shed light on the differential impact of firm shutdowns across the distribution of workers, we adopt the wage determination framework of Bonhomme, Lamadon, and Manresa (2019) to uncover workers' unobserved types. Worker types relate to workers' position in the wage distribution: all else equal, a higher type implies higher wages. We use the universe of social security records of Italy's Veneto region, one of the leading Italian regional economies, from 1975 to 2001. Based on this rich matched employer-employee data, we measure wage losses after firm shutdowns for different worker types using an event-study framework. Aggregate wage losses directly after a shutdown are 4.5% of the daily wage, which almost halves after six years. This aggregate trajectory masks stark heterogeneity: top-type workers face initial losses of 12.4%, which remain persistent even after six years. Conversely, initial losses for bottom-type workers are 2.6%, which become statistically insignificant after six years. We identify losses in firm tenure as the main source of wage reductions following the shutdown of a worker's firm. Finally, we show that the AKM model (Abowd, Kramarz, and Margolis, 1999), the current workhorse model of wage determination, does not capture this heterogeneity and can lead to misleading conclusions regarding the sources of wage losses.

Keywords firm shutdowns, wage losses, unobserved heterogeneity

JEL Classification J63, J65, J31

*Acknowledgements: The authors extend their gratitude to Elena Manresa and Stéphane Bonhomme for valuable insights into BLM. Further we wish to thank Martyn J. Andrews for his insightful comments. We thank Robin Jessen, Carsten Schröder as well as the participants of the AIEL 2023, the IAAE 2023, the EEA/ESEM 2023, the EALE 2023, the VfS 2023, and of the IAB workshop on "Identifying Worker and Firm Quality" for helpful comments and insights.

[†] DIW Berlin, Mohrenstr. 58, 10117 Berlin, Germany (jseebauer@diw.de), corresponding author

[‡] Roma Tre University, Via Ostiense, 161/165, 00154 Roma RM, Italy (matteo.targa@uniroma3.it)

[§] DIW Berlin, Mohrenstr. 58, 10117 Berlin, Germany (jkoenig@diw.de)

[¶] Stone Center / Graduate Center @ City University of New York (mlongmuir@gc.cuny.edu)

1 Introduction

When firms shut down and lay off their workforce, the consequences can be dire for the individual worker. The employment prospects after a shutdown vary significantly between individuals: some workers will be sought after by firms, while others will not. The literature on firm closures and mass layoffs has been informative about the aggregate consequences of job loss but has had less to say about who suffers the most and who the least (see, e.g., Bertheau et al., 2023; Fackler, Mueller, and Stegmaier, 2021; Farber, 2017; Helm, Kügler, and Schönberg, 2022; Jarosch, 2021; Lachowska, Mas, and Woodbury, 2020; Schmieder, Von Wachter, and Heining, 2023). Besides observable characteristics, be it an individual's education or occupational position, this will hinge in no small part on unobservable factors, such as their ability. A very capable craftsperson stuck in a failing company will likely be quickly picked up by another firm. However, an inept manager in the same circumstances may struggle to find a new placement.

This paper takes a distributional perspective on worker's wage trajectories after firm shutdowns. The main objective of the paper is to quantify the extent to which workers suffer unequally after their layoff. We focus on the differences in the severity and the persistence of wage losses and unpack the underlying forces that drive them.

However, meeting our objective poses a distinct challenge. We would commit classification errors if we simply used workers' positions in the wage distribution at a given point in time as a reference point, as we may assign a worker to a part of the distribution that she would *usually* not belong to. The solution to this problem comes in the form of the wage determination framework developed by Bonhomme, Lamadon, and Manresa (2019), henceforth BLM. It is a dynamic model with two-sided unobserved heterogeneity. Firms are sorted into classes and workers into types. The firm classes result from k-means clustering based on how similar the firms' wage distributions are, and the worker types result from a maximum likelihood estimation of the wage determination model, which accounts for both wages and mobility patterns. BLM, thus, uncovers worker's unobserved types, which determine their wages. Accordingly, we sort workers into a latent wage distribution that would be inaccessible without the use of BLM. Figure 1 demonstrates how the BLM worker type matters for an individual's position in the wage distribution: top worker types sort into the top deciles of the wage distribution, while bottom types generally sort into the bottom of the wage distribution.

We apply BLM to a long-run administrative social security data set from Italy's Veneto region (see, e.g., Card, Devicienti, and Maida, 2014; Bartolucci, Devicienti, and Monzón, 2018). The panel data covers work histories with detailed information on workers' wages, employment spells, and firm identifiers to describe worker mobility between 1975 and 2001. We

Figure 1: Worker Types along the Wage Distribution

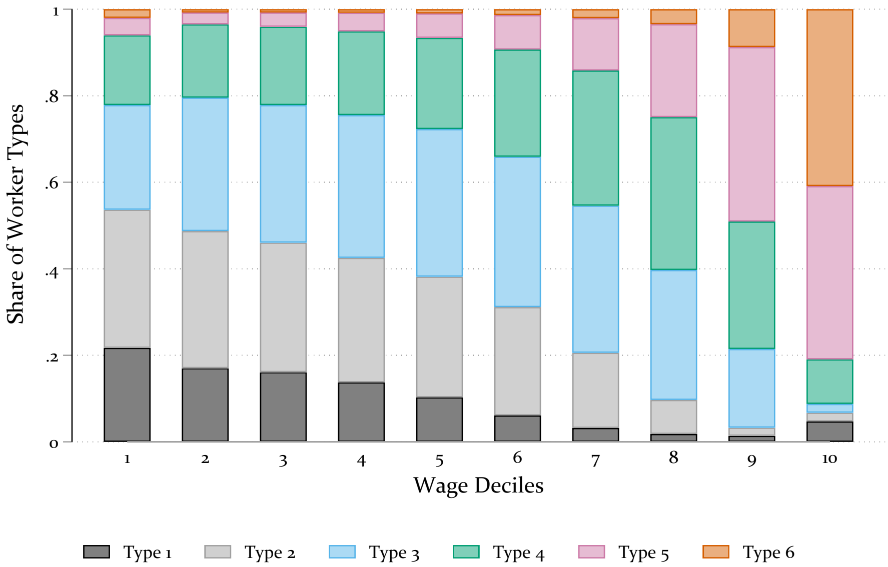

Note: Compiled by authors based on the VWH. The figure shows what share of workers sort into the year-specific deciles of the wage distribution.

follow Bonhomme, Lamadon, and Manresa (2019) and specify ten different firm classes and six different worker types. Sorting among workers and firms is considerable: On average, more than one-fourth of workers in top firms (class 10) are of the highest type (type 6) and more than a third are of the second highest type (type 5). Bottom-tier firms (class 1) hire less than 2% of the top workers.

We then analyze the impact of firm shutdowns on wages in an event-study framework. Based on the estimated worker types, we are able to disaggregate the overall wage losses. On average, one year post-shutdown, wage losses are 4.5%, which almost halve after six years. This aggregate effect, however, masks substantial heterogeneity. Investigating losses by worker type, we find that top types face initial losses of 12.4%, which do not show any sign of recovery over the following six years. Conversely, initial losses for bottom types are 2.6% and statistically insignificant six years later.

We also delve into the heterogeneous consequences of job losses for workers' employment trajectories. On average, one year post-shutdown, displaced workers are 7.1 percentage points less likely to be in employment than comparable non-displaced workers. Six years after displacement, the relative employment losses of treated workers decrease to less than 1 percentage point, implying an almost complete recovery.

Again, there is substantial heterogeneity across worker types. While low-type workers suffer the most substantial losses in the year after the shutdown (more than 11 percentage points), their recovery is the steepest. High-type workers, instead, face slightly lower initial employment losses, which are, however, more persistent throughout the following years. Workers in the middle of the type distribution generally face lower initial losses of about 6-8 percentage points, and recover fully to pre-shutdown levels of employment.

To investigate the mechanisms underlying the observed wage losses, we decompose them into the contributions from general labor market experience, firm tenure, firm classes, and complementarities between workers and firms.

Remarkably, treated workers tend to *upgrade* in firm class, finding jobs in firms paying higher wages than than their previous employer. This shields workers from larger wages losses that they would have otherwise incurred. According to our estimates, the contribution of firm effects to wage losses is rather modest, ranging from about -20% (upgrading) to 15%, against the 40% to 80% found in the literature (e.g., Bertheau et al., 2023; Helm, Kügler, and Schönberg, 2022).

Instead, we find that for both high- and low-type workers, the loss of *firm tenure* plays the dominant role in explaining wage losses. Its contribution ranges from almost 90% to roughly 70% of the aggregated wage loss within the first four years after shutdown and then diminishes to 41% in the sixth period. This suggests that it is the surplus accumulated over time in a given worker-firm match that explains most of the observed wage losses rather than the loss of a job at a particularly well-paying firm.

To reconcile our findings with the literature, we replicate our analysis using the model by Abowd, Kramarz, and Margolis (1999), henceforth AKM. Exploiting firm-to-firm mobility of job movers to identify fixed worker and firm wage components, the authors propose a simple two-way fixed effect linear model. Based on these fixed effects, it is then possible to classify firms and workers into low and high types, similarly to the BLM wage determination model. AKM has been extensively employed in inequality research (e.g., Card, Heining, and Kline, 2013; Devicienti, Fanfani, and Maida, 2019; Mueller, Ouimet, and Simintzi, 2017; Song et al., 2019), and typically serves as the baseline model in the displacement literature (e.g., Bertheau et al., 2023; Fackler, Mueller, and Stegmaier, 2021; Helm, Kügler, and Schönberg, 2022; Lachowska, Mas, and Woodbury, 2020; Schmieder, Von Wachter, and Heining, 2023).

It turns out that AKM paints a fundamentally different picture of wage losses and their sources. First, wage losses are fairly similar in magnitude across all AKM-based worker fixed effects groups. Second, and in line with previous literature relying on AKM, firm effects, or *firm premia*, are a key driver of observed losses, particularly for low-type workers.

We posit that this discrepancy relates to how the two methodologies identify unobserved heterogeneity. If firm-to-firm worker mobility is limited, then estimation error in AKM parameter estimates can be particularly large. This shortcoming of AKM is well-documented in the literature and leads to an *underestimation* of assortative matching and an *overestimation* of the role of firm premia in variance decompositions (Andrews et al., 2008, 2012; Abowd, McKinney, and Schmutte, 2019; Kline, Saggio, and Sølvsten, 2020; Di Addario et al., 2023; Lachowska et al., 2023; Bonhomme et al., 2023). This so-called *limited mobility bias* is particularly pronounced when panels are short and there are many small employers. In our data these concerns are warranted: treated firms are typically small in size and average worker mobility per firm is limited. Thus, in the presence of limited worker mobility, large firm fixed effects will be found among workers with small worker fixed effects and vice versa. By contrast, BLM is by design robust to limited mobility bias as it relies on transitions between clusters of firms rather than between individual firms.

This, however, has important implications for the estimation of *heterogeneity* in wage losses. Our findings suggest that the distribution of AKM fixed effects insufficiently captures heterogeneous patterns in the impact of firm shutdowns on workers' wages in the region of Veneto. Moreover, decomposition exercises based on AKM likely overestimate the contribution of firm effects to wage losses since firms with large fixed effects are falsely associated with workers that have small fixed effects. Thus, mechanically, the wage losses of low fixed effects workers will be ascribed in large part to the loss of the disproportionately large firm effects attributed to them. This is precisely what we find.

Our paper contributes to several strands of labor economics. Although there is an extensive literature on the general association between job displacements and wage losses (Fackler, Mueller, and Stegmaier, 2021; Farber, Hall, and Pencavel, 1993; Farber, Haltiwanger, and Abraham, 1997; Farber, 2017; Huckfeldt, 2022; Jacobson, LaLonde, and Sullivan, 1993; Jarosch, 2021; Lachowska, Mas, and Woodbury, 2020; Topel, 1990) efforts to describe effect heterogeneity are still scarce (Bertheau et al., 2023; Fackler, Mueller, and Stegmaier, 2021; Helm, Kügler, and Schönberg, 2022; Schmieder, Von Wachter, and Heining, 2023). Therefore, we contribute to the growing understanding of the distributional consequences of job displacements, and to a better description of who is most affected by such events. Our study also makes a methodological contribution by demonstrating that the choice of the wage determination model can significantly affect the estimation of wage losses and the understanding of what causes them.

Policymakers can use our results to target those affected by firm shutdowns more effectively. Our results on job matching post-layoff point to a large contribution of decreased sorting, which costs higher types dearly in terms of wages. Since high-type workers lose up to more than 12% of their wage, may want to consider to provide tailored assistance to help these workers search for a better match.

The article is structured as follows: Section 2 discusses recent related literature. Section 3 describes the methods applied, and Section 4 describes the data. Section 5 presents the results, Section 6 details several robustness checks, and Section 7 concludes.

2 Literature Review

There is a long-standing literature analyzing the effect of job displacement on workers' wages. Job displacements 'scar' workers, that is, wage growth trajectories are often permanently altered for the worse (Farber, Hall, and Pencavel, 1993; Farber, Haltiwanger, and Abraham, 1997; Farber, 2017; Jacobson, LaLonde, and Sullivan, 1993; Jarosch, 2021; Lachowska, Mas, and Woodbury, 2020; Topel, 1990). This finding holds internationally, even though the magnitude of the effect differs (Bertheau et al., 2023; Escudero, 2018; Gathmann, Helm, and Schönberg, 2020; Raposo, Portugal, and Carneiro, 2021).

Several explanations have been put forth. One is based on human capital theory, describing the devaluation of firm or industry-specific human capital, e.g., by Topel (1990); Jacobson, LaLonde, and Sullivan (1993), and more recently by Huckfeldt (2022). Another strain of the literature focuses on matching theory, emphasizing that the dissolution of a favorable worker-employer match contributes to reduced workers' wages (Burdett, Carrillo-Tudela, and Coles, 2020; Lachowska, Mas, and Woodbury, 2020; Jarosch, 2021; Raposo, Portugal, and Carneiro, 2021). Recent contributions support this matching-centered perspective, with many studies focusing on the relevance of firm wage premium effects (Bertheau et al., 2023; Brandily, Hémet, and Malgouyres, 2022; Fackler, Mueller, and Stegmaier, 2021; Guerrico and Tojerow, 2021; Helm, Kügler, and Schönberg, 2022; Schmieder, Von Wachter, and Heining, 2023). Following a firm shutdown, displaced workers may lose a favorable match and the associated wage premium, i.e., they fall down the *wage premium ladder*.

The BLM framework allows us to simultaneously study mobility effects related to the firm, like the wage premium ladder, while accounting for the considerable share of wage variation due to individual heterogeneity, and for the complementarities between the two. BLM and its advantages over other wage determination models have recently been discussed (Alvarez et al., 2018; Bagger and Lentz, 2019; Card et al., 2018; Engbom and Moser, 2022). However, only a few empirical applications exist. Bassier, Dube, and Naidu (2022) use BLM to estimate individual separation elasticities in the presence of monopsonistic competition. Lamadon, Mogstad, and Setzler (2022) apply the BLM classification to analyze imperfect competition by estimating the size of labor market rents realized by firms and workers in the US. Palladino, Roulet, and Stabile (2021) analyze the gender wage gap in firm pay premia in France. Applying the BLM framework in the context of an event study is a central contribution of this paper. Specifically, it allows us to explore heterogeneous dynamics after a firm shutdown, and to identify the drivers of the observed wage losses.

3 Method

Our empirical strategy proceeds in four steps. First, we estimate firm classes and worker types using BLM. The worker types are the essential building block for the heterogeneity analysis that follows. Second, we identify workers subject to a firm shutdown and match treated to control workers based on a number of worker and firm characteristics as well as the worker types and firm classes retrieved in the first step. Third, we estimate wage losses of displaced workers compared to matched non-displaced workers in an event study framework; both in aggregate and for each worker type separately. Fourth, we use an auxiliary wage determination model to decompose the observed wage losses into individual contributions of several underlying drivers.

3.1 Recovering Firm Classes and Worker Types

We use the dynamic model proposed by Bonhomme, Lamadon, and Manresa (2019) to categorize workers and firms into latent types and classes. Unlike the workhorse model in labor economics developed by Abowd, Kramarz, and Margolis (1999), BLM allows worker mobility between firms to depend on previous wages realizations conditional on worker and firm attributes. In addition, wages after a move are allowed to depend on the previous firm's characteristics. Finally, BLM includes interactions between worker and firm attributes (match effects), allowing for rich complementarity patterns in wages. Thus, BLM discards three of the major limiting assumptions of the AKM model: 1) exogenous mobility; 2) the absence of state dependence; and 3) the non-complementarity of workers and firms.

Given the improved modeling of the wage determination process, the worker types we recover more credibly reflect unobserved differences across workers. Following Bonhomme, Lamadon, and Manresa (2019), we specify the model with ten firm classes and six worker types in our primary analysis.

Defining unobservable firm heterogeneity at the level of firm classes rather than at the individual firm level has two main advantages: first, the cluster aggregation allows for the inclusion of small firms. In the AKM world, this is typically not possible since all firms need to form a connected set of worker-firm movements. In the BLM framework, such a connected set only needs to exist at the cluster level. Second, since identification requires worker mobility only at the level of firm classes, limited mobility bias is considerably less likely to be a concern.[1]

[1] See also Palladino, Roulet, and Stabile (2021) for a detailed discussion.

We estimate the dynamic version of BLM, which requires wage information for five-year intervals. The model groups workers into two categories: stayers remain in the same firm the entire time, while movers are in the same firm in years one and two and move to a new firm in year three, where they remain employed in years four and five.

The estimation proceeds in two steps[2]: first, firms are classified according to their wage distributions by solving a weighted k-means problem. In particular, firms are grouped into clusters based on how similar their wage distributions are. This step only includes the wages of stayers in the sample. Second, we use the firm classes to retrieve parameter vectors of wage densities and worker-type distributions for stayers and movers in a maximum likelihood procedure. The wage-setting processes in these functions are defined for stayers and movers and expressed as means of wages conditional on the respective period, firm classes, and worker types.

3.2 Matching

Individual workers are considered treated if they were displaced following a firm shutdown. Potential control units are all workers who have never faced a shutdown event. In order to ensure that workers in the control group do not systematically differ from those in the treated group, we apply Coarsened Exact Matching (Iacus, King, and Porro, 2012). We match on the following characteristics: worker type, firm class, sex, year, occupational group (manager, white collar, or blue collar worker), one-digit-industry classification, wage decile, ten-year age and three-year establishment tenure brackets, as well as firm size category. Our matching procedure ensures that we compare only workers within the cells created by the matching characteristics.

We match displaced workers to controls between three and four years before the firm shutdown (at $k = -4$) since wage losses for displaced workers may already occur prior to the actual closure of the firm (see also Helm, Kügler, and Schönberg, 2022; Couch and Placzek, 2010). Our matching procedure works well as Table 1 in Section 4 illustrates.

3.3 Event Study

Since our empirical strategy combines the BLM framework with an event study design, we need to ensure that the classification of workers and firms is unaffected by the event itself, i.e., the firm shutdown. Thus, we estimate worker types and firm classes on five-year-rolling windows. In this way, we obtain year-specific firm class and worker type classifications, ensuring that for a displaced worker, and the respective control units, an up-to-date classification at the time of

[2] A more detailed discussion is provided in Appendix A.

matching is available.[3]

We perform the event study analyses using linear regression and control for individual and year fixed effects as well as a set of dummies for the pre- and post-shock relative periods and their interactions with the treatment dummy. The reference period is the time of matching ($k = -4$). The post-shock interaction dummies measure the average treatment effect on the treated of a firm shutting down. The regression equations take the form

$$Y_{it} = \sum_{k=-6,k\neq-4}^{6} \gamma_k P_{it}^k + \sum_{k=-6,k\neq-4}^{6} \delta_k P_{it}^k \times T_i + \nu_i + \tau_t + \varepsilon_{it} \ , \tag{1}$$

where Y_{it} is the outcome of interest for person i in year t, e.g., log daily wages. $\{P_{it}^k\}_{k=-6,k\neq-4}^{6}$ is a set of relative period-dummies running from -6 to 6, but excluding $k = -4$, the period of matching[4], with a shutdown occurring between period $k = -1$ and $k = 0$ such that $k = 0$ corresponds to the first wage observation post displacement if the displaced worker finds a job directly in the year following the shutdown.[5] T_i is the treatment group dummy, ν_i is an individual fixed effect, τ_t is the year dummy, and ε_{it} is an idiosyncratic error. The coefficients of interest are the δ_k's.

Identification relies on the common trends assumption. This requires that in the absence of the shutdown, the wage dynamics for those who were displaced would have been the same as for the control group.[6] If this assumption holds, the differences in outcomes between control and treatment following the shutdown give the causal effect of the shutdown on wages. A common way to investigate the plausibility of the common trends assumption is to show pre-trends, that is, outcome differences prior to the shutdown. In our case, it is plausible that there are relevant anticipation effects and that wages may change in the run-up to the shutdown. Therefore, it is the pre-trends before matching ($k < -4$) are informative about the common trend assumption.

We do not impose any restrictions on the size of firms. Therefore, our approach differs in this regard from analyses focusing on mass layoffs. The reason for this is twofold: first, the Veneto region includes many small businesses and firms. Focusing on mass layoffs would provide an incomplete picture of the wage impact of firm closures. Second, recent findings suggest that the effects of firm shutdowns highly differ by firm size, highlighting the importance

[3]We investigate the stability of worker types and firm classes over time in Appendix B.

[4]We exclude this period to avoid perfect multicollinearity in the dummy set.

[5]Since we only keep the main spell in a given year, short-lived employment spells occurring between the shutdown and the following main spell are disregarded. Note as well that the first wage observation for a given worker may occur well after $k = 0$ if it takes the worker longer to find a new job.

[6]See, for example, Sun and Abraham (2021) and Goodman-Bacon (2021) for recent expositions on the topic.

of including smaller firms (Fackler, Mueller, and Stegmaier, 2021).[7]

3.4 AKM

To understand how our BLM results compare to the those obtained using the workhorse AKM model, we estimate the following AKM specification:

$$w_{it} = \alpha + \mathbf{X}_{it}^{w}\beta^{w} + \mathbf{X}_{it}^{f}\beta^{f} + \theta_i + \psi_j + \varepsilon_{it}, \tag{2}$$

where w_{it} are log daily wages, $\mathbf{X}^{w}$ and $\mathbf{X}^{f}$ are time-varying worker and firm characteristics, respectively,, θ_i is the worker fixed effect, ψ_j is the firm fixed effect, and ε_{it} is an idiosyncratic error.

Identification of firm and workers effects in AKM is possible within a given connected set of employers linked by worker mobility (Abowd et al., 2002). Such a connected set contains all the workers who have ever worked for any of the firms in the group and all the firms at which any of the workers were ever employed. Within each group, all the parameters in Equation 3.4 can be estimated by OLS and *J-1* firm fixed effects, ψ_j, will be identified.

In our implementation of AKM, we include a cubic of worker age and year-fixed effects as the observed covariates. Further, to mirror the related literature and our implementation of BLM, we estimate a five-year rolling window AKM.

We then construct worker types and firm classes based on the estimated θ_i's and ψ_j's, respectively. In particular, to make the AKM and BLM analyses comparable, we generate ten quantile groups based on the estimated firm fixed effects, and six quantile groups based on the estimated worker fixed effects. This gives us the AKM firm classes and worker types.

3.5 Decomposition framework for wage losses

Our analysis also sheds light on the sources of wage losses. To decompose wage losses from firm shutdowns, we adopt a framework that integrates labor market experience and firm tenure in addition to firm classes and worker effects. The decomposition is a multi-step procedure that builds on the estimation of the wage determination model. When we decompose wage losses after the estimation of BLM, we proceed in three steps. The first step consists of the following auxiliary regression:

[7] In Section H.2, we present an additional analysis where we restrict the size of firms to maximize comparability with the literature on mass layoffs.

$$
\begin{aligned}
w_{it} = & \sum_{k=1}^{K} \zeta_k F_{it}^k + \sum_{l=1}^{L} o_l W_{it}^l + \sum_{k=1}^{K} \sum_{l=1}^{L} \eta_{kl} (F_{it}^k \times W_{it}^l) \\
& + \sum_{l=1}^{L} \lambda_{l1} W_{it}^l \times \text{experience}_{it} + \lambda_{l2} W_{it}^l \times \text{experience}_{it}^2 \\
& + \sum_{l=1}^{L} \lambda_{l3} W_{it}^l \times \text{tenure\_firm}_{it} + \lambda_{l4} W_{it}^l \times \text{tenure\_firm}_{it}^2 \\
& + \nu_i + \tau_t + \varepsilon_{it} \ , \qquad (3)
\end{aligned}
$$

where w_{it} is the log daily wage for person i in year t. ζ_k are the regression coefficients of the indicator functions F_{it}^k. The latter are equal to one when the worker's BLM firm class is k and, analogously, o_l are the regression coefficients of indicator functions W_{it}^l. The latter are equal to one when the worker's BLM type is l. Accordingly, by interacting these indicators, we can estimate η_{kl}, which measures complementarities between firms and workers.[8] The variable experience$_{it}$ measures a worker's labor market experience.[9] The variable tenure_firm$_{it}$ measures the length of the employment spell at the current firm. We estimate separate quadratic labor market experience and firm tenure functions interacted with the indicator function for BLM worker types, W_{it}^l. This gives the experience coefficients, λ_{l1} and λ_{l2}, and the firm tenure coefficients, λ_{l3} and λ_{l4}, respectively. Finally, we include individual fixed effects, ν_i, year fixed effects, τ_t, and the idiosyncratic error ε_{it}.

The second step consists of predicting several components from this auxiliary regression. In total, we calculate four components: 1) the experience component, 2) the firm tenure component, 3) the firm (class) component, and 4) the complementarities component. We use the coefficients ζ_k to predict the firm component, the coefficients λ_{l1} and λ_{l2} to predict the experience component, the coefficients λ_{l3} and λ_{l4} coefficients to predict the firm tenure component, and the coefficients η_{kl} to predict the complementarities component.

The third step consists of estimating the event study coefficients for wage losses but using only one of the wage components as the outcome. The ratio of the component-specific event study coefficient to the event study coefficient on the total wage determines the contribution of

[8] This is similar in spirit to the AKM-style variance decomposition in Bonhomme, Lamadon, and Manresa (2019), where complementarities are measured by the interaction of worker types and firm classes. By contrast, Helm, Kügler, and Schönberg (2022), who decompose wage losses after mass layoffs in the German manufacturing sector in an AKM framework, interpret the regression residual as the match effect. Lachowska, Mas, and Woodbury (2020) estimate match quality by explicitly decomposing the AKM residual.

[9] Because of data restrictions, this variable is measured with error for a subset of the workers in the dataset. Since we we do not observe the date of labor market entry for individuals entering the labor market before 1975 and thus cannot count all spells of employment, we estimate separate experience and firm tenure coefficients for individuals born before 1962. See Section 4 for a description of the variable.

the component to the total wage loss. Note that this ratio can be larger than 1 and can also be negative.

For the AKM decomposition, the procedure is adapted slightly and involves four steps. The first step consists of removing the previously estimated AKM firm effect, $\hat{\psi}_{it}$, from the log daily wage w_{it}:

$$\tilde{w}_{it} = w_{it} - \hat{\psi}_{it} \ .$$

The second step then proceeds with these residualized wages, and we estimate

$$\begin{aligned}
\tilde{w}_{it} = & \sum_{l=1}^{L} \lambda_{l1} \tilde{W}_{it}^{l} \times \text{experience}_{it} + \lambda_{l2} \tilde{W}_{it}^{l} \times \text{experience}_{it}^{2} \\
& + \sum_{l=1}^{L} \lambda_{l3} \tilde{W}_{it}^{l} \times \text{tenure_firm}_{it} + \lambda_{l4} \tilde{W}_{it}^{l} \times \text{tenure_firm}_{it}^{2} \\
& + \nu_i + \tau_t + \varepsilon_{it} \ ,
\end{aligned} \tag{4}$$

where $\tilde{W}_{it}^{l}$ are dummy variables equal to one if a worker belongs to one of L worker fixed effect quantile groups formed by the estimated AKM worker fixed effects. In our case we have $L = 6$, as discussed above. Just as with BLM, we use an auxiliary regression to predict the different components driving the wage loss. However, in the AKM decomposition, firm effects are already estimated outside of the auxiliary regression, and by virtue of the AKM design, there are no terms capturing complementarities. Note that complementarities, therefore, will be contained in the residual.

Steps 3 and 4 are analogous to steps 2 and 3 in the BLM decomposition, although the firm component of the wage loss is directly measured by the AKM firm fixed effects.

4 Data

Our analysis is based on the Veneto Worker History (VWH) file. This is an administrative, matched employer-employee panel data set, which follows workers over time and across different employers. It is obtained from the administrative records of the Italian Social Security System, and it includes the universe of workers employed in the private sector in Veneto[10] from

[10]Veneto is one of the 20 administrative regions in Italy. It represents around 10 percent of the national GDP, and its economy relies on a well-developed manufacturing sector, has a close-to-natural unemployment rate and limited out-migration, making it comparable to other well-developed Western economies (Devicienti, Fanfani, and Maida, 2019).

1975 through 2001.[11] The dataset provides extensive information on labor market characteristics and has been utilized by a large number of researchers in labor economics and related fields (Card, Devicienti, and Maida, 2014; Bartolucci, Devicienti, and Monzón, 2018; Devicienti, Fanfani, and Maida, 2019; Serafinelli, 2019; Kline, Saggio, and Sølvsten, 2020; Fanfani, 2022). For each year in the sample, the data includes information on the job spells of each worker ever employed in Veneto, providing detailed information on the worker's wages and income, job spell length, occupation, age, and gender, all alongside basic information on the employing firm. Importantly, the VWH includes the job spells of those workers who moved outside the Veneto region throughout their careers as long as they remain employed in the private sector. Workers employed in agriculture, public administration, public services (e.g., the health system and railway transportation), and the self-employed are excluded from the sample.

Employment, Earnings, and Wage Information In the analysis, we rely on the primary job spell per year between 1982 and 2001 (Card, Devicienti, and Maida, 2014; Devicienti, Fanfani, and Maida, 2019; Fanfani, 2022).[12] In cases where a worker has multiple job spells in the same year, we select the longest spell in terms of weeks worked. If this criterion is not sufficient, we prioritize the job spell with the highest earnings. In the VWH, earnings are defined pre-tax, including all in-cash benefits, but excluding all in-kind ones. For the main analysis, we rely on (log) *daily* wages from the main job expressed in 2003 euro prices. Daily wages are constructed by dividing the earnings by the number of days in a given main job spell within a year. We drop part-time workers from the analysis.

Residualized Wages The BLM estimation of worker types is based on residualized wages in order to control for basic observable heterogeneity. To do so, we regress log daily wages on a gender dummy, a cubic polynomial of age, their interactions, and a set of yearly fixed effects.

Firm Shutdowns Due to the longitudinal design of the dataset, we can trace the employment trajectories of firms over time. A firm is deemed to have ceased operations when it no longer appears in social security records, as indicated by the absence of any employment affiliations (job spells) linked to the employer in question. Note that within the context of the VWH, the term 'firm' is a broad concept, referring to any unit functioning as an employer. These

[11] See Tattara and Valentini (2010) for details.

[12] Following Devicienti, Fanfani, and Maida (2019), information on earnings is only used from 1982 onwards, that is, we exclude earnings from years in which the wage indexation system (*ScalaMobile*) was strongest. See Leonardi, Pellizzari, and Tabasso (2019) for details. This implies that variables such as tenure and labor market experience are calculated using the entire sample period, whereas BLM model parameters are estimated beginning in 1982.

units may range from single establishments to expansive conglomerates, contingent on how payroll and social security contributions are administered. Each employer unit is allocated a distinct identifier, which is subject to change in the event of a change in ownership or if the employer moves to a new municipality. Therefore, mergers and acquisitions might mask as firm shutdowns.[13] The VWH history archive is, however, complemented with auxiliary information reconstructing the starting and ending date (if applicable) of any job spell recorded, independently from mergers and acquisitions, addressing the issue of potential distortions in firm mortality rates.[14] Ultimately, we confine our analysis of shutdowns solely to firms located in Veneto.[15]

Labor Market Experience and Tenure We calculate tenure as the cumulative sum of months in paid employment registered in the VWH archive. Accordingly, we calculate a worker's occupational and firm tenure as the cumulative sum of months worked in the same occupation, i.e., blue collar worker, white collar worker or manger, and in the same firm, respectively. For these reasons tenure is left-censored at the year 1975, the earliest year in the dataset. We address this censoring as described in the previous section.

Working Sample The working sample consists of treated workers, i.e., workers who have been displaced following the closure of their employing firm, and matched control workers who have never been affected by a firm shutdown.[16] The sample is restricted to male workers between the ages of 18 to 65 and female workers between the ages of 18 to 60, accounting for differences in the eligibility criteria for old-age pension.[17]

Descriptive statistics for the treatment and control group are provided in Table 1. The

[13]Such cases are, however, very limited and occur in only 2.3% of the firms in our sample.

[14]Such information has been previously used in Card, Devicienti, and Maida (2014) and in Bartolucci, Devicienti, and Monzón (2018) to adjust workers' tenure in each job, and it is provided in an auxiliary dataset called *spell.dta* (see Tattara and Valentini, 2010). Because workers are followed across firms that change their identification number due to false mortality episodes, we are able to assign a unique firm ID to all the firms involved. Specifically, the assigned identification number corresponds to the ID of the largest firm among those that are connected by false mortality cases.

[15]Since workers' employment histories are tracked even if they leave the Veneto labor market, the VWH dataset contains information on firms located in all Italian regions. However, the job spells of the entire workforce in these firms are not matched to the data, and reliable information on firm mortality cannot be assessed.

[16]In case a worker experiences multiple shutdowns, we only consider the first one.

[17]Before the 2011 pension reform, the eligibility threshold for old-age pension for private sector workers was a minimum age of 65 years with at least 20 years of paid social security contributions for men, and women had to be at least 60 years old with 20 years of social security contribution. Earlier retirement was possible at any age for workers who had paid social security contributions for at least 35 years (Bertoni and Brunello, 2021). Before 1992, the age cutoff was 60 years for both male and female private sector employees, conditional on 15 years of paid social security contributions. In Appendix H.3, we show that our results are robust to restricting our sample to workers between the age of 25 and 55.

	Treated	Control	Δ	S.E.
Daily Wage	59.479	59.560	-0.081	0.120
Share of Women	0.327	0.327	0.000	0.003
Age	34.800	34.845	-0.046	0.052
Tenure in Months				
in same Firm	114.722	114.312	0.410	0.270
in same Occupation	125.975	124.796	1.179	0.269
Firm Size (share)				
<50	0.731	0.664	0.068	0.003
50-100	0.095	0.140	-0.045	0.002
100-250	0.127	0.147	-0.020	0.002
250-500	0.030	0.032	-0.002	0.001
>500	0.017	0.017	-0.000	0.001
Occupation (share)				
Blue Collar	0.831	0.831		
Manager	0.003	0.003		
White Collar	0.166	0.166		
Worker Type (share)				
1	0.075	0.075		
2	0.188	0.188		
3	0.360	0.360		
4	0.281	0.281		
5	0.074	0.074		
6	0.022	0.022		
Firm Class (share)				
1	0.057	0.057		
2	0.135	0.135		
3	0.173	0.173		
4	0.133	0.133		
5	0.150	0.150		
6	0.096	0.096		
7	0.097	0.097		
8	0.101	0.101		
9	0.048	0.048		
10	0.011	0.011		
N of workers	19,846	101,253		
N of firms at time of shutdown	8,866	34,481		

Notes: Compiled by the authors based on the VWH. The table compares treated and control groups at the time of matching ($k = -4$). Means test results are reported in the third and fourth columns, where we include differences and standard errors respectively.

Table 1: Descriptive Statistics

treatment and control group consists of 19,846 and 101,253 workers, respectively. In total, 8,866 firm shutdowns between 1989 and 1994 are considered.[18] The similarity in all covariates between the two groups demonstrates that our matching procedure works well.

5 Results

This section provides the results of our analyses: first, we start by discussing some descriptive evidence relating to the BLM estimation. Second, we turn to the event study, providing results for wages and employment, both aggregated and by type. Moreover, we discuss transitions between firm classes after firm shutdowns. Third, we decompose the observed wage losses to identify their main drivers for different worker type aggregations. Fourth, we benchmark our results using firm classes and worker types retrieved from the AKM model.

5.1 BLM Descriptives: Worker Types and Firm Classes

BLM uncovers a considerable amount of sorting. In Figure 2, we describe the BLM classification, where the left panel shows the average of residualized log daily wages of the six worker types along the ten firm classes, and the right panel shows the proportions of worker types within firm classes.

[18]This ensures that we have sufficient years pre- and post-shutdown available to track the wage and employment trajectories of affected workers.

Figure 2: Worker Types and Firm Classes

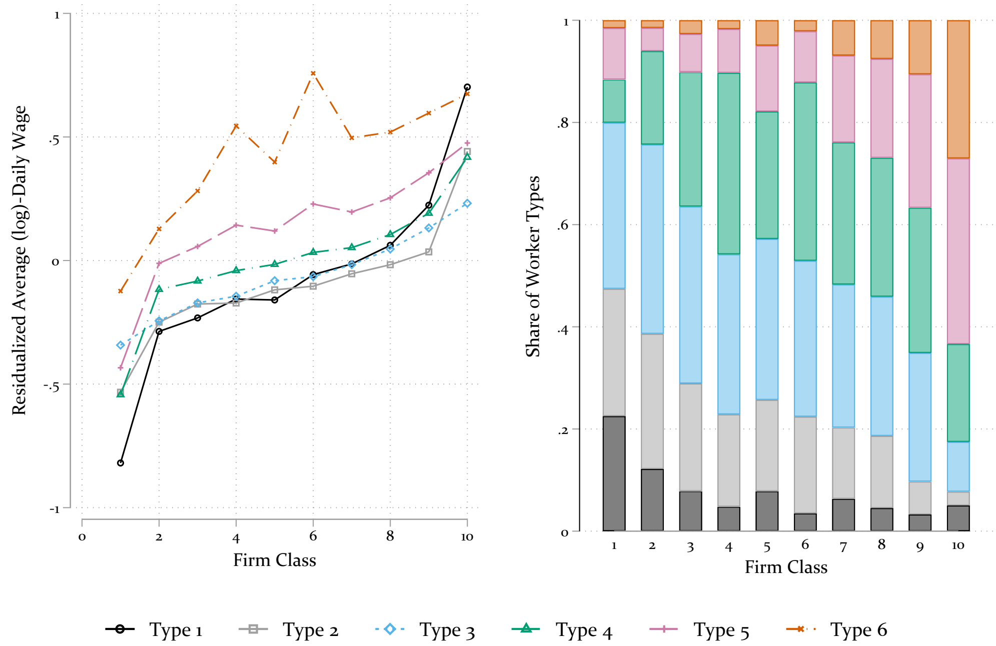

Note: Compiled by authors based on the VWH. The left panel shows the mean log residual wages of worker types along firm classes. The right panel displays the proportion of worker types in different firm classes. The two panels are based on the BLM specification.

The left panel shows that for classes 3 to 9, the worker types are mostly stratified along the wage distribution. For firm classes 1, 2, and 10, the residualized mean log wages of the worker types partly overlap. Two points merit attention: First, workers of type 1 mark an exception in that they manage to overtake worker type 2 to 4 in terms of their wages if they are in firm class 7 to 9 and even earn higher wages than top type workers if they are in the highest firm class. Second, observe that the slope of wages for workers of type 6 changes notably when moving, for example, from firm class 5 to 6. These patterns illustrate the relevance of worker-firm complementarities.

The right panel reveals strong sorting between workers and firms: high-type workers are employed in high-class firms and vice versa. The workforce of top firms (class 10) consists, on average, of 27% top workers (type 6) and 36% of the second highest type (type 5). Bottom-tier firms (class 1) hire less than 2% of top-type workers.

Table 2 in Appendix C describes the BLM worker types along observable characteristics. While top-type workers tend to work in larger firms and are more likely to be managers or white collar workers further down the type distribution, the discrepancy with respect to, for example, age or tenure is rather small and not systematic. Overall, Figure 2 and Table 2 suggest that the classification of worker types goes beyond merely reflecting differences in wages and a number of observable characteristics that are typically deemed important in the wage determination process.

5.2 Event Study

Aggregate Wage Losses We estimate the aggregate effects of firm shutdowns on wage losses based on equation (1). We find that job displacement after firm shutdowns leads to significant wage losses for treated workers compared to the control group. Figure 3 plots six years before and after the firm shutdown, and shows the average wage difference between the two groups. Our estimates show a significant wage drop in the period directly after the firm shuts down ($k = 0$). We find that displaced workers earn, on average, 4.5% lower wages after the shutdown of their firm. These losses shrink to 2.5% after six years.

The wage losses are smaller than what most of the recent literature has found, where they differ both in size and recovery rate. Based on German data, Helm, Kügler, and Schönberg (2022) and Gathmann, Helm, and Schönberg (2020) find around 8% to 9% initial losses with a small recovery. Analyzing US recessions, Huckfeldt (2022) finds losses of around 7% for workers who do not switch occupations, and considerably larger ones, of 18.1%, for those who do. Analyzing Italy, Bertheau et al. (2023) estimate initial losses of around 7%. Only two recent studies, however, are directly comparable with our specification, as they also do not focus on *mass* layoffs. Relating firm closures to the business cycle in Germany, Schmieder, Von Wachter, and Heining (2023) observe wage losses of 9%, which reduce to 6% over the following periods. Also using German data, Fackler, Mueller, and Stegmaier (2021) find initial wage losses of around 5.3%, which fall to 2.9% after five years. They specifically include smaller firms, which makes their estimates highly comparable to our analysis.[19]

In the years prior to the shutdown, the difference between the treatment and control group is close to zero. Importantly, this also holds for the periods pre-matching ($k = -5, -6$), which confirms the quality and reliability of our matching procedure. Note that for the wage observation just before the shutdown, the coefficient is slightly larger than zero. Since this is a pattern that will hold for several of the type-specific event studies, it is useful to explain the mechanics behind this: workers laid off due to an imminent shutdown may receive severance payments. Since workers' wages are simply income divided by days worked in the main job,

[19]To increase comparability of our results to the literature, we exclude firms that have fewer than 30 or more than 500 employees in Appendix H.2. The wage losses are close to the primary analysis. However, the recovery rate is slightly smaller.

Figure 3: Wage Effects of Job Displacement after Firm Shutdown

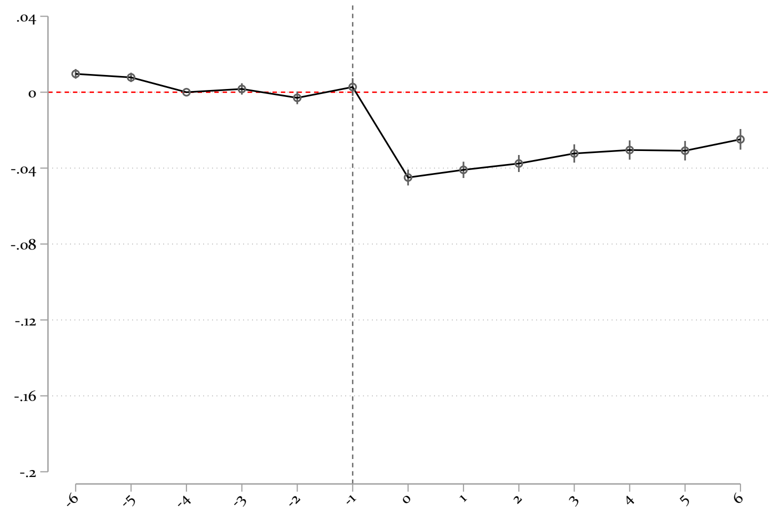

Note: Compiled by authors based on the VWH. The figure displays event study coefficients of log wage losses after a firm shutdown conditional on employment based on the BLM specification. 95% confidence intervals are represented by the horizontal whiskers in the figure.

such severance payments can mask as raises in pay in $k = -1$.

Type-specific Wage Losses Wage losses differ substantially across the six BLM worker types. Figure 4 depicts the event study estimates for each type. Compared to the aggregate, type 1 workers face lower initial wage losses of 2.6%, and the difference to the control group is non-significant four years after the shutdown. Type 2 workers suffer slightly larger initial losses of 3.4%, but follow a similar recovery trajectory thereafter. The initial wage losses of worker types 3 and 4 amount to more than 4%, yet they reduce to about 2% after six years. Strikingly, and by contrast to all other types, the losses of types 5 and 6 show no sign of recovery even after six years. Type 5 workers face initial losses of 8.3%, which hardly change thereafter. The losses of type 6 workers initially lie at 12.4% and continue to hover around 10-12% with no sign of recovery. Finally, for both type 5 and 6 workers initial losses are statistically indistinguishable from the losses six years after the shutdown.

The heterogeneous effects for different worker types have several implications. Even though all worker types are negatively affected by a firm shutdown, the magnitude of losses and the recovery over time are considerably different. As wages increase in worker types, a shutdown appears to be a transitory shock for low type workers, while high-type workers face a persistent

Figure 4: Wage Effects of Job Displacement by Type

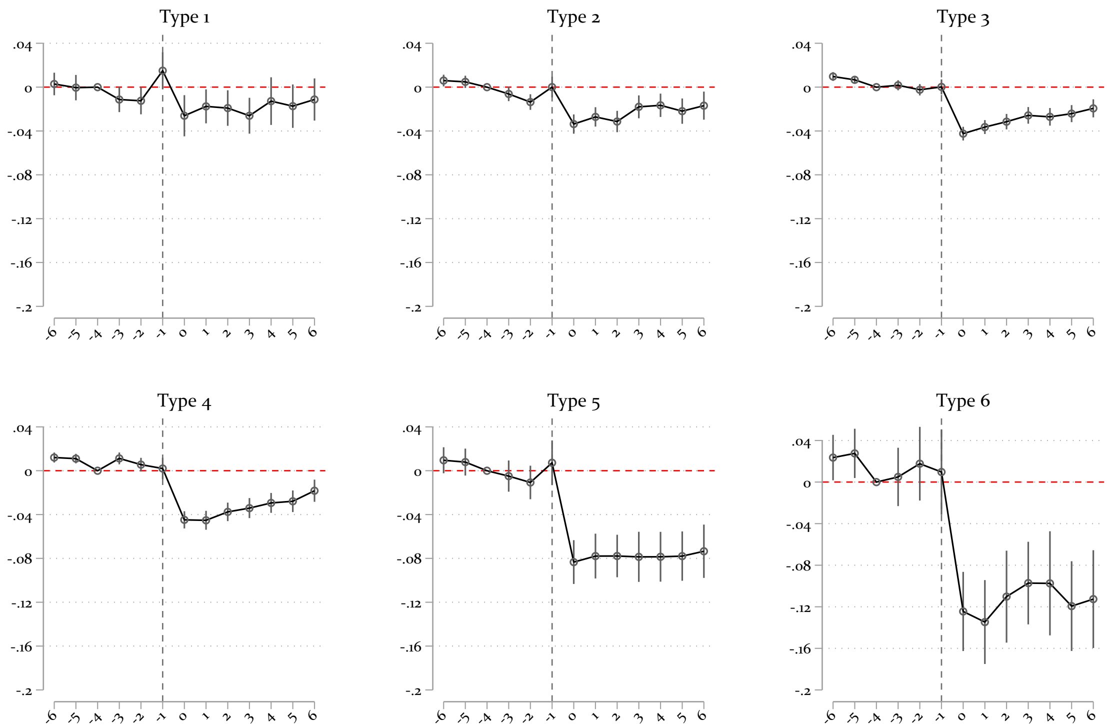

Note: Compiled by authors based on the VWH. The panels display event study estimates for each worker type of the effects of job displacement after a firm shutdown on log-wages, conditional on employment. results in Figure 4 are based on the BLM classification. 95% confidence intervals are represented by the horizontal whiskers in the figure.

drop in their wages.

Worker mobility A reason for the differential impact on wages could be that higher types fall down the wage premium ladder (Bertheau et al., 2023; Brandily, Hémet, and Malgouyres, 2022; Fackler, Mueller, and Stegmaier, 2021; Guerrico and Tojerow, 2021; Helm, Kügler, and Schönberg, 2022; Schmieder, Von Wachter, and Heining, 2023). Since it is difficult and time-consuming to 'climb the ladder' again, we observe a persistent reduction in wages for top types. We examine the relevance of mobility patterns across firm classes post-shutdown in Figure 5. The figure compares changes between the time of the matching ($k = -4$) and the first period post-shutdown ($k = 0$). The panels on the left show regular mobility patterns in the control group; the panels on the right show mobility of displaced workers across firm classes after shutdowns. The upper two panels display transitions of type 1 workers, and the lower two those of type 6 workers.

Figure 5: Transitions Across Firm Classes by Worker Type

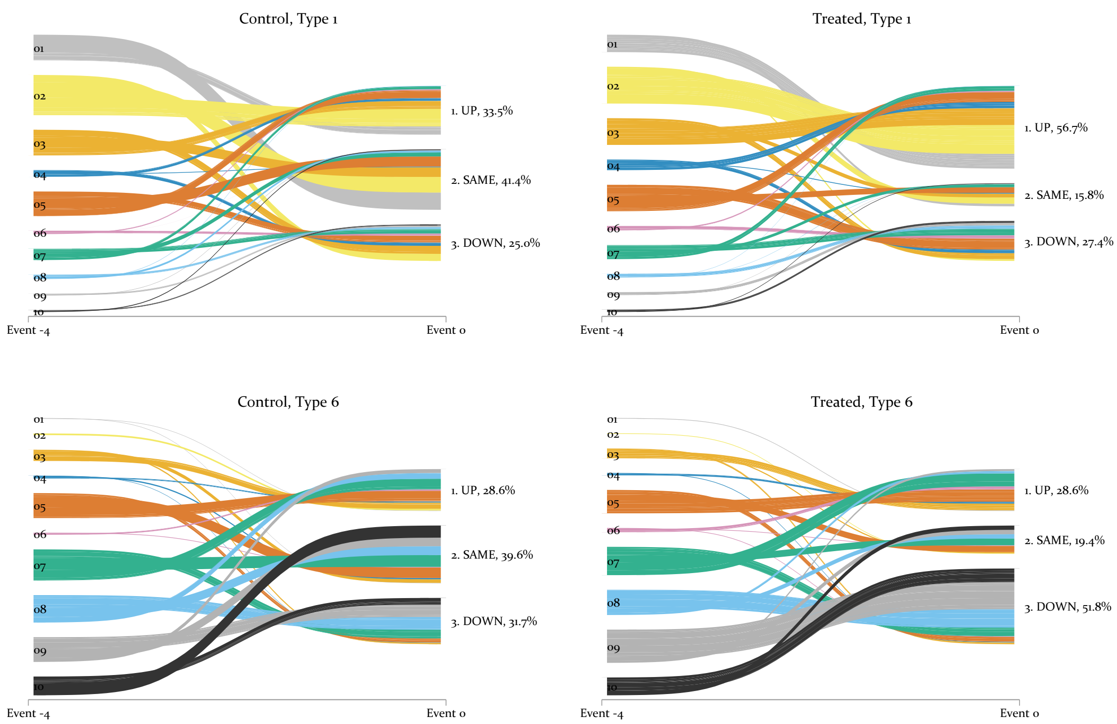

Note: Compiled by authors based on the VWH. The figure displays worker-type specific transition frequencies within and between firm classes for treated and control observations. The figure shows whether workers upgrade, downgrade, or remain in the same firm class.

Type 1 workers are typically found in lower firm classes, while the opposite is true for high-type workers. Interestingly, however, absent a shutdown event, type 1 and type 6 workers up- and downgrade with fairly similar probabilities (left panels). Thus, higher types are not disproportionately more likely to downgrade than low types as one would expect if strong regression-to-the-mean dynamics were present.

Rather, there are significant differences between types within the treatment group. Displaced type 1 workers are significantly less likely to remain in the same firm class than their counterparts in the control group who are not forced to move, which is almost entirely driven by a higher upward mobility. Hence, at least in terms of sorting across firms, these workers appear to benefit from the firm shutdown. One potential explanation is that they were stuck in subpar matches. This may occur, for instance, when workers are not fully informed about their outside options (Jäger et al., 2023).

For displaced type 6 workers the opposite is true: more than half of them downgrade in firm class post-shutdown. In total, they are almost twice as likely to downgrade than displaced type 1 workers. A simple difference-in-differences calculation suggests that the effect of a shutdown on the probability of downgrading is roughly ten times higher for a type 6 than for a type 1 worker. Therefore, firm shutdowns decrease the degree of positive assortative matching.

Figure 20 in Appendix D summarizes the mobility results across all worker types. The figures confirm the findings above: displaced workers of a lower type upgrade, while the opposite is true for high-type workers.

Similarly to Figure 5, Figure 6 compares the mobility of treated and control workers across industries, spanning the time from the matching period ($k = -4$) to the initial post-shutdown phase ($k = 0$). The figure shows outcomes for worker types 1 and 6 in the upper and lower panels, respectively.[20] Two notable observations emerge from this analysis. First, there appears to be little discernible sorting of worker types across sectors, as evidenced by the similar distribution of worker types 1 and 6 over industries. Second, while movements across sectors appear to be the exception rather than the norm, its likelihood increases significantly for workers affected by a firm closure. Specifically, 19.6% of treated type 1 workers and 29.4% of treated type 6 workers transition to different industries, compared to 6.1% and 11.3%, respectively, within the corresponding control groups.

[20] A summary of industry mobility patterns for all worker types can be found in Figure 21 of Appendix D.

Figure 6: Transitions Across Industry by Worker Type

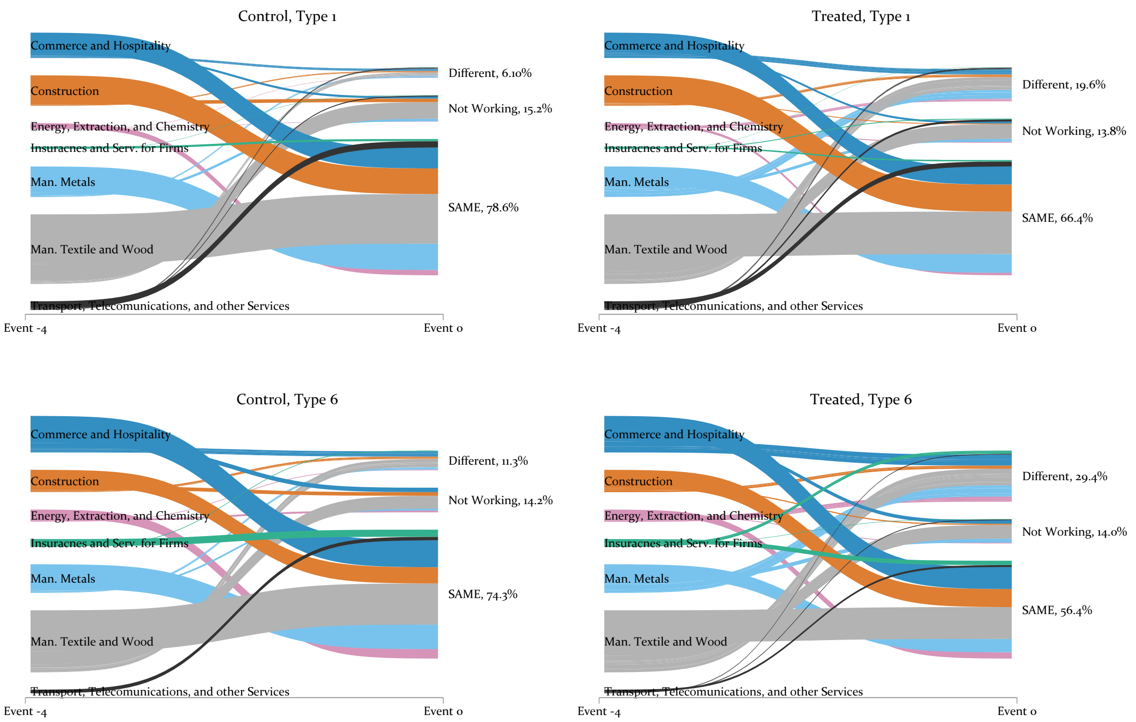

Note: Compiled by authors based on the VWH. The figure displays worker-type specific transition frequencies within and between the industrial sector for treated and control observations. The figure shows whether workers change or remain in the same industry.

Aggregate employment effects Next, we examine differences in employment effects of firm closures in order to complement our analysis of wage losses. Figure 7 shows the employment trajectory of displaced workers relative to non-displaced workers up to six years post-shutdown. We restrict the sample to those individuals who will regain employment at some point during their labor market career.[21] This implies that the observed employment effects constitute a lower-bound estimate.

The figure reveals that displaced worker are, on average, 7.1 percentage points less likely to be in employment in the year following the shutdown event. Thereafter, however, there is a fairly steep increase in the probability of reemployment. Six years after displacement,

[21] We observe employment spells until the year 2001, which is the last year in our data. Displacement events considered in our analysis occur between 1986 and 1991. Therefore, individuals are observed for a maximum of 15 years post-shutdown. If they have not regained employment within that time frame, we do not consider them here. We opted for this approach since we do not want to classify individuals as unemployed who may have gone into early retirement, entered self-employment, or left the labor market in favor of non-employment.

the relative employment losses of treated workers decrease to less than 1 percentage point, implying an almost complete recovery.

Figure 7: Employment effects of Job Displacement after Firm Shutdown

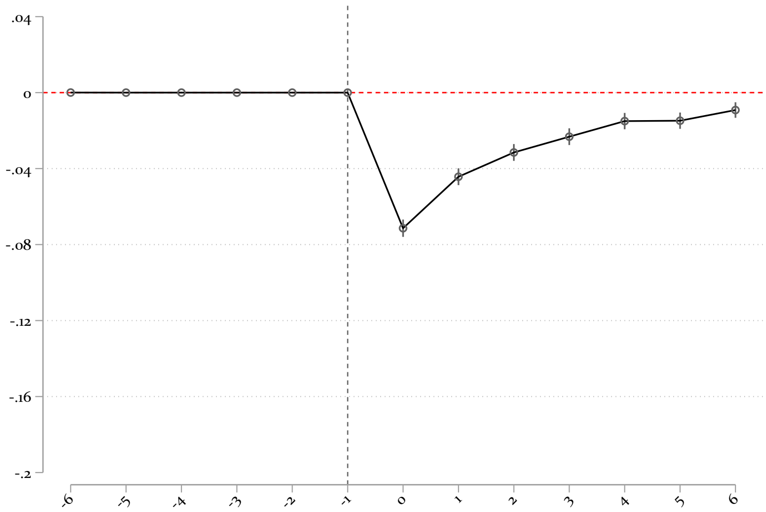

Note: Compiled by authors based on the VWH. Figure 7 displays event study estimates employment losses after a firm shutdown conditional on employment based on the BLM specification. 95% confidence intervals are represented by the horizontal whiskers in the figure.

The effects on employment are smaller in size than, for example, in Bertheau et al. (2023) or Helm, Kügler, and Schönberg (2022), and similar in size to Fackler, Mueller, and Stegmaier (2021) or Schmieder, Von Wachter, and Heining (2023). Interestingly, the employment recovery we observe is typically not seen in studies focusing on mass layoffs. Displaced workers from smaller firms seem to return faster to the workforce, potentially due to less regional competition and other spillover effects implied by mass layoffs.

Type-specific employment effects Figure 8 disaggregates the employment losses by type. While type 1 workers suffer the largest losses in terms of employment - their employment probability drops by more than 11 percentage points in the year after displacement - their recovery trajectory is steep: four periods post-shutdown, the difference in employment between displaced and non-displaced type 1 workers is indistinguishable from zero. Worker types 2 to 5 experience a notably smaller drop in their employment probability (around 6-8 percentage points), and, similar to type 1 workers, they recover quickly over the subsequent years. Displaced workers of type 6, however, are not only almost 10 percentage points less likely to be in employment in the year after the shutdown than their non-displaced counterparts, their drop in employment is also quite persistent. Even six years after the shutdown, the difference in the employment probability between type 6 workers in the treatment and control group is close to 4 percentage points. This implies that the substantive wage losses of top-type workers cannot be solely explained by displaced high-skilled workers trading off wages and employment. While they do tend to find reemployment somewhat quicker directly after the shutdown than type 1 workers, they are considerably less likely to be employed than their non-displaced counterparts even years after displacement; and those who do find a job, earn, on average, significantly lower wages.

While we focus on wage losses and employment effects in our primary analysis, we also discuss the results for yearly earnings in Appendix E. In contrast to daily wages, yearly earnings account for lost income during intra-year non-employment spells. Displaced workers' earnings drop by about 24% in the year following displacement and stay roughly 4-5% lower six years post-shutdown. In line with the combined results for wages and employment, initial earnings losses are particularly pronounced for low (type 1) and high types (types 5 and 6). While the former recover from their earnings losses over time, the latter do not.

Figure 8: Employment Effects of Job Displacement by Type

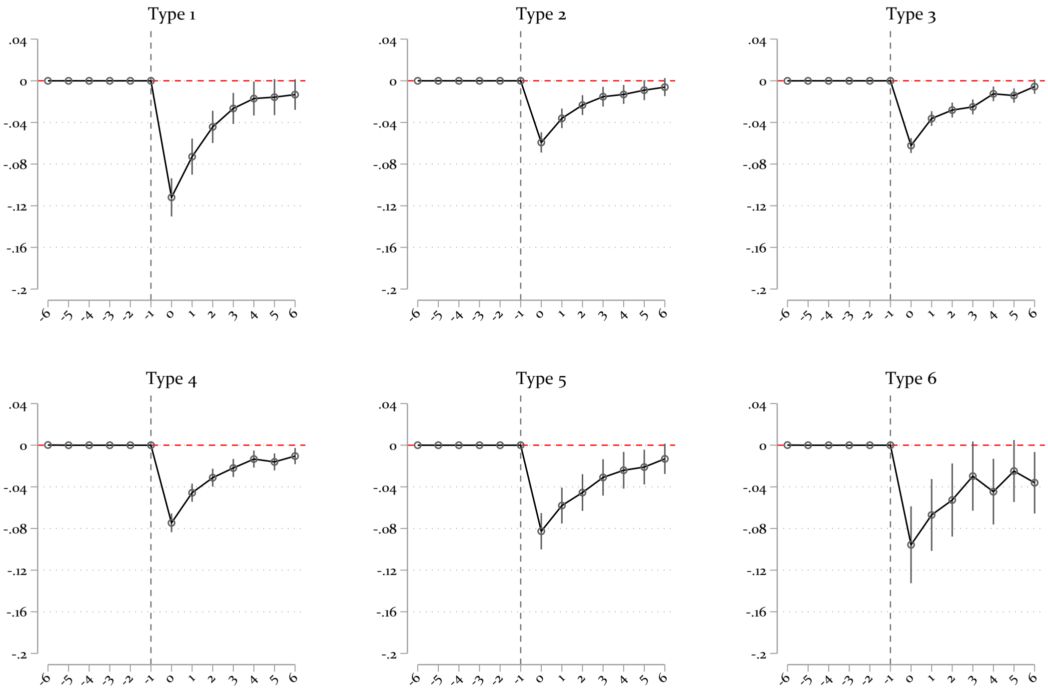

Note: Compiled by authors based on the VWH. The panels display event study estimates for each worker type of the effects of job displacement after a firm shutdown on employment status. Results in Figure 8 are based on the BLM classification. 95% confidence intervals are represented by the horizontal whiskers in the figure.

5.3 Decomposition of Wage Losses

What are the underlying drivers of workers wage trajectories after they have been displaced because of a firm shutdown? To answer this question, we rely on the decomposition framework described in Section 3. We focus on the following sources of wage loss: 1) general labor market experience; 2) firm tenure; 3) firm effects; and 4) complementarities between worker and firm.[22] Firm effects are captured by a dummy set for the BLM firm classes and complementarities are captured by the interaction of firm class and worker type dummies.

Aggregate Losses Figure 9 shows the decomposition of the aggregate wage loss. Almost all of the wage loss can be attributed to the loss of firm tenure for the first four periods. The contribution ranges from almost 90% to roughly 70%. After the first four periods, the contribution of firm tenure diminishes rapidly, from 62% in the fourth year to 41% in the sixth period. Loss

[22] Since our information on occupation is limited to a few broad categories, introducing occupational tenure into the decomposition is not feasible.

of labor market experience plays a minor role for the first four periods as contributions slowly increase from about 6% to about 14%. In period six, the contribution grows to approximately 22%.

These off-setting dynamics can be understood considering what we have shown regarding employment. For most worker types, the bulk of unemployed workers have rejoined the workforce after the fourth period post-shutdown. Workers who rejoin in year 0, that is, immediately after being laid off, hardly lose any labor market experience, but they do lose all of their firm tenure. Workers who rejoin later have lost both their firm tenure and some labor market experience. Therefore, the relevance of labor market experience gradually increases as workers trickle back into employment.

We can understand the contribution of firm effects through the same lens: Workers who find a new job early are likely to accept it from firms that are at least of the same quality as their old firm, that is, they are more likely to upgrade to a higher firm class, while workers who rejoin the labor market later are more likely to accept the offer of a lower pay firm.[23] Therefore, we see a positive effect of the firm class of between 7-10% in the first four periods and a reversal after period 4.

Our findings show that time-invariant complementarities beteween firm classes and worker types do not seem to be relevant for wage losses after firm shutdowns. Lastly, the residual, that is, the part of wage loss we cannot explain using the sum of the components, is fairly small: it amounts to between 10% and 27%. It only exhibits a consistent time trend for the last three years as it grows from about 11% to 27%.

Firm tenure clearly explains a very large share of the overall wage loss. This component, however, picks up time-variant match effects. Surplus stemming from a specific worker-firm match may not materialize immediately, but evolve over time as both parties in the employment relationship adapt more and more to one another. This also suggests that at least some of the wage loss results from workers changing their occupations and industries as they rejoin the workforce. In Appendix F, Figure 24, we augment the decomposition equation with year-specific industry and occupation dummies and make this component explicit. The results confirm that part of the contribution attributed to the loss of firm tenure stems from changes in industry and occupation. For example, the contribution of firm tenure in period 3 is reduced from 73% in the baseline analysis to about 60%.

[23]Note that these dynamics are unlikely to be driven by re-classification of firms within the BLM framework as firms may, in principle, switch classes every year because of our rolling-window implementation of BLM. Hence, if these effects were driven by re-classification, we should be able to see them even earlier after the event. Moreover, we discuss the stability of firm classes in Appendix B, which are found to be quite persistent over time.

Figure 9: Decomposition of Aggregate Wage Losses

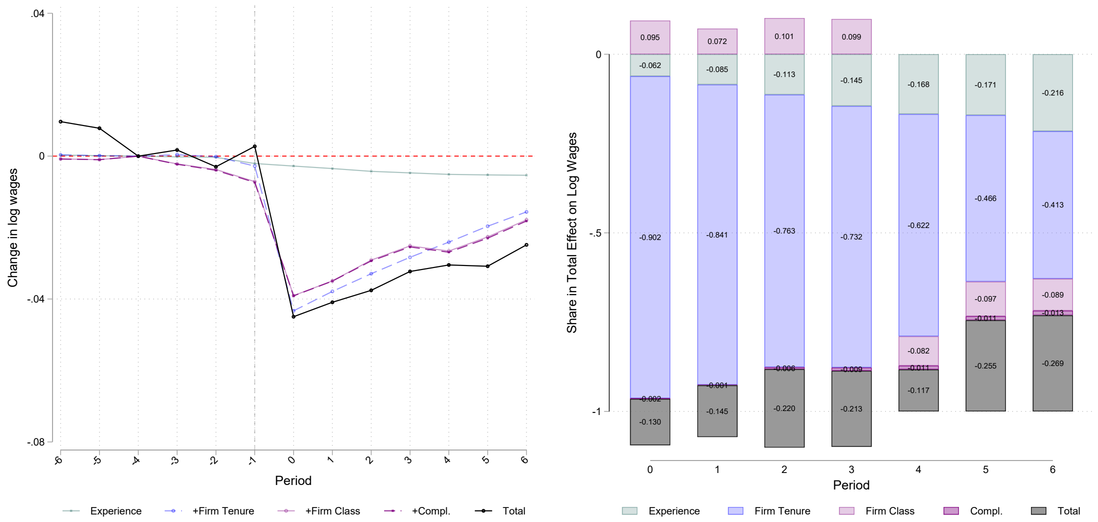

Note: Compiled by authors based on the VWH. The left panel shows the change in log wages pre- and post-event and how much of the wage loss can be attributed to labor market experience, firm tenure, firm effects, and complementarities. The right panel shows the share of the total effect for the six periods after the shutdown. The height of the bars and the numbers within them correspond to minus the source-specific effect divided by the total effect. The total effect, thus, is normalized to -1.

Losses for Low and High Types To reveal whether decomposition dynamics vary across types, we split the distribution of types into two groups: low types and high types. Low types consist of types 1, 2, and 3, while high types are 4, 5, and 6.

Figure 10 displays the decomposition results for low types. The pattern mirrors that of the overall decomposition: there is a positive effect of the firm class for the first four periods and the bulk of the wage loss is attributed to the loss of firm tenure. The role of losses in labor market experience gradually becomes more relevant. An important difference is the magnitude of the positive effect of changes in firm class: it ranges from 18% to 21% for the first four years. This is to be expected, as upgrading in firm class is concentrated among lower types (Figure 5). The decomposition fits the pattern of losses particularly well for low types: at most, the residual contributes about 19% to losses, and in several periods, it is very close to zero.

The pattern shown for high types in Figure 11 is quite distinct from the aggregate decomposition. First, although losses in firm tenure remain the primary contributor to wage loss, their role is less prominent in size. In period 0, they make up about 70% of total losses and contribute only about 26% in period 6. Loss in general labor market experience plays a more prominent role, making up about 7% of total loss in period 0 and growing to roughly 27% in period 6. Most importantly however, although small, firm effects are negative, that is, high types must, on average, work in firms with a lower class compared to their firm prior to the shutdown. There is, however, still a separation between the first four and the final three periods: losses due to firm effects are fairly small in the first four periods (between 3% and 5%) and considerably larger in the final three periods (between 10% and 14%). This aligns well with our previous finding that higher worker types tend to downgrade in firm class (cf. Figure 5). Time-invariant complementarities are, again, an almost negligible component of wage loss, contributing at most roughly 3% to total loss. Finally, the unexplained residual is quite a bit larger for high types than for low types, but it does not exhibit a clear trend over time. The contribution ranges from 18% to 36%.[24]

In sum, the decomposition analysis shows that the loss of firm tenure and, over time, the loss of labor market experience are the premier contributors to wage loss both for high- and low-type workers. The contribution of firm tenure picks up the loss of firm-specific human capital, but also changes in workers' occupations and industry, as Appendix F illustrates. However, our finding that firm tenure, and thus firm-specific human capital, is by far the largest contributor to wage losses, is in stark contrast to the related literature, which emphasizes the loss of firm effects (Fackler, Mueller, and Stegmaier, 2021; Bertheau et al., 2023; Helm, Kügler, and Schönberg, 2022). Especially for low wage workers, several related papers on mass layoffs document a large contribution due to the loss of firm premia that often ranges from 40% to 80% (Bertheau et al., 2023; Helm, Kügler, and Schönberg, 2022). Our context is Italy, where Bertheau et al. (2023) show the contribution of firm effects to be about 40-50%. Instead, our findings suggest that it is not the loss of a job at a particularly well-paying company that explains workers' wage losses but rather the loss of a specific worker-firm relationship, in which benefits have accrued over time.

The question is whether this stark difference in attribution is due to differences in the data or due to differences in the methodology. A clear difference in terms of methodology between the current paper and the literature is our use of BLM, while the rest of the literature generally uses AKM.

A first indication on whether data or methods drive the differences in results can be gleaned by comparing the aggregate wage loss from our sample and that reported by Bertheau et al. (2023) for Italy as a whole. They report immediate losses of about 6-7% in the aggregate, which is fairly close to the 4.5% we find for Veneto. Hence, a first glance suggests that methodological differences do matter for explaining the differences in firm effects. In the following, we explicitly test this hypothesis, and adopt the design of other studies that use AKM. It turns out that we can indeed replicate the large contribution of firm effects to overall wage losses, in

[24]In Helm, Kügler, and Schönberg (2022), the AKM-based decomposition of wages losses for high-wage workers attributes up to more than half of the overall wage losses to the residual component.

line with the findings of related literature using an AKM decomposition framework. We also offer potential explanations for the discrepancy in findings resulting from the two approaches.

Figure 10: Decomposition of Wage Losses, Low Types

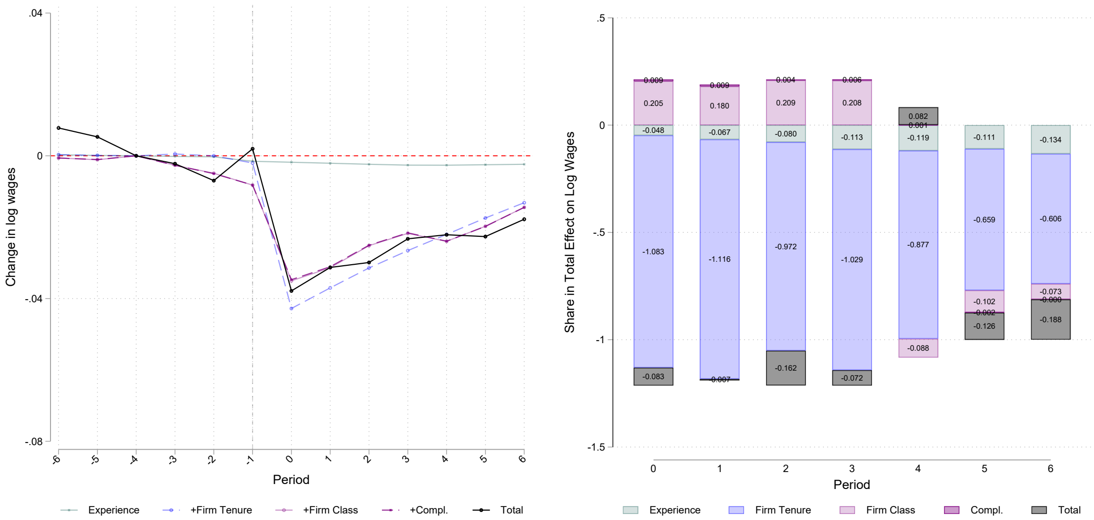

Figure 11: Decomposition of Wage Losses, High Types

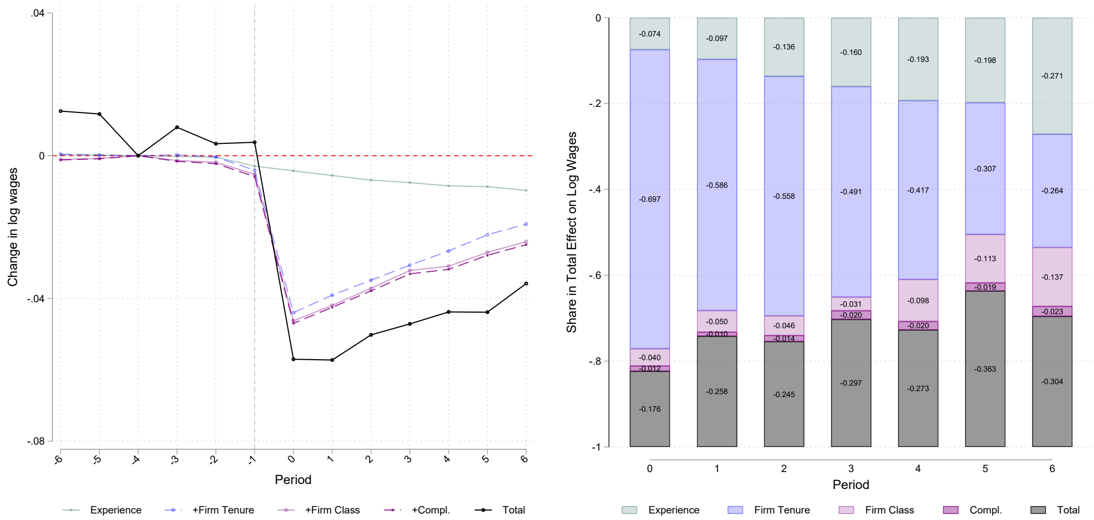

Note: Compiled by authors based on the VWH. The left panel shows the change in log wages pre- and post-event and how much of the wage loss can be attributed to labor market experience, firm tenure, firm effects, and complementarities. The right panel shows the share of the total effect for the six periods after the shutdown. The height of the bars and the numbers within them correspond to minus the source-specific effect divided by the total effect. The total effect, thus, is normalized to -1.

5.4 Comparison to AKM Results

To facilitate the comparison to BLM, we group workers and firms into a class and type structure that resembles BLM. We sort firms into ten quantile-based classes and workers into six quantile-based types.[25]

Our first comparison is with respect to the sorting between worker and firms. Figure 12, analogously to Figure 2, shows how workers' residual mean log wages stratify along the firm class distribution and how workers and firms sort according to AKM. Note on the left panel of the figure that, unlike for BLM, all wage differences are level shifts that do not overlap. This is by construction since AKM does not allow for worker-firm complementarities. That is, type 1 workers cannot earn more at a form in class 10 than type 2 workers, which is indeed the case for BLM, as shown in Figure 2. In the right panel, we see that sorting is fairly weak compared to BLM, and in some cases, there are indications of negative assortative matching. Notably, type 1 workers make up over 20% of the workers in the top firm class, which is a higher proportion than in any other class. Mirroring this at the other end of the distribution, top type workers make up more than 40% of workers in the bottom firm class, which is also more than in any other class. An important takeaway from this is that workers with a large worker fixed effect will often have a low firm fixed effect and vice versa.

Next, we directly compare the association between AKM and BLM worker types. Figure 13 shows how the quantiles of the AKM fixed effects distribution sort across BLM types. Generally, there is quite some, but definitely no perfect, agreement between the two classifications. The BLM types 2 to 5 do not correspond well with their AKM counterparts. Moreover, types 1 and 6 also align only for 40%, and 49%, respectively and 11% of the bottom BLM types are classified as top types in the AKM.

Limited Mobility Bias What is the reason for the discrepancy in sorting patterns obtained from AKM and BLM, respectively? The most plausible explanation is the well-known limited mobility bias. The AKM model identifies the firm and worker fixed effects based on worker mobility across firms. However, Andrews et al. (2008) prove and Andrews et al. (2012) empirically demonstrate that when the number of firm-to-firm job movers is limited, the vector of estimated firm and worker fixed effects suffers from standard least squares estimation error. Specifically, worker and firm fixed effects are negatively and linearly related via the two-way fixed effects specification: if a worker fixed effect is estimated with a positive and large estimation error, then the corresponding firm fixed effect must be estimated with an oppositely signed

[25] We could also use k-means clustering to arrive at the AKM classes and types. However, since the related literature generally explores heterogeneity by sorting on quantiles of the firm and worker fixed effects, we find it most instructive to do the same.

Figure 12: Worker Types and Firm Classes (AKM)

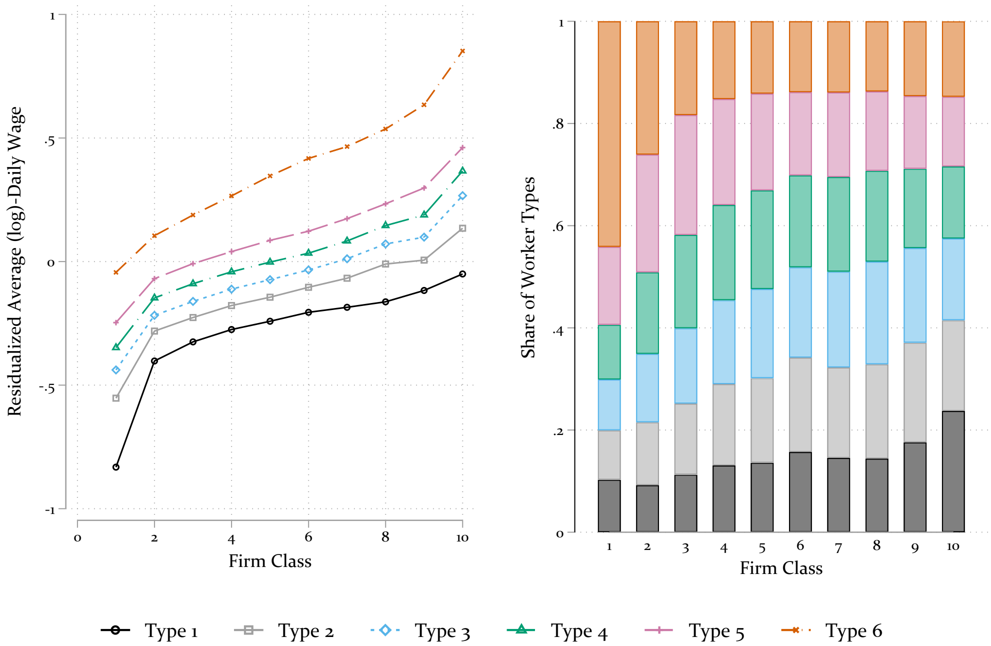

Note: Compiled by authors based on the VWH. The left panel shows mean log wages of worker types along firm classes. The right panel displays the proportion of worker types in different firm classes. The two panels are based on the AKM specification.

Figure 13: Shares of Worker Types: BLM and AKM

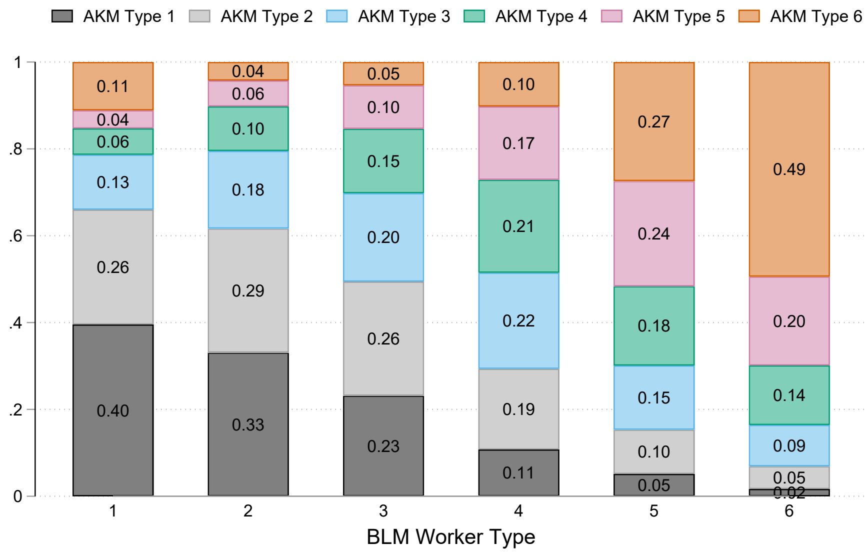

Note: Compiled by authors based on the VWH. Figure 13 displays the six worker types according to BLM (x-axis) and AKM (y-axis).

error.[26] This finding has an important implication: In presence of limited worker mobility, we will most likely find large firm fixed effects for workers with small worker fixed effects and vice versa. AKM, therefore, risks to *underestimate* sorting (highly productive workers selecting into highly remunerative firms) within the economy. This shortcoming of AKM is well documented in the literature (Abowd, McKinney, and Schmutte, 2019; Kline, Saggio, and Sølvsten, 2020; Lachowska et al., 2023; Bonhomme et al., 2023; Di Addario et al., 2023).

For our empirical example, the limited mobility bias is particularly relevant. Firms in Veneto are typically small and characterized by low mobility.[27] Specifically, about two thirds of workers are employed in firms with less than 50 employees (Table 1), while the average number of job movers per firm ranges from 3.3 in the interval 1983-1987 to 4.1 in the interval 1996-2000.[28] Note that this is well below the number of movers that Andrews et al. (2012) indicate (>25) for acceptable AKM estimates of the covariance between the fixed effects. We implement the bias correction proposed in Kline, Saggio, and Sølvsten (2020) and report the

[26]For a concise and intuitive explanation, see section 2 in Andrews et al. (2012).

[27]Lachowska et al. (2023) (p.390) explicitly address the fact that Italian firms are characterized by lower mobility rates than Washington State in the US, suggesting extra care in the application of bias correction methodologies within the AKM framework.

[28]Figure 25 in Appendix G reports the average number of job movers per firm in each rolling AKM sample.

corresponding estimates in Figure 26 of Appendix G. The results indeed show large biases in the AKM estimates. We will also discuss how these can help explain the differential role of firm effects in the decomposition of wage losses.

Wage Losses With these qualifications in mind, we can examine wage losses within the AKM framework. We find that wage losses differ considerably when disaggregating them according to the AKM classification compared to the findings based on BLM. Figure 14 shows the wage losses for six worker types based on AKM. It is striking that across all groups, wage losses are fairly similar in magnitude, as for all groups initial losses are close to about 4%. Moreover, all worker types share a very similar (partial) recovery pattern.

Figure 14: Wage Effects of Job Displacement by Type (AKM)

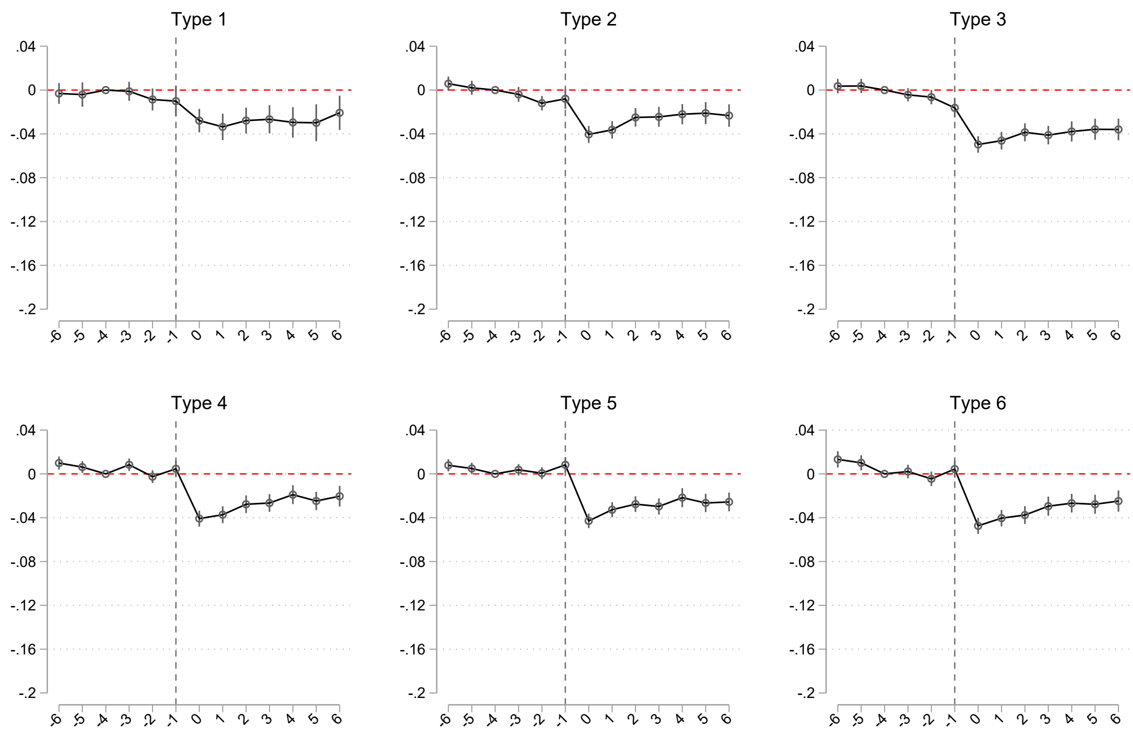

Note: Compiled by authors based on the VWH. Both figures display event study estimates for each worker type of the effects of job displacement after a firm shutdown on log-wages, conditional on employment. Results in Figure 4 are based on the AKM specification respectively. 95% confidence intervals are represented by the horizontal whiskers in the figure.

As discussed above and seen in Figure 13, AKM types mix many workers that are of distinct

BLM types. Hence, the AKM results are some convex combination of the BLM results. Let us consider type 6 according to AKM: it makes up about half of BLM type 6, but the rest sorts into other BLM types with more muted wage loss patterns (Figure 13). Accordingly, wages losses of AKM type 6 will be far less severe than those of BLM type 6. Analogous arguments extend to the rest of the AKM types. In all cases, less than half of a given BLM type consists of the corresponding AKM type.

Decomposition Results Analogous to section 5.3, we show decompositions for the AKM-based wage loss results. To measure the contribution of firm effects to wage losses, we use the AKM firm fixed effects. Note as well that the decomposition based on the AKM framework does not allow for explicitly computing the contribution of complementarities. Instead, these will be part of the residual.

Figure 15: Decomposition of Wage Losses (AKM)

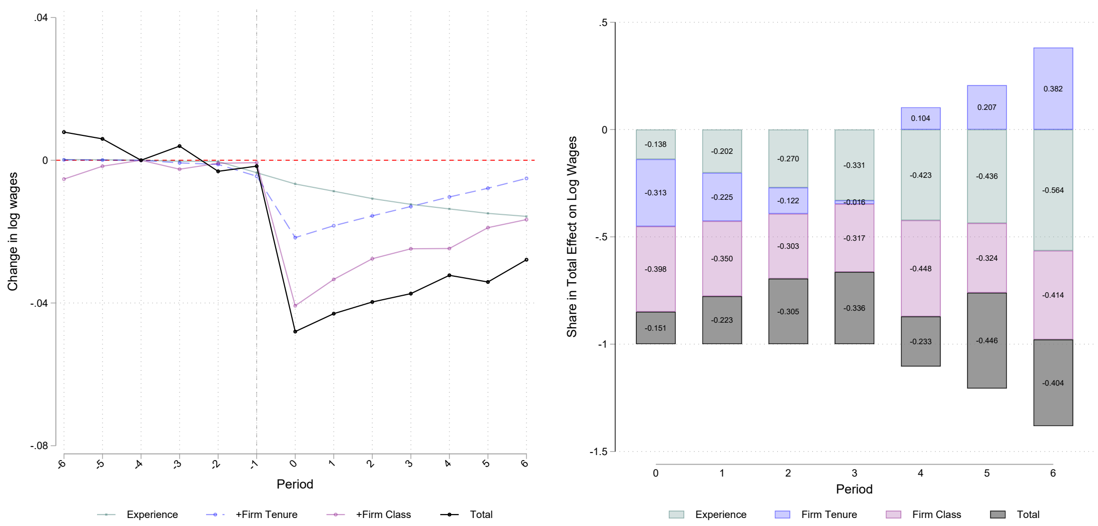

Note: Compiled by authors based on the VWH. The left panel shows the change in log wages pre- and post-event and how much of the wage loss can be attributed to labor market experience, firm tenure, and firm effects. The right panel shows the share of the total effect for the six periods after the shutdown. The height of the bars and the numbers within them correspond to minus the source-specific effect divided by the total effect. The total effect, thus, is normalized to -1.

Figure 15 shows decomposition results with respect to the aggregate wage loss. The most important difference to the BLM results are that firm effects now have a very prominent role in contributing to losses. Firm effects explain between roughly 30% and 45% of total wage loss. This contribution is also fairly stable over time. The role of firm tenure loss is also much smaller compared to the BLM results: while the loss of firm tenure explains about 31% of total wage losses in the period directly following the shutdown, its contribution rapidly shrinks thereafter, and even changes sign in period 4.

By contrast, the role of general labor market experience in explaining the observed wage losses is much more prominent. Its contribution grows from roughly 14% in period 0 to about 56% in period 6.

Figures 16 and 17 show decomposition results for low and high types. Most importantly, we find that firm effects play an even larger role for low types (roughly 50-70%), while for high types firm effects are almost negligible, especially in later periods.

Implications The above analysis provides two key insights:

1) The BLM framework allows for identifying latent worker types based on earnings and mobility patterns. Our event study estimates reveal that this heterogeneity in types translates into substantial heterogeneity in wage losses. By contrast, relying on the AKM fixed effects distribution to identify unobserved heterogeneity is insufficient to capture meaningful differences in the impact of firm shutdowns on workers' wages in the region of Veneto.

2) The methodological differences between BLM and AKM are of particular relevance for the decomposition of observed wage losses. While firm effects play a minor role in the BLM decomposition, they are of utmost importance in the AKM framework. First, it is important to note that we are able to match the large contribution of firm effects to aggregate wage losses that is documented in related literature simply by switching the wage determination framework to AKM. Bertheau et al. (2023) find that firm effects contribute between 43% to 49% to aggregate wage loss in Italy, which is very close to what our AKM decomposition suggests for the Veneto region. However, the separate decompositions for low and high AKM types provide the key piece of evidence for the central role of limited mobility bias in driving the contribution of firm effects to wage losses. Limited mobility bias causes low fixed effects workers to be disproportionately often misclassified as working in firms with large fixed effects, while the opposite is true for high fixed effects workers. Thus, given these mechanics, one would reasonably expect that much of the wage losses of low fixed effects workers is ascribed to the large losses in firm fixed effects, which the AKM framework falsely attributes to them. By contrast, disproportionately small firm fixed effects are attributed to high fixed effects workers, so their potential role in explaining wage losses is limited. This is exactly what we find: firm effects are the main driver of wage losses for low AKM types, while they hardly matter for the losses of high AKM types.

Figure 16: Decomposition of Wage Losses (AKM), Low Types

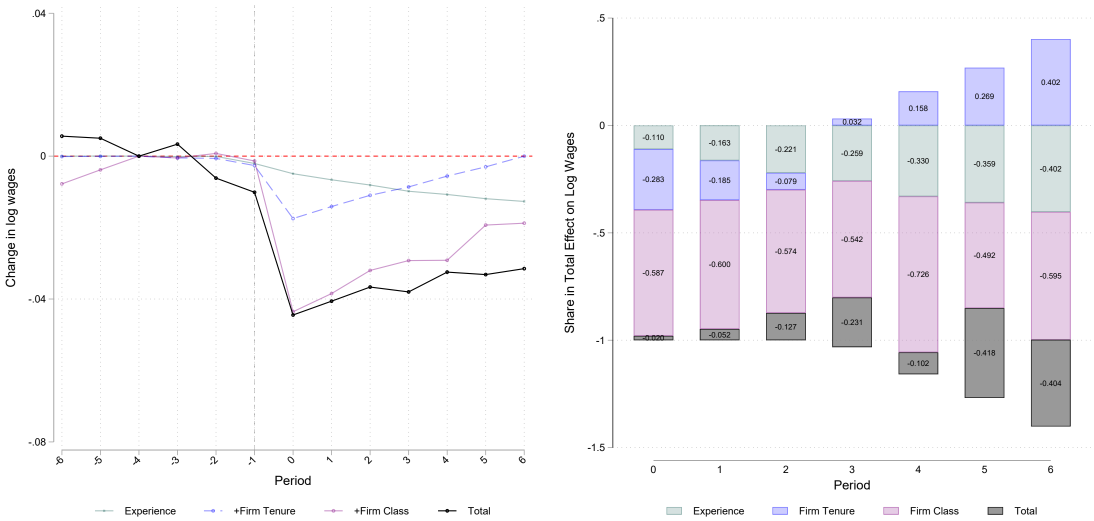

Figure 17: Decomposition of Wage Losses (AKM), High Types

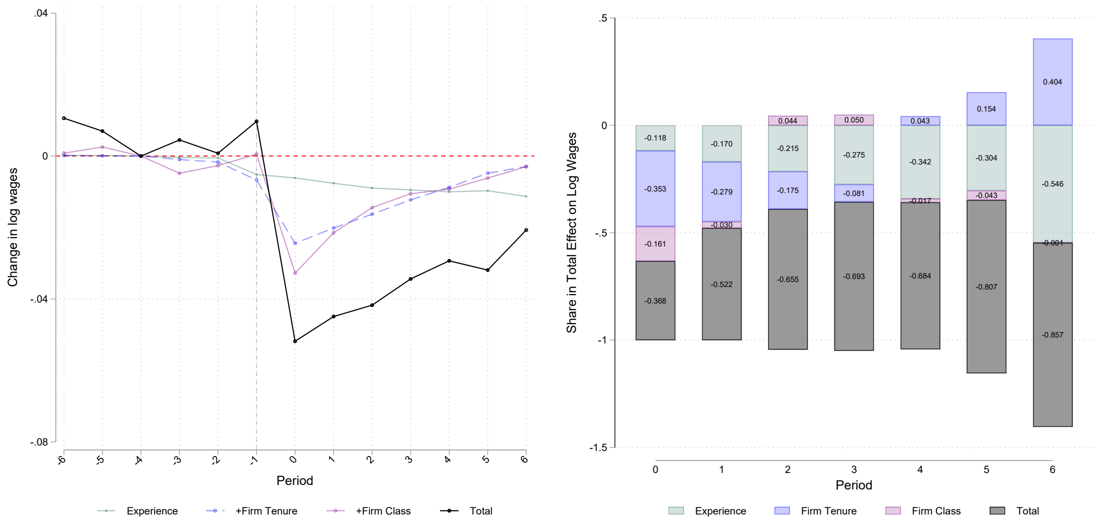

Note: Compiled by authors based on the VWH. The left panel shows the change in log wages pre- and post-event and how much of the wage loss can be attributed to labor market experience, firm tenure, and firm effects. The right panel shows the share of the total effect for the six periods after the shutdown. The height of the bars and the numbers within them correspond to minus the source-specific effect divided by the total effect. The total effect, thus, is normalized to -1.

6 Robustness Checks

To buttress the generalizability of our results, we run several robustness checks.

Number of Firm Classes Our results indicate that the role of firm effects is much smaller than the literature typically suggests, which attributes a large share of the losses to workers falling down the wage premium ladder. By contrast, we find that low-type workers are actually shielded from larger losses by *upgrading* in firm class. It may be conceivable that clustering into ten firm classes insufficiently captures firm heterogeneity and thus ignores important intra-class differences between firms. We pursue this question in detail in Appendix H.1, where we estimate BLM with 20 firm classes instead.

Overall, we find very similar wage loss patterns to our main results. Most notably, the event study estimates for top-type workers are now very noisy. However, this is to be expected since doubling the number of firm classes significantly decreases the size of the worker-firm cells in a group that has the fewest number of observations to begin with.

Reassuringly, the decomposition of aggregate wage losses shows patterns that are both qualitatively and quantitatively very close to the main results. Distinguishing between low and high types, we observe that the importance of firm effects has increased, although in opposite directions. This leads to a dilution of effects when viewed in aggregate. For low types, for example, firm effects counteract 28% (vs. 21% in the main specification) of the wage loss in period 3, while for high types firm effects contribute 16% (vs. 3% in the main specification) to the overall wage loss in that period. Hence, while the role of firm effects becomes more pronounced, as one would expect given the increase in the number of classes, the direction of their contribution is oppositely signed for low and high types.

Firm Size To align our results more with the related literature on mass layoffs and to further scrutinize the credibility to the exogeneity of the firm shutdown, we restrict the sample of firms we consider. We reestimate both the BLM model and the event study with a sample that only features treated firms with at least 30 but no more than 500 employees before the shutdown. Given that the vast majority of treated workers are employed in small firms 1, this reduces the sample size by about two-thirds. Yet, the wage trajectories across worker types are remarkably similar to our main findings. The full results are shown in Appendix H.2.

Prime-Age Workers It is common practice in the literature to focus labor market analyses like ours on workers in the prime of their labor market career. While we put minimal restrictions on the age of workers in our main analysis, we test the robustness of our results to excluding workers younger than 25 and older than 55, respectively. As shown in Appendix H.3, our results hardly differ when only considering prime-age workers.

Worker Type Assignment After obtaining the BLM estimates for a given five-year interval, different options exist for recovering worker types. In our primary analysis, we assign workers the type with the largest likelihood. As shown in Appendix B, this approach provides reasonable estimates of class and type stability over time. However, as this is based on a relatively short time frame, an immediate concern is whether this assignment is relatively noisy and not representative of the worker's 'true' latent type. An alternative is to randomly draw from the estimated type distribution for a given worker. We explore this option in Appendix H.4 and find no relevant differences from our main results. We also check whether a type-stabilization procedure leads to divergent results in Appendix H.5. Instead of fixing the worker type at the time of matching, we assign workers the type to which they were assigned most frequently in the past five years. Results obtained from this procedure are also consistent with our main results. Given that there is some (expected) worker mobility across types (see Appendix B), it is very reassuring that this type stabilization exercise does not meaningfully alter our results.

7 Conclusion

In this paper, we provide a distributional perspective on the effects of firm shutdowns on workers' wages. We use the wage determination model by Bonhomme, Lamadon, and Manresa (2019) to sort workers into distributionally relevant, latent types, and integrate these into an event study framework. We show that wage losses increase in worker types. While lower-type workers experience transitory losses of a smaller magnitude, high-type workers' wages drop precipitously and permanently. As a consequence of these heterogenous losses, firm shutdowns give rise to wage compression.

As our decomposition of wage losses reveals, lower-type workers tend to upgrade in firm class, and benefit from this upgrading in terms of their post-shutdown wage. However, other factors, particularly the loss of firm tenure, far outweigh the effect of upgrading in firm class. For high-type workers, it is considerably more likely to downgrade after a shutdown, yet also for them the loss of firm tenure, as well as the loss of general labor market experience play a more dominant role in explaining their sizable wage losses.

We also demonstrate that we would not have recovered these findings using the more prominent alternative model of wage determination: the AKM model. We show that AKM delivers both a more muted picture of the heterogeneity in wage losses and that it attributes a sizable portion of those losses to the loss of firm effects, especially for low AKM worker types. We argue that this finding is partly a consequence of limited mobility bias as described in Andrews et al. (2008, 2012) and empirically demonstrated in Abowd, McKinney, and Schmutte (2019), Kline, Saggio, and Sølvsten (2020), and Bonhomme et al. (2023), among others.

Our study has revealed for whom a forced separation from their current position is particularly impactful: high BLM worker types, that is, workers that are high up on the job ladder. Their way back to a comparable position at a new firm takes more time and, as the wage loss results indicate, is not always successful. Our decomposition results further show that, in our framework, what matters for losses is not the overall wage policy of the firm, but rather the dynamic evolution of the quality of the new match as indicated by the strong contribution of the firm tenure component. This leaves ample scope for policy-makers to address important questions concerning mismatches, both before and after job loss. This may include the provision of tailored assistance for job search, but also the facilitation of effective training on the job.

References

Abowd, John M, Robert H Creecy, Francis Kramarz et al. 2002. "Computing person and firm effects using linked longitudinal employer-employee data." Tech. rep., Center for Economic Studies, US Census Bureau.

Abowd, John M, Francis Kramarz, and David N Margolis. 1999. "High wage workers and high wage firms." *Econometrica* 67 (2):251–333.

Abowd, John M., Kevin L. McKinney, and Ian M. Schmutte. 2019. "Modeling Endogenous Mobility in Earnings Determination." *Journal of Business & Economic Statistics* 37 (3):405–418.

Alvarez, Jorge, Felipe Benguria, Niklas Engbom, and Christian Moser. 2018. "Firms and the decline in earnings inequality in brazil." *American Economic Journal: Macroeconomics* 10 (1):149–89.

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. Series A (Statistics in Society)* 171 (3):673–697. URL http://www.jstor.org/stable/30135090.

Andrews, M.J., L. Gill, T. Schank, and R. Upward. 2012. "High wage workers match with high wage firms: Clear evidence of the effects of limited mobility bias." *Economics Letters* 117 (3):824–827. URL https://www.sciencedirect.com/science/article/pii/S0165176512004272.

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

Bartolucci, Cristian, Francesco Devicienti, and Ignacio Monzón. 2018. "Identifying sorting in practice." *American Economic Journal: Applied Economics* 10 (4):408–38.

Bassier, Ihsaan, Arindrajit Dube, and Suresh Naidu. 2022. "Monopsony in Movers The Elasticity of Labor Supply to Firm Wage Policies." *Journal of Human Resources* 57 (S):S50–s86.

Bertheau, Antoine, Edoardo Maria Acabbi, Cristina Barceló, Andreas Gulyas, Stefano Lombardi, and Raffaele Saggio. 2023. "The Unequal Consequences of Job Loss across Countries." Tech. Rep. 3.

Bertoni, Marco and Giorgio Brunello. 2021. "Does a higher retirement age reduce youth employment?" *Economic Policy* 36 (106):325–372.

Bonhomme, Stéphane, Kerstin Holzheu, Thibaut Lamadon, Elena Manresa, Magne Mogstad, and Bradley Setzler. 2023. "How Much Should We Trust Estimates of Firm Effects and Worker Sorting?" *Journal of Labor Economics* 41 (2):291–322.

Bonhomme, Stéphane, Thibaut Lamadon, and Elena Manresa. 2019. "A distributional framework for matched employer employee data." *Econometrica* 87 (3):699–739.

Brandily, Paul, Camille Hémet, and Clément Malgouyres. 2022. "Understanding the Reallocation of Displaced Workers to Firms." .

Burdett, Kenneth, Carlos Carrillo-Tudela, and Melvyn Coles. 2020. "The cost of job loss." *The Review of Economic Studies* 87 (4):1757–1798.

Card, David, Ana Rute Cardoso, Joerg Heining, and Patrick Kline. 2018. "Firms and labor market inequality: Evidence and some theory." *Journal of Labor Economics* 36 (S1):S13–S70.

Card, David, Francesco Devicienti, and Agata Maida. 2014. "Rent-sharing, holdup, and wages: Evidence from matched panel data." *Review of Economic Studies* 81 (1):84–111.

Card, David, Jörg Heining, and Patrick Kline. 2013. "Workplace heterogeneity and the rise of West German wage inequality." *The Quarterly journal of economics* 128 (3):967–1015.

Couch, Kenneth A and Dana W Placzek. 2010. "Earnings losses of displaced workers revisited." *American Economic Review* 100 (1):572–589.

Davis, Steven J and Till M Von Wachter. 2011. "Recessions and the cost of job loss." Tech. rep., National Bureau of Economic Research.

Devicienti, Francesco, Bernardo Fanfani, and Agata Maida. 2019. "Collective bargaining and the evolution of wage inequality in Italy." *British Journal of Industrial Relations* 57 (2):377–407.

Di Addario, Sabrina, Patrick Kline, Raffaele Saggio, and Mikkel Sølvsten. 2023. "It ain't where you're from, it's where you're at: hiring origins, firm heterogeneity, and wages." *Journal of Econometrics* 233 (2):340–374.

Engbom, Niklas and Christian Moser. 2022. "Earnings inequality and the minimum wage: Evidence from Brazil." *American Economic Review* 112 (12):3803–47.

Escudero, Verónica. 2018. “Are active labour market policies effective in activating and integrating low-skilled individuals? An international comparison.” *IZA Journal of Labor Policy* 7 (1):1–26.

Fackler, Daniel, Steffen Mueller, and Jens Stegmaier. 2021. “Explaining wage losses after job displacement: Employer size and lost firm wage premiums.” *Journal of the European Economic Association* 19 (5):2695–2736.

Fanfani, Bernardo. 2022. “Tastes for discrimination in monopsonistic labour markets.” *Labour Economics* 75:102107.

Farber, Henry S. 2017. “Employment, hours, and earnings consequences of job loss: US evidence from the displaced workers survey.” *Journal of Labor Economics* 35 (S1):S235–S272.

Farber, Henry S, Robert Hall, and John Pencavel. 1993. “The incidence and costs of job loss: 1982-91.” *Brookings papers on economic activity. Microeconomics* 1993 (1):73–132.

Farber, Henry S, John Haltiwanger, and Katharine G Abraham. 1997. “The changing face of job loss in the United States, 1981-1995.” *Brookings papers on economic activity. Microeconomics* 1997:55–142.

Gathmann, Christina, Ines Helm, and Uta Schönberg. 2020. “Spillover effects of mass layoffs.” *Journal of the European Economic Association* 18 (1):427–468.

Goodman-Bacon, Andrew. 2021. “Difference-in-differences with variation in treatment timing.” *Journal of Econometrics* 225 (2):254–277. URL https://www.sciencedirect.com/science/article/pii/S0304407621001445.

Guerrico, Sofía and Ilan Tojerow. 2021. “Flexible Wages and the Costs of Job Displacement.” .

Helm, Ines, Alice Kügler, and Uta Schönberg. 2022. “Displacement effects in manufacturing.” Tech. rep., Technical Report, mimeo UCL.

Huckfeldt, Christopher. 2022. “Understanding the scarring effect of recessions.” *American Economic Review* 112 (4):1273–1310.

Iacus, Stefano M, Gary King, and Giuseppe Porro. 2012. “Causal inference without balance checking: Coarsened exact matching.” *Political analysis* 20 (1):1–24.

Jacobson, Louis S, Robert J LaLonde, and Daniel G Sullivan. 1993. "Earnings losses of displaced workers." *The American economic review* :685–709.

Jäger, Simon, Christopher Roth, Nina Roussille, and Benjamin Schoefer. 2023. "Worker beliefs about outside options." *Quarterly Journal of Economics* .

Jarosch, Gregor. 2021. "Searching for job security and the consequences of job loss." Tech. rep., National Bureau of Economic Research.

Kline, Patrick, Raffaele Saggio, and Mikkel Sølvsten. 2020. "Leave-out estimation of variance components." *Econometrica* 88 (5):1859–1898.

Lachowska, Marta, Alexandre Mas, Raffaele Saggio, and Stephen A Woodbury. 2023. "Do firm effects drift? Evidence from Washington administrative data." *Journal of Econometrics* 233 (2):375–395.

Lachowska, Marta, Alexandre Mas, and Stephen A Woodbury. 2020. "Sources of displaced workers' long-term earnings losses." *American Economic Review* 110 (10):3231–66.

Lamadon, Thibaut, Magne Mogstad, and Bradley Setzler. 2022. "Imperfect competition, compensating differentials, and rent sharing in the US labor market." *American Economic Review* 112 (1):169–212.

Leonardi, Marco, Michele Pellizzari, and Domenico Tabasso. 2019. "Wage compression within the firm: Evidence from an indexation scheme." *The Economic Journal* 129 (624):3256–3291.

Mueller, Holger M, Paige P Ouimet, and Elena Simintzi. 2017. "Wage inequality and firm growth." *American Economic Review* 107 (5):379–83.

Palladino, Marco Guido, Alexandra Roulet, and Mark Stabile. 2021. "Understanding the role of firms in the gender wage gap over time, over the life cycle, and across worker types." .

Raposo, Pedro, Pedro Portugal, and Anabela Carneiro. 2021. "The Sources of the Wage Losses of Displaced Workers The Role of the Reallocation of Workers into Firms, Matches, and Job Titles." *Journal of Human Resources* 56 (3):786–820.

Schmieder, Johannes F, Till Von Wachter, and Jörg Heining. 2023. "The costs of job displacement over the business cycle and its sources: evidence from Germany." *American Economic Review* 113 (5):1208–1254.

Serafinelli, Michel. 2019. "“Good” firms, worker flows, and local productivity." *Journal of Labor Economics* 37 (3):747–792.

Song, Jae, David J Price, Fatih Guvenen, Nicholas Bloom, and Till Von Wachter. 2019. "Firming up inequality." *The Quarterly journal of economics* 134 (1):1–50.

Sun, Liyang and Sarah Abraham. 2021. "Estimating dynamic treatment effects in event studies with heterogeneous treatment effects." *Journal of Econometrics* 225 (2):175–199.

Tattara, Giuseppe and Marco Valentini. 2010. "Turnover and excess worker reallocation. The Veneto labour market between 1982 and 1996." *Labour* 24 (4):474–500.

Topel, Robert. 1990. "Specific capital and unemployment: Measuring the costs and consequences of job loss." In *Carnegie-Rochester conference series on public policy*, vol. 33. Elsevier, 181–214.

Appendix

A The BLM Estimation

To recover worker types and firm classes, we follow Bonhomme, Lamadon, and Manresa (2019) in using a two-step grouped fixed effects estimator. In the first step, firms are classified according to their wages distributions. This step only includes the wages of stayers in the sample. Including J firms in the sample, the firm classes are estimated by solving the following weighted k-means problem:

$$\min_{k(1),\dots,k(J),H_1,\dots,H_k} \sum_{j=1}^{J} n_j \int (\hat{F}_j(y) - H_{k(j)}(y))^2 d\mu(y), \tag{5}$$

where k(1), ...,k(J), denotes a partition of firms into K classes, $H_1,\dots,H_k$ are cdfs, n_j is the number of workers in firm j, $\hat{F}_j$ represents the cdf of log-wages in firm j, and μ denotes a discrete or continuous measure. In the main analysis, K is set to 10. Equation (5) is minimized with respect to all classes and cdfs. The estimates of $\hat{k}(j)$ are then used to identify worker types α with the maximum number of types $L = 6$ in the next step.

We estimate the parameter vectors of the first two years' wage density θ_{ff}, the last two years' wage density of movers θ_{f^m} and stayers θ_{f^s}, the wage density of movers between classes θ_{f^b}, the density of worker-type proportions θ_q, and worker type proportions for job movers θ_p.

To estimate these parameters, the BLM specifies the wage-setting process with the conditional means in the first period, $Y_{i,1}$, given workers and firm heterogeneity, and $Y_{i,2}$ is defined as $Y_{i,1} = \mu_{1,k} + \rho_{1|2}Y_{i2}$ and $Y_{i,3} = \mu_{k',\alpha} + \rho_{4|3}Y_{i3}$, with $\mu_{k',\alpha}$ as a k',α-specific intercept. The parameters $\rho_{1|2}$ and $\rho_{4|3}$ describe the persistence of log-wages within job.

For the categorization, two attributes of the dynamic BLM model are crucial: First, the presence of endogenous mobility, i.e., a worker's decision to move, and to which firm she moves, depends on the current wage. Second, the presence of state dependence, i.e., for movers, the previous firm has a direct effect on log wages. These attributes are represented in the conditional means, $Y_{i,2} = \mu_{2,k,\alpha} + \xi_2(k')$, and $Y_{i,3} = \mu_{3,k',\alpha} + \xi_3(k)$, for movers between k and k'. Hence, $\xi_2(k')$ ($\xi_3(k)$) defines endogenous mobility (state dependence). This leads to the second step, where wages and mobility parameters are estimated based on both movers' and stayers' wages. The wage density parameters are estimated successively by first estimating $\hat{\theta}_p$. $\hat{\theta}_{ff}$, $\hat{\theta}_{f^m}$, $\hat{\theta}_{f^b}$ with the log-likelihood function:

$$\sum_{i=1}^{N_m}\sum_{k=1}^{K}\sum_{k'=1}^{K}\mathbf{1}\{\hat{k}_{i2}=k\}1\{\hat{k}_{i3}=k'\}\times ... \tag{6}$$
$$\times ln\left(\sum_{\alpha=1}^{L} p_{kk'}(\alpha;\theta_p) f^f_{Y_{i2}k\alpha}(Y_{i1};\hat{\rho}_{1|2},\theta_{f^f}) f^m_{Y_{kk'}\alpha}(Y_{i2},Y_{i3};\theta_{f^m}) f^b_{Y_{i3}k'\alpha}(Y_{i4},\hat{\rho}_{1|2};\theta_{f^b})\right),$$

and then parameters $\hat{\theta}_q$ and $\hat{\theta_{f^s}}$ with

$$\sum_{i=1}^{N}\sum_{k=1}^{K} 1\{\hat{k}_{i2}=k\} \tag{7}$$
$$\times\, ln\left(\sum_{\alpha=1}^{L} q_k(\alpha;\theta_q) f^f_{Y_{i2}k\alpha}(Y_{i1};\hat{\rho}_{1|2},\hat{\theta}_{f^f}) f^s_{Y_{k\alpha}}(Y_{i2},Y_{i3};\theta_{f^s}) f^b_{Y_{i3}k'\alpha}(Y_{i4},\hat{\rho}_{1|2};\hat{\theta}_{f^b})\right),$$

using the EM-algorithm. To reduce the computational burden, we estimate $\hat{\rho}_{1|2}$ and $\hat{\rho}_{5|4}$ in an initial step restricting the covariance structure. Specifically, this means that the effect of worker types on mean log-wages is constant over time within the firm.

B Descriptive Statistics on Type and Class Stability Over Time

For a consistent interpretation of type-specific wage losses some stability of worker types and firm classes over time is helpful. Large fluctuations in the classification would also negatively affect the matching quality. At the same time, however, we follow workers and firms over almost two decades, and thus, changes in the classification are to be expected. In fact, the ability to accommodate such variation over time constitutes an advantage of our framework. Assuming worker or firm characteristics, such as ability or productivity, to be invariant over a long period of time may often be undesirable. However, given the inherent limited mobility bias, studies relying on AKM typically estimate their two-way fixed effects models over extended time periods.

We find that our BLM classification remains reasonably stable over time. Figure 18 shows the share of firms that remain within or switch their class compared to the previous year. We observe that between around 30% and 55% of firms stay in the same class. This proportion is stable over the sample period. The figure also shows that the majority of firms for which the class changes either up or downgrade by one class. Overall, between around 70% and 80% of firms change their class either not at all or by a maximum of one class. Hence, the frequency of changes in firm class is limited.

Figure 19 provides the shares of workers categorized into the same type or switching type over time, in the same fashion as above. We observe more worker than firm mobility over time. Yet again, most types are categorized into the same class, upgraded, or downgraded by one. An exception is the early 1990s, where we observe some larger shares of up- and downgrading by two types. To rule out that such fluctuations drive our results, we 1) use an alternative approach to retrieving worker types from the BLM, and 2) fix types over time, as detailed in Section 6 and Appendix H.4 and H.5.

Figure 18: Stability of Classes Over Time

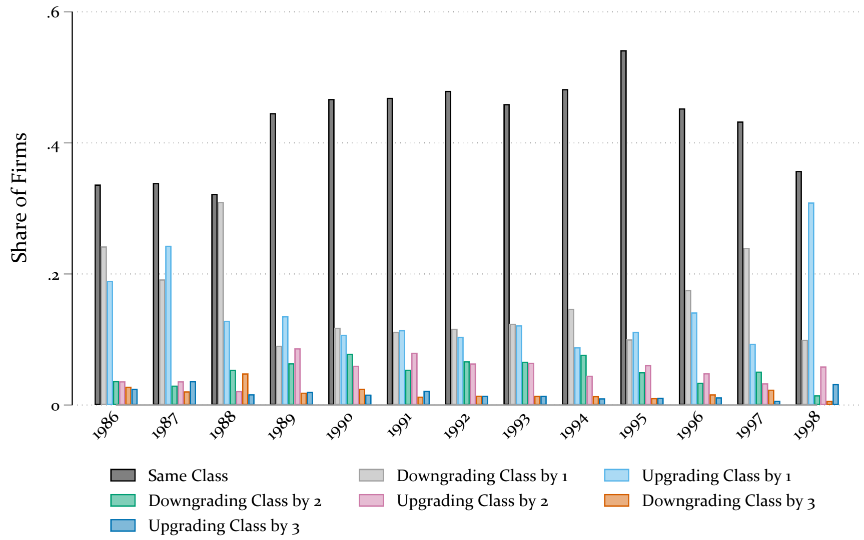

Note: Compiled by authors based on the VWH. Figure 18 provides the share of firms remaining in the same class compared to the previous period, and the share of firm up- or downgrading by one, two, or three or more classes from year to the next.

Figure 19: Stability of Types Over Time

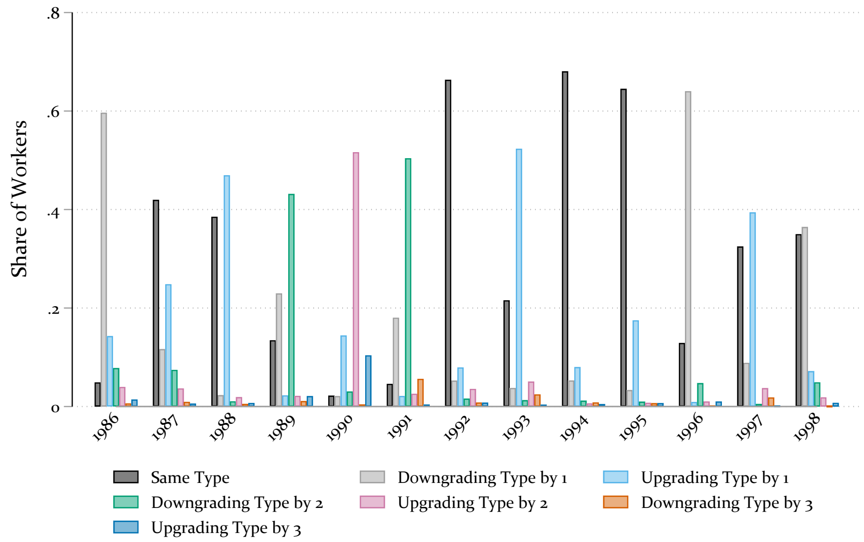

Note: Compiled by authors based on the VWH. Figure 19 provides the share of workers categorized as the same type compared to the previous period, and the share of workers up- or downgrading by one, two, or three or more types from year to the next.

C Descriptive Statistics by Worker Types

Variables			Types					
			1	2	3	4	5	6
log wage		mean	4.069	4.061	4.113	4.178	4.399	4.744
		s.e.	0.003	0.001	0.001	0.001	0.001	0.002
female		mean	0.223	0.225	0.231	0.324	0.305	0.239
		s.e.	0.002	0.001	0.001	0.001	0.001	0.002
age		mean	40.406	38.936	37.996	34.623	34.491	37.788
		s.e.	0.047	0.032	0.023	0.023	0.029	0.046
tenure in months:	firm	mean	110.031	110.836	111.466	105.384	102.671	109.107
		s.e.	0.272	0.160	0.113	0.117	0.163	0.250
	occupation	mean	132.388	128.086	126.607	115.543	112.638	118.568
		s.e.	0.252	0.147	0.106	0.115	0.157	0.234
firm size:	<50	mean	0.475	0.489	0.461	0.457	0.335	0.262
		s.e.	0.002	0.002	0.001	0.001	0.001	0.002
	50-100	mean	0.121	0.116	0.111	0.103	0.104	0.095
		s.e.	0.002	0.001	0.001	0.001	0.001	0.001
	100-250	mean	0.146	0.128	0.120	0.114	0.116	0.121
		s.e.	0.002	0.001	0.001	0.001	0.001	0.002
	250-500	mean	0.086	0.086	0.089	0.084	0.092	0.100
		s.e.	0.001	0.001	0.001	0.001	0.001	0.001
	500+	mean	0.171	0.182	0.218	0.241	0.354	0.422
		s.e.	0.002	0.001	0.001	0.001	0.002	0.002
occupation:	apprentice	mean	0.003	0.001	0.002	0.001	0.001	0.000
		s.e.	0.000	0.000	0.000	0.000	0.000	0.000
	blue collar	mean	0.772	0.826	0.783	0.688	0.424	0.221
		s.e.	0.002	0.001	0.001	0.001	0.002	0.002
	white collar	mean	0.190	0.170	0.215	0.304	0.570	0.671
		s.e.	0.002	0.001	0.001	0.001	0.002	0.002
	manager	mean	0.035	0.002	0.000	0.007	0.005	0.108
		s.e.	0.001	0.000	0.000	0.000	0.000	0.001
Obs.			41029	97776	179114	167486	99439	43275

Notes: Compiled by the authors based on the VWH. The table compares the different worker types. both treated and controls. at the time of matching ($k = -4$).

Table 2: Descriptive Statistics - Worker Types

D Additional Mobility Results

Figure 20 and Figure 21 report for each type the share of workers that change firm class and industry between the time of the matching (event = -4) and the first period post-shutdown (event = 0), distinguishing between treated (red) and control workers (black).

Figure 20: Transitions Across Firm Classes by Worker Type

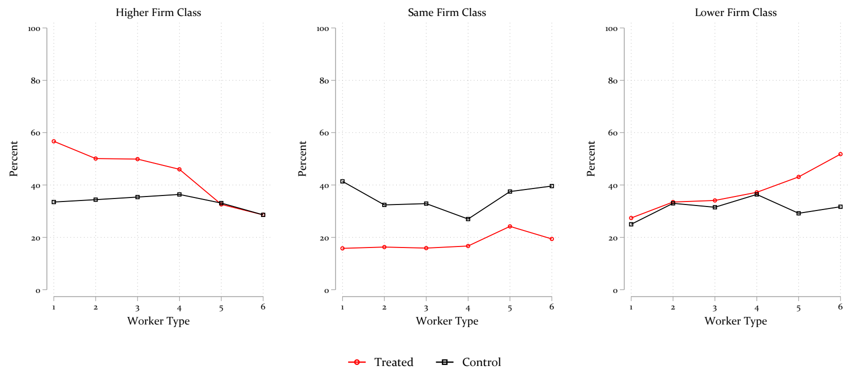

Note: Compiled by authors based on the VWH. The figure displays worker-type specific transition frequencies within and between firm classes for treated and control observations.

Figure 21: Transitions Across Industry by Worker Type

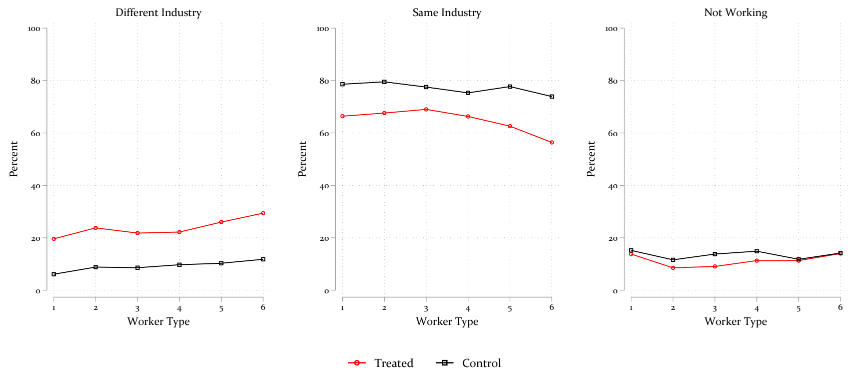

Note: Compiled by authors based on the VWH. The figure displays worker-type specific transition frequencies within and between firm classes for treated and control observations.

E Aggregate Earnings Losses

This section complements our findings on wages and employment with an analysis of yearly earnings, shown in Figure 22. Unsurprisingly, the drop in yearly earnings is significantly more pronounced than the corresponding wage losses.[29] Overall, the yearly earnings of displaced workers drop by approximately 24% in the year following displacement. Thereafter, earnings losses quickly attenuate to about 9%, after which they hover around 4-7%.[30]

Figure 22: Earnings Effects of Job Displacement after Firm Shutdown

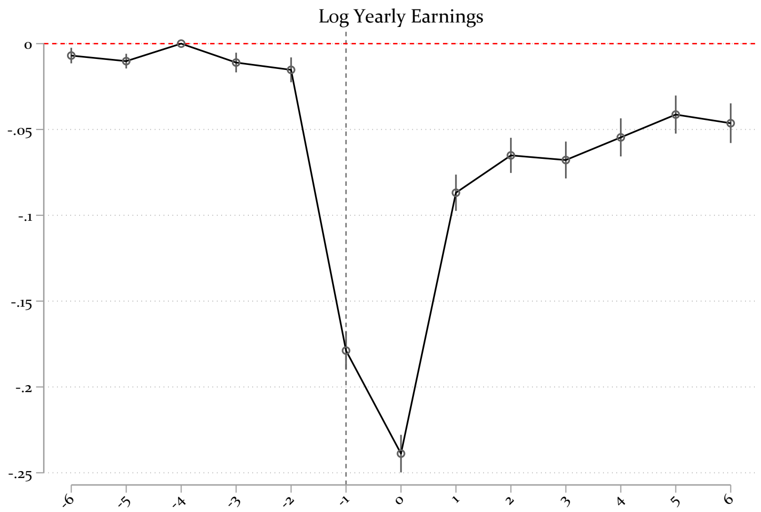

Note: Compiled by authors based on the VWH. The figure displays event study coefficients of log yearly earnings losses after a firm shutdown, conditional on being employed in the respective calendar year. Results are are based on the BLM specification. 95% confidence intervals are represented by the horizontal whiskers in the figure.

Type-specific Earnings Losses Figure 23 shows the corresponding results by worker type. These mirror our findings for wages and employment. The lowest types suffer the least in terms of their wage losses and recover fairly quickly from their initial earnings losses. Yearly earnings of worker types 5 and 6 drop precipitously in the year of the shutdown (about 30%

[29]Consider the following example: If a worker loses her job in November and finds a new one in April the following year at the same wage rate, her daily wage, that is, yearly earnings divided by days worked, will be unchanged. However, her total earnings over the entire calendar year will be substantially lower. See also the discussion based on US administrative data in Farber (2017).

[30]Note that yearly earnings already drop noticeably in $k = -1$. Take again the example of a worker who loses her job in November and finds a new job in April of the following year. Since firms of treated workers shut down between $k = -1$ and $k = 0$, this implies that she will have lower yearly earnings in both periods.

to 35%) and never recover fully (about 10% to 15% after 6 years). Accordingly, differences in earnings losses become more pronounced in the years following displacement, where high worker types cannot catch up to pre-shutdown levels.

Figure 23: Earnings Effects of Job Displacement by Type

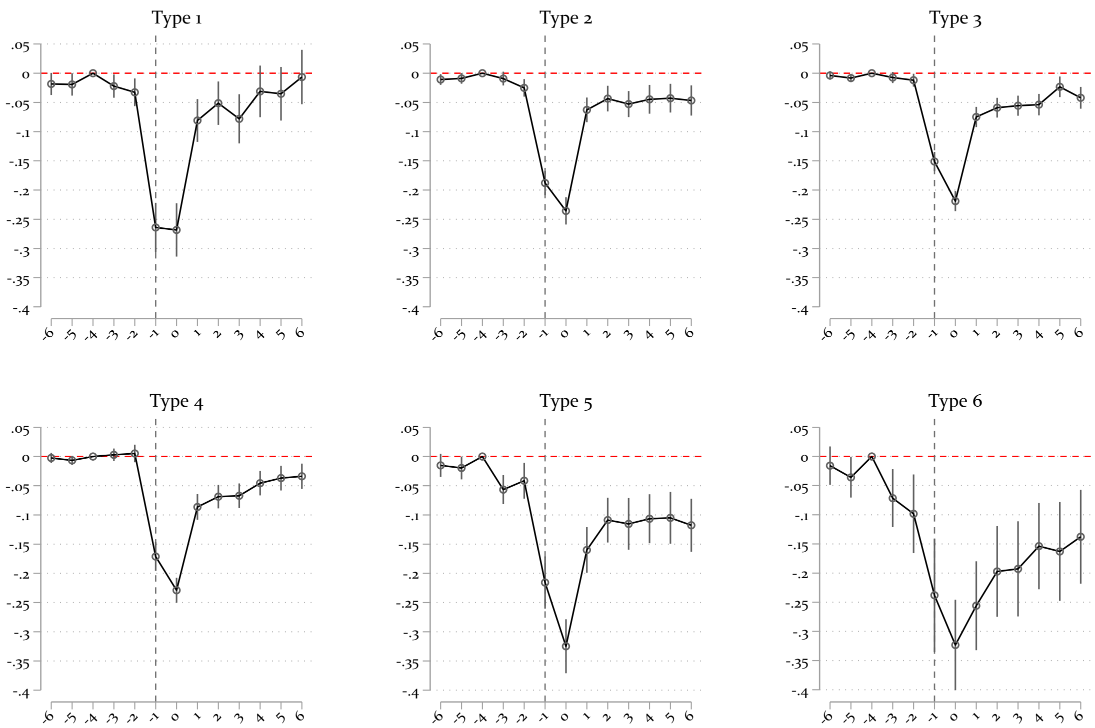

Note: Compiled by authors based on the VWH. The panels display event study estimates for each worker type of the effects of job displacement after a firm shutdown on log yearly earnings, conditional on being employed in the respective calendar year. Results are based on the BLM classification. 95% confidence intervals are represented by the horizontal whiskers in the figure.

Comparing earnings losses in the literature presents some challenges, particularly in terms of accounting for the earnings of unemployed workers. The literature does not establish a standardized method for addressing this issue, but it does outline three primary strategies. The first strategy involves assigning a value of zero to the earnings of unemployed workers and incorporating these values into the regression analysis without applying a logarithmic transformation on the dependent variable (Bertheau et al., 2023; Lachowska, Mas, and Woodbury, 2020). By construction, this method tends to amplify the total earnings losses. A second strategy involves excluding unemployed workers from the analysis, resulting in an underestimation of the overall earnings losses (Couch and Placzek, 2010; Jacobson, LaLonde, and Sullivan, 1993; Lachowska, Mas, and Woodbury, 2020). The third strategy compensates for the lack of earnings by imputing income based on unemployment benefits, thereby acknowledging the role of social security mechanisms (Davis and Von Wachter, 2011). Due to the insufficient information on the unemployment status in our dataset, we are unable to employ the third strategy and thus opt for the second approach. In alignment with the existing literature, our findings reveal substantial initial earnings losses followed by a robust recovery.

Our estimated losses in yearly earnings for Veneto are in the lower range of estimates found in the literature. In part due to the methodological differences outlined above, findings across studies differ considerably, even within the most frequently analyzed context of the United States. Analyzing mass layoffs, estimated initial earnings losses range from 33%, recovering to 15%, in Connecticut (Jacobson, LaLonde, and Sullivan, 1993), to 40%, recovering to 25%, in Pennsylvania (Couch and Placzek, 2010), to 45%, recovering to 23%, in Washington State (Lachowska, Mas, and Woodbury, 2020). Focusing on mass layoffs during recessions in the United States overall, Davis and Von Wachter (2011) find initial earnings losses of around 40% with a subsequent recovery to 20%. Diverging from the recession-centric perspective, Huckfeldt (2022) finds slightly smaller initial losses and a faster recovery for the United States. In addition, Huckfeldt (2022) finds that analyzing shutdowns during a time of expansion or recession plays a crucial role for working hours, leading to more variation in estimates of yearly earnings losses. Lastly, Bertheau et al. (2023) find earnings losses after firm shutdowns in Italy of around 40% initially, and 28% after 5 years.

F Additional Decomposition Results

Figure 24: Decomposition of Wage Losses Considering Industry and Occupation

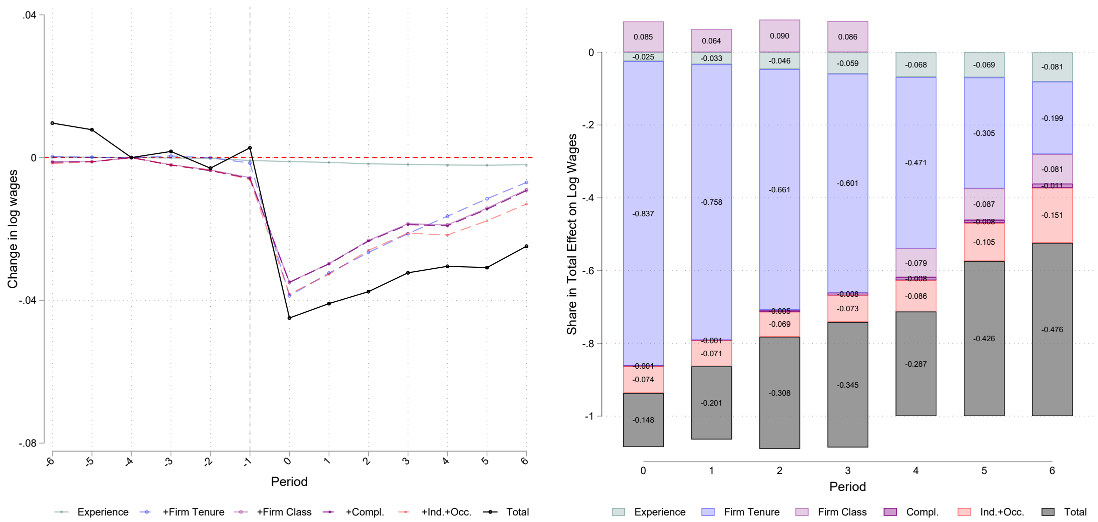

Note: Compiled by authors based on the VWH. The left panel shows the change in log wages pre- and post-event and how much of the wage loss can be attributed to labor market experience, firm tenure, firm effects, industry, occupation, and complementarities. The right panel shows the share of the total effect for the six periods after the shutdown. The height of the bars and the numbers within them correspond to minus the source-specific effect divided by the total effect. The total effect, thus, is normalized to -1.

G Limited Mobility Bias and KSS Bias Correction

Figure 25: Average number of movers per firm

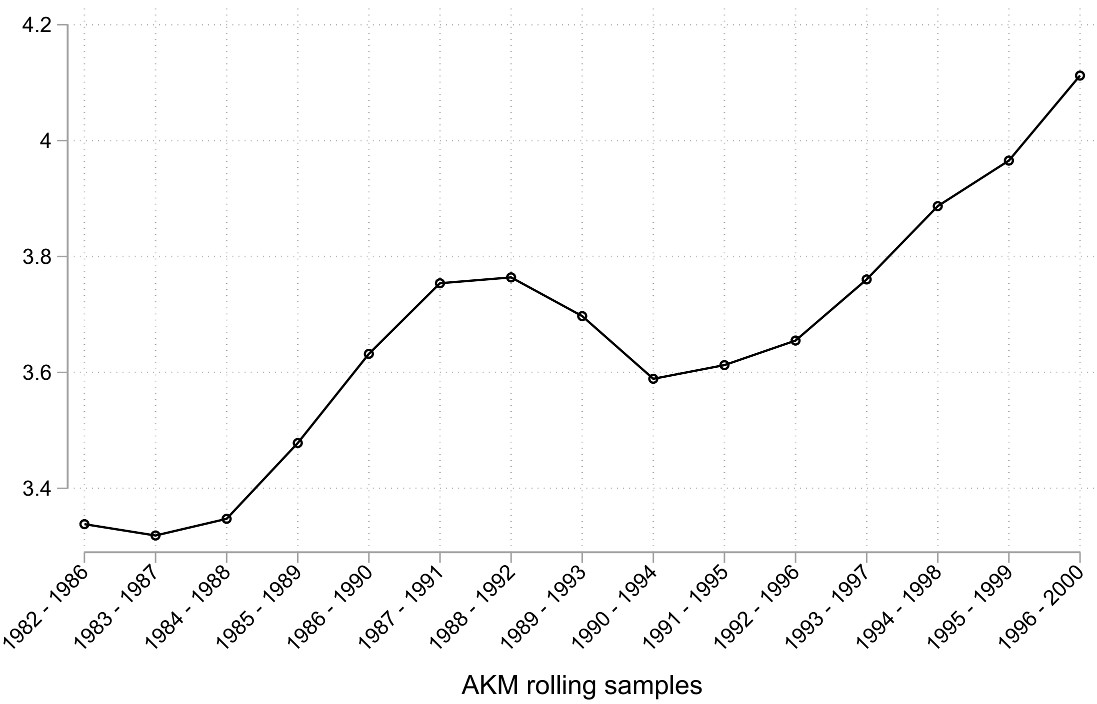

Note: Compiled by authors based on the VWH. The figures reports the average number of firm-to firm movers in each rolloing KSS leave-one-out sample.

Figure 26: Comparison AKM vs. KSS estimation

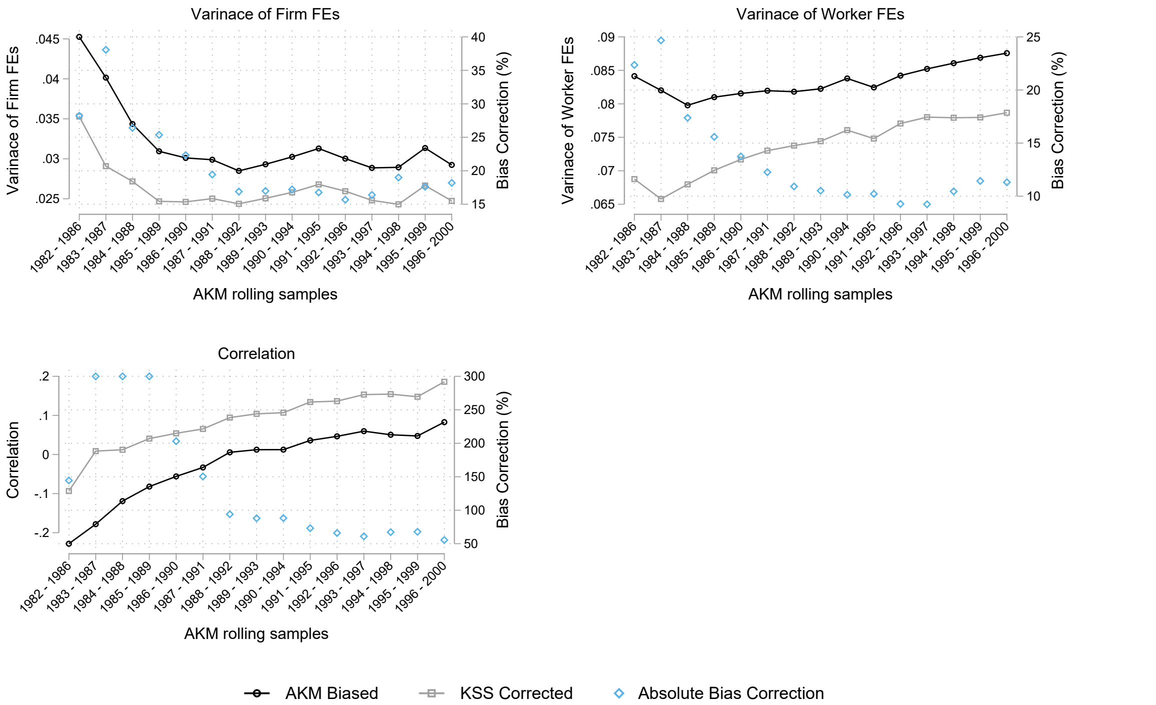

Note: Compiled by authors based on the VWH. The figure compares the AKM (in black) and KSS-corrected (in gray) estimates of $Var(\psi_j)$, $Var(\theta_i)$, and $Corr(\psi_j, \theta_i)$ in each rolling sample. The figure also reports with light blue squared markers the absoulte relative bias correction caluclautaed as $|\frac{\sigma^2_{KSS} - \sigma^2_{AKM}}{\sigma^2_{KSS}}| * 100$ where σ^2_{KSS} and σ^2_{AKM} are respectively the bias-corrected, and uncorrected second moments. Note that for $Corr(\psi_j, \theta_i)$ we cap the absolute relative bias correction to 300% to ease the readability of the graph.

H Robustness Checks

H.1 Increasing the Number of Firm Classes

In our baseline specification, we have clustered firms into ten distinct classes based on their wage distributions. This dimensionality reduction enables many of the desirable features of the BLM framework, such as endogenous mobility and wage persistence. In principle, however, BLM can nest the workhorse AKM framework, and can accommodate as many firm clusters as there are firms. As discussed in Section 5.3, the literature using AKM decomposition techniques tends to attribute considerably larger fractions of observed wage losses to firm effects than our findings suggest. To verify that our results are not merely driven by disregarding important intra-class heterogeneities, we replicate our main analysis, doubling the number of firm classes to 20. Figures 27 and 28 show the event study estimates and Figures 29 to 31 the corresponding decompositions of wage losses.

The event study graphs confirm the pattern we observe in the main analysis. Most notable are the large confidence bands around the estimates for worker type 6. However, recall that top-type workers are by far the smallest group in our population. Since we match treated workers only with control units in the same cell of observed covariates, doubling the number of firm classes significantly reduces the observation number within cells, thus decreasing the precision of the estimates.

As Figure 29 reveals, the overall decomposition results are largely unaffected by increasing the number of firm classes. Only once we differentiate between high and low types, the contribution of firm effects changes. Figure 30 shows that low types continue to move to firms in higher firm classes over the first three years post-shutdown, and the contribution of this upgrading mechanism to overall wage losses even slightly increases compared to what we obtain in the main analysis. However, while the firm class contribution to wage losses switches sign after the third year in our main specification, this is no longer the case. Instead, it vanishes almost entirely, or, if anything, remains positive.

For high types, by contrast, the relevance of firm effects changes noticeably only in the first three years post-shutdown. Thereafter, the firm effects are quantitatively very close to our main results. As discussed in Section 5.2, displaced high-type workers are more likely to downgrade in firm class after displacement. This manifests itself more clearly in Figure 31, where the contribution of firm effects to the losses of high-type workers increases by roughly 10 percentage points in the first three years after their firm shuts down compared to the main analysis. Instead, the contribution of firm tenure becomes smaller. However, firm tenure remains by far the most important contributor to wage losses.

Figure 27: Wage Effects of Job Displacement after Firm Shutdown - 20 Firm Classes

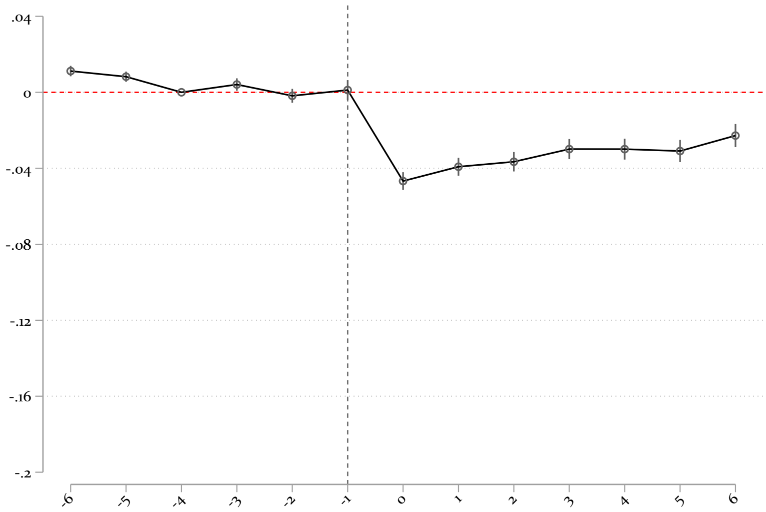

Note: Compiled by authors based on the VWH. The figure displays event study coefficients of log-wage losses after a firm shutdown conditional on employment based on the BLM specification using 20 firm classes. 95% confidence intervals are represented by the horizontal whiskers in the figure.

Figure 28: Wage Effects of Job Displacement by Type - 20 Firm Classes

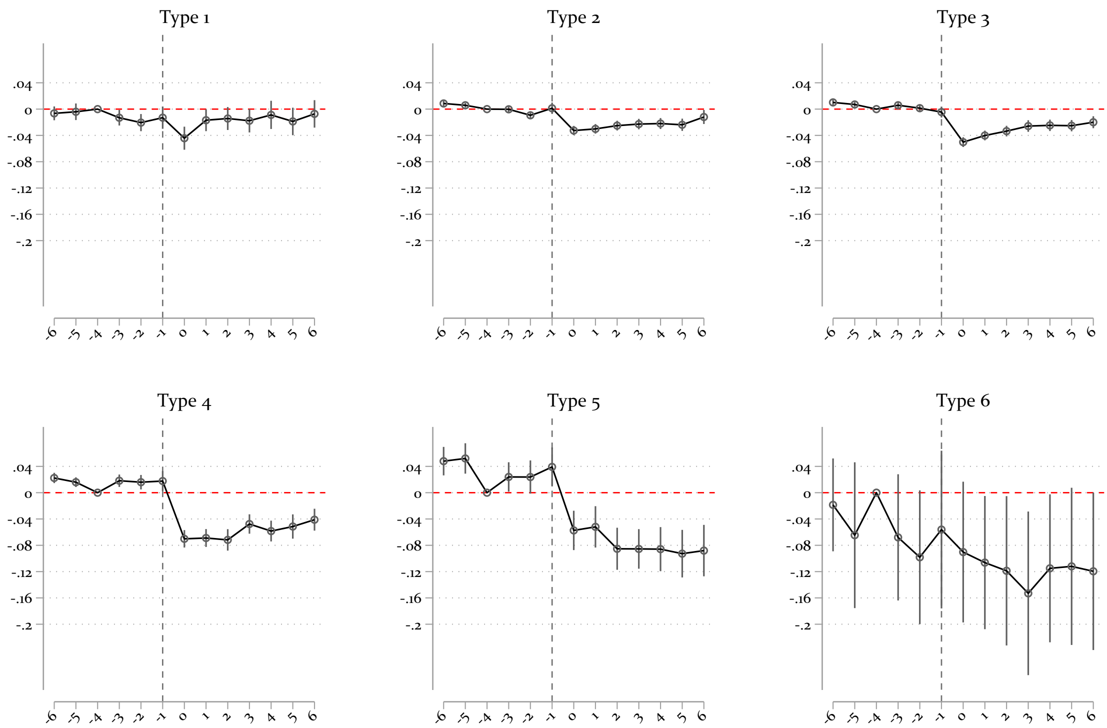

Note: Compiled by authors based on the VWH. The panels display event study estimates for each worker type of the effects of job displacement after a firm shutdown on log-wages, conditional on employment. Results are based on the BLM classification. 95% confidence intervals are represented by the horizontal whiskers in the figure.

Figure 29: Decomposition of Aggregate Wage Losses - 20 Firm Classes

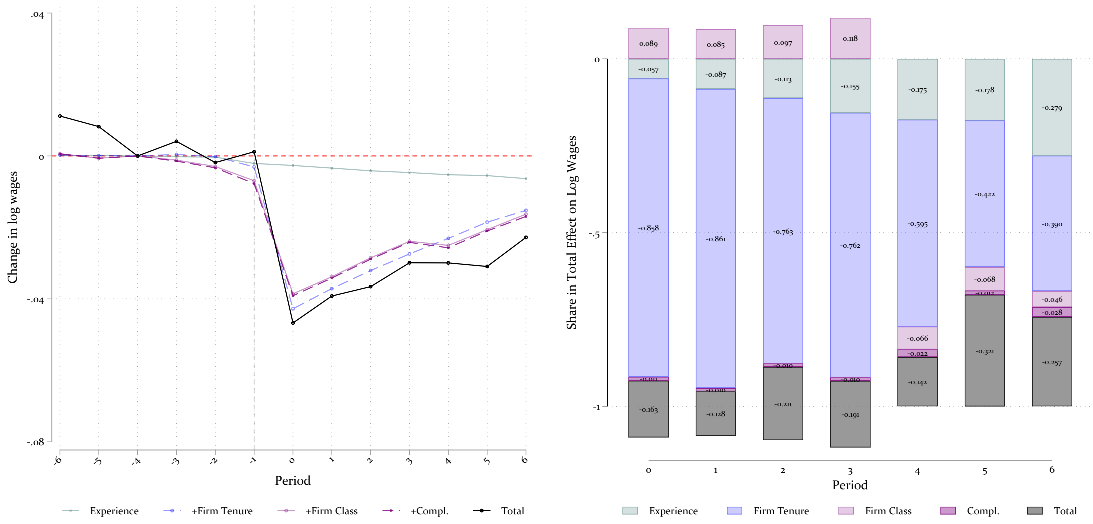

Note: Compiled by authors based on the VWH. The left panel shows the change in log wages pre- and post-event and how much of the wage loss can be attributed to labor market experience, firm tenure, firm effects, and complementarities. The right panel shows the share of the total effect for the six periods after the shutdown. The height of the bars and the numbers within them correspond to minus the source-specific effect divided by the total effect. The total effect, thus, is normalized to -1.

Figure 30: Decomposition of Wage Losses, Low Types - 20 Firm Classes

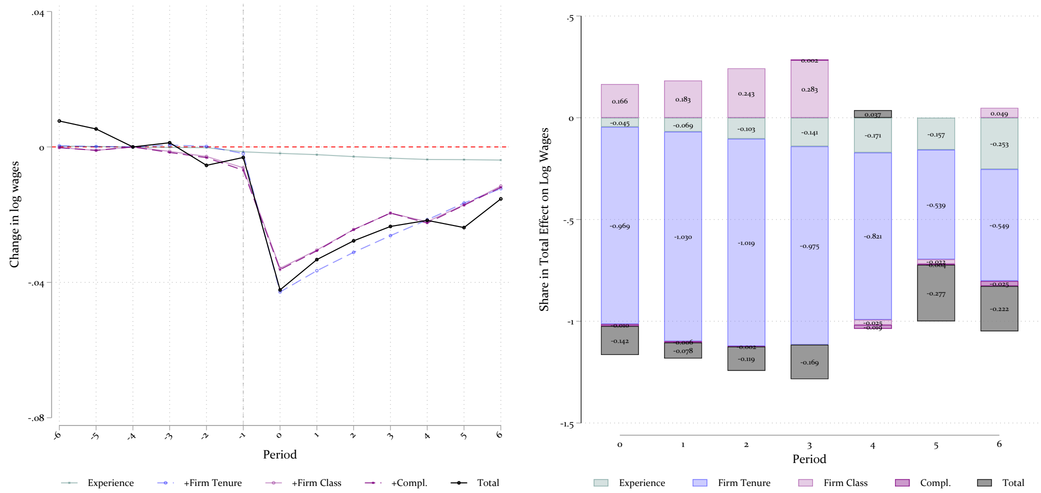

Figure 31: Decomposition of Wage Losses, High Types - 20 Firm Classes

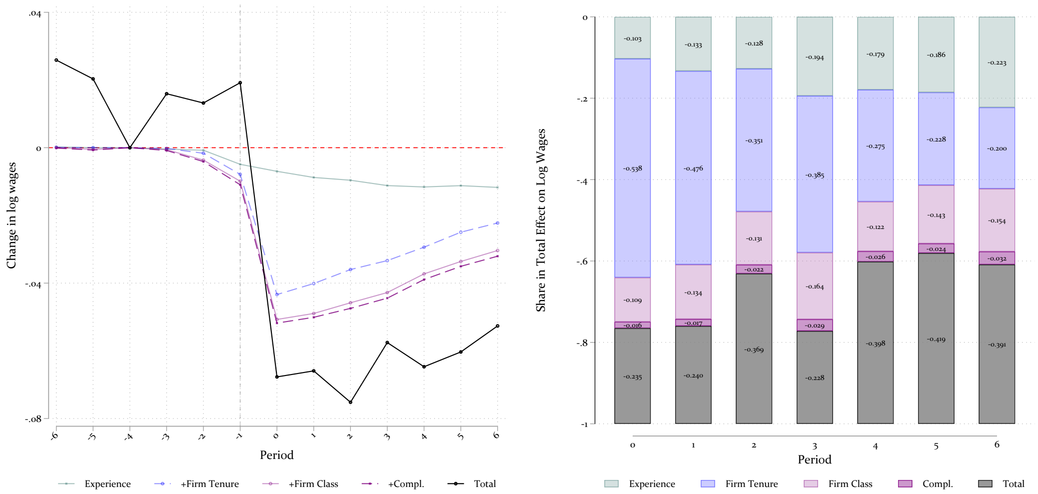

Note: Compiled by authors based on the VWH. The left panel shows the change in log wages pre- and post-event and how much of the wage loss can be attributed to labor market experience, firm tenure, firm effects, and complementarities. The right panel shows the share of the total effect for the six periods after the shutdown. The height of the bars and the numbers within them correspond to minus the source-specific effect divided by the total effect. The total effect, thus, is normalized to -1.

H.2 Restricting the Size of Firms

One distinct advantage of the BLM framework is that it allows for the inclusion of small firms since it requires mobility only between firm classes rather than the individual firm level as is the case in AKM. In this section, however, we follow the convention in the literature to restrict the size of treated firms. In particular, we exclude firms that have fewer than 30 or more than 500 employees before the shutdown (see, e.g., Helm, Kügler, and Schönberg, 2022). This offers two advantages: first, we maximize comparability with the literature in terms of estimated effect sizes. Second, restricting the size of treated firms strengthens the plausibility of the exogeneity assumption. From the perspective of the individual worker, the closure of the firm she works in is more likely to be exogenous to her own actions if the firm is larger and her individual contribution to the firm's success is smaller. Limiting the firm size from above ensures that spillover effects on other firms in the region, industry, or economy overall do not confound the estimates.

Restricting the firm size in this way significantly reduces the number of firm closures under consideration. This becomes evident in Figure 32. in the main sample, around 75% of treated workers are employed in firms with fewer than 30 employees. These firms make up about 95% of the firm closures under consideration. Overall, the number of treated workers is reduced from 19,846 in the main sample to 5,715 in the restricted sample. The number of firm closures considered drops from 8,866 to a mere 193. Figures 33 and 34 show the corresponding event study estimates. Despite the reduction in sample size, the magnitude of wage losses mirrors the results obtained from the main specification, both in aggregate and by worker type.

Figure 32: Distribution of Worker and Firms

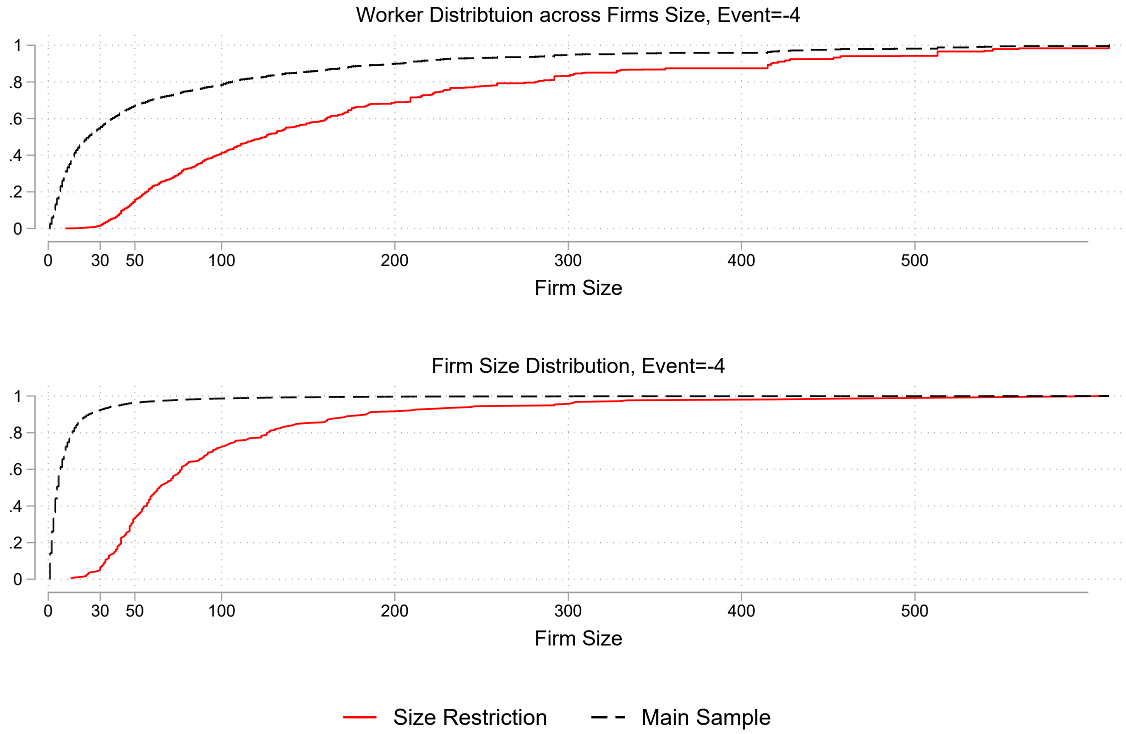

Note: Compiled by authors based on the VWH. The top panel displays the distribution of treated workers across firms of different sizes for the overall sample and the sample where we restrict the size of firms that shut down to between 30 and 500 workers. The bottom panel shows the size distribution of treated firms for both samples.

Figure 33: Wage Effects of Job Displacement after Firm Shutdown - Restricted Firm Size

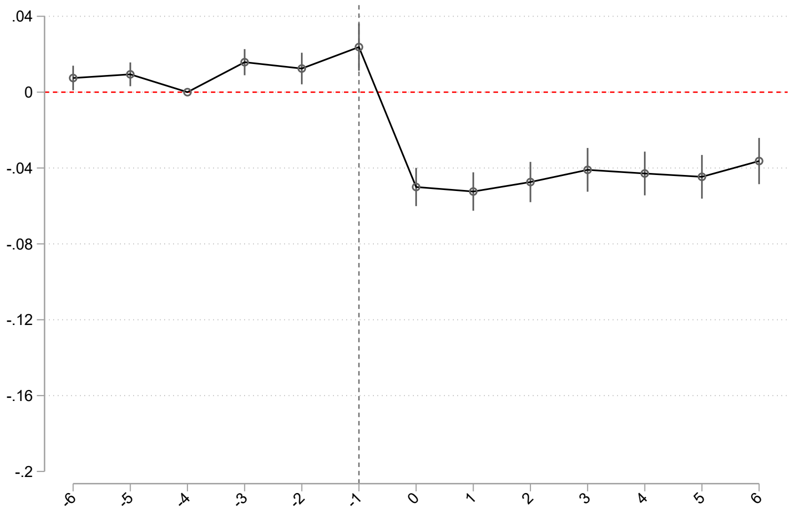

Note: Compiled by authors based on the VWH. The figure displays event study coefficients of log-wage losses after a firm shutdown conditional on employment based on the BLM specification. 95% confidence intervals are represented by the horizontal whiskers in the figure.

Figure 34: Wage Effects of Job Displacement by Type - Restricted Firm Size

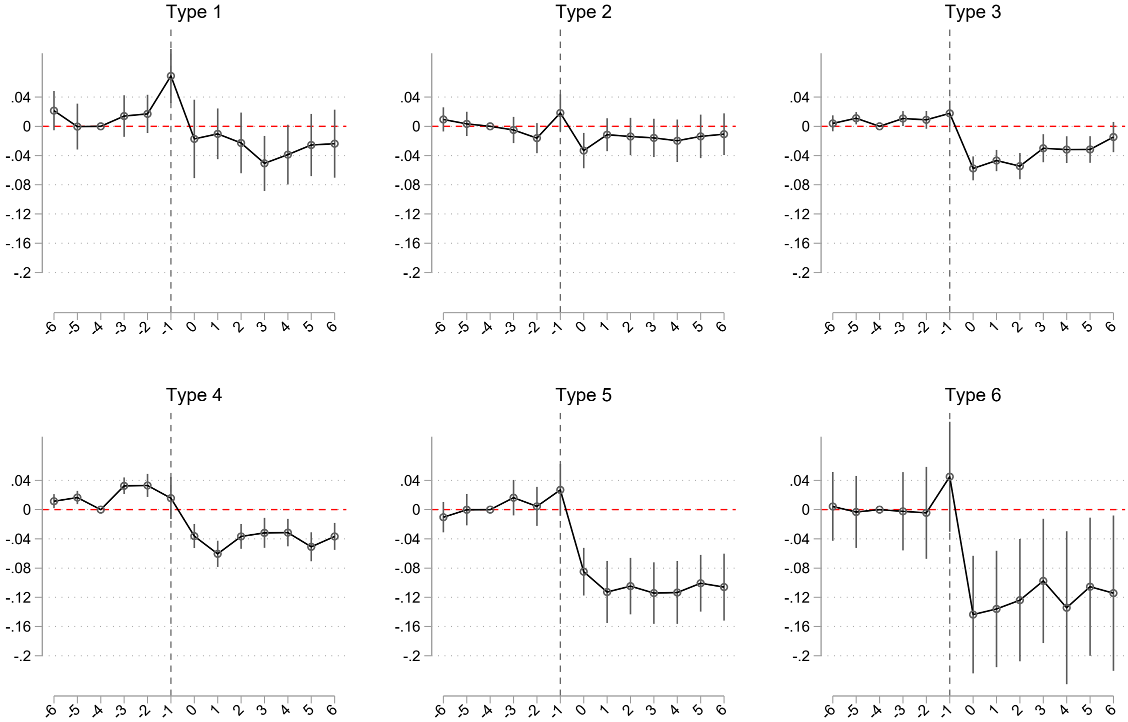

Note: Compiled by authors based on the VWH. The panels display event study estimates for each worker type of the effects of job displacement after a firm shutdown on log-wages, conditional on employment. Results are based on the BLM classification. 95% confidence intervals are represented by the horizontal whiskers in the figure.

H.3 Prime-Age Worker Sample

In our main analysis, we consider male workers between the ages of 18 and 65 and female workers between the ages of 18 and 60. In this section, we reproduce our main results excluding both very young workers and those close to retirement, focusing on prime-age workers instead as is often done in the literature. Figure 35 shows event study estimates for workers aged 25 to 55, and Figure 36 the corresponding results by worker type. The results hardly differ from our main analysis.

Figure 35: Wage Effects of Job Displacement after Firm Shutdown - Core Age Workers

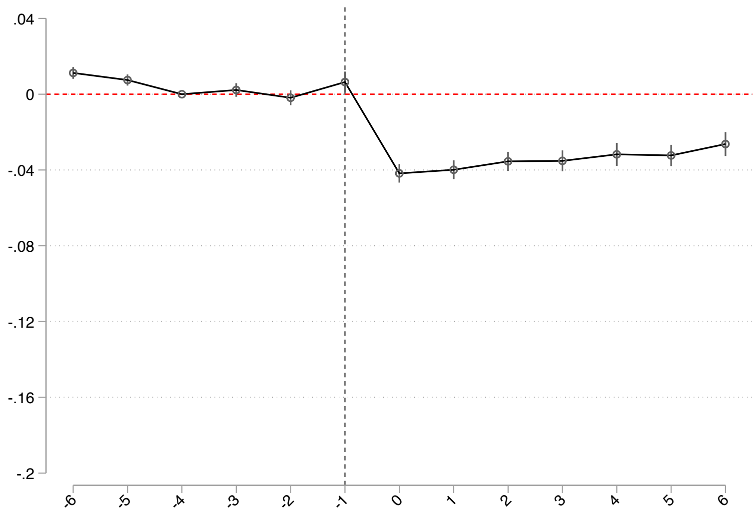

Note: Compiled by authors based on the VWH. The figure displays event study coefficients of log-wage losses after a firm shutdown conditional on employment based on the BLM specification. 95% confidence intervals are represented by the horizontal whiskers in the figure.

Figure 36: Wage Effects of Job Displacement by Type - Core Age Workers

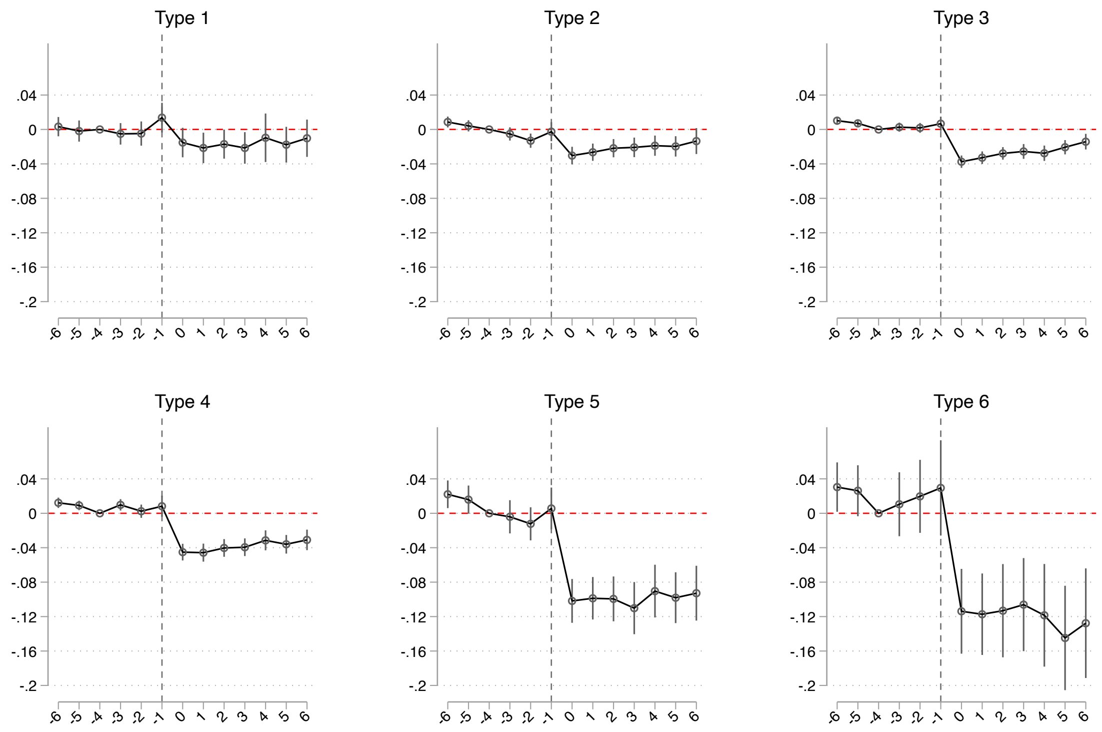

Note: Compiled by authors based on the VWH. The panels display event study estimates for each worker type of the effects of job displacement after a firm shutdown on log-wages, conditional on employment. Results are based on the BLM classification. 95% confidence intervals are represented by the horizontal whiskers in the figure.

H.4 Probabilistic Assignment of Worker Types

In BLM, workers belong to a certain worker type with some probability. In our main analysis, we simply assign each worker to the type with the largest likelihood. An alternative to this hard assignment is to randomly draw from the estimated type distribution for a given worker.

We proceeds as follows: First, we recover the individual-specific CDF of the worker type from our BLM estimation. Next, we draw a uniform random number for each worker. Finally, we assign a type to each worker by backing out the type from the inverted individual-specific CDF.

As the central point of comparison, Figure 37 shows the type-disaggregated wage losses, analogous to Figure 4. Overall, there are hardly any noticeable differences between the two. The wage trajectories are both quantitatively and qualitatively very similar. Therefore, we are content with the hard assignment of worker types chosen for the main results.

Figure 37: Wage Effects of Job Displacement by Type - Probabilistically Assigned Worker Types

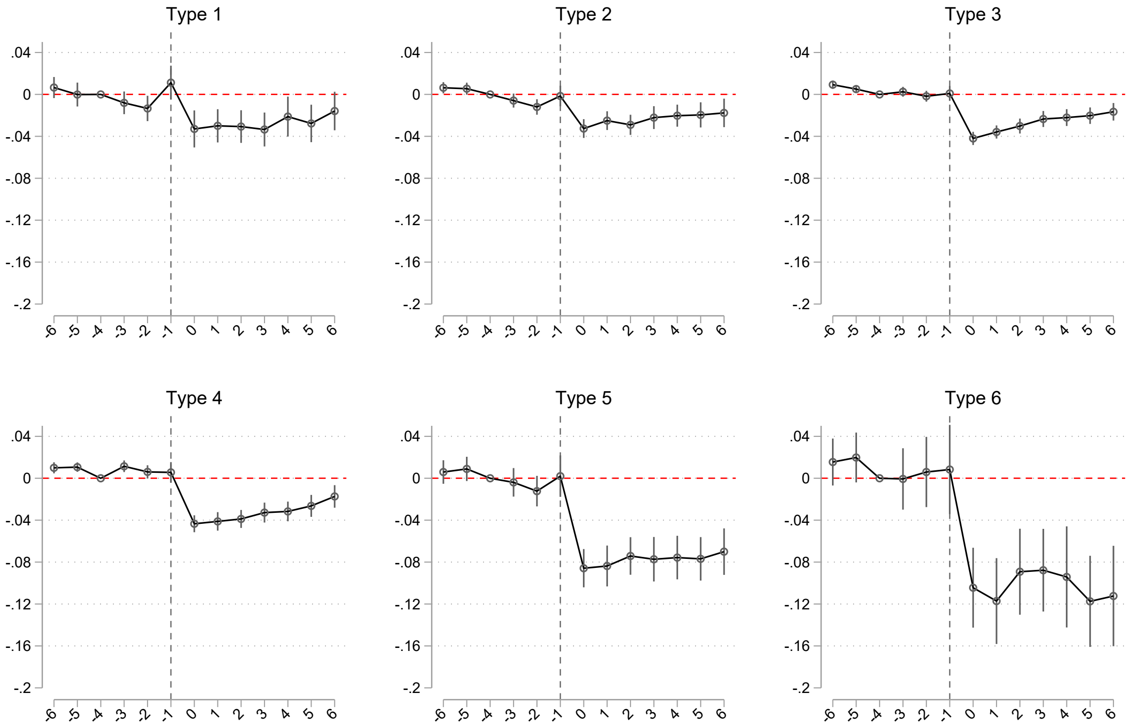

Note: Compiled by authors based on the VWH. The panels display event study estimates for each worker type of the effects of job displacement after a firm shutdown on log-wages, conditional on employment. Results in Figure 37 are based on the BLM classification and the probabilistic selection of the worker type. 95% confidence intervals are represented by the horizontal whiskers in the figure.

H.5 Stabilization of Worker Types over Time

To check whether mobility of workers across types leads to notably different results, we re-estimate the dissagregated wage losses using a time-stabilized worker type.

In particular, instead of fixing the worker type at the time of matching, we consider all types observed for a given worker from period -8 until period -4, that is, 5 years in total. For each worker, we then generate the mode of this individual-specific type distribution and use it as the assigned worker type. In the case of a tie, we select the higher type, which leads to lower types being less populated.

In Figure 38 we show the event study estimates for the stabilized types. The figure shows that results for types 2-5 are very similar to those in our main specification. For type 6, results are also qualitatively very similar, although the size of the initial wage loss is somewhat smaller. For type 1 there are more noticeable changes: Point estimates are now mostly positive, but the confidence bands are very large; so large that they generally contain the original negative effects we see in the main results. Since the type 1 group has become much smaller due to nature of the mode-assignment, it is unsurprising that the estimates have become less precise. Overall, we do not find meaningful differences to our main results.

Figure 38: Wage Effects of Job Displacement by Type - Stabilized Worker Types

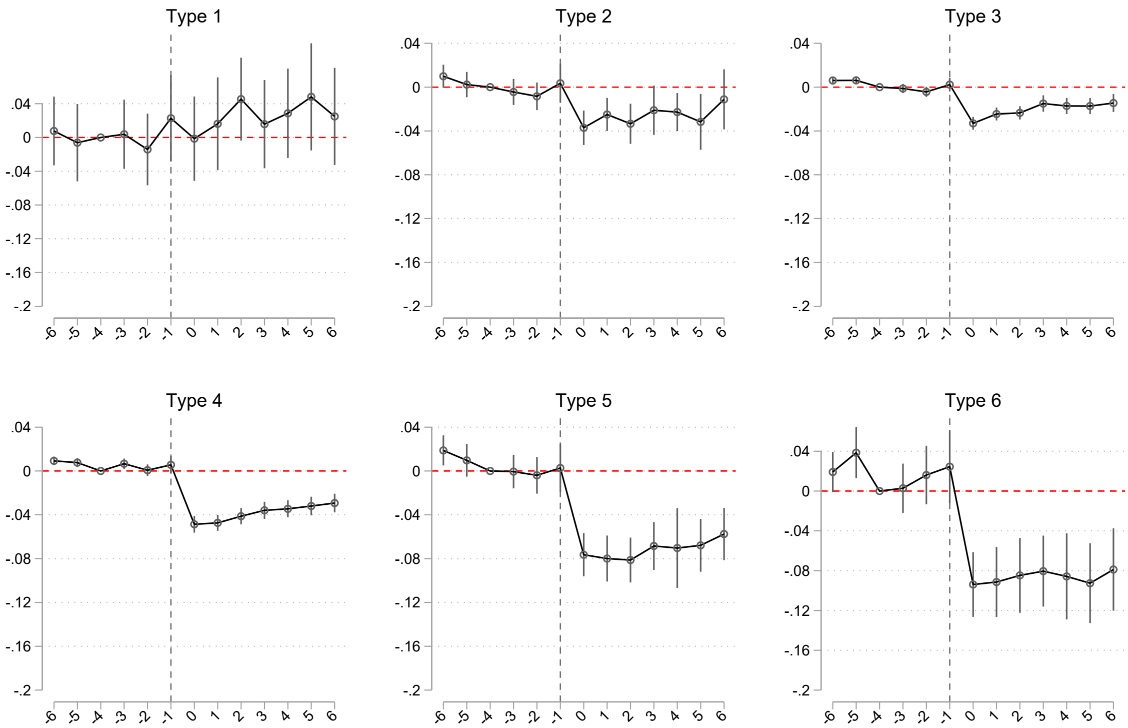

Note: Compiled by authors based on the VWH. The panels display event study estimates for each worker type of the effects of job displacement after a firm shutdown on log-wages, conditional on employment. Results in Figure 38 are based on the BLM classification and the probabilistic selection of the worker type. 95% confidence intervals are represented by the horizontal whiskers in the figure.