Separations Revisited: Do Layoffs or Quits Drive Lower Separation Rates in High-Productivity Firms? *

Cauê Dobbin† Daniel Fernandez‡ Tom Zohar§

April 30, 2026

Abstract

A well-known empirical regularity is that high-productivity firms have lower worker separation rates, but it is unclear whether this pattern reflects quits or layoffs. Using matched employer-employee data from Brazil that distinguish the reason for each separation, we show that the productivity–separation gradient is driven primarily by layoffs rather than quits. We then propose and test a mechanism in which downward wage rigidity prevents firms from adjusting wages in response to adverse shocks, causing those shocks to translate into layoffs. High-productivity firms are less exposed because their larger markdowns provide a buffer between productivity and wages. Consistent with this mechanism, we find that firms with higher markdowns have lower layoff rates, and that markets with stronger wage rigidity exhibit both higher layoff rates and a steeper productivity–layoff gradient. These findings suggest that productivity differences across firms shape not only wages, but also workers' exposure to job loss.

Keywords: Layoffs, Wage Rigidity, Firm Productivity.
JEL Codes: J63, J31.

*We thank Edoardo Acabbi, Nezih Guner, David Green, Simon Jäger, Kurt Mitman, Chris Moser, Pascuel Plotkin, Josep Pijoan-Mas, Jason Sockin, and Isaac Sorkin, as well as participants in various seminars and conferences, for helpful comments and suggestions. We are grateful to Carles Pare Ogg, David Herskovits, and Moritz Osterhuber for excellent research assistance. We gratefully acknowledge financial support from Georgetown University, the Juan de la Cierva program, the European Union, CEMFI, and Fundación Carolina. All errors are our own.

†Georgetown University. Email: caue.dobbin@georgetown.edu
‡CEMFI. Email: daniel.fernandez@cemfi.edu.es
§CEMFI. Email: tom.zohar@cemfi.es

Introduction

A well-known stylized fact in labor economics is that high-productivity firms exhibit lower worker separation rates.[1] Yet no existing work establishes whether this negative productivity–separation gradient reflects quits or layoffs, and beliefs about this decomposition are highly dispersed.[2] The distinction matters because quits often reflect workers moving to better jobs, whereas layoffs impose large and persistent earnings losses. If the gradient is driven by layoffs, low-productivity firms are a systematic source of job destruction that is costly for workers.

Using Brazilian matched employer-employee data that records the reason for every separation, we provide the first direct decomposition of this gradient into quit and layoff components. We find that layoffs explain 93% of the productivity–separation gradient. We then propose and empirically validate a mechanism in which adverse shocks trigger layoffs when wages cannot adjust downward, while high-productivity firms are less affected because their larger markdowns provide a cushion.

Brazil is a particularly informative setting to study the relative roles of quits and layoffs because dismissal regulations are relatively permissive, making layoff decisions more responsive to underlying economic forces. Layoff rates are broadly comparable to those in the United States, which also has limited employment protection. By contrast, layoff rates are lower in many countries with stricter regulations, where these forces may be more muted.[3]

We study this setting using matched Brazilian administrative employment records combined with data on value added per worker and markdowns for the manufacturing sector. A key advantage of these data is that we observe the stated reason for every separation. These classifications are likely to be informative because workers benefit from layoff status through severance payments and government benefits, while firms bear severance costs and a government fine when dismissing workers, creating strong opposing incentives against misreporting (Gerard and Naritomi, 2021). Consistent with

[1]See, for example, Topel and Ward (1992), Davis et al. (2013), Haltiwanger et al. (2018a), and Haltiwanger et al. (2018b).

[2]In a forecast we fielded on the Social Science Prediction Platform (DellaVigna et al., 2020), beliefs were highly dispersed: the interquartile range for the layoff share was 31–70%, and most respondents reported low confidence.

[3]See Supplemental Appendix E for a comparison of layoff rates and dismissal regulations across Brazil, the United States, and other countries.

this interpretation, post-separation outcomes are substantially worse for laid-off workers than for workers who quit.

These data reveal a sharp asymmetry: annual layoff rates decline steeply with firm productivity, while quit rates vary much less. As a result, most of the overall separation gradient is accounted for by layoffs, with quits explaining only a modest share.

Our interpretation of the productivity–layoff gradient is that wage rigidity prevents firms from adjusting wages in response to adverse shocks, causing those shocks to translate into layoffs. High-productivity firms are less exposed because their larger markdowns provide a buffer between productivity and wages. To assess this mechanism, we exploit another advantage of the Brazilian data, which separately reports contract wages and variable pay. We first show that contract wages exhibit strong downward rigidity, while variable pay adjusts more flexibly. We then proxy firm-level wage rigidity by the share of compensation paid through contract wages.[4] Using this measure, we find that firms with more rigid wages exhibit higher layoff rates, consistent with the proposed mechanism.

Two alternative explanations for these results are worker sorting and differential exposure to shocks across high- and low-productivity firms. We find that worker sorting accounts for about half of the productivity–layoff gradient and about half of the association between layoff rates and wage rigidity. Yet, even after controlling for worker composition, layoffs remain the main driver of the productivity–separation gradient, and wage rigidity remains strongly associated with higher layoff rates. Controlling for exposure to shocks leaves the results essentially unchanged. The findings are also robust to alternative measures of productivity and to different sample definitions.

To organize these empirical findings, we present a wage-posting framework with wage rigidity and match-specific shocks. In the model, wage rigidity gives rise to endogenous layoffs: firms set wages before productivity is fully realized, so sufficiently adverse shocks make the match unprofitable and trigger layoffs.

Our main theoretical result is that more productive firms have larger markdowns, even though they pay higher wages. These larger markdowns buffer adverse shocks and explain why more productive firms have lower layoff rates. The intuition is simple.

[4] Several previous papers also show that variable pay is more flexible than contract pay (Altonji and Devereux, 1999; Messina et al., 2010; Anger, 2011; Grigsby et al., 2021). In addition, several recent papers use related measures of wage rigidity based on compensation structure, including Makridis and Gittleman (2022), Reizer (2022), and Sockin and Sockin (2025).

Firms choose wages by trading off worker retention, which rises with wages, against markdowns, which fall with wages. Expected profit per worker is the product of the retention rate and the expected markdown, so the two are complementary: a higher markdown raises the value of retaining workers, and higher retention raises the value of a larger markdown. As a result, more productive firms optimally choose both higher retention and larger expected markdowns.

The framework delivers a set of testable predictions that we take to the data.

First, we show that, consistent with the model, high-productivity firms have larger markdowns and that firms with larger markdowns have lower layoff rates. Moreover, a mediation analysis shows that, after accounting for worker sorting, markdowns explain essentially all of the productivity–layoff gradient.

Second, the model highlights that the relationship between wages, markdowns, and layoffs depends on the underlying sources of wage dispersion across firms. If wages vary across firms for reasons unrelated to productivity, markdowns will be mechanically negatively correlated with wages, implying that high-wage firms would have more layoffs because of their lower markdowns. We find that markdowns rise and layoff rates fall with firm pay premiums, consistent with productivity variation being a central driver of wage dispersion across firms.

Third, the model predicts that greater wage rigidity increases layoffs. It further implies that, when markdowns rise steeply with productivity—as they do in our empirical setting—wage rigidity amplifies the productivity–layoff gradient. The intuition is that high-productivity firms' larger markdowns already buffer negative shocks, making wage flexibility less important for them. Exploiting cross-market variation in the share of rigid-wage jobs, we find support for both implications: markets with greater wage rigidity exhibit both higher layoff rates and a steeper productivity–layoff gradient.

Our work contributes to the empirical literature on wage rigidity and its role in generating layoffs. Schmieder and von Wachter (2010) show that persistent wages can amplify employment fluctuations, Jäger et al. (2023) provide direct evidence of inefficient separations, Ehrlich and Montes (2024) show that firms with more rigid wages have higher layoff rates, and Davis and Krolikowski (2025) demonstrate that a substantial share of layoffs violate bilateral efficiency. We complement these findings by showing that the effects of wage rigidity vary systematically with firm productivity:

high-productivity firms are more insulated because larger markdowns buffer shocks, generating the negative productivity–layoff gradient.

Our findings also speak to research on the importance of firms for worker outcomes. The AKM tradition shows that the firm a worker is employed at plays a central role in determining wages and earnings trajectories (Abowd et al., 1999; Card et al., 2013; Song et al., 2019; Gerard et al., 2021). Closer to our focus, Pinheiro and Visschers (2015) and Jarosch (2023) emphasize that firm-level differences in layoff risk are an important source of unemployment persistence and worker inequality. We connect these two strands by showing that workers at lower-productivity firms face both lower wages and higher layoff risk, so the disadvantage of working at such firms is larger than wage differences alone would suggest.[5]

More broadly, our mechanism relates to recent work examining how wage-setting frictions interact with firm-level shocks to generate endogenous separations (Carlsson and Westermark, 2022; Blanco et al., 2024).[6] We show that these forces also help explain persistent cross-firm differences in layoff risk.

The remainder of the paper is organized as follows. Section 1 describes the institutional setting and presents evidence supporting the reliability of the quit-layoff distinction in our data. Section 2 documents our main empirical finding: the negative relationship between firm productivity and separation rates is primarily driven by layoffs. Section 3 examines the empirical relationship between wage rigidity and layoffs. Section 4 develops a conceptual framework to interpret our findings. Section 5 evaluates the framework's mechanisms in the data. Section 6 concludes.

1 Distinguishing Quits and Layoffs Empirically

1.1 Data and Sample

Employer-Employee Data from Brazil. The *Relação Anual de Informações Sociais* (RAIS) is a comprehensive administrative dataset that records all formal employment relation-

[5] A related literature shows that layoffs impose large and persistent earnings losses on laid-off workers (Jacobson et al., 1993; Couch and Placzek, 2010; Bertheau et al., 2023).

[6] Related work on the interaction between shocks, wage-setting, and separation dynamics includes Hopenhayn and Rogerson (1993), Mortensen and Pissarides (1994), Mueller (2017), and Acabbi et al. (2024).

ships in Brazil. Each year, firms submit RAIS filings covering their workforce in the previous year, including worker characteristics such as gender, date of birth, and education, as well as contract information such as earnings, contracted hours, and detailed occupation (2,638 unique occupation codes). Importantly, RAIS requires firms to report both the date and reason for each separation, allowing us to distinguish between quits and layoffs. The data also report contract wages separately from variable pay components, which include bonuses, performance pay, and overtime.

Manufacturing Survey. The Annual Manufacturing Survey (*Pesquisa Industrial Anual*, PIA), conducted by the *Instituto Brasileiro de Geografia e Estatística* (IBGE), provides detailed information on production, employment, and costs in the manufacturing sector. The PIA defines value added (VA) per worker as the value of industrial transformation, computed as the difference between the gross value of industrial production and the costs of industrial operations (excluding labor costs). The data also report firm size, which we use to construct value added per worker, and total labor costs, including wages, benefits, and mandatory social security contributions. The survey is conducted at the firm level; however, we use the publicly available version of the data, which is aggregated at the 3-digit industry $\times$ state level.[7]

Firm Finances. The Orbis dataset is a proprietary firm-level database compiled by Bureau van Dijk and widely used in Economics and Finance research. It aggregates and harmonizes administrative and financial information on firms worldwide. In Brazil, Orbis covers a large but selected subset of firms, with greater representation of larger firms that report financial statements, while still spanning a broad range of industries beyond manufacturing. The data report total firm revenue and employment for each year, which we use to construct revenue per worker.

Estimating firm wage premiums. AKM firm pay premiums are estimated following Abowd et al. (1999). To improve estimation precision, we classify firms into 100 clusters using a k-means clustering algorithm, as recommended by Bonhomme et al. (2019). Supplemental Appendix C details the estimation procedures and validates the AKM model assumptions in our sample.

Sample: Private-Sector Jobs. Our sample spans 2010–2017, beginning after the Great

[7] There are 285 industries at the 3-digit level and 4 states in the Southeast region of Brazil, yielding 1,140 industry–state cells. This level of disaggregation provides substantial within-state variation in value added across industries. The firm-level microdata are accessible only through secure data rooms at IBGE headquarters.

Recession and ending before Brazil's 2018 labor market reform. To focus on workers with substantial labor market attachment, we apply the following restrictions. We restrict the sample to Brazilian men and women born between 1959 and 1987 with at least one year of potential labor market experience.[8] We further restrict the sample to private-sector, nonfarm, open-ended contracts paid monthly at wages above the minimum wage, and to firms with at least five workers. We retain jobs that are active on December 31 of each year and have at least one month of tenure. When a worker holds multiple jobs in a given year, we select the job with the highest contracted hours or, in the case of a tie, the highest hourly wage.

Given the significant presence of informal employment in Brazil, which is not captured in our data, we follow Gerard et al. (2021) and restrict our sample to the Southeast region, comprising the states of Espírito Santo, Minas Gerais, Rio de Janeiro, and São Paulo. This region accounts for nearly half of the country's formal employment and has relatively lower informality rates (33.8% in 2017, versus 56.2% in the Northeast and a national average of 40.8%).[9] Additionally, its minimum-to-median wage ratio is comparable to those in other developing and developed economies.[10] We then apply standard wage-trimming rules, dropping observations with year-over-year log-wage changes exceeding one log point or above the 99th state-year percentile of log-hourly wages. In addition to these sample restrictions, most of our analysis focuses on the manufacturing sample, where we can measure productivity using PIA.

Our data span 2010–2017, but the analysis uses the 2011–2016 period. This restriction has two sources. First, we classify a worker as laid off in year t if the worker is employed at a firm in year t and is laid off in year $t+1$, so any analysis of layoffs must end in 2016. Second, throughout the paper, we construct measures of firm productivity and wage rigidity using data from the first year each firm appears in the sample and then exclude that year from the analysis. Since 2010 is the first sample year for all firms, it is excluded throughout. For consistency, we use the 2011–2016 period in all results, even

[8]"Potential labor market experience" is defined as age minus years of education minus 6. "Years of education" is inferred from the highest reported degree. For example, completing high school corresponds to 12 years of education, and completing college to 16.

[9]Informality rates are from IBGE (2018), based on PNAD Contínua 2017. An informal worker is defined as a private-sector employee without a signed work contract, a domestic worker without contract, a self-employed worker not contributing to social security, or an unpaid family worker.

[10]Using U.S. data, Davis and Krolikowski (2025) finds that the minimum wage is not a primary driver of wage rigidity.

in specifications for which dropping 2010 or 2017 would not otherwise be necessary.

Table I reports descriptive statistics for the manufacturing sector, comparing our estimation sample with the broader national and regional samples. Table B.I presents analogous statistics for all sectors, in which manufacturing accounts for 18.65% of workers. Workers across samples have similar profiles in terms of age, education, and tenure. Importantly, layoff rates are also similar across samples, and layoffs account for most separations in all of them. In our main analytical sample, the average annual layoff rate is 16.79%, compared with a quit rate of 2.74%.

Table I
Descriptive Statistics: Manufacturing Sector

	Brazil	Southeast region	Sample	
			2010–2017	2011–2016
Number of firms	343,513	209,591	98,065	76,354
Average firm size	13.6	14.8	29.4	31.3
Number of workers	5,735,901	3,851,462	3,570,717	3,005,657
Average age (years)	37.4	37.4	37.3	37.4
Average log-hourly wage	2.423	2.477	2.485	2.509
Median log-hourly wage	2.237	2.283	2.313	2.338
Average tenure (months)	57.9	59.3	60.0	61.8
Average schooling (years)	10.2	10.4	10.4	10.4
Share non-white (%)	27.34	31.35	31.73	31.65
Share female (%)	34.75	33.76	33.55	33.57
Average annual layoff rate (%)	18.01	17.53	15.56	16.79
Average annual quit rate (%)	2.92	2.50	2.68	2.74
Average log-value added per worker	11.326	11.385	11.227	11.230

Notes. Summary statistics from RAIS restricted to the manufacturing sector (CNAE divisions 10–33). Columns progressively restrict the sample: "Brazil" applies demographic, contract, and job-selection restrictions nationwide; "Southeast region" further restricts to the states of Minas Gerais, Espírito Santo, Rio de Janeiro, and São Paulo; "Sample" incorporates the additional restrictions described in Section 1. Value added per worker is from the *Pesquisa Industrial Anual* (PIA), matched at the 3-digit industry × state level. Table B.I reports analogous statistics for all sectors.

1.2 Context: Quits and Layoffs in Brazil

How layoffs are identified in the data. The RAIS dataset distinguishes between quits and layoffs, a critical distinction given that the government uses this data for administrative purposes. In the case of a layoff, the firm must pay a fine to the government and provide severance pay to the worker. Additionally, the worker becomes eligible for unemployment benefits and gains access to their public pension fund, which is typically reserved for retirement. Given the low incidence of quits in the data, a natural concern is whether these policies create incentives to misclassify quits as layoffs. Supplemental Appendix D provides more details on these policies, and below we discuss why such incentives are unlikely to result in systematic misreporting.

Incentives for accurately reporting layoffs. If a separation is reported as a quit, it benefits the firm; if reported as a layoff, it benefits the worker. Consequently, both parties have strong incentives to ensure that the separation is accurately reported. However, there is a potential issue: in the case of a layoff, the firm incurs a cost by paying a fine to the government, while the worker benefits from unemployment payments and gains early access to their pension funds. If a worker highly values immediate liquidity—such as accessing their pension funds early—the total benefits received from the government could outweigh the costs to the firm. This scenario might create an incentive for collusion between the worker and the firm, where they agree to misclassify the separation as a layoff in exchange for side payments that leave both parties better off.

Nevertheless, such collusion is unlikely in practice. When a separation is classified as a layoff, the firm must make substantial payments to both the government and the worker. For collusion to succeed, the firm would need to trust that the worker will return a portion of these payments after accessing their pension funds, an arrangement that is difficult to enforce given its illegal nature.

Evidence on reporting accuracy. Empirical evidence suggests that collusion agreements to misreport the cause of separation are rare. Since 2018, firms and workers in Brazil have had the option to terminate contracts by mutual agreement. Under this arrangement, the worker receives severance pay and can access 80% of their pension funds, but the firm avoids the government fine. If early access to pension funds were a strong motivator for misreporting quits as layoffs, mutual agreement separations would be more common. However, they account for only 0.5% of all separations. Another potential motive for misreporting is access to unemployment benefits. Using the same RAIS data, Van Doornik et al. (2023) find that workers eligible for unemployment insurance are 11% more likely to be laid off. However, their analysis shows that these excess layoffs are not merely misclassified quits, further suggesting that misreporting is uncommon.

We conducted further validation by comparing the post-separation outcomes of workers who quit versus those who were laid off. Figure B.I presents compelling evidence:

workers who quit are significantly more likely to secure employment within a year compared to those who were laid off—51% versus 25%, respectively. Moreover, among those who found jobs, quitting workers tended to secure new positions more quickly (46% found immediate employment, compared to 25% of those laid off) and experienced more favorable wage growth, with an average increase of 11% in wages compared to a 2% decrease among those laid off. These patterns align with the hypothesis that separations categorized as quits are indeed voluntary and initiated by the workers, while those labeled as layoffs are not, further substantiating the accuracy of the reporting in our data.

Layoff rates in Brazil in comparative perspective. Brazil exhibits relatively high layoff rates by international standards, though it is not unusual among countries with flexible labor markets. Brazil's annual layoff rate is broadly comparable to that of the United States (16.9% and 17.3%, respectively), while many other countries display substantially lower rates.[11] A key factor underlying this pattern is the relatively low cost of employer-initiated separations in both Brazil and the United States. According to the EPLex dataset from the International Labor Organization (ILO), which measures the strictness of employment protection legislation, both countries rank among those with the least restrictive dismissal regulations.[12]

Supplemental Appendix E provides additional cross-country comparisons, documenting a negative relationship between layoff protection and layoff rates. It also reports estimates of layoff rates in Brazil using an alternative dataset, confirming that Brazil's rates are close to those in the United States and higher than in most other countries for which comparable data are available.

These comparisons suggest that the mechanisms we study are most relevant in labor markets with relatively permissive dismissal regulations, such as Brazil and the United States, where observed separation decisions more closely reflect underlying economic forces rather than institutional frictions. As a result, these settings provide a particularly informative benchmark for isolating and quantifying such forces. While the magnitudes

[11] For example, 6.6% in Canada, 3.5% in Australia, 1.4% in Japan, and 5.9% in Latin America excluding Brazil.

[12] The EPLex index summarizes multiple dimensions of employment protection, including notice periods, severance payments, and administrative requirements for dismissal. The index ranges from 0 to 1, with higher values indicating stronger protection. Brazil has a score of 0.23 and the United States 0.22, compared with higher values in other countries, such as 0.29 in Canada, 0.40 in Australia, 0.36 in Japan, and 0.41 in Latin America excluding Brazil.

may differ in more regulated labor markets, the same underlying mechanisms are likely to operate more broadly, albeit partially muted by policy-induced constraints.

2 The Productivity–Layoff Gradient

In this section, we investigate the layoff–separation gradient. We first show that the negative productivity–separation gradient is primarily driven by lower layoff rates at more productive firms. We then demonstrate that this gradient is not explained by differential worker sorting or by differential exposure to shocks across firms. Finally, we show that the results are robust to a range of alternative specifications and sample restrictions.

2.1 Layoffs Drive the Productivity–Separation Gradient

This section presents our main results: the productivity–separation gradient, and its decomposition into quits and layoffs. We use value added per worker as our preferred measure of productivity.[13] This measure reflects both output prices and production efficiency: firms may exhibit high value added due to strong product-market demand (high prices), capital intensity, or high TFP. We abstract from these distinctions. This interpretation is consistent with the notion of productivity in our theoretical framework (Section 4). Finally, to capture cross-sectional differences, we measure value added in the first period of the sample and keep it fixed throughout the analysis.

Figure I presents our key finding: layoff rates decline sharply with productivity. The magnitude of this gradient is substantial. Firms in the bottom 5% of the value-added distribution have layoff rates of 22.3%, compared to 8.8% among firms in the top 5%. Total separations and quits also decline with productivity. However, since quits are relatively infrequent, they account for only a small share of the overall productivity–separation gradient.

To formally quantify the relative contributions of layoffs and quits to the overall productivity–separation gradient, we estimate the following regression:

$$Y_{jt} = \beta_Y \text{VA}_j + \theta_t + \epsilon_{jt}, \tag{1}$$

[13] As explained in Section 1, value added is observed at the state–3-digit-industry level. Section 2.4 shows that the results are robust to using alternative firm-level measures.

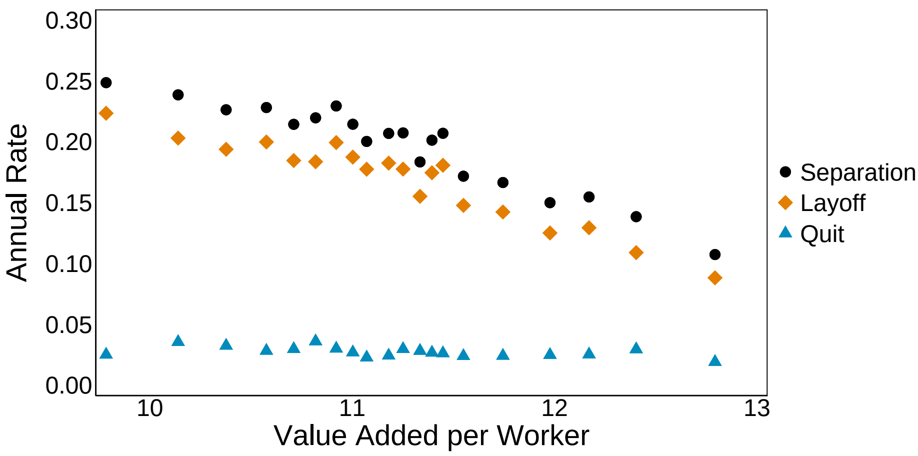

Figure I
Productivity and Separation Rates

Notes. Binscatter with 20 equal-sized bins of annual separation, layoff, and quit rates on firm productivity (Value Added). Data at the firm-year level, weighted by firm size; year fixed effects absorbed. Sample as in Section 1.

where t represents a year, j is a firm, and VA_j is the firm's value added. The dependent variable Y_{jt} represents the firm's quit, layoff, or overall separation rate, θ_t are year fixed effects, and ϵ_{jt} captures residuals. The parameter of interest, β_Y, measures the slope of separations, layoffs, or quits with respect to productivity. Since the overall separation rate is the sum of the layoff and quit rates, it follows that:

$$\beta_{Separation} = \beta_{Layoff} + \beta_{Quit}. \tag{2}$$

Motivated by the decomposition in Equation (2), we assess the role of layoffs in the productivity–separation gradient using the following ratio:

$$\text{Role of layoffs in productivity–separation gradient} \equiv \frac{\beta_{Layoff}}{\beta_{Separation}}.$$

The first column of Table II reports OLS estimates of β_{Layoff} and $\beta_{\text{Separation}}$, along with the implied ratio $\frac{\beta_{\text{Layoff}}}{\beta_{\text{Separation}}} = 0.934$. These results confirm the visual evidence in Figure I: the negative productivity–separation gradient is driven predominantly by layoffs.

2.2 The Role of Worker Sorting

A possible explanation for the patterns observed in Figure I is worker sorting. There is substantial evidence showing that higher-skilled workers tend to sort into higher-productivity firms (Card et al., 2013; Gerard et al., 2021). As a result, the lower layoff rates observed in high-productivity firms could simply reflect the higher skill levels of their employees rather than firm productivity itself.

To assess the role of worker sorting, we extend Equation (1) to the individual level:

$$Y_{it} = \beta_Y \text{VA}_{j(i,t)t} + \theta_t + \gamma X_{it} + \epsilon_{it}, \tag{3}$$

where i indexes workers, t years, and $j(i,t)$ denotes the firm employing worker i at time t. We estimate two versions of (3): (i) Y_{it} is an indicator for a layoff in year t, and (ii) Y_{it} indicates any separation. The regressor $\text{VA}_{j(i,t)t}$ measures firm productivity, X_{it} is a vector of worker characteristics, and ϵ_{it} captures unobservables. The coefficient of interest, β_Y, captures the productivity–separation or productivity–layoff gradient, net of worker heterogeneity.

Without controls ($X_{it} = \emptyset$), the OLS estimate $\hat{\beta}_Y$ from (3) is algebraically equivalent to the firm-level slope from Equation (1), weighted by firm size. This follows because all workers within a firm share the same VA_j, while $\mathbb{E}[Y_{it} \mid j, t]$ equals the firm's separation (or layoff) rate. Hence, Column (1) of Table II can be obtained from either specification and yields the same estimate.

Subsequent columns of Table II introduce a rich set of worker controls to isolate the firm-level component of the gradient. We include tenure and tenure squared, given their central role in separation dynamics (Jovanovic, 1979; Topel and Ward, 1992; Ureta, 1993). We control for age and age squared, interacted with gender and education fixed effects, to capture heterogeneous career trajectories. We add race fixed effects to account for potential discrimination and occupation fixed effects to capture occupation-level differences, such as unionization rates. Finally, we include AKM worker effects to proxy for unobserved ability.[14]

Three results emerge. First, adding worker controls substantially reduces both $\hat{\beta}^{\text{Layoff}}$ and $\hat{\beta}^{\text{Separation}}$. This indicates that worker sorting accounts for part of the productivity–separation gradient: higher-skill workers, who are less likely to be laid off, sort into higher-productivity firms. Second, the strong negative layoff–productivity gradient remains even after flexibly controlling for worker composition, suggesting that sorting is

[14] Because AKM effects are estimated with error, specifications including them as controls should be interpreted with caution.

not the only force behind this pattern and that firm-specific differences in layoff risk play an important role. Third, the ratio $\frac{\hat{\beta}^{\text{Layoff}}}{\hat{\beta}^{\text{Separation}}}$ remains remarkably stable across specifications, declining only from 0.934 in the baseline to 0.862 in the richest specification.

Taken together, these results show that the component of the overall productivity–separation gradient that remains after controlling for worker composition is still driven predominantly by layoffs. Our main conclusion from this exercise is that accounting for sorting leaves the central finding unchanged: firm-level variation in separation risk is accounted for mostly by layoffs.

Table II
Productivity–Layoff Gradient Drives the Productivity–Separation Gradient

Dependent variables: Annual layoff and separation rates				
	(1)	(2)	(3)	(4)
β^{Layoff}	-0.042***	-0.022***	-0.024***	-0.018***
	(0.0041)	(0.0037)	(0.0027)	(0.0025)
$\beta^{\text{Separation}}$	-0.045***	-0.025***	-0.026***	-0.021***
	(0.0041)	(0.0038)	(0.0027)	(0.0026)
$\frac{\beta^{\text{Layoff}}}{\beta^{\text{Separation}}}$	0.934***	0.851**	0.911***	0.862***
	(0.1758)	(0.2743)	(0.1949)	(0.2255)
Observations	9,196,989	9,196,989	9,196,346	9,196,346
Worker covariates		✓		✓
Worker AKM Effect			✓	✓

Notes. OLS estimates of Equation (3). The dependent variable is a worker-level layoff or separation indicator. β^{Layoff} and $\beta^{\text{Separation}}$ denote the productivity–layoff and productivity–separation slopes; the third row reports their ratio. Worker controls include AKM worker effects, race, occupation, tenure, and flexible age × gender × education interactions. Year fixed effects absorbed; standard errors clustered at the 3-digit industry × state level. Sample as in Section 1. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

2.3 Accounting for Time-Varying Shocks

Another potential explanation for the productivity–layoff gradient is that cross-sectional variation in productivity covaries with time-varying shocks. At least two stories are plausible. First, firms that are more productive at baseline may also experience the most growth. If so, these firms could exhibit lower layoff rates not because of their higher baseline productivity, but because of their faster growth. Second, more productive firms might be more stable—that is, face smaller shocks—which could also explain their lower layoff rates.

We construct two measures of growth: average annual employment growth from RAIS and average annual sales growth from PIA. In both cases, these averages are computed across years—at the firm level in RAIS and at the 3-digit industry–state level in PIA. We also construct corresponding measures of volatility, defined as the within-firm (or within-industry–state) standard deviation of each growth measure across years. To assess whether time-varying shocks drive our results, we regress layoff rates on value added while controlling for each of these measures. We use a piecewise linear specification in the growth measures, allowing the slope to differ above and below zero. This choice is motivated by the hockey-stick pattern documented by Davis and Haltiwanger (1999) and replicated in many contexts: shrinking firms exhibit much higher separation rates than the median firm, while rapidly expanding firms also have higher separation rates than the median firm, though to a lesser extent.

The results are reported in Table III. Column (1) presents the baseline specification without controls. Columns (2)–(3) add the growth controls and recover the well-known hockey-stick pattern: layoffs decline sharply with growth among shrinking firms and rise modestly among expanding firms. Columns (4)–(5) incorporate our measures of volatility and show that higher volatility is associated with higher layoff rates.

Across all specifications, one common pattern emerges: the coefficient on baseline productivity is virtually unchanged. These results indicate that the productivity–separation gradient is not driven by high-productivity firms facing a systematically different distribution of shocks.

2.4 Alternative Sample Definitions and Productivity Measures

There are at least two limitations with our baseline productivity measure, value added sourced from the PIA. First, it is aggregated at the industry–state level rather than observed at the firm level. Second, it is restricted to manufacturing, so our findings might be specific to this sector.

To assess whether these limitations affect our conclusions, we replicate the analysis using two alternative measures observed at the firm level: revenue per worker from ORBIS and log firm size. Both are informative proxies for firm productivity—revenue per worker captures firm-level output directly, while firm size reflects equilibrium predictions from canonical models that more productive firms are larger (e.g., Burdett and

Table III
The Productivity–Layoff Gradient Is Not Explained by Firm-Level Shocks

	Dependent variable: Annual layoff rate				
	(1)	(2)	(3)	(4)	(5)
Value Added	-0.042***	-0.039***	-0.043***	-0.036***	-0.042***
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)
Emp. Growth (positive)		0.155***			
		(0.019)			
Emp. Growth (negative)		-0.315***			
		(0.024)			
Sales Growth (positive)			-0.031		
			(0.062)		
Sales Growth (negative)			-0.283***		
			(0.098)		
SD(Emp. Growth)				0.145***	
				(0.013)	
SD(Sales Growth)					0.027
					(0.017)
Observations	290,444	288,730	278,173	277,136	278,043

Notes. OLS estimates of the productivity–layoff gradient controlling for firm-level growth and volatility. Firm-year level, weighted by employment; year fixed effects absorbed; standard errors clustered at the 3-digit industry $\times$ state level. Column (1): baseline. Columns (2)–(3): piecewise linear employment and sales growth, following Davis and Haltiwanger (1999). Columns (4)–(5): add within-firm (or within-industry–state) standard deviation of each growth measure. Sample as in Section 1. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Mortensen, 1998; Hopenhayn, 1992; Moscarini and Postel-Vinay, 2012). Additionally, they cover different samples than the PIA data. As discussed in Section 1.1, the ORBIS data are not representative but cover a large share of the Brazilian labor market across different sectors. Firm size is computed from RAIS and, hence, is measured for all firms in the country.[15]

The results are reported in Table B.II—using revenue per worker—and Table B.III—using firm size. The estimates are very similar to those obtained using value added. This consistency suggests that the industry–state aggregation of the PIA data does not substantially affect our findings, likely because the level of aggregation is sufficiently granular. It also suggests that our findings extend beyond the manufacturing sector.

The relationship between firm size and productivity raises a potential confounder: more productive firms may exhibit lower layoff rates not because of higher productivity per se, but because they are larger. For instance, larger firms may face higher bureau-

[15] A potential concern is that larger firms have more within-firm transfers, biasing our estimates. However, transfers between establishments of the same firm are not recorded as separations in RAIS and therefore do not affect our measures of separation, quit, or layoff rates.

cratic or organizational costs of laying off workers. Table B.IV reports our main specifications using value added per worker as the productivity measure while controlling for firm size. The results are similar to our baseline findings, indicating that the lower layoff rates in more productive firms are not driven by their larger size.

Another concern is that small firms may add considerable noise to the estimates, since our main outcomes are layoff rates computed as a share of total employment. To address this issue, Table B.V reports results from a sample restricted to firms with at least 20 workers in every year. The estimates are very similar to those in the main specification.

In sum, our results are robust across alternative productivity measures and sample definitions. Across all specifications, more productive firms exhibit both lower separation and lower layoff rates. Moreover, the productivity–separation gradient is mostly driven by layoffs.

3 Wage Rigidity and the Productivity–Layoff Gradient

The central empirical pattern of this paper is that high-productivity firms lay off fewer workers. In this section, we show that downward wage rigidity is one of the key mechanisms behind this finding.

The intuition is as follows. When a firm-worker match experiences a negative productivity shock, the firm's markdown for that worker may turn negative. In the absence of wage rigidity, the firm and worker could renegotiate to a lower wage and preserve the match. But if the contract wage cannot be cut, the firm must lay off the worker instead, even when a mutually beneficial wage exists. Higher-productivity firms are better insulated against such shocks: their larger average markdowns provide a buffer that makes it less likely any given shock turns the markdown negative. This mechanism therefore predicts that the productivity–layoff gradient should be steeper in markets where wages are more rigid. Importantly, the relevant form of rigidity here is not about the level at which wages are initially set—as with minimum wages or collectively bargained wage floors—but rather about the inability to adjust wages downward after a match is formed.

We develop this intuition into a formal model in Section 4. In this section, we provide empirical evidence of its relevance. First, we document the extent of downward wage rigidity in our setting. Second, we show that wage rigidity is associated with higher layoff rates and a steeper productivity–layoff gradient.

3.1 Documenting Wage Rigidity

To study wage rigidity, we exploit the fact that the RAIS data report total compensation separately as contract wages and variable pay.[16] Variable pay encompasses bonuses, performance pay, and overtime. Firms face substantial rigidity in adjusting contract wages, as the Brazilian Constitution prohibits wage reductions unless authorized by a collective bargaining agreement.[17] In contrast, variable pay is not constrained by these regulations. Figure II presents the distribution of yearly wage changes for workers who remain in the same firm across two consecutive years. Consistent with these regulations, Panel (A) shows that only 1.34% of workers experience a reduction in their contract wage, whereas Panel (B) reveals that 9.28% see a reduction in their total compensation. Additionally, contract wage changes cluster around zero, whereas total compensation exhibits no such bunching.[18] These patterns highlight the substantial rigidity of contract wages relative to the flexibility of variable pay, a phenomenon well-documented in other contexts (Altonji and Devereux, 1999; Messina et al., 2010; Anger, 2011; Grigsby et al., 2021).

Motivated by the patterns in Figure II, we construct a firm-level measure of wage rigidity based on its reliance on contract wages versus variable pay, an approach similar to Makridis and Gittleman (2022), Reizer (2022), and Sockin and Sockin (2025). Specifically, we proxy wage rigidity using the average share of contract wages in total compensation:

$$\text{ContractShare}_j = \frac{1}{N_j} \sum_{i|j(i,t_{j0})=j} \frac{\text{ContractWage}_i}{\text{VariablePay}_i + \text{ContractWage}_i}, \tag{4}$$

where $j(i,t)$ denotes the firm that employs worker i in year t, t_{j0} is the first year firm

[16]Throughout the paper, "wage" refers to total compensation—the sum of the contract and variable components—unless noted otherwise.

[17]Title II, Chapter I, Article 7, Paragraph VI of the 1988 Constitution.

[18]High inflation can relax nominal wage rigidity, since firms can reduce real wages by holding nominal wages fixed. Yet the pronounced bunching at zero in Panel (A) of Figure II indicates that this channel is often insufficient: many firms appear constrained from making nominal wage cuts they would otherwise choose.

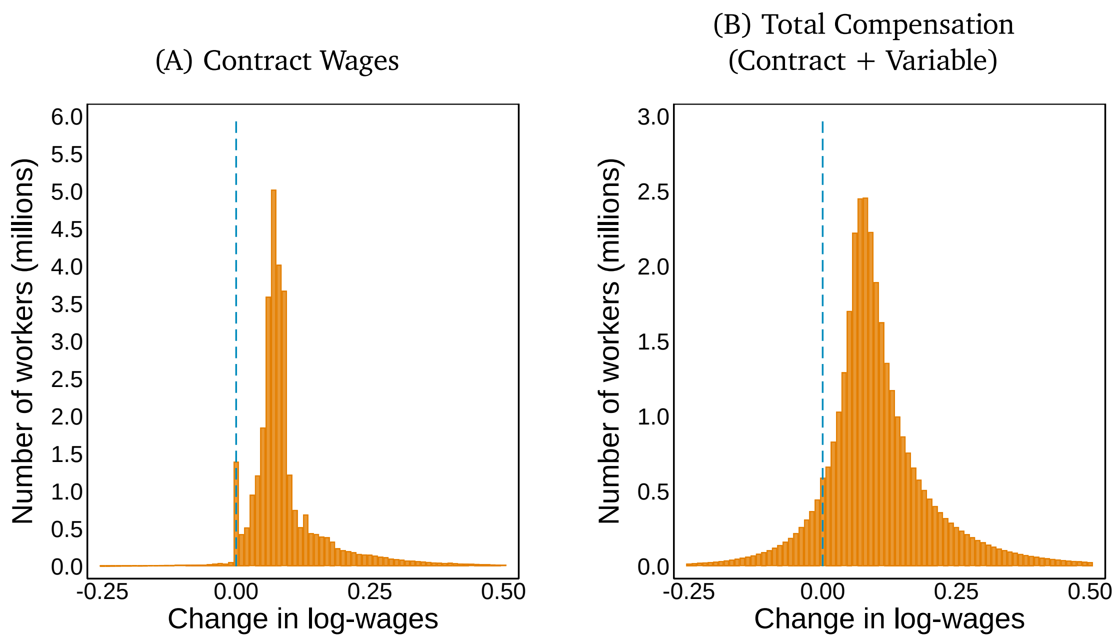

Figure II
Distribution of Wage Changes for Stayers

Notes. Distribution of nominal wage changes for workers who remain at the same firm for two consecutive years (stayers). Panel (A): contract wages; Panel (B): total compensation (contract wages plus variable pay). Wages are not inflation-adjusted. Sample as in Section 1.

j appears in the sample, and N_j is the size of firm j in that year. Since variable pay can be adjusted while contract wages cannot, higher ContractShare indicates stronger wage rigidity. This measure captures rigidity in wage *adjustment*—the degree to which firms can reduce compensation in response to negative shocks—rather than rigidity in wage *setting*. The distinction is deliberate: our mechanism operates through the firm's inability to cut an existing wage, not through constraints on initial wage offers.[19] To address concerns about endogeneity, such as ContractShare responding to productivity shocks, we compute ContractShare using the first year each firm appears in the sample and hold it fixed throughout the analysis. Furthermore, we exclude the year used to define ContractShare from subsequent analysis. Figure B.II presents the distribution of ContractShare across firms and reveals substantial variation: the median share is 0.92, the 25th percentile is 0.83, and the 75th percentile is 0.98.

To validate ContractShare as a proxy for wage rigidity, we examine its relationship with wage changes among workers who remain at the same firm for two consecutive

[19]Importantly, collective bargaining agreements (CBAs) in this context serve to *relax* rather than tighten this constraint: the default rule prohibits wage cuts unilaterally, but a CBA can authorize them for a given sector or firm. Therefore, ContractShare should be seen as an upper bound for the legally imposed wage rigidity.

years. The results, reported in Figure B.III, support this interpretation: firms with higher ContractShare are less likely to reduce workers' total compensation.

3.2 Wage Rigidity, Productivity, and Layoffs

Having established that ContractShare is a valid proxy for wage rigidity, we now examine how it relates to layoff rates. A natural concern is that any relationship between ContractShare and layoffs could reflect productivity differences across firms rather than wage rigidity per se. More productive firms may differ from less productive ones across many dimensions, including the composition of their compensation packages. Indeed, Figure B.IV shows that ContractShare is negatively related to value added: more productive firms tend to rely less on contract wages. To disentangle these effects, we estimate the following regression:

$$\text{Layoff}_{it} = \alpha + \beta_{\text{CS}}\text{ContractShare}_{j(i,t)} + \beta_{\text{VA}}\text{VA}_{j(i,t)t} + \theta_t + \beta_X X_{it} + \beta_H H_{j(i,t)t} + \varepsilon_{it}, \quad (5)$$

where $\text{VA}_{j(i,t)t}$ measures firm productivity (value added), θ_t denotes year fixed effects, X_{it} is a vector of worker characteristics, $H_{j(i,t)t}$ is a vector of time-varying firm-level shocks, and ε_{it} is an error term. Our coefficient of interest is β_{CS}, which captures the relationship between ContractShare and layoffs conditional on productivity, worker composition, and both nationwide and firm-specific time-varying factors.

Table IV reports OLS estimates of Equation (5). Column (1) regresses layoff rates on ContractShare alone. The coefficient is positive and highly significant, confirming that firms with more rigid wages exhibit higher layoff rates. This is consistent with the mechanism outlined above: when contract wages cannot be cut, firms must rely on layoffs to absorb negative productivity shocks. Column (2) adds value added as a control. The coefficient on ContractShare declines but remains large and highly significant, confirming that wage rigidity predicts layoffs beyond what can be explained by productivity differences alone.

Part of this relationship could reflect worker sorting: high-skill workers are less likely to be laid off and may also receive more variable pay, generating a spurious positive relationship between ContractShare and layoffs. To address this concern, Column (3) controls for worker composition. As with the productivity–layoff gradient itself (Section 2), worker sorting explains part of the relationship: the coefficient on ContractShare attenuates once we add controls. However, the main result remains. Even after conditioning on a rich set of worker characteristics—race, occupation, tenure, AKM worker effects, and flexible interactions among gender, age, and education—firms with more rigid wages continue to exhibit significantly higher layoff rates.

A related concern is that a firm's own ContractShare is endogenous: a low ContractShare could reflect unusually high bonus payments in a particularly good year for the firm. This concern is partly mitigated by our measurement strategy, since we compute ContractShare in the first year each firm appears in the sample and exclude that year from the analysis. Still, if shocks are persistent, this may not fully resolve the issue. To address it, we implement an additional robustness check that controls for time-varying shocks, using the same set of controls as in Table III. The results, reported in Column (4) of Table IV, show that including these controls has little impact on our estimates.

In summary, the results in this section indicate that wage rigidity plays an important role in explaining variation in layoff rates across firms. We return to this question in Section 5.3, where we exploit variation in wage rigidity across markets. This market-level analysis advances the results in two ways. First, variation across markets helps alleviate concerns that ContractShare is endogenously chosen by firms. Second, it allows us to study not only the role of wage rigidity in layoff levels, but also its role in the productivity–layoff gradient.

Table IV
Wage Rigidity Is Associated with Higher Layoff Rates

	Dependent variable: Annual layoff rate			
	(1)	(2)	(3)	(4)
Contract Share	0.191***	0.093***	0.048***	0.048***
	(0.026)	(0.019)	(0.014)	(0.014)
Value Added		-0.038***	-0.017***	-0.013***
		(0.004)	(0.002)	(0.002)
Worker covariates			✓	✓
Time-varying shocks				✓
Observations	8,962,115	8,962,115	8,961,480	8,552,280

Notes. OLS estimates of Equation (5). The dependent variable is a worker-level layoff indicator. ContractShare defined in Equation (4). Year fixed effects absorbed; standard errors clustered at the 3-digit industry $\times$ state level. Sample as in Section 1. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

4 Conceptual Framework

This section introduces a simple model to explain our empirical finding that higher-productivity firms have lower layoff rates. Building on the evidence in Section 3, which documents a link between wage rigidity and layoffs, our framework treats wage rigidity as the source of endogenous layoffs. Firms set wages before match productivity is realized, so a sufficiently bad productivity shock makes retention unprofitable and triggers a layoff. Apart from these features, the model is kept as parsimonious as possible.

We consider a partial-equilibrium, one-period random-search model with homogeneous workers.[20] The homogeneity assumption can be interpreted as focusing on a specific labor market segment, for instance defined by education or occupation. Random search implies that firms are matched with an exogenously given type and number of workers. The partial-equilibrium environment treats the distribution of workers' outside options as exogenous: firms take as given the determinants of these outside options, such as competing offers, government policies, and market-level conditions.

The rest of this section proceeds as follows. We first describe the economy and the timing of agents' decisions (Section 4.1). We then analyze the model in three steps. First, we characterize equilibrium across firms that differ only in productivity (Section 4.2). Our central theoretical result is that higher-productivity firms pay higher wages, accumulate larger markdowns, and lay off workers less frequently. We then introduce heterogeneity in wage-setting power (Section 4.3) and wage rigidity (Section 4.4).

4.1 Setup

There is a single firm with productivity ψ.[21] The firm meets a unit mass of ex-ante homogeneous workers, each with productivity α. Each worker-firm match realizes a mean-zero idiosyncratic productivity shock η_i. Total match productivity is therefore $\alpha + \eta_i$, with expected value α. The total revenue the firm receives from worker i is $\psi + \alpha + \eta_i$.

[20] The one-period formulation is sufficient to capture the key patterns in the data. Extending the model to multiple periods leaves the qualitative results unchanged; we adopt the static formulation for expositional clarity.

[21] We interpret ψ as value added per worker, consistent with our empirical productivity measure (Section 2). This measure conflates output prices and production efficiency. Since the model's predictions depend only on total match value, it is immaterial whether variation in ψ reflects either source.

If employed, workers derive utility equal to the wage, and if not, they receive a stochastic, exogenously determined outside option b_i. Workers quit if the wage they receive is below b_i. Both shocks, η_i and b_i, are idiosyncratic and follow known distributions, F_η and F_b, respectively, with $\mathbb{E}[\eta_i] = 0$ normalized without loss of generality.

The firm's per-worker profit is given by the wage markdown $\mu_i(w)$, the difference between worker-specific productivity and the wage:

$$\overbrace{\mu_i(w)}^{\text{markdown}} \equiv \overbrace{\psi}^{\text{firm productivity}} + \overbrace{\alpha + \eta_i}^{\text{match productivity}} - \overbrace{w}^{\text{wage}}. \tag{6}$$

Since $\mathbb{E}[\eta_i] = 0$, the *average* per-worker markdown at wage w is $\mu_\psi(w) \equiv \psi + \alpha - w$.

The timing of the model is illustrated in Figure III. The firm first posts a wage w to all workers it meets. Match-specific productivity shocks η_i and outside option shocks b_i are then realized. The firm then decides which workers to lay off, after which the remaining workers choose whether to quit. Payoffs are realized at the end of the period.

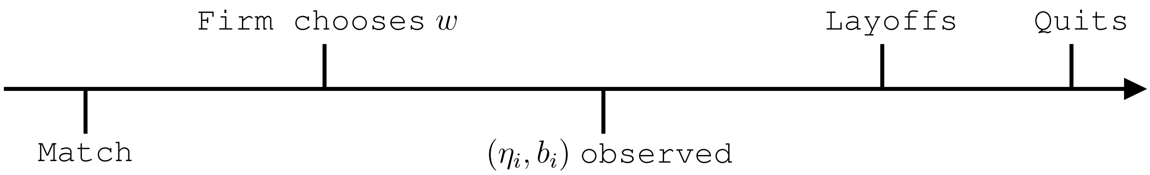

Figure III
Model Timeline

Notes. This figure shows the timeline of the model presented in Section 4.

The assumption that wages are set before match-specific productivity is realized reflects that the quality of a worker–firm match is not fully observed ex ante: firms post wages in advance and, in our framework, do not condition them on realized match productivity. This timing is standard in wage-posting models (e.g., Burdett and Mortensen, 1998) and is consistent with the literature on match-quality learning (Jovanovic, 1979). In practice, however, firms may partially adjust compensation ex post through variable components such as bonuses, as discussed in Section 3. We therefore present an extension in Section 4.4 that allows for partial wage rigidity.

If the realized markdown is negative ($\mu_i(w) < 0$), the worker is laid off. Since firms cannot adjust wages after observing the shocks, layoffs occur even in cases where an alternative wage would make both the worker and the firm better off—the source of inefficient layoffs in the model.[22]

An equilibrium is defined by the optimality of three decisions: layoffs, quits, and wages. First, the firm lays off a worker if their markdown is negative. Second, workers quit if the outside option is higher than the wage. Third, the firm chooses wages to maximize expected profits per meeting. Appendix A.1 defines an equilibrium formally.

4.2 Drivers of the Layoff–Separation Gradient

In Section 2, we show that high-productivity firms have lower layoff rates. We now use the model to shed light on the mechanism underlying this gradient. We begin with Proposition 1, which characterizes how equilibrium outcomes—including layoff rates—vary with firm productivity.

Proposition 1 *Assume that outside option and match-productivity shocks are distributed such that both F_b and $1 - F_\eta$ are log-concave. Then, in equilibrium:*

(I) Firm size is weakly increasing in firm productivity;[23]
(II) Wages are weakly increasing in firm productivity;
(III) The separation rate is weakly decreasing in firm productivity;
(IV) The layoff rate is weakly decreasing in firm productivity;
(V) The average markdown is weakly increasing in firm productivity.

Formal statement: *Theorem 1 and Remark 2 in Appendix A.1.*

Proof: *Appendix A.2.*

Remark 1 *The log-concavity assumptions of Proposition 1 are satisfied by Normal and Uniform distributions for all parameter values.*

Consistent with prior theoretical work (Postel-Vinay and Robin, 2002; Elsby and Gottfries, 2022), the model predicts that wages and employment comove with firm productivity, while total separation rates decline with it. Our framework builds on this literature by additionally characterizing the role of layoffs.[24]

[22] Inefficient quits may also arise when the firm would prefer to raise wages to retain a worker but cannot. While the model accommodates both margins, we focus the discussion on layoffs given their dominant role in our empirical results (Section 2).

[23] Firm size is defined as the expected number of workers employed by the firm at the end of the period, after layoff and quit decisions are made.

[24] Since wages increase with productivity, high-productivity firms mechanically have lower quit probabilities among workers who are not laid off. Consistent with the focus of the paper, we emphasize the layoff margin.

Our main contribution is to explain why higher-productivity firms lay off workers less frequently. A layoff occurs when a worker's individual markdown is negative, which happens when the realized match-specific shock is sufficiently low to offset the firm's average markdown. Layoff rates therefore decline as markdowns increase, and the key question is whether higher-productivity firms systematically exhibit larger markdowns. While this is intuitive, the result is not immediate: higher-productivity firms also pay higher wages to increase worker retention, and if the retention incentives were sufficiently strong they could end up with *lower* markdowns—as in Burdett and Mortensen (1998).

The firm's tradeoff: retention versus markdown. We can gain further insight into why markdowns increase with firm productivity by reformulating the firm's problem as a choice over *retention* (ρ), defined as the share of workers who do not quit, and the *average markdown* (μ). This choice is subject to the existence of a wage w that supports the pair $\{\rho, \mu\}$. We refer to this constraint as the production possibility frontier (PPF).[25] Since higher wages increase retention but reduce markdowns, the PPF is downward sloping.

In this formulation, the firm solves:

$$\max_{\rho,\mu} \overbrace{\rho \cdot \left[1 - \delta(\mu)\right]}^{\text{retention} \times \text{survival}} \cdot \overbrace{\pi(\mu)}^{\text{profit per retained worker}} \quad \text{subject to PPF,} \tag{7}$$

where $\delta(\mu)$ is the layoff rate (decreasing in μ), and $\pi(\mu)$ is per-worker profit (increasing in μ). The objective is expected profit per match.[26]

Equation (7) makes explicit the trade-off firms face between retention and markdowns. Expected profits are determined by three components: retention, survival in the layoff stage, and average profit per retained worker.[27] The key tension is that higher markdowns (i.e., lower wages) increase both survival and profit per retained worker,

[25] Formally, the PPF requires a wage w such that $\mu = \psi + \alpha - w$ and $\rho = P(b_i \leq w)$.

[26] Maximizing expected profit per match is equivalent to maximizing total profits because the mass of matches is exogenously determined and independent of the firm's choices.

[27] Profits per retained worker and average markdown are distinct: the average markdown is defined over all workers the firm meets, whereas profit per retained worker conditions on those still employed at the end of the period, when profits are realized. Formally, $\pi(\mu) = \mu + \mathbb{E}[\eta_i \mid \mu + \eta_i > 0]$. Under the assumptions of Proposition 1, $\pi(\mu)$ is increasing.

but reduce retention.

Illustrative example. To see why higher-productivity firms choose larger markdowns, consider a simple calibration in which $\eta_i \sim U[-\sigma_\eta, \sigma_\eta]$ and $b_i \sim U[0, \sigma_b]$. Appendix A.2 shows that, under this parameterization, the firm's problem in Equation (7) becomes:

$$\max_{\rho,\mu} \rho \cdot (\mu+1)^2 \quad \text{subject to: } \mu + \sigma_b \rho = \psi + \alpha. \tag{8}$$

Figure IV illustrates Equation (8). The solid lines depict the PPF for a high-productivity firm (orange) and a low-productivity firm (blue); the dashed lines are isoprofit curves. A higher ψ shifts the PPF outward. Crucially, the firm's objective depends on the *product* of ρ and μ,[28] so these are *complementary inputs*: higher markdowns strengthen the incentive to retain workers, and vice versa. When the PPF expands, the firm therefore optimally raises *both* ρ and μ—higher-productivity firms end up with higher markdowns even though they pay higher wages, which delivers lower layoff rates.

In sum, the model implies that high-productivity firms have lower layoff rates because they sustain larger markdowns. We take this prediction to the data in Section 5.1.

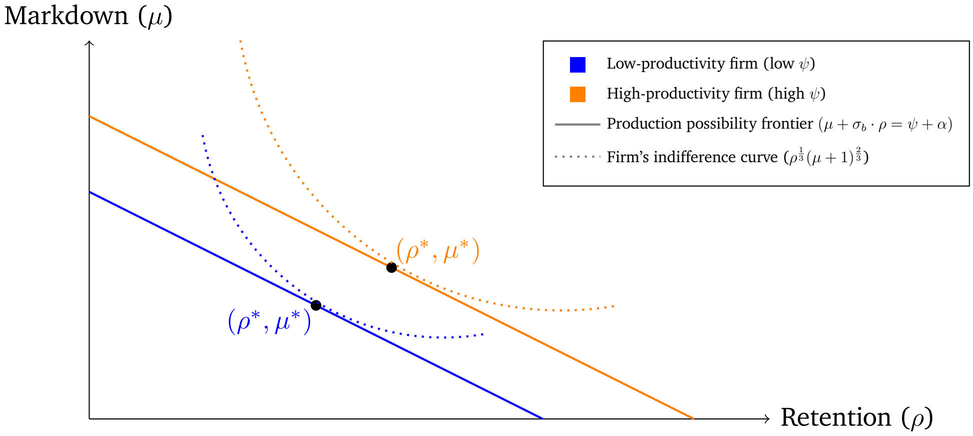

Figure IV
High-Productivity Firm Has *Both* Higher Markdown and Retention

Notes. This figure illustrates the model presented in Equation (8). Solid lines represent the production possibility frontier and dotted lines represent firms' indifference curves. Dots denote equilibrium outcomes. Two firms are represented in the figure: high-productivity (orange) and low-productivity (blue).

[28] The exponent on $(\mu + 1)$ is twice that on ρ in Equation (8) because a higher markdown influences both per-worker profit *and* the layoff rate.

4.3 The Wage–Layoff Gradient and Wage-Setting Power

Proposition 1 holds wage-setting power—captured by the outside option distribution F_b—fixed across firms. But wage-setting power can vary across firms for reasons such as differences in local labor market competition (i.e., monopsony power) or firm-specific amenities. In our model, both of these non-productivity factors are captured by the distribution of outside options F_b: monopsony power is reflected in the shape of F_b (the elasticity of labor supply to the firm is $\varepsilon(w) = wf_b(w)/F_b(w)$), while firm-specific amenities shift its location.[29]

Consider next the case in which ψ is held fixed and F_b varies. Since $\mu^* = \psi + \alpha - w^*$, each additional dollar of wages compresses markdowns one-for-one. Lower markdowns, in turn, imply more layoffs. This contrast with the productivity channel yields a sharp empirical test: when productivity drives wage variation across firms, higher wages are associated with fewer layoffs; when wage-setting power drives wage variation, higher wages are associated with more layoffs. The *sign* of the cross-sectional wage–layoff relationship therefore reveals which channel dominates.

Proposition 2 *Assume that outside option and match-productivity shocks are distributed such that both F_b and $1 - F_\eta$ are log-concave. Further assume an interior equilibrium. Then:*

> *(I) Productivity channel (ψ varies, outside option distribution fixed): Markdown is weakly increasing in wages; firms paying higher wages have layoff rates that are weakly lower.*
>
> *(II) Wage-setting power channel (ψ fixed, outside option distribution varies across firms): markdown is decreasing in wages; firms paying higher wages have higher layoff rates.*

Formal statement: *Theorem 2 and Remark 3 in Appendix A.1.*

Proof: *Appendix A.2.*

Part (I) follows directly from Proposition 1: since both w^* and μ^* are increasing in ψ, they move together along the equilibrium path. Part (II) is a mechanical identity: holding ψ fixed, $\mu^* = \psi + \alpha - w^*$ is strictly decreasing in w^*. Any force that raises wages without changing productivity therefore compresses markdowns, increasing the incidence of negative-markdown matches and hence layoffs.

[29] A firm offering amenity value a retains workers whenever $b \leq w + a$, equivalent to facing a shifted distribution $\tilde{F}_b(w) = F_b(w + a)$.

In practice, both channels operate simultaneously: wages reflect both productivity differences and wage-setting power differences. A negative wage–layoff relationship in the cross-section implies the productivity channel dominates; a positive relationship would implicate wage-setting power. We take these predictions to the data in Section 5.2.

4.4 Heterogeneous Wage Rigidity and Layoff Rates

The results in Propositions 1 and 2 take the degree of wage rigidity as given: wages are set before the match-specific shock, and if the realized markdown turns negative, the worker is laid off. In practice, however, firms differ in the extent of wage rigidity. As documented in Section 3, wage flexibility varies substantially across firms and markets. In this section, we extend the model to incorporate this heterogeneity and derive its implications for layoffs and the productivity–layoff gradient.[30]

Following Gertler and Trigari (2009), we introduce a Calvo-style friction: each match is rigid (wage cannot be adjusted) with probability ϕ, and flexible with probability $1-\phi$. The baseline model analyzed in Sections 4.2 and 4.3 corresponds to fully rigid wages ($\phi = 1$).

The timing in the extended model is as follows. The firm first posts a wage w. Match-specific productivity and outside option shocks, η and b, are then realized. The match is subsequently determined to be rigid (with probability ϕ) or flexible (with probability $1-\phi$). In rigid matches, quit and layoff decisions follow the baseline model of Section 4.1. In flexible matches, if there exists a wage that makes both the firm and the worker better off than separating, the wage is renegotiated to some w' within this feasible set; we impose no restrictions on bargaining power at this stage.[31] If no such wage exists, the wage remains at the posted level w and the baseline layoff/quit rules apply: the worker is laid off if $\mu_i(w) < 0$, and otherwise quits if $b > w$.

Proposition 3 characterizes how the degree of wage rigidity affects layoffs and the productivity–layoff relationship. Before stating it, we define the key object: a *savable layoff* is a match that is laid off under a rigid wage but would be retained under a flexible wage. Equivalently, it is a match in which (i) the markdown at the posted wage is

[30]Propositions 1 and 2 continue to hold in the extended model under additional assumptions; see Theorems 1 and 2 in Appendix A for details.

[31]We assume that the renegotiated wage w_i' may depend flexibly on b_i, ψ, α, and η_i, but not on the pre-posted wage w.

negative—so the firm would want to lay off—and (ii) the total match surplus is positive—so the worker could be retained at some Pareto-improving wage.

Proposition 3 *Assume that the distributions of outside options and match-specific productivity shocks are such that both F_b and $1 - F_\eta$ are log-concave. Further assume an interior equilibrium. Then:*

(I) The layoff rate is increasing in wage rigidity.

(II) The effect of wage rigidity on the productivity–layoff gradient depends on how the savable-layoff rate varies with firm productivity:

(II.i) If the savable-layoff rate is decreasing in firm productivity, wage rigidity amplifies the productivity–layoff gradient.

(II.ii) If the savable-layoff rate is increasing in firm productivity, wage rigidity weakens the productivity–layoff gradient.

Formal statement: *Theorem 3 in Appendix A.1.*

Proof: *Appendix A.2.*

The first result in Proposition 3 formalizes the intuition that guided our empirical analysis in Section 3: stronger wage rigidity is associated with more layoffs. The logic is straightforward. In flexible matches, a layoff requires both a negative markdown *and* a negative surplus (so that no Pareto-improving wage exists). In rigid matches, a negative markdown alone suffices. Hence more flexible matches (lower ϕ) must mean fewer layoffs.

The relationship between wage rigidity and the productivity–layoff gradient is more complex. As stated in Proposition 3, it depends on whether the savable-layoff rate is decreasing or increasing in firm productivity—and this is theoretically ambiguous because two opposing forces are at work. A savable layoff requires two conditions to hold simultaneously: the markdown at the posted wage is negative, and the total surplus of the match is positive. Higher productivity affects both:

- *Markdown-buffer effect.* Higher productivity raises the expected markdown, so fewer matches have a negative markdown. This lowers the savable-layoff rate and pushes toward amplification.
- *Surplus-composition effect.* Higher productivity raises total match value, so among matches with a negative markdown, a larger share remains surplus-positive, i.e., satisfies $\psi + \alpha + \eta - b \geq 0$. This increases the savable-layoff rate and pushes toward de-amplification.

Which force dominates depends, among other factors, on the slope of the relationship between productivity and markdowns.[32] When this relationship is sufficiently steep, even a small increase in productivity substantially reduces the mass of matches at risk of layoff. In this case, the markdown-buffer effect dominates: the savable-layoff rate declines with productivity, and wage rigidity amplifies the productivity–layoff gradient.

In sum, Proposition 3 delivers two lessons for the empirical analysis. First, more rigid wages generate more layoffs—an unambiguous prediction that we can test directly. Second, whether rigidity amplifies or attenuates the productivity–layoff gradient depends on the strength of the markdown–productivity relationship, which can also be assessed empirically. We provided initial evidence on these predictions in Section 3 and present additional evidence in Section 5.3.

5 Empirical Evidence on the Mechanism

In this section, we take the model's predictions to the data. Section 5.1 tests the central mechanism: higher-productivity firms exhibit larger markdowns, which are in turn associated with lower layoff rates (Proposition 1). Section 5.2 examines the wage–layoff relationship and shows that its sign is consistent with productivity being the dominant driver of wage variation across firms (Proposition 2). Section 5.3 shows that greater wage rigidity is associated with higher layoff rates (Proposition 3).

5.1 Proposition 1: Productivity, Markdowns, and Layoff Rates

In this subsection, we describe the empirical relationship between markdowns, productivity, and layoffs, and show that it aligns with the predictions of Proposition 1. We measure markdowns using the PIA dataset: our baseline definition is Markdown $= \log\left(\frac{\text{VA}-\text{Labor Costs}}{\text{workers}}\right)$, i.e., log net value added per worker.[33] We fix markdowns to each firm's first year in the sample, consistent with our treatment of value added.

Figure V presents two facts. Panel (A) shows a strong positive gradient between markdowns and firm productivity (value added), consistent with the prediction of our

[32] A formal statement is provided in Equation (A.3) in Appendix A.1.

[33] We consider three alternative definitions as robustness checks: $\log(\text{VA/workers}) - \log(\text{Labor Costs/workers})$; $(\text{VA} - \text{Labor Costs})/\text{VA}$; and $(\text{VA} - \text{Labor Costs})/\text{Labor Costs}$. Table B.VI shows that our findings in this section are robust to these alternative definitions.

model that higher-productivity firms retain a larger surplus. Panel (B) shows that annual layoff rates fall sharply with markdowns: firms with larger wage cushions lay off fewer workers. Panels A and B of Table B.VI show that these results are robust to alternative definitions of markdowns.

In our theoretical framework, which holds worker type fixed, the negative productivity–layoff gradient is entirely driven by the larger markdowns of more productive firms. To evaluate this mechanism empirically, we implement a mediation analysis following Gelbach (2016). Specifically, we decompose the observed layoff–productivity gradient into three components: (i) worker sorting, (ii) markdowns, and (iii) a residual term.[34] The results indicate that markdowns account for 59% of the gradient, with worker sorting accounting for essentially all of the remainder and leaving only a small and statistically insignificant residual term. Thus, consistent with the model, once worker type is held fixed, markdowns explain essentially all of the remaining productivity–layoff gradient.

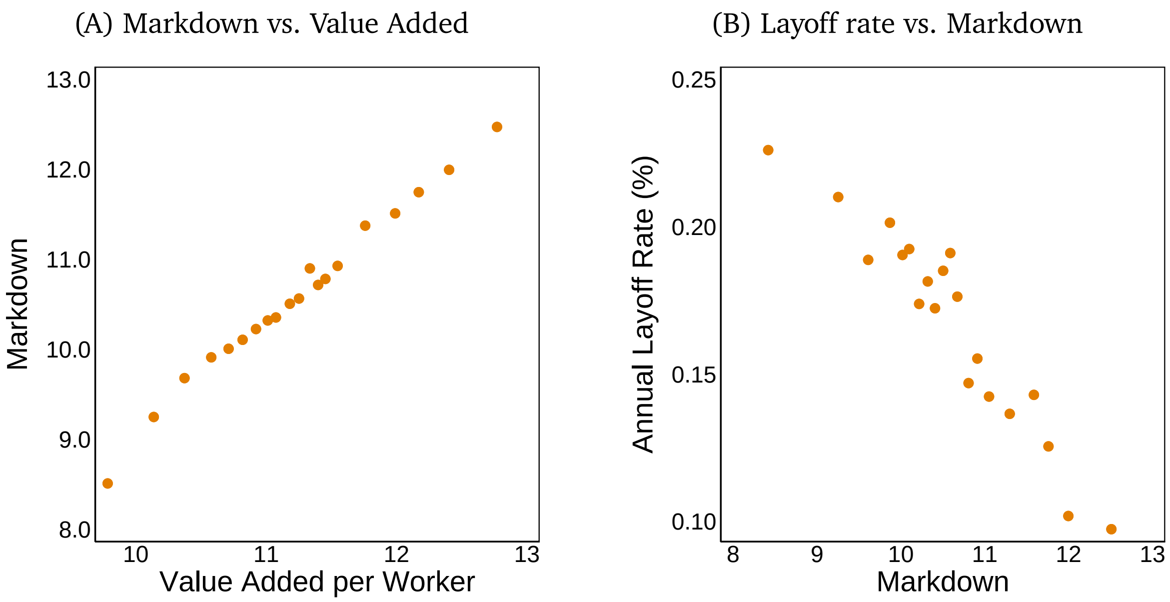

Figure V
Markdowns, Firm Productivity, and Layoff Rates

Notes. Binscatters with 20 equal-sized bins at the firm-year level. Markdowns defined as $\log((\text{VA} - \text{Labor Costs})/\text{workers})$, fixed to each firm's first sample year. Year fixed effects absorbed; weighted by firm size. Sample as in Section 1.

[34] Appendix F details the decomposition.

5.2 Proposition 2: Wages, Markdowns, and Layoff Rates

Proposition 2 yields the following empirical test: if productivity is the dominant source of wage variation across firms, higher wages should predict larger markdowns and lower layoff rates. If variation in wage-setting power dominates instead, the signs reverse. In this section, we take this prediction to the data using the AKM firm pay premium—the firm component of wages estimated from a two-way fixed effects decomposition (Supplemental Appendix C)—as our measure of "wages." This measure captures variation in what firms pay to observationally identical workers, which corresponds directly to the wage w in the model.[35]

Figure VI presents the results. Panel (A) plots markdowns against the firm pay premium. Consistent with the productivity channel, higher-paying firms exhibit *larger* markdowns. This confirms that, in the cross-section, the higher wages paid by high-pay-premium firms are more than offset by their higher productivity, resulting in larger wage cushions. Panel C of Table B.VI shows that this result is robust to alternative definitions of markdowns.

Panel (B) plots layoff rates against the firm pay premium. The gradient is strongly negative: firms that pay higher wages to identical workers lay off fewer of them. This is exactly the sign predicted by the productivity channel in Proposition 2, and the opposite of what would arise if heterogeneity in wage-setting power were the dominant source of wage dispersion.

Note that our results do not imply that firms have weak wage-setting power or that such power is unimportant. Rather, they indicate that wage differences across firms are driven primarily by variation in productivity rather than by *variation* in wage-setting power.

5.3 Proposition 3: Wage Rigidity, Productivity, and Layoff Rates

We now turn to an empirical assessment of the model's predictions on wage rigidity (Proposition 3). While Section 3.2 provided initial evidence, here we revisit the issue through the lens of the model. In particular, we examine not only the role of wage rigidity in layoff levels, but also its role in the productivity–layoff gradient.

[35]Figure B.V shows that the firm pay premium is positively related to value added, confirming that higher-paying firms are also more productive, as predicted by Proposition 1.

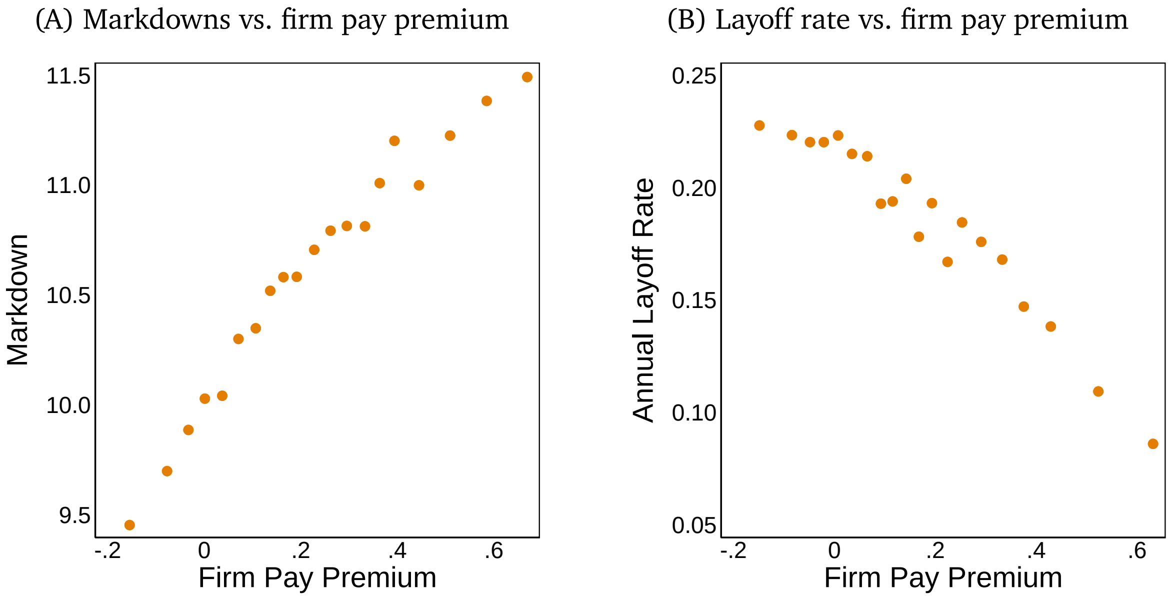

Figure VI
The Wage–Markdown and Wage–Layoff Gradients

Notes. Binscatters with 20 equal-sized bins. The firm pay premium is the AKM firm fixed effect (Supplemental Appendix C). Markdowns defined as $\log((\text{VA} - \text{Labor Costs})/\text{workers})$, fixed to each firm's first sample year. Year fixed effects absorbed; weighted by firm size. Sample as in Section 1.

We exploit cross-market variation in wage rigidity to test whether the productivity–layoff gradient is systematically steeper in more rigid labor markets. We define each market m as a 2-digit industry–state pair.[36] Following the framework in Section 4.4, we measure wage rigidity in market m as the share of jobs with rigid wages in that market. We classify a job as having a rigid wage if its contract share is at least 0.8. We assess robustness using a 0.9 threshold, as well as a continuous alternative that measures market-level wage rigidity by the average contract share in the market.

A natural concern is that the realized compensation structure in a market is itself an equilibrium outcome, reflecting endogenous firm choices rather than exogenous variation in wage-setting institutions. To mitigate this concern, we construct a leave-one-out (LOO) measure defined as the share of jobs with rigid wages across all other markets in the same 2-digit industry as m, excluding market m itself. We measure this object in 2010 and exclude that year from the subsequent analysis. We denote the resulting measure by $\widehat{\phi}_m$.

To quantify the productivity–layoff gradient within each market, we estimate the

[36] Value added is observed at the more disaggregated 3-digit industry–state level, so there remains within-market variation in value added.

following regression, analogous to Equation (3), separately by market:

$$\text{Layoff}_{it} = C_m + \beta_m \, \text{VA}_{j(i,t)t} + \theta_{mt} + \gamma_m^X X_{it} + \gamma_m^H H_{j(i,t)t} + \epsilon_{ijt}^M, \tag{9}$$

where $j(i,t)$ denotes the firm employing worker i at time t, and $m = m(i,t)$ denotes that worker's market. The terms X_{it} and $H_{j(i,t)t}$ control for worker characteristics and time-varying firm shocks, respectively; θ_{mt} denotes market-year fixed effects that capture market-wide time-varying shocks; and ϵ_{ijt}^M is an error term. The intercept C_m and coefficients β_m, γ_m^X, and γ_m^H are estimated separately for each market. Our coefficient of interest is β_m, which captures the market-specific productivity–layoff gradient.

We next test whether markets with greater wage rigidity exhibit more layoffs and a steeper productivity–layoff gradient by estimating:

$$\text{outcome}_m = \chi \, \widehat{\phi}_m + \epsilon_m^\beta, \tag{10}$$

where outcome_m denotes either the layoff rate in market m or $\widehat{\beta}_m$, χ is the coefficient of interest, and ϵ_m^β is an error term.[37]

OLS estimates of Equation (10) are reported in Table V. Column (1) shows that greater wage rigidity is associated with higher layoff rates, consistent with Proposition 3 and with the firm-level evidence in Section 3. This provides additional support for the mechanism through which wage rigidity leads to layoffs.

The relationship between wage rigidity and the productivity–layoff gradient is theoretically ambiguous (Proposition 3). The model predicts amplification when the markdown–productivity relationship is sufficiently steep,[38] in which case the markdown-buffer effect dominates. In our setting, the markdown–productivity gradient is strong, as shown in Panel (A) of Figure V, suggesting that the markdown-buffer effect may dominate and that greater wage rigidity may therefore be associated with a steeper productivity–layoff gradient.

To evaluate this prediction empirically, Columns (2) through (4) of Table V report OLS estimates of Equation (10), using the market-specific productivity–layoff gradient, β_m, as the outcome. Column (2) uses β_m estimated controlling only for time fixed effects.

[37]The term $\widehat{\beta}_m$ is estimated with error. Our main object of interest, however, is Equation (10), which uses $\widehat{\beta}_m$ as an outcome and therefore yields consistent estimates despite this measurement error.

[38]See Equation (A.3) in Appendix A.2.

Table V
Productivity–Layoff Gradient Is Steeper in Markets with More Wage Rigidity

	Average Layoff Rate	β_m		
	(1)	(2)	(3)	(4)
$\widehat{\phi}_m$	0.186***	-0.055**	-0.083***	-0.093***
	(0.059)	(0.022)	(0.019)	(0.028)
Worker covariates			✓	✓
Time-varying shocks				✓
Markets	75	72	72	72

Notes. OLS estimates of Equation (10). A market is a 2-digit industry $\times$ state pair. $\widehat{\phi}_m$ is the leave-one-out share of jobs with contract wages $\geq$ 80% of total pay. Column (1): market-level layoff rate on $\widehat{\phi}_m$, weighted by market employment. Columns (2)–(4): market-specific productivity–layoff slope $\widehat{\beta}_m$ on $\widehat{\phi}_m$, weighted by inverse squared SE of $\widehat{\beta}_m$. Each $\widehat{\beta}_m$ is estimated market-by-market with market$\times$year fixed effects, Equation (9). Checkmarks refer to controls in the estimation of $\widehat{\beta}_m$, not in the market-level regression. Worker covariates are as in Table II; time-varying shocks are as in Table III. Robust standard errors. Sample as in Section 1. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Column (3) uses the version estimated with worker controls, and Column (4) further adds controls for time-varying firm shocks. Across all specifications, we find that greater wage rigidity is associated with a steeper productivity–layoff gradient.

Note that controlling for worker composition does not weaken the relationship between wage rigidity and the productivity–layoff gradient. This contrasts with the results in Table IV, where controlling for worker composition substantially attenuates the association between wage rigidity and layoff rates. The reason is that, in Equation (10), the market-specific intercepts and time effects already absorb much of the relevant variation in worker composition, even in the specification without worker controls.

Table B.VII shows that these results are robust to alternative measures of wage rigidity. Estimates using an alternative threshold to classify a job as having a rigid wage—0.9 instead of 0.8—are very similar to the baseline results and are not statistically distinguishable from them. When we use the continuous measure, the magnitudes are not directly comparable, but the results point in the same direction: greater wage rigidity is associated with both higher layoff rates and a steeper productivity–layoff gradient.

Taken together, these results provide empirical support for the mechanism in Proposition 3. Markets with greater wage rigidity exhibit not only higher average layoff rates, but also a steeper productivity–layoff gradient. This pattern is consistent with wage rigidity operating through the markdown-buffer channel to shape cross-firm variation in layoff rates.

6 Final Remarks

This paper revisits a central empirical regularity in labor economics: high-productivity firms have lower worker separation rates. Using matched employer-employee data from Brazil that distinguishes quits from layoffs, we show that this pattern is driven primarily by layoffs rather than quits. The productivity–separation gradient therefore reflects systematic differences in job stability across firms, rather than mainly voluntary worker reallocation.

We develop a parsimonious framework in which downward wage rigidity and productivity uncertainty generate inefficient layoffs. Higher-productivity firms are better able to absorb adverse shocks because they operate with larger markdown buffers, making layoffs less likely. Consistent with this mechanism, we show that markets with stronger wage rigidity exhibit both higher layoff rates and a steeper productivity–layoff gradient, and that firms with higher estimated markdowns lay off fewer workers.

These findings have broader implications for labor market models and for worker welfare. If low-productivity firms disproportionately generate layoffs, then productivity differences across firms shape not only wages, but also exposure to costly job loss. More generally, wage-setting frictions may be central to understanding how employment risk is distributed across workers and firms.

References

Abowd, John, Francis Kramarz, and David Margolis, "High Wage Workers and High Wage Firms," *Econometrica*, 1999, *67* (2), 251–333.

Acabbi, Edoardo Maria, Andrea Alati, and Luca Mazzone, "Human Capital Ladders, Cyclical Sorting, and Hysteresis," 2024.

Altonji, Joseph G and Paul J Devereux, "The Extent and Consequences of Downward Nominal Wage Rigidity," *NBER Working Paper Series*, 1999.

Anger, Silke, "The Cyclicality of Effective Wages within Employer–Employee Matches in a Rigid Labor Market," *Labour Economics*, December 2011, *18* (6), 786–797.

Bertheau, Antoine, Edoardo Maria Acabbi, Cristina Barcelo, Andreas Gulyas, Stefano Lombardi, and Raffaele Saggio, "The Unequal Consequences of Job Loss across Countries," *American Economic Review: Insights*, 2023, *5* (3), 393–408.

Blanco, Andrés, Andrés Drenik, Christian Moser, and Emilio Zaratiegui, "A Theory of Labor Markets with Inefficient Turnover," 2024.

Bonhomme, Stephane, Thibaut Lamadon, and Elena Manresa, "A Distributional Framework for Matched Employer Employee Data," *Econometrica*, 2019, *87* (3), 699–739.

Burdett, Kenneth and Dale T. Mortensen, "Wage Differentials, Employer Size, and Unemployment," *International Economic Review*, May 1998, *39* (2), 257.

Card, David, Jörg Heining, and Patrick Kline, "Workplace Heterogeneity and the Rise of West German Wage Inequality," *The Quarterly Journal of Economics*, 2013, *128* (3), 967–1015.

Carlsson, Mikael and Andreas Westermark, "Endogenous Separations, Wage Rigidities, and Unemployment Volatility," *American Economic Journal: Macroeconomics*, January 2022, *14* (1), 332–354.

Couch, Kenneth A. and Dana W. Placzek, "Earnings Losses of Displaced Workers Revisited," *American Economic Review*, March 2010, *100* (1), 572–589.

Davis, Steven J. and John Haltiwanger, "Chapter 41 Gross Job Flows," in "Handbook of Labor Economics," Vol. 3, Elsevier, 1999, pp. 2711–2805.

__ **and Pawel M. Krolikowski**, "Sticky Wages on the Layoff Margin," *American Economic Review*, February 2025, *115* (2), 491–524.

__ **, R. Jason Faberman, and John C. Haltiwanger**, "The Establishment-Level Behavior of Vacancies and Hiring*," *The Quarterly Journal of Economics*, May 2013, *128* (2), 581–622.

DellaVigna, Stefano, Nicholas Otis, and Eva Vivalt, "Forecasting the Results of Experiments," *AEA Papers and Proceedings*, 2020, *110.*

Doornik, Bernardus Van, David Schoenherr, and Janis Skrastins, "Strategic Formal Layoffs: Unemployment Insurance and Informal Labor Markets," *American Economic Journal: Applied Economics*, January 2023, *15* (1), 292–318.

Ehrlich, Gabriel and Joshua Montes, "Wage Rigidity and Employment Outcomes: Evidence from Administrative Data," *American Economic Journal: Macroeconomics*, January 2024, *16* (1), 147–206.

Elsby, Michael W L and Axel Gottfries, "Firm Dynamics, On-the-Job Search, and Labor Market Fluctuations," *The Review of Economic Studies*, May 2022, *89* (3), 1370–1419.

Gelbach, Jonah B., "When Do Covariates Matter? And Which Ones, and How Much?," *Journal of Labor Economics*, April 2016, *34* (2), 509–543.

Gerard, Francois and Joana Naritomi, "Job Displacement Insurance and (the Lack of) Consumption-Smoothing," *American Economic Review*, 2021, (3).

Gerard, François, Lorenzo Lagos, Edson Severnini, and David Card, "Assortative Matching or Exclusionary Hiring? The Impact of Employment and Pay Policies on Racial Wage Differences in Brazil," *American Economic Review*, 2021, *111* (10), 40.

Gertler, Mark and Antonella Trigari, "Unemployment Fluctuations with Staggered Nash Wage Bargaining," *Journal of Political Economy*, 2009, *117* (1), 38–86.

Grigsby, John, Erik Hurst, and Ahu Yildirmaz, "Aggregate Nominal Wage Adjustments: New Evidence from Administrative Payroll Data," *American Economic Review*, February 2021, *111* (2), 428–471.

Haltiwanger, John C., Henry R. Hyatt, Lisa B. Kahn, and Erika McEntarfer, "Cyclical Job Ladders by Firm Size and Firm Wage," *American Economic Journal: Macroeconomics*, April 2018, *10* (2), 52–85.

Haltiwanger, John, Henry Hyatt, and Erika McEntarfer, "Who Moves Up the Job Ladder?," *Journal of Labor Economics*, 2018, *36.*

Hopenhayn, Hugo A., "Entry, Exit, and Firm Dynamics in Long Run Equilibrium," *Econometrica*, September 1992, *60* (5), 1127.

Hopenhayn, Hugo and Richard Rogerson, "Job Turnover and Policy Evaluation: A General Equilibrium Analysis," *Journal of Political Economy*, October 1993, *101* (5), 915–938.

IBGE, *Síntese de indicadores sociais* number 39. In 'Síntese de Indicadores Sociais.', Rio de Janeiro: Instituto Brasileiro de Geografia e Estatistica (IBGE), 2018.

Jacobson, Louis, Robert Lalonde, and Daniel Sullivan, "Earning Losses of Displaced Workers," *The American Economic Review*, 1993.

Jäger, Simon, Benjamin Schoefer, and Josef Zweimüller, "Marginal Jobs and Job Surplus: A Test of the Efficiency of Separations," *The Review of Economic Studies*, May 2023, *90* (3), 1265–1303.

Jarosch, Gregor, "Searching for Job Security and the Consequences of Job Loss," *Econometrica*, 2023, *91* (3), 903–942.

Jovanovic, Boyan, "Job Matching and the Theory of Turnover," *Journal of Political Economy*, 1979, *87* (5), 972–990.

Makridis, Christos A and Maury Gittleman, "On the Cyclicality of Real Wages and Employment: New Evidence and Stylized Facts from Performance Pay and Fixed Wage Jobs," *The Journal of Law, Economics, and Organization*, October 2022, *38* (3), 889–920.

Messina, Julián, Cláudia Filipa Duarte, and Mario Izquierdo, "The Incidence of Nominal and Real Wage Rigidity: An Individual-Based Sectoral Approach," *Journal of the European Economic Association*, 2010.

Mortensen, D. T. and C. A. Pissarides, "Job Creation and Job Destruction in the Theory of Unemployment," *The Review of Economic Studies*, July 1994, *61* (3), 397–415.

Moscarini, Giuseppe and Fabien Postel-Vinay, "The Contribution of Large and Small Employers to Job Creation in Times of High and Low Unemployment," *American Economic Review*, May 2012, *102* (6), 2509–2539.

Mueller, Andreas I., "Separations, Sorting, and Cyclical Unemployment," *American Economic Review*, July 2017, *107* (7), 2081–2107.

Pinheiro, Roberto and Ludo Visschers, "Unemployment Risk and Wage Differentials," *Journal of Economic Theory*, 2015, *157* (C), 397–424.

Postel-Vinay, Fabien and Jean-Marc Robin, "Equilibrium Wage Dispersion with Worker and Employer Heterogeneity," *Econometrica*, 2002, *70* (6), 2295–2350.

Reizer, Balázs, "Employment and Wage Consequences of Flexible Wage Components," *Labour Economics*, October 2022, *78*, 102256.

Schmieder, Johannes F. and Till von Wachter, "Does Wage Persistence Matter for Employment Fluctuations? Evidence from Displaced Workers," *American Economic Journal: Applied Economics*, July 2010, *2* (3), 1–21.

Sockin, Jason and Michael Sockin, "Variable Pay and Risk Sharing between Firms and Workers," *SSRN Electronic Journal*, 2025.

Song, Jae, David J. Price, Fatih Guvenen, Nicholas Bloom, and Till von Wachter, "Firming Up Inequality," *The Quarterly Journal of Economics*, February 2019, *134* (1), 1–50.

Sorkin, Isaac, "Ranking Firms Using Revealed Preference," *The Quarterly Journal of Economics*, August 2018, *133* (3), 1331–1393.

Topel, Robert and Michael Ward, "Job Mobility and the Careers of Young Men," *The Quarterly Journal of Economics*, 1992, *107.*

Ureta, Manuelita, "The Importance of Lifetime Jobs in the U.S. Economy, Revisited," *The American Economic Review*, 1993, *82.*

We gratefully acknowledge financial support from: Georgetown University; Juan de la Cierva 2022; Grant FJC2022-049259, funded by MICIU/AEI/10.13039/501100011033 and by the European Union Next Generation EU/PRTR; the Maria de Maeztu Unit of Excellence grant CEMFI MDM-2016-0684 and CEMFI CEX2020-001104-M, funded by MICIU/AEI/10.13039/501100011033; PRE2022-101381 grant, funded by MICIU/AEI/10.13039/501100011033 and by ESF +; the graduate scholarship funded by CEMFI; Fundación Carolina and CEMFI.

Appendix

A Model Details

A.1 Definitions

In this appendix, we formalize the model introduced in Section 4. Relative to the baseline framework, we introduce a wage rigidity parameter $\phi \in (0,1]$ that governs the fraction of worker-firm matches subject to downward wage rigidity. The model discussed in Sections 4.1, 4.2, and 4.3 corresponds to $\phi = 1$ (fully rigid wages), but all results extend to any $\phi \in (0,1]$, as shown in this appendix.[39] The model presented here coincides with the extension described in Section 4.4.

Following the Calvo-style friction of Gertler and Trigari (2009), we assume that each match is rigid with probability ϕ, in which case the wage cannot be adjusted from the posted wage w, and flexible with probability $1 - \phi$, in which case the wage may be renegotiated after shocks are realized. When renegotiation is possible, we impose only that the new wage leaves both parties weakly better off than separation. If no such wage exists, the wage remains at w and the baseline layoff and quit rules apply.

Assumption 1 (Pareto-improving renegotiation) *In a flexible match, after the productivity shock η_i and the worker's outside option b_i are observed:*

- *If the renegotiation set $[b_i,\ \psi+\alpha+\eta_i]$ is non-empty (equivalently, $b_i \leq \psi+\alpha+\eta_i$), the wage is reset to some w_i' in this interval and the match continues. The renegotiated wage w_i' may depend flexibly on b_i, ψ, α, and η_i, but not on the pre-posted wage w.*
- *Otherwise, the wage remains at the posted level w and the baseline layoff/quit rules apply: the worker is laid off if $\mu_\psi(w) + \eta_i < 0$, and otherwise quits if $b_i > w$.*

Assumption 1 is deliberately weak: it requires only that renegotiation be Pareto-improving relative to separation, without specifying a particular bargaining protocol. The specific rule for selecting w_i' within the renegotiation set $[b_i, \psi + \alpha + \eta_i]$ does not affect any of our results.

The key analytic implication of Assumption 1 is that the firm's expected profit from flexible matches depends on the primitives $(\psi, \alpha, F_b, F_\eta)$ and the renegotiation rule, but

[39] One result requires an additional assumption when $\phi < 1$, as discussed in Remark 3 of this appendix.

not on the posted wage w. Matches that fall back on the baseline rules yield zero profit regardless of w (the worker either is laid off or quits). Consequently, the posted wage w enters the firm's objective only through the rigid-match term, and the firm's first-order condition, optimal wage w^*_ψ, and average markdown μ^*_ψ are all independent of ϕ.

Throughout, the match-specific productivity shock η_i and the outside option shock b_i are mutually independent.

The timing of the model is:

1. The firm posts a wage w to offer to all workers it meets.
2. Match-specific productivity shocks η_i and outside option shocks b_i are realized.
3. The rigidity shock is realized: with probability ϕ the match is rigid; with probability $1-\phi$ the match is flexible.
4. Separation decisions are made:
   - In rigid matches, the firm lays off the worker if $\mu_\psi(w)+\eta_i<0$; among non-laid-off workers, the worker quits if $b_i>w$.
   - In flexible matches, the wage is renegotiated according to Assumption 1: if $b_i \le \psi+\alpha+\eta_i$, the match continues at some $w'_i \in [b_i,\ \psi+\alpha+\eta_i]$; otherwise, the wage remains at w, and the layoff and quit rules for rigid matches apply.
5. Payoffs are realized.

In rigid matches, the layoff condition $\mu_\psi(w)+\eta_i<0$ depends only on η_i (not on b_i); by mutual independence of η_i and b_i, the layoff rate in rigid matches is therefore simply $F_\eta(-\mu_\psi(w))$, without a retention-composition factor.

Definition 1 *An equilibrium given wage rigidity $\phi \in (0,1]$ is defined by wages w^*_ψ, retention function $\rho(w)$, and layoff function $\delta_\psi(w)$, such that conditions (I)–(III) below hold:*

(I) In rigid matches, workers quit if $b>w$. Among non-laid-off rigid-match workers, the retention function is:

$$\rho(w) = P_b(b \le w) = F_b(w).$$

(II) In rigid matches, the firm lays off the worker if the realized markdown is negative. In flexible matches where renegotiation fails, the worker is laid off under the same rule. The layoff function is:

$$\delta_\psi(w) = \phi \cdot F_\eta\big[-\mu_\psi(w)\big] + (1-\phi)\cdot L_{\mathit{flex}}(\psi, w),$$

where the flexible-match layoff rate is

$$L_{\textit{flex}}(\psi, w) \equiv \int_{-\infty}^{-\mu_\psi(w)} \big(1 - F_b(\psi + \alpha + \eta)\big)\, f_\eta(\eta)\, d\eta.$$

(III) The firm chooses the posted wage to maximize expected profits per meeting:

$$w_\psi^* = \arg\max_w V_\psi(w). \tag{A.1}$$

Where $\mu_\psi(w)$ and $V_\psi(w)$ are defined as follows.
Since $\mathbb{E}_\eta[\eta] = 0$, the average markdown at the posted wage is:

$$\mu_\psi(w) \equiv \psi + \alpha - w.$$

The expected profit per meeting is:

$$V_\psi(w) \equiv \phi \cdot \underbrace{F_b(w) \cdot \big(1 - F_\eta[-\mu_\psi(w)]\big) \cdot \Big\{\mu_\psi(w) + \mathbb{E}_\eta[\eta \mid \mu_\psi(w) + \eta \geq 0]\Big\}}_{V_{\textit{rigid}}(w)\text{: depends on } w} + (1-\phi) \cdot V_{\textit{flex}},$$

where $V_{\textit{flex}} \equiv \mathbb{E}_{\eta,b}\big[(\psi + \alpha + \eta - w'(b, \eta)) \cdot \mathbf{1}\{b \leq \psi + \alpha + \eta\}\big]$ does not depend on the posted wage w.

The expected employment per meeting at posted wage w is:

$$\sigma_\psi(w) \equiv \phi \cdot \big(1 - F_\eta[-\mu_\psi(w)]\big) F_b(w) + (1 - \phi) \cdot \mathbb{E}_\eta\big[F_b(\psi + \alpha + \eta)\big].$$

And $\mu_\psi^* \equiv \mu_\psi(w^*)$, $V_\psi^* \equiv V_\psi(w^*)$, $\rho_\psi^* \equiv \rho(w^*)$, $\delta_\psi^* \equiv \delta_\psi(w^*)$, $\sigma_\psi^* \equiv \sigma_\psi(w^*)$.

V_{flex} is independent of w because $w'(b, \eta)$ depends on the realized shocks only, and matches that fall back on the baseline rules yield zero profit—either the worker is laid off ($\mu_\psi(w) + \eta < 0$, profit zero) or the worker quits ($b > w$, which holds because $b > \psi + \alpha + \eta \geq w$ when $\mu_\psi(w) + \eta \geq 0$). Consequently, only $V_{\text{rigid}}(w)$ enters the first-order condition for w, and the equilibrium wage w_ψ^* and average markdown μ_ψ^* are independent of ϕ.

We now state the main theoretical results.

Theorem 1 *Assume both F_b and $1 - F_\eta$ are log-concave, and that the ranges of H_η and H_b overlap:* $\inf H_b < \sup H_\eta$ *and* $\inf H_\eta < \sup H_b$.[40] *Then, for any $\phi \in (0, 1]$, there is a unique equilibrium and:*

(I) Expected employment per meeting is weakly increasing in firm productivity;
(II) Wages are weakly increasing in firm productivity $\left(\frac{dw_\psi^*}{d\psi} \geq 0\right)$;
(III) The separation rate is weakly decreasing in firm productivity;
(IV) The layoff rate is weakly decreasing in firm productivity $\left(\frac{d\delta_\psi(w_\psi^*)}{d\psi} \leq 0\right)$;

[40]Here $H_\eta(x) \equiv \mathbb{E}_\eta[\eta - x \mid \eta \geq x]$ is the mean residual life and $H_b(w) \equiv F_b(w)/f_b(w)$.

(V) The average markdown is weakly increasing in firm productivity $\left(\frac{d\mu^*_\psi}{d\psi} \geq 0\right)$.

Proof: *Appendix A.2.*

Remark 2 *The range-overlap condition and the implicit non-degeneracy requirement* $\psi + \alpha + \bar{\eta} > \underline{b}$ *are needed* only *for the existence and uniqueness of equilibrium; the comparative statics of Parts (I)–(V) hold under log-concavity alone, as stated in Proposition 1.*

The range-overlap condition rules out pathological distribution pairs in which the mean residual life of η *is bounded below by a value that* H_b *never reaches. This occurs, for example, when* η *has a constant hazard rate (such as the exponential) and* b *has a lower tail thick enough that* $\inf H_b > 0$ *(such as the logistic). For such pairs the first-order condition* $H_\eta(-\mu^*_\psi) = H_b(w^*)$ *has no finite solution and no equilibrium exists.*

For Normal distributions (in either role), $\inf H_b = 0$*,* $\sup H_\eta = \infty$*,* $\inf H_\eta = 0$*, and* $\sup H_b = \infty$*, so the range-overlap condition is satisfied trivially under any parameterization. For Uniform distributions the condition holds whenever* $\psi + \alpha + \bar{\eta} > \underline{b}$*, a parameter restriction satisfied by all economically relevant calibrations.*

Theorem 1 characterizes equilibrium when productivity ψ varies across firms. The following theorem considers the complementary question of how the equilibrium changes when wage-setting power varies across firms.

Definition 2 (Interior equilibrium) *An equilibrium is* interior *if* $F_\eta(-\mu^*_\psi) > 0$*, i.e., layoffs occur with strictly positive probability. At any interior equilibrium, the first-order condition for* w^*_ψ *(derived in the proof of Theorem 1) gives* $H_\eta(-\mu^*_\psi) = H_b(w^*_\psi) > 0$*—since* $w^*_\psi > \underline{b}$ *at an interior wage optimum—which is equivalent to* $F_\eta(-\mu^*_\psi) < 1$*. Hence at an interior equilibrium both layoffs and retention occur with strictly positive probability.*

Theorem 2 *Under the assumptions of Theorem 1:*

(I) (Productivity channel.) Holding the outside option distribution fixed, as ψ *increases,* w^* *and* μ^* *move weakly in the same direction and* w^* *and* δ^* *move weakly in opposite directions along the equilibrium path. Formally, whenever* $\frac{dw^*}{d\psi} > 0$*:* $\frac{d\mu^*}{dw^*} \geq 0$ *and* $\frac{d\delta^*}{dw^*} \leq 0$*. When* $\frac{dw^*}{d\psi} = 0$*, wages are flat in* ψ *and there is no cross-sectional wage variation along the equilibrium path.*

(II) (Wage-setting power channel.) Holding ψ *fixed, if equilibrium wages differ across firms: further assume an interior equilibrium (*$F_\eta(-\mu^*_\psi) > 0$*) and that cross-firm variation in* F_b *is ordered by first-order stochastic dominance—firms paying higher wages face workers with first-order stochastically higher outside options. Then the*

firm paying higher wages has strictly lower markdowns $\left(\frac{d\mu^*}{dw^*} = -1 < 0\right)$ *and strictly higher layoff rates* $\left(\frac{d\delta^*}{dw^*} > 0\right)$.

Proof: *Appendix A.2.*

Remark 3 *The first-order stochastic dominance assumption (FOSD) in Part (II) of Theorem 2 is required only when* $\phi < 1$. *When* $\phi = 1$ *(fully rigid wages), the layoff function reduces to* $\delta_\psi(w) = F_\eta(-\mu_\psi(w))$, *which does not depend on* F_b. *The comparison* $\delta^*_B > \delta^*_A$ *therefore follows directly from* $w^*_B > w^*_A$*—which implies* $\mu^*_B < \mu^*_A$ *and hence* $F_\eta(-\mu^*_B) > F_\eta(-\mu^*_A)$, *given the interior equilibrium condition* $F_\eta(-\mu^*_\psi) > 0$*—without any restriction on how* F_b *varies across firms. Proposition 2 in the main text, which implicitly corresponds to the* $\phi = 1$ *case of Theorem 2, is accordingly valid without the FOSD assumption.*

The final theorem establishes the role of the wage rigidity parameter ϕ. It uses the *savable-layoff rate*, the per-meeting probability that a match is laid off under a rigid wage but would be retained under a flexible wage:

$$L_{\text{savable}}(\psi) \equiv \int_{-\infty}^{-\mu^*_\psi} F_b(\psi + \alpha + \eta)\, f_\eta(\eta)\, d\eta.$$

This is the probability of the event $\{\eta < -\mu^*_\psi\} \cap \{b \leq \psi + \alpha + \eta\}$: negative markdown combined with positive surplus.

Theorem 3 (Effect of wage rigidity) *Under the assumptions of Theorem 1, assuming an interior equilibrium (Definition 2):*

(I) (Layoff level.) The equilibrium layoff rate is strictly increasing in wage rigidity. The marginal effect equals the savable-layoff rate:

$$\frac{\partial \delta^*_\psi}{\partial \phi} = L_{savable}(\psi) > 0.$$

(II) (Effect on the productivity–layoff gradient.) The cross-derivative of the layoff rate with respect to productivity and rigidity equals the derivative of the savable-layoff rate:

$$\frac{\partial^2 \delta^*_\psi}{\partial \psi\, \partial \phi} = \frac{dL_{savable}(\psi)}{d\psi}.$$

This admits the decomposition

$$\frac{\partial^2 \delta^*_\psi}{\partial \psi\, \partial \phi} = \underbrace{-F_b(w^*_\psi) \cdot f_\eta[-\mu^*_\psi] \cdot \frac{d\mu^*_\psi}{d\psi}}_{\textit{markdown-buffer effect}(\leq 0)} + \underbrace{\int_{-\infty}^{-\mu^*_\psi} f_b(\psi + \alpha + \eta)\, f_\eta(\eta)\, d\eta}_{\textit{surplus-composition effect}(> 0)}. \tag{A.2}$$

Equivalently, factoring out $F_\eta(-\mu^*_\psi) \cdot \rho^*_\psi \cdot r_\eta(-\mu^*_\psi)$, where $r_\eta(x) \equiv f_\eta(x)/F_\eta(x)$:

$$\frac{\partial^2 \delta^*_\psi}{\partial\psi\,\partial\phi} = F_\eta(-\mu^*_\psi) \cdot \rho^*_\psi \cdot r_\eta(-\mu^*_\psi) \cdot \left[\underbrace{-\frac{d\mu^*_\psi}{d\psi}}_{\textit{markdown-buffer}} + \underbrace{\frac{\mathbb{E}_\eta\big[f_b(\psi+\alpha+\eta) \mid \eta < -\mu^*_\psi\big]}{\rho^*_\psi \cdot r_\eta(-\mu^*_\psi)}}_{\textit{surplus-composition}} \right]. \tag{A.3}$$

The sign depends on which force dominates:

(II.i) If the markdown-buffer effect dominates, $\frac{\partial^2 \delta^*_\psi}{\partial\psi\,\partial\phi} < 0$: wage rigidity amplifies *the productivity–layoff gradient.*

(II.ii) If the surplus-composition effect dominates, $\frac{\partial^2 \delta^*_\psi}{\partial\psi\,\partial\phi} > 0$: wage rigidity weakens *the gradient.*

Proof: *Appendix A.2.*

A.2 Proofs

Existence, Uniqueness, and FOC Independence Lemma.

Under the assumptions of Theorem 1 (log-concavity and range overlap), for each ψ there is a unique equilibrium. Moreover, the equilibrium wage w^*_ψ and average markdown μ^*_ψ are independent of ϕ.

Proof. From Definition 1, the firm's objective is $V_\psi(w) = \phi \cdot V_{\text{rigid}}(w) + (1-\phi) \cdot V_{\text{flex}}$, where $V_{\text{rigid}}(w) = F_b(w)(1 - F_\eta[-\mu_\psi(w)])\{\mu_\psi(w) + \mathbb{E}_\eta[\eta \mid \eta \geq -\mu_\psi(w)]\}$ and $V_{\text{flex}} = \mathbb{E}_{\eta,b}[(\psi+\alpha+\eta - w'(b,\eta)) \cdot \mathbf{1}\{b \leq \psi+\alpha+\eta\}]$. Under Assumption 1, $w'(b,\eta)$ depends on b and η only; therefore V_{flex} does not depend on w. Since $\phi > 0$, the FOC $\partial V_\psi/\partial w = 0$ reduces to $\partial V_{\text{rigid}}/\partial w = 0$, yielding:

$$H_\eta(-\mu^*_\psi) = H_b(w^*), \tag{A.4}$$

where $H_\eta(x) \equiv \mathbb{E}_\eta[\eta \mid \eta \geq x] - x$ and $H_b(x) \equiv 1/\frac{\partial \ln F_b(x)}{\partial x} = F_b(x)/f_b(x)$. Since $\mu_\psi(w) = \psi + \alpha - w$ does not depend on ϕ, equation (A.4) is also independent of ϕ.

Existence and uniqueness. Write (A.4) as $L(w) = R(w)$, where $L(w) \equiv H_\eta(w - \psi - \alpha)$ and $R(w) \equiv H_b(w)$.

L is decreasing. Log-concavity of $1 - F_\eta$ implies that F_η has an increasing hazard rate. A classical result in reliability theory states that an increasing hazard rate implies a decreasing mean residual life, i.e. $H_\eta(x) = \mathbb{E}[\eta - x \mid \eta \geq x]$ is decreasing in x. Since $w - \psi - \alpha$ is strictly increasing in w, $L(w)$ is decreasing.

R is increasing. Log-concavity of F_b implies f_b/F_b is decreasing, so $R(w) = F_b(w)/f_b(w)$ is increasing.

Existence. We apply the intermediate value theorem. Since L is decreasing and R is increasing, $L(w) - R(w)$ is decreasing. The range-overlap condition $\inf H_b < \sup H_\eta$ guarantees that $L(w) > R(w)$ at the lower boundary of the relevant domain (or in the limit as $w \to \underline{b}$), and $\inf H_\eta < \sup H_b$ guarantees that $L(w) < R(w)$ at the upper boundary (or in the limit as $w \to \psi+\alpha+\overline{\eta}$). By continuity and the intermediate value theorem, a crossing exists.

Uniqueness and global maximum. Since L is decreasing and R increasing, they cross at most once, giving uniqueness. To confirm the crossing is a global maximum, differentiate $V_{\text{rigid}}(w) = F_b(w)(1 - F_\eta(-\mu_\psi(w)))H_\eta(-\mu_\psi(w))$ and use $H'_\eta(x) = r_\eta(x)H_\eta(x) - 1$ (where r_η is the hazard rate of F_η) to show that the hazard-rate terms cancel in the derivative, yielding $\text{sign}(\partial V_{\text{rigid}}/\partial w) = \text{sign}(1/R(w) - 1/L(w))$. This quantity is decreasing in w: positive for $w < w^*$ and negative for $w > w^*$. Hence w^*_ψ is the unique global maximizer, and the pair (w^*_ψ, μ^*_ψ) is uniquely determined by (A.4) independently of ϕ. □

Theorem 1

We prove the parts in order of logical dependence: (II) and (V) first (from the FOC), then (IV), (I), (III).

(II) Wages are increasing in firm productivity: Replacing $\mu^*_\psi = \psi+\alpha-w$ in Equation (A.4), taking the total derivative with respect to ψ, and isolating $\frac{\mathrm{d}w^*}{\mathrm{d}\psi}$:

$$\frac{\mathrm{d}w^*}{\mathrm{d}\psi} = \frac{H'_\eta(-\mu^*_\psi)}{H'_\eta(-\mu^*_\psi) - H'_b(w^*_\psi)}. \tag{A.5}$$

Log-concavity of F_b implies $H'_b(w^*_\psi) \geq 0$.[41] Log-concavity of $1 - F_\eta$ implies $H'_\eta(-\mu^*_\psi) \leq 0$.[42] Therefore $\frac{\mathrm{d}w^*}{\mathrm{d}\psi} \geq 0$.[43]

(V) Average markdown is increasing in firm productivity: Replacing $w = \psi + \alpha - \mu^*$ in

[41]Since F_b is log concave, $\partial \ln F_b(w)/\partial w$ is decreasing, hence H_b is increasing: $H'_b \geq 0$.

[42]Log-concavity of $1 - F_\eta$ implies an increasing hazard rate, which gives $\partial \mathbb{E}_\eta[\eta|\eta \geq x]/\partial x \leq 1$, so $H'_\eta \leq 0$. For distributions with strictly increasing hazard rate, the inequality is strict.

[43]If both $H'_\eta(-\mu^*_\psi) = 0$ and $H'_b(w^*_\psi) = 0$, the first-order condition holds for all w and the profit function is flat; one may then select $w^*(\psi)$ to be weakly increasing in ψ.

Equation (A.4) and differentiating with respect to ψ:

$$\frac{\mathrm{d}\mu^*}{\mathrm{d}\psi} = \frac{H_b'(w_\psi^*)}{H_b'(w_\psi^*) - H_\eta'(-\mu_\psi^*)}. \tag{A.6}$$

Since $H_b'(w_\psi^*) \geq 0$ and $H_\eta'(-\mu_\psi^*) \leq 0$, we have $\frac{\mathrm{d}\mu^*}{\mathrm{d}\psi} \geq 0$.

(IV) Layoff rate is decreasing in firm productivity: From Definition 1, $\delta_\psi^* = \phi \cdot F_\eta(-\mu_\psi^*) + (1-\phi) \cdot L_{\text{flex}}(\psi, w^*(\psi))$. Differentiating:

$$\frac{\mathrm{d}\delta_\psi^*}{\mathrm{d}\psi} = \phi \cdot \big(- f_\eta(-\mu^*) \cdot \tfrac{\mathrm{d}\mu^*}{\mathrm{d}\psi} \big) + (1-\phi) \cdot \frac{\mathrm{d}L_{\text{flex}}}{\mathrm{d}\psi}.$$

The first term is non-positive. For the second, applying Leibniz's rule (using $\psi + \alpha - \mu^*(\psi) = w^*(\psi)$):

$$\frac{\mathrm{d}L_{\text{flex}}}{\mathrm{d}\psi} = -\big(1 - F_b(w^*)\big) \cdot f_\eta(-\mu^*) \cdot \frac{\mathrm{d}\mu^*}{\mathrm{d}\psi} - \int_{-\infty}^{-\mu^*} f_b(\psi + \alpha + \eta)\, f_\eta(\eta)\, d\eta \leq 0. \tag{A.7}$$

Both terms are non-positive (since $\mathrm{d}\mu^*/\mathrm{d}\psi \geq 0$ and $f_b, f_\eta \geq 0$). Hence $\mathrm{d}\delta_\psi^*/\mathrm{d}\psi \leq 0$.

(I) Expected employment per meeting is increasing in firm productivity: A meeting results in employment if the worker is neither laid off nor quit. In a rigid match, this requires $\eta \geq -\mu^*$ and $b \leq w^*$. In a flexible match, this requires renegotiation to succeed ($b \leq \psi + \alpha + \eta$); if renegotiation fails, the worker separates—either laid off (if $\mu + \eta < 0$) or quit (if $\mu + \eta \geq 0$, which gives $\psi + \alpha + \eta \geq w^*$ and hence $b > \psi + \alpha + \eta$ implies $b > w^*$). Therefore:

$$\sigma_\psi(w^*) = \phi \cdot \big(1 - F_\eta[-\mu_\psi^*]\big) F_b(w^*) + (1-\phi) \cdot \mathbb{E}_\eta\big[F_b(\psi + \alpha + \eta)\big].$$

As ψ increases: by Parts (II) and (V), both w^* and μ^* rise, so both $1 - F_\eta(-\mu^*)$ and $F_b(w^*)$ rise; and $\mathbb{E}_\eta[F_b(\psi + \alpha + \eta)]$ rises because F_b is increasing. So $\sigma_\psi(w^*)$ is increasing in ψ.

(III) The separation rate is decreasing in firm productivity: The separation rate is $1 - \sigma_\psi(w^*)$. Since $\sigma_\psi(w^*)$ is increasing in ψ (Part I), the separation rate is decreasing. □

Theorem 2

Part (I) — Productivity channel: By Theorem 1 (Parts II and V), both w_ψ^* and μ_ψ^* are weakly increasing in ψ. We consider two cases.

Case (a): $dw^*/d\psi > 0$. This holds when $H'_\eta(-\mu^*_\psi) < 0$ (strictly increasing hazard rate of $1-F_\eta$). Then $\mathrm{d}\mu^*/\mathrm{d}w^* = (\mathrm{d}\mu^*/\mathrm{d}\psi)/(\mathrm{d}w^*/\mathrm{d}\psi) \geq 0$ and $\mathrm{d}\delta^*/\mathrm{d}w^* = (\mathrm{d}\delta^*/\mathrm{d}\psi)/(\mathrm{d}w^*/\mathrm{d}\psi) \leq 0$.

Case (b): $dw^*/d\psi = 0$. From (A.5), this occurs when $H'_\eta(-\mu^*_\psi) = 0$ (constant hazard rate of $1 - F_\eta$). In this case w^* is flat in ψ while $\mathrm{d}\mu^*/\mathrm{d}\psi = 1 > 0$ (from (A.6)). There is therefore no cross-sectional wage variation along the equilibrium path, the antecedent $\mathrm{d}w^*/\mathrm{d}\psi > 0$ is false, and the conditional claims $\mathrm{d}\mu^*/\mathrm{d}w^* \geq 0$ and $\mathrm{d}\delta^*/\mathrm{d}w^* \leq 0$ are not invoked. $\square$

Part (II) — Wage-setting power channel: Fix ψ and α. Since $\mu^* = \psi + \alpha - w^*$, any two firms A and B with the same ψ satisfy $\mu^*_B - \mu^*_A = -(w^*_B - w^*_A)$, so $\mathrm{d}\mu^*/\mathrm{d}w^* = -1 < 0$. For the layoff rate, define $g_j(\eta) \equiv 1 - (1-\phi)F_{b,j}(\psi + \alpha + \eta)$, so $g_j \geq \phi > 0$ and $\delta^*_j = \int_{-\infty}^{-\mu^*_j} g_j(\eta)\, f_\eta(\eta)\, d\eta$. For $w^*_B > w^*_A$ (so $-\mu^*_B > -\mu^*_A$), decompose:

$$\delta^*_B - \delta^*_A = \underbrace{\int_{-\infty}^{-\mu^*_A} \big[g_B(\eta) - g_A(\eta)\big]\, f_\eta(\eta)\, d\eta}_{\geq 0} + \underbrace{\int_{-\mu^*_A}^{-\mu^*_B} g_B(\eta)\, f_\eta(\eta)\, d\eta}_{> 0}.$$

First integral (≥ 0): $g_B(\eta) - g_A(\eta) = (1-\phi)\big[F_{b,A}(\psi + \alpha + \eta) - F_{b,B}(\psi + \alpha + \eta)\big] \geq 0$ pointwise, by the FOSD assumption ($F_{b,B} \leq F_{b,A}$ everywhere).

Second integral (> 0): $g_B \geq \phi > 0$ uniformly. The mass $\int_{-\mu^*_A}^{-\mu^*_B} f_\eta\, d\eta = F_\eta(-\mu^*_B) - F_\eta(-\mu^*_A) > 0$ because $-\mu^*_B > -\mu^*_A$, $F_\eta(-\mu^*_B) > 0$ (interior equilibrium), and F_η is atomless, so it is strictly increasing at any point with positive density.

Hence $\delta^*_B > \delta^*_A$. $\square$

Theorem 3

Part (I) — Layoff level: From Definition 1, $\delta^*_\psi = \phi \cdot F_\eta(-\mu^*_\psi) + (1-\phi) \cdot L_{\text{flex}}(\psi, w^*_\psi)$. By the Existence, Uniqueness, and FOC Independence Lemma, w^*_ψ and μ^*_ψ are independent of ϕ, so L_{flex} is also independent of ϕ. Therefore:

$$\begin{aligned}\frac{\partial \delta^*_\psi}{\partial \phi} &= F_\eta(-\mu^*_\psi) - L_{\text{flex}}(\psi, w^*_\psi) \\ &= \int_{-\infty}^{-\mu^*} f_\eta(\eta)\, d\eta - \int_{-\infty}^{-\mu^*} (1 - F_b(\psi + \alpha + \eta))\, f_\eta(\eta)\, d\eta \\ &= \int_{-\infty}^{-\mu^*} F_b(\psi + \alpha + \eta)\, f_\eta(\eta)\, d\eta \;=\; L_{\text{savable}}(\psi) > 0,\end{aligned}$$

where strict positivity holds because: (i) $F_\eta(-\mu^*_\psi) > 0$ (interior equilibrium), so f_η places positive mass on $(-\infty, -\mu^*_\psi)$; and (ii) as $\eta \nearrow -\mu^*_\psi$, the factor $F_b(\psi+\alpha+\eta) \to F_b(w^*_\psi) = \rho^*_\psi$; by the FOC (A.4), $H_b(w^*_\psi) = H_\eta(-\mu^*_\psi) > 0$ (since $F_\eta(-\mu^*_\psi) < 1$ at any interior equilibrium, by Definition 2), so $\rho^*_\psi = H_b(w^*_\psi) \cdot f_b(w^*_\psi) > 0$ (as $f_b(w^*_\psi) > 0$ by log-concavity of F_b); by continuity the integrand is strictly positive on a neighborhood of $-\mu^*_\psi$. Together these give a set of positive f_η-measure on which the integrand is strictly positive. □

Part (II) — Effect on the productivity–layoff gradient: From Part (I), $\partial\delta^*/\partial\phi = L_{\text{savable}}(\psi)$. Since L_{savable} depends on ψ but not on ϕ:

$$\frac{\partial^2 \delta^*_\psi}{\partial\psi\,\partial\phi} = \frac{\mathrm{d}L_{\text{savable}}(\psi)}{\mathrm{d}\psi}.$$

Applying Leibniz's rule to $L_{\text{savable}}(\psi) = \int_{-\infty}^{-\mu^*(\psi)} F_b(\psi+\alpha+\eta) f_\eta(\eta)\, d\eta$ (using $\psi+\alpha-\mu^*(\psi) = w^*(\psi)$) gives Equation (A.2):

$$\frac{\mathrm{d}L_{\text{savable}}}{\mathrm{d}\psi} = \underbrace{-F_b(w^*) \cdot f_\eta(-\mu^*) \cdot \frac{\mathrm{d}\mu^*}{\mathrm{d}\psi}}_{\text{markdown-buffer } (\le 0)} + \underbrace{\int_{-\infty}^{-\mu^*} f_b(\psi + \alpha + \eta)\, f_\eta(\eta)\, d\eta}_{\text{surplus-composition } (>0)}.$$

The markdown-buffer term is non-positive ($\mathrm{d}\mu^*/\mathrm{d}\psi \geq 0$ by Theorem 1 Part V). The surplus-composition term is positive ($f_b, f_\eta > 0$).

The factored form in Equation (A.3) follows by writing $F_b(w^*) = \rho^*_\psi$ and $f_\eta(-\mu^*) = r_\eta(-\mu^*_\psi) \cdot F_\eta(-\mu^*_\psi)$:

$$\begin{aligned}\text{markdown-buffer} &= -\rho^*_\psi \cdot r_\eta(-\mu^*_\psi) \cdot F_\eta(-\mu^*_\psi) \cdot \frac{\mathrm{d}\mu^*}{\mathrm{d}\psi}, \\ \text{surplus-composition} &= F_\eta(-\mu^*_\psi) \cdot \mathbb{E}_\eta\big[f_b(\psi + \alpha + \eta) \mid \eta < -\mu^*_\psi\big].\end{aligned}$$

Factoring out $F_\eta(-\mu_\psi^*) \cdot \rho_\psi^* \cdot r_\eta(-\mu_\psi^*)$ yields Equation (A.3). The sign depends on which force dominates. $\square$

Derivation of Equations (7) and (8)

Step 1: Firm problem in retention–markdown space (Equation (7)).

From Definition 1, the firm's expected profit per meeting is

$$V_\psi(w) = \phi \cdot V_{\text{rigid}}(w) + (1-\phi) \cdot V_{\text{flex}},$$

where V_{flex} does not depend on w (Assumption 1). Since $\phi > 0$, maximizing $V_\psi(w)$ over w is equivalent to maximizing $V_{\text{rigid}}(w)$. Expanding from Definition 1:

$$V_{\text{rigid}}(w) = F_b(w) \cdot \big(1 - F_\eta\big[-\mu_\psi(w)\big]\big) \cdot \Big\{\mu_\psi(w) + \mathbb{E}_\eta\big[\eta \mid \eta \geq -\mu_\psi(w)\big]\Big\}.$$

Define retention $\rho \equiv F_b(w)$, average markdown $\mu \equiv \mu_\psi(w) = \psi + \alpha - w$, layoff rate $\delta(\mu) \equiv F_\eta(-\mu)$, and per-worker profit $\pi(\mu) \equiv \mu + \mathbb{E}_\eta[\eta \mid \eta \geq -\mu]$. Then:

$$V_{\text{rigid}}(w) = \rho \cdot \big[1 - \delta(\mu)\big] \cdot \pi(\mu).$$

A pair (ρ, μ) is feasible if and only if there exists a wage w such that $\rho = F_b(w)$ and $\mu = \psi + \alpha - w$ simultaneously; this is the production possibility frontier (PPF). Since F_b is strictly increasing and continuous, the map $w \mapsto (\rho, \mu)$ is one-to-one (w is uniquely pinned down by either coordinate), so maximizing $V_\psi(w)$ over w is equivalent to:

$$\max_{\rho,\mu} \; \rho \cdot \big[1 - \delta(\mu)\big] \cdot \pi(\mu) \quad \text{subject to PPF},$$

which is Equation (7).

Step 2: Functional-form specialization (Equation (8)).

Assume $\eta_i \sim U[-\sigma_\eta, \sigma_\eta]$ and $b_i \sim U[0, \sigma_b]$, with $\sigma_\eta = 1$ (normalization without loss of generality).

PPF. Since $b_i \sim U[0, \sigma_b]$, we have $F_b(w) = w/\sigma_b$ for $w \in [0, \sigma_b]$, so $\rho = w/\sigma_b$, i.e.

$w = \sigma_b \rho$. Substituting into $\mu = \psi + \alpha - w$ gives the linear PPF constraint:

$$\mu + \sigma_b\, \rho = \psi + \alpha.$$

Survival rate. Since $\eta_i \sim U[-1, 1]$, for $\mu \in (-1, 1)$:

$$\delta(\mu) = F_\eta(-\mu) = \frac{1-\mu}{2}, \qquad \text{so} \qquad 1 - \delta(\mu) = \frac{1+\mu}{2}.$$

Per-worker profit. For $\eta_i \sim U[-1, 1]$, the conditional mean $\mathbb{E}_\eta[\eta \mid \eta \geq -\mu] = (1-\mu)/2$, so:

$$\pi(\mu) = \mu + \frac{1-\mu}{2} = \frac{1+\mu}{2}.$$

Objective. Substituting into Equation (7):

$$\rho \cdot \big[1 - \delta(\mu)\big] \cdot \pi(\mu) = \rho \cdot \frac{1+\mu}{2} \cdot \frac{1+\mu}{2} = \frac{\rho\,(\mu+1)^2}{4}.$$

Since $1/4 > 0$ is a constant, maximizing this is equivalent to maximizing $\rho(\mu + 1)^2$. The firm's problem therefore reduces to:

$$\max_{\rho,\mu}\ \rho\,(\mu+1)^2 \quad \text{subject to: } \mu + \sigma_b\, \rho = \psi + \alpha,$$

which is Equation (8). $\square$

Supplemental Appendix for:

Separations Revisited: Do Layoffs or Quits Drive Lower Separation Rates in High-Productivity Firms?

Cauê Dobbin   Daniel Fernandez   Tom Zohar

B Supplemental Figures and Tables

Table B.I
Descriptive Statistics: All Sectors

	Brazil	Southeast region	Sample 2010–2017	Sample 2011–2016
Number of firms	4,348,254	2,119,052	684,762	506,419
Average firm size	8.0	8.6	23.4	25.4
Number of workers	36,341,616	19,448,425	17,293,138	15,095,581
Average age (years)	37.3	37.6	37.5	37.6
Average log-hourly wage	2.238	2.341	2.373	2.397
Median log-hourly wage	2.067	2.149	2.201	2.225
Average tenure (months)	46.3	47.8	48.5	50.3
Average schooling (years)	10.3	10.5	10.5	10.5
Share non-white (%)	44.73	37.82	39.02	38.97
Share female (%)	40.94	42.36	41.58	41.78
Average annual layoff rate (%)	20.45	19.66	17.08	18.46
Average annual quit rate (%)	3.58	3.73	3.93	4.11
Share of firms in manufacturing (%)	6.76	8.36	12.58	13.68
Share of workers in manufacturing (%)	13.68	17.08	18.65	18.99

Notes. Summary statistics from RAIS across all sectors. Columns progressively restrict the sample: "Brazil" applies demographic, contract, and job-selection restrictions nationwide; "Southeast region" further restricts to the states of Minas Gerais, Espírito Santo, Rio de Janeiro, and São Paulo; "Sample" incorporates the additional restrictions described in Section 1.

Table B.II
Productivity–Layoff Gradient Using Log Revenue per Worker (ORBIS) as Proxy for Productivity

	Dependent variables: Annual layoff and separation rates			
	(1)	(2)	(3)	(4)
β^{Layoff}	-0.017***	-0.011***	-0.010***	-0.009***
	(0.0020)	(0.0019)	(0.0014)	(0.0014)
$\beta^{\text{Separation}}$	-0.018***	-0.012***	-0.012***	-0.011***
	(0.0022)	(0.0020)	(0.0014)	(0.0014)
$\frac{\beta^{\text{Layoff}}}{\beta^{\text{Separation}}}$	0.946***	0.936**	0.877***	0.856***
	(0.2206)	(0.3286)	(0.2239)	(0.2496)
Observations	6,253,327	6,253,327	6,252,603	6,252,603
Worker covariates		✓		✓
Worker AKM Effect			✓	✓

Notes. Replicates Table II using log revenue per worker from ORBIS as the productivity measure. Standard errors clustered at the firm level, since the ORBIS data are observed at the firm level. Worker controls and specification as in Table II. Sample as described in Section 1, further restricted to firms in ORBIS. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B.III
Productivity–Layoff Gradient Using Log Firm Size as Proxy for Productivity

	Dependent variables: Annual layoff and separation rates			
	(1)	(2)	(3)	(4)
β^{Layoff}	-0.013***	-0.009***	-0.008***	-0.006***
	(0.0018)	(0.0017)	(0.0011)	(0.0010)
$\beta^{\text{Separation}}$	-0.013***	-0.009***	-0.007***	-0.006***
	(0.0018)	(0.0016)	(0.0011)	(0.0011)
$\frac{\beta^{\text{Layoff}}}{\beta^{\text{Separation}}}$	1.048***	1.093**	1.127***	1.158**
	(0.2967)	(0.3904)	(0.3276)	(0.4120)
Observations	48,430,694	48,430,694	48,424,427	48,424,427
Worker covariates		✓		✓
Worker AKM Effect			✓	✓

Notes. Replicates Table II using log firm size as the productivity measure. Firm size is total employment in the firm's first sample year. Standard errors are clustered at the firm level, since firm size is measured at the firm level. Worker controls and specification as in Table II. Sample as in Section 1. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B.IV
Productivity–Layoff Gradient Controlling for Log Firm Size

	Dependent variables: Annual layoff and separation rates			
	(1)	(2)	(3)	(4)
β^{Layoff}	-0.028***	-0.012***	-0.017***	-0.013***
	(0.0039)	(0.0035)	(0.0023)	(0.0022)
$\beta^{\text{Separation}}$	-0.030***	-0.015***	-0.019***	-0.015***
	(0.0038)	(0.0035)	(0.0023)	(0.0023)
$\frac{\beta^{\text{Layoff}}}{\beta^{\text{Separation}}}$	0.934***	0.815	0.886***	0.827**
	(0.2502)	(0.4227)	(0.2297)	(0.2668)
Observations	9,196,989	9,196,989	9,196,346	9,196,346
Worker covariates		✓		✓
Worker AKM Effect			✓	✓

Notes. Replicates Table II adding log firm size (total employment in first sample year) as a control. All other specifications as in Table II. Sample as in Section 1. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B.V
Productivity–Layoff Gradient: Firms with at Least 20 Workers

Dependent variables: Annual layoff and separation rates				
	(1)	(2)	(3)	(4)
β^{Layoff}	-0.035***	-0.017***	-0.020***	-0.015***
	(0.0037)	(0.0034)	(0.0026)	(0.0025)
$\beta^{\text{Separation}}$	-0.038***	-0.021***	-0.023***	-0.019***
	(0.0039)	(0.0036)	(0.0026)	(0.0026)
$\frac{\beta^{\text{Layoff}}}{\beta^{\text{Separation}}}$	0.920***	0.791**	0.868***	0.801**
	(0.1898)	(0.2990)	(0.2134)	(0.2496)
Observations	6,468,165	6,468,165	6,467,757	6,467,757
Worker covariates		✓		✓
Worker AKM Effect			✓	✓

Notes. Replicates Table II restricting to firms with ≥20 workers in every sample year. All other specifications as in Table II. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B.VI
Markdown Results Are Robust to Alternative Markdown Definitions

	(1) M1	(2) M2	(3) M3	(4) M4
Panel A: Markdown $\sim$ Value Added				
Value Added	1.2472***	0.3284***	0.1316***	1.0832***
	(0.0297)	(0.0290)	(0.0111)	(0.1935)
Observations	277,210	277,800	277,800	277,800
Panel B: Layoff Rate $\sim$ Markdown				
Markdown	-0.0330***	-0.0605***	-0.1055***	-0.0116***
	(0.0037)	(0.0096)	(0.0311)	(0.0036)
Observations	277,210	277,800	277,800	277,800
Panel C: Markdown $\sim$ Firm Pay Premium				
Firm Pay Premium	2.4656***	0.3286***	0.1444***	0.8731***
	(0.3331)	(0.1098)	(0.0513)	(0.2935)
Observations	277,210	277,800	277,800	277,800

Notes. OLS estimates across four markdown definitions: M1 = $\log((\text{VA}-\text{LC})/N)$; M2 = $\log(\text{VA}/N) - \log(\text{LC}/N)$; M3 = $(\text{VA} - \text{LC})/\text{VA}$; M4 = $(\text{VA} - \text{LC})/\text{LC}$. LC denotes total labor cost, VA denotes value added, and N denotes the number of workers. All definitions fixed to each firm's first sample year. Panel A: markdown on VA; Panel B: layoff rate on markdown; Panel C: markdown on AKM pay premium. Firm-year level, weighted by firm size; year fixed effects; standard errors clustered at the 3-digit industry-state level. Sample as in Section 1. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B.VII

Wage Rigidity Amplifies the Productivity–Layoff Gradient: Alternative Rigidity Measures

	Average Layoff Rate			β_m		
	(1)	(2)	(3)	(4)	(5)	(6)
$\widehat{\phi}_m^{80}$	0.186***			-0.093***		
	(0.059)			(0.028)		
$\widehat{\phi}_m^{90}$		0.157***			-0.077*	
		(0.043)			(0.039)	
Ave. Contract Share			0.568***			-0.256**
			(0.160)			(0.103)
Worker covariates				✓	✓	✓
Time-varying shocks				✓	✓	✓
Markets	75	75	75	72	72	72

Notes. Replicates Table V using alternative measures of wage rigidity. Columns (1)–(3): market-level layoff rate on alternative rigidity measures, weighted by market employment. Columns (4)–(6): $\widehat{\beta}_m$ on alternative rigidity measures, weighted by inverse squared SE of $\widehat{\beta}_m$. $\widehat{\phi}_m^{80}$ and $\widehat{\phi}_m^{90}$ denote the leave-one-out share of jobs with contract wages $\geq$ 80% and $\geq$ 90% of total pay, respectively. "Ave. Contract Share" is the leave-one-out average contract share in market m. Checkmarks refer to controls in the estimation of $\widehat{\beta}_m$, not in the market-level regression. Robust standard errors. Sample as in Section 1. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

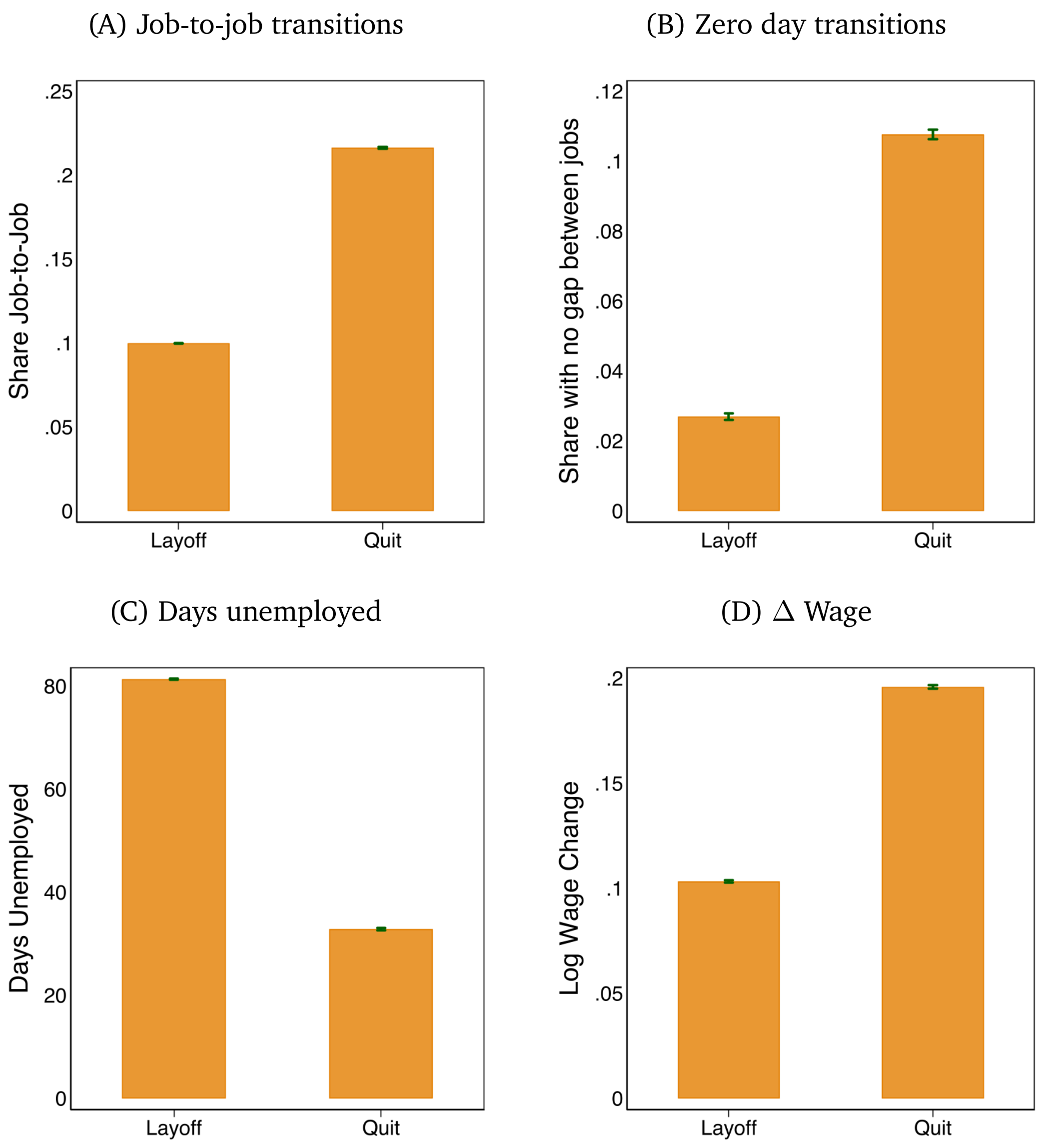

Figure B.I
Quitting Workers Make Better Moves than Laid-Off Ones

Notes. Comparison of post-separation outcomes for laid-off and quitting workers: (A) the share of separated workers who make a job-to-job transition, defined as finding a new job in the same calendar year as the separation, which allows for a period of nonemployment between jobs; (B) the share whose next job begins with no gap between the separation date from the previous job and the hiring date of the new job; (C) the average number of days between jobs, conditional on making a job-to-job transition; and (D) differential wage growth for workers making a job-to-job transition. The sample includes only jobs satisfying the restrictions described in Section 1.

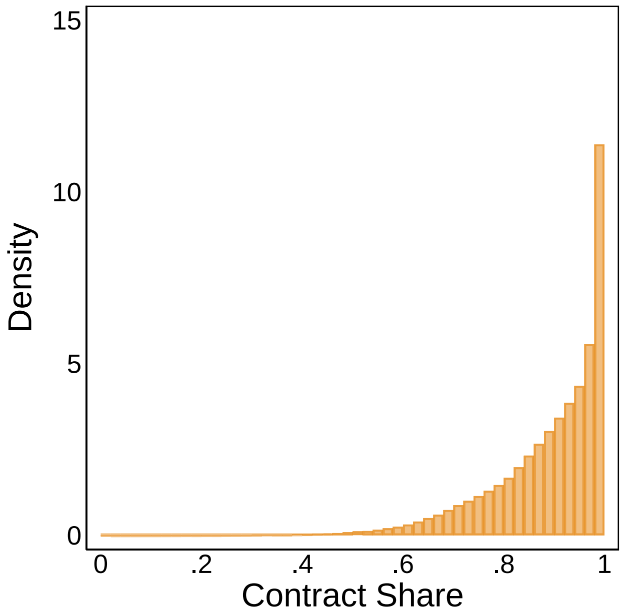

Figure B.II
Distribution of Wage Rigidity Proxy (ContractShare) across Firms

Notes. Distribution of ContractShare, Equation (4), across firms, fixed to each firm's first sample year. Sample as in Section 1.

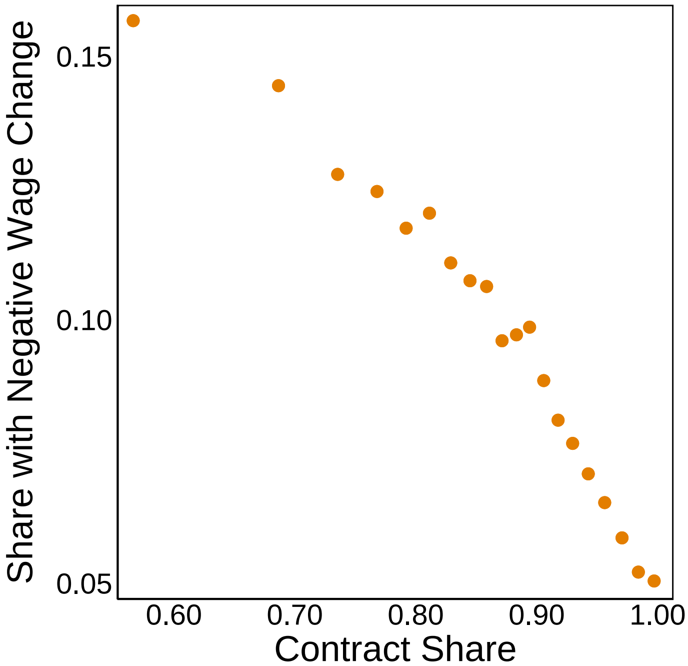

Figure B.III
Higher ContractShare Is Associated with Fewer Wage Cuts

Notes. Binscatter with 20 equal-sized bins of ContractShare, Equation (4), against the share of stayers with negative changes in total compensation. Firm-year level, weighted by firm size; year fixed effects absorbed. Sample as in Section 1. ContractShare is fixed at each firm's first sample year, so firms with ContractShare $= 1$ may still include workers who receive variable pay in later years, which explains the nonzero share of negative wage changes for those firms.

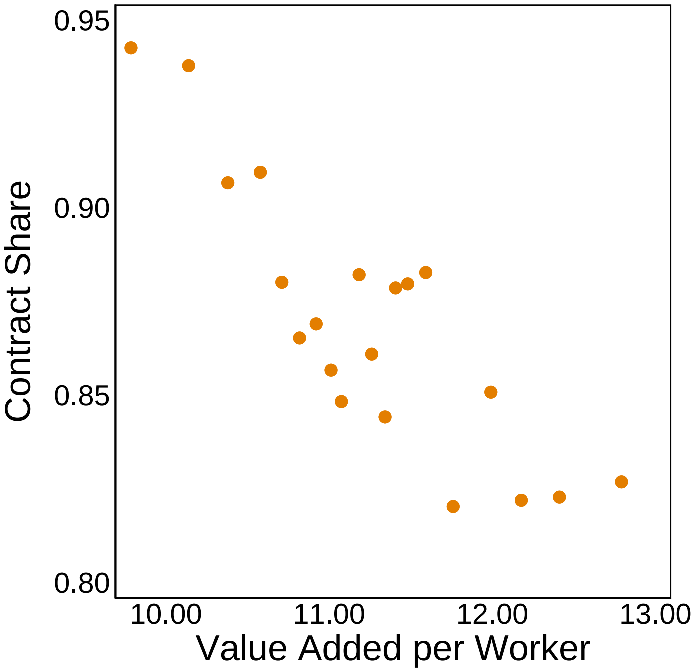

Figure B.IV
ContractShare Is Negatively Related to Value Added

Notes. Binscatter with 20 equal-sized bins of ContractShare, Equation (4), on Value Added. Firm level, weighted by firm size. Sample as in Section 1.

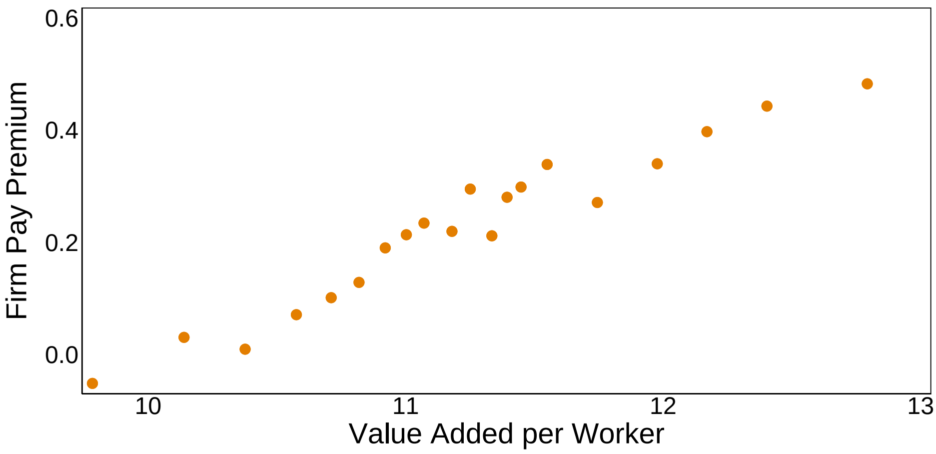

Figure B.V
Firm Pay Premium and Productivity

Notes. Binscatter with 20 equal-sized bins of the AKM firm pay premium (Supplemental Appendix C) against Value Added. Firm level, weighted by firm size. Sample as in Section 1.

C AKM Estimation

C.1 Estimation

In this appendix, we detail the estimation procedure for firm wage premiums.

To improve estimation precision, we first classify firms into 100 clusters using the method proposed by Bonhomme et al. (2019), which groups firms with similar wage distributions. Formally, we partition firms by solving the following weighted k-means problem:

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

where $\hat{F}_j(y)$ is the empirical cumulative distribution function (CDF) of log-wage in firm j, n_j is the number of workers in firm j, μ is a measure supported on a finite grid, $k(1), ..., k(J)$ denotes a partition of J firms into K clusters, and $H_1, ..., H_K$ are cluster-specific wage CDFs.

We then estimate firm wage premiums using the methodology of Abowd et al. (1999), but with cluster effects instead of firm effects. Specifically, we assume that the log of hourly wages for worker i of gender g in year t follows:

$$\log Y_{it} = \alpha_i + \psi_{gK(i,t)} + x'_{it}\beta^x_g + r_{it}, \tag{C.2}$$

where α_i is a worker fixed effect capturing the portable component of individual wages, x_{it} is a set of time-varying controls (including year fixed effects and a polynomial of age interacted with education), ψ_k is a wage premium paid at cluster k, $K(i,t)$ is an index function indicating the cluster of worker i's workplace in year t, and r_{it} is an error component capturing all other factors. We allow all parameters to vary by gender, including wage premiums.

We estimate Equation (C.2) by OLS. Table C.I presents the resulting variance decomposition. The results are broadly consistent with Gerard et al. (2021), who analyze the same setting. However, our estimates of cluster effects are lower than the firm effects reported in their study. This highlights the trade-off between firm effects—which may overestimate pay premiums due to small-sample bias—and cluster effects, which may underestimate them by ignoring within-cluster variation.

Table C.I
AKM Variance Decomposition

	All	Female	Male
Standard deviation of log-wages	0.664	0.654	0.658
AKM decomposition			
SD of worker effects	0.504	0.510	0.489
SD of cluster effects	0.213	0.200	0.220
SD of covariates	0.098	0.095	0.100
Corr. of worker and cluster effects	0.598	0.607	0.585
Percentage of variance of log wages due to:			
Worker effect	57.5	60.8	55.4
Cluster effect	10.2	9.3	11.2
Cov. of worker and cluster effects	29.0	28.9	29.1

Notes. AKM variance decomposition of wages. Firm clusters constructed via k-means (see text above). Sample as in Section 1.

C.2 Model Fit

We now test the restrictions imposed by the AKM framework. In particular, the restriction that wages follow a log-linear structure and that the job moving probability is uncorrelated with the error term. We test these restrictions with the approach proposed by Sorkin (2018).

From Equation (C.2), we have:

$$\log Y_{i,t} = \alpha_i + \psi_{gK(i,t)} + x'_{i,t}\beta_g^x + r_{i,t},$$
$$\log Y_{i,t+1} = \alpha_i + \psi_{gK(i,t+1)} + x'_{i,t+1}\beta_g^x + r_{i,t+1}.$$

Taking first differences:

$$\Delta \log Y_{i,t} - \Delta x'_{i,t}\beta_g^x = \Delta\psi_{gK(i,t)} + \Delta r_{i,t}.$$

We now take expectations, conditional on moving:

$$\mathbb{E}\big[\Delta \log Y_{i,t} - \Delta x'_{i,t}\beta^x_g | M_{i,t} = 1\big] = \Delta\mathbb{E}\big[\psi_{gK(i,t)}|M_{i,t} = 1\big] + \mathbb{E}\big[\Delta r_{i,t}|M_{i,t} = 1\big],$$

where $M_{i,t}$ indicates whether worker i moved between clusters in year t.

The key assumption to estimate Equation (C.2) by OLS is that the probability of moving is uncorrelated with the error term, that is $\mathbb{E}\big[\Delta r_{i,t}|M_{i,t} = 1\big] = 0$. Under this assumption:

$$\mathbb{E}\big[\Delta \log Y_{i,t} - \Delta x'_{i,t}\beta^x_g | M_{i,t} = 1\big] = \Delta\mathbb{E}\big[\psi_{gK(i,t)}|M_{i,t} = 1\big].$$

We take this restriction to the data by focusing on job switchers and comparing their residualized wage changes against their firm-effect changes. The results are in Figure C.I. The solid blue line plots the best-fitting line. The dashed line plots the 45 degree line. We find that wage changes closely follow changes in firm premiums, showing that the AKM framework fits the data well.

Figure C.I
Wage Change Corresponds to Firm Fixed-Effect Change

Notes. Residualized log wage changes for cluster switchers plotted against changes in cluster pay premiums. Dots are bin means; solid line is the best-fit line on micro-data; dashed line is the 45-degree line.

D Setting: Details

The Public Pension Fund: FGTS. All formally employed workers in the private sector are required to have an account at *Caixa*, a public bank. This account is known as FGTS (Fundo de Garantia do Tempo de Serviço). Employers must deposit 8% of each worker's gross monthly salary into this account. Furthermore, if a worker is laid off, they receive a severance payment amounting to 40% of the total balance accrued in their FGTS account. Workers can access these funds if they are laid off or upon reaching retirement age.

Layoff Fine. In the event of a layoff, firms are required to pay a government fine equivalent to 10% of the worker's total FGTS balance. This is in addition to the 40% severance payment made directly to the worker.

Unemployment Benefits. Workers who are laid off are eligible for unemployment benefits, which are contingent upon the length of their formal employment. The benefits are structured as follows:

- Workers employed for 6 to 11 months within the last 36 months receive three months of benefits.
- Workers employed for 12 to 23 months within the last 36 months receive four months of benefits.
- Workers employed for 24 months or more within the last 36 months receive five months of benefits.

→ In 2015, the monthly unemployment payment ranged from one to 1.76 times the minimum wage, dependent on the worker's average salary prior to being laid off.

E Layoff Rates and Layoff Costs across Countries

Layoff rates vary substantially across countries. This appendix documents the cross-country comparisons referenced in the main text, describes the data sources used to construct each country's layoff rate, and examines the relationship between layoff rates and the strictness of employment protection legislation.

Measuring layoff rates across countries. Comparing layoff rates internationally is challenging for two reasons. First, many countries do not distinguish between quits and layoffs in official statistics, reporting only total separations or unemployment rates. Second, statistical agencies that do decompose separation rates often rely on definitions that differ in periodicity (monthly, quarterly, or annual) and in the types of employer-initiated separations they include. To ensure comparability, we focus on annual layoff rates, which provide the broadest coverage across the countries for which data are available.

For the *United States*, we use the Job Openings and Labor Turnover Survey (JOLTS), published by the Bureau of Labor Statistics. JOLTS reports monthly layoffs and discharges alongside total nonfarm employment. We compute the annual layoff rate as total layoffs divided by average employment over the year. For 2010, this yields an annual layoff rate of 17.3%.

For *Brazil*, we compute the annual layoff rate directly from the RAIS transitions matrix used in the main analysis. Workers are classified as employed if they hold a formal job, and a layoff is defined as an employer-initiated separation (separation cause codes 10–19 in RAIS). The resulting annual layoff rate is 16.9% in 2010. As an independent check, we also estimate Brazil's layoff rate using the ECAF survey (described below), restricting to São Paulo—one of the cities in the Southeastern region covered by our main sample. The ECAF-based estimate closely matches the RAIS figure, providing reassurance that the administrative records and survey data yield consistent layoff rates.

For *Canada*, we use the Canadian Longitudinal Worker File, which reports annual layoff rates for permanent layoffs (excluding recalls). The annual layoff rate is approximately 6.6% in 2010.

For *Australia*, we use the Job Mobility Survey from the Australian Bureau of Statistics, which reports an annual separation rate of 8.8% in 2010. Applying the survey's estimate that approximately 40% of separations are involuntary yields an annual layoff rate of 3.5%.

For *Japan*, we use the Survey of Employment Trends from the Ministry of Health, Labour and Welfare, which reports a 1.4% annual rate of separations initiated by employers in 2010.

For *Latin American countries*, we use the ECAF (Encuesta CAF), a survey covering 11 large cities: Buenos Aires (Argentina), La Paz (Bolivia), São Paulo (Brazil), Santiago (Chile), Bogotá (Colombia), Quito (Ecuador), México City (Mexico), Panamá City (Panama), Lima (Peru), Montevideo (Uruguay), and Caracas (Venezuela). The ECAF asks respondents about their most recent separation from a previous employer, conditional on having experienced at least one separation. To convert these conditional separation shares into unconditional layoff rates, we estimate the hazard rate of separation at each tenure level using the pooled sample across all cities, and then compute the unconditional layoff rate as $\mathbb{E}[\Pr(\text{Layoff} \mid \text{Separation}) \times h(t)]$, where $h(t)$ is the empirical hazard rate of separation at tenure t. This approach ensures that cross-country differences in the distribution of tenure at separation do not mechanically drive differences in estimated layoff rates. The sample is restricted to private-sector, urban, formal employees on open-ended contracts. We report both individual country estimates and an aggregate for Latin America excluding Brazil, which yields an average annual layoff rate of 5.9%.

Employment protection and layoff rates. A key institutional factor underlying the cross-country variation in layoff rates is the cost of employer-initiated separations. We measure this using the EPLex dataset from the International Labor Organization (ILO), which constructs a composite indicator of the strictness of employment protection legislation across 102 countries. The index summarizes multiple dimensions of protection against individual dismissals, including valid and prohibited grounds for dismissal, the maximum length of probationary periods, notice periods, severance and redundancy payments, redress mechanisms, and administrative requirements. The composite indicator ranges from 0 to 1, with higher values indicating stronger protection.

Brazil (0.23) and the United States (0.22) rank among the countries with the least restrictive dismissal regulations. Other countries for which we have layoff rate data exhibit higher EPLex scores: Canada (0.29), Japan (0.36), Australia (0.40), and Latin

America excluding Brazil (0.41). Although the index varies somewhat over time, the relative ranking of countries remains stable between 2000 and 2020.

Figure E.I plots annual layoff rates against EPLex scores for all countries and regions described above. The figure reveals a clear negative relationship: countries with lower employment protection costs exhibit higher layoff rates. Brazil and the United States stand out as having both the lowest layoff costs and the highest layoff rates. At the other extreme, Japan and Australia combine strict employment protection with very low layoff rates. Most Latin American countries fall in between, with moderate protection scores and layoff rates well below those of Brazil.

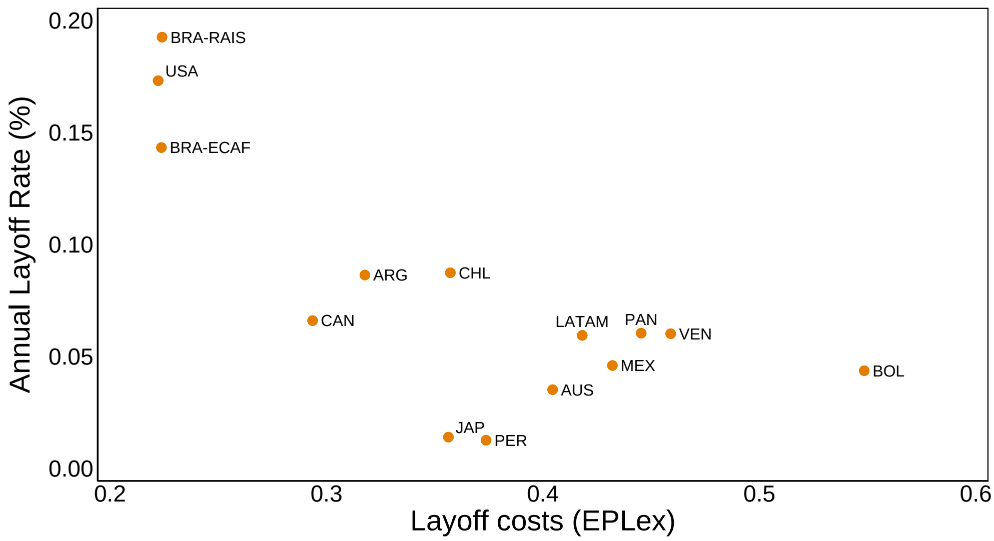

Figure E.I

Layoff Rates and Employment Protection across Countries

Notes. Annual layoff rate vs. EPLex employment protection index (ILO). Sources: JOLTS (US), RAIS (Brazil), ECAF (Latin America), Canadian Longitudinal Worker File, Australian Bureau of Statistics, Japanese Ministry of Health. All rates circa 2010.

These patterns support the interpretation that the relatively high layoff rates observed in Brazil—and studied in the main text—reflect institutional features, particularly permissive dismissal regulations, rather than anomalies specific to the Brazilian context. The United States, with similarly low employment protection, exhibits comparable layoff rates. The mechanisms we study are therefore most relevant in labor markets where observed separation decisions closely reflect underlying economic forces, rather than being heavily constrained by institutional frictions.

F Gelbach Decomposition

We decompose the unconditional productivity–layoff slope using the method of Gelbach (2016) to quantify the contributions of worker sorting and markdowns. The Gelbach decomposition avoids the well-known problem that sequentially adding controls yields path-dependent results: the contribution of each group of covariates depends on the order in which they are introduced. Instead, it uses the coefficients from the fully specified model to obtain a unique, order-invariant decomposition.

Let the base specification be

$$\text{Layoff}_i = \alpha + \beta^{\text{base}} \cdot \text{VA}_j + \varepsilon_i, \tag{F.1}$$

and the full specification be

$$\text{Layoff}_i = \tilde{\alpha} + \beta^{\text{full}} \cdot \text{VA}_j + \hat{\gamma}' X_i^{\text{sort}} + \hat{\delta} \cdot \text{Markdown}_j + u_i, \tag{F.2}$$

where X_i^{sort} denotes the vector of worker sorting controls (race and occupation fixed effects, tenure, AKM worker effects, and flexible interactions of gender, age, and education) and Markdown_j is the firm-level markdown measure, both as defined in Section 5.1. Gelbach (2016) shows that the change in the coefficient of interest can be decomposed exactly as

$$\beta^{\text{base}} - \beta^{\text{full}} = \underbrace{\frac{\text{Cov}(\hat{\gamma}' X_i^{\text{sort}}, \text{VA}_j)}{\text{Var}(\text{VA}_j)}}_{\text{worker sorting}} + \underbrace{\frac{\text{Cov}\left(\hat{\delta} \cdot \text{Markdown}_j, \text{VA}_j\right)}{\text{Var}(\text{VA}_j)}}_{\text{markdowns}}, \tag{F.3}$$

where the predicted components $\hat{\gamma}' X_i^{\text{sort}}$ and $\hat{\delta} \cdot \text{Markdown}_j$ are computed from the full model. Both models are estimated on the same sample (the estimation sample of the full specification in Column 3 of Table F.I).

The results of the decomposition indicate that sorting accounts for 57% of the gradient and markdowns for 59%. Taken together, these two components slightly more than account for the full gradient, implying a small residual term of -16%.

Table F.I
Markdowns Explain the Productivity–Layoff Gradient

	Dependent variable: Annual layoff rate		
	(1)	(2)	(3)
ValueAdded	-0.043***	-0.019***	0.007
	(0.001)	(0.001)	(0.006)
Markdown			-0.021***
			(0.005)
Worker controls	No	Yes	Yes
N	8,879,352	8,878,736	8,855,880

Notes. OLS estimates underlying the Gelbach decomposition. Column (1): layoff indicator on Value Added. Column (2): adds worker controls (as in Table II). Column (3): adds markdown ($\log((\text{VA}-\text{LC})/N)$, fixed to first sample year). Year fixed effects absorbed; standard errors clustered at the 3-digit industry × state level. Sample as in Section 1. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.