The Rise of Inaction: Job Reallocation and the Declining Labor Share*

Byunghee Choi[†]
Mount St. Mary's University

Choongryul Yang[‡]
Federal Reserve Bank of Atlanta

Abstract

Since 2000, the U.S. labor share and job reallocation rate have both fallen. We ask whether rising labor adjustment frictions can help explain this joint decline. We document a rise in establishment inaction and a shift toward smaller conditional employment changes, patterns consistent with kinked adjustment costs. In state-industry data, larger reallocation declines are associated with larger payroll-share declines, and the same cells with larger increases in establishment-level inaction also show larger declines in firm-level job reallocation and payroll shares. A calibrated firm-dynamics model disciplined by inaction can account for a meaningful share of both aggregate declines over the sample period.

Keywords: Labor Share, Job Reallocation, Business Dynamism, Labor Market Frictions, Kinked Adjustment Costs
JEL Codes: E24, E25, J23, J42

---

*The views expressed here are those of the authors and do not necessarily reflect those of the Federal Reserve Bank of Atlanta or the Federal Reserve System. This version: April 2026

†16300 Old Emmitsburg Road, Emmitsburg, MD 21727, U.S.A. Email: b.choi@msmary.edu.

‡1000 Peachtree St NE, Atlanta, GA 30309, U.S.A. Email: cryang1224@gmail.com.

1 Introduction

Two of the most significant secular transformations in the U.S. economy since 2000 are the persistent decline in the labor share of income and the simultaneous erosion of business dynamism.[1] These phenomena are central puzzles in modern macroeconomics, yet they have been studied largely in isolation, each with its own set of proposed causes. This paper asks whether rising firm-level labor adjustment frictions can help explain both trends.

The literature has advanced several leading explanations for the labor share decline (see Elsby, Hobijn and Şahin, 2013, for an early comprehensive account). A prominent class of explanations centers on capital-labor substitution: a fall in the relative price of investment goods has induced firms to replace workers with capital (*e.g.*, Karabarbounis and Neiman, 2014). However, this channel is contested; Glover and Short (2020) and Oberfield and Raval (2021) estimate aggregate elasticities of substitution at or below one, concluding that capital deepening alone cannot account for the decline. An increasingly influential set of explanations emphasizes rising product market power. The "superstar firm" hypothesis of Autor, Dorn, Katz, Patterson and Van Reenen (2020) argues that a between-firm reallocation of economic activity to highly productive, low-labor-share firms drives the aggregate decline. Relatedly, De Loecker, Eeckhout and Unger (2020) and Barkai (2020) document a broad rise in firm-level markups and pure profits at the expense of both labor and capital shares. Yet Kehrig and Vincent (2021) show that roughly half of the manufacturing labor share decline occurs *within* firms, not just between them—leaving a substantial within-firm component that the superstar story does not address. Baqaee and Farhi (2020) develop a general equilibrium framework showing that misallocation of factors across heterogeneous firms can have first-order effects on aggregate output and factor shares—a channel that our model formalizes through the inaction wedge. Grossman and Oberfield (2022) provide a comprehensive survey and note that no single mechanism has proven fully satisfactory.

The decline in business dynamism has attracted a parallel but largely separate literature. Davis, Faberman and Haltiwanger (2012) and Hyatt and Spletzer (2013) document falling rates of job creation, destruction, and worker turnover. Molloy, Smith, Trezzi and Wozniak (2016) explore potential explanations ranging from demographic shifts to occupational licensing. More recently, Akcigit and Ates (2023) link declining dynamism to reduced knowledge diffusion from frontier firms, while Davis and Haltiwanger (2024) emphasize the role of housing markets and credit conditions. Pries and Rogerson (2022) attribute declining worker turnover to a reduction in short-duration employment spells. Despite the concurrent timing, Albrecht and Decker (2026) find no systematic industry-level link between rising markups and declining dynamism, suggesting that

[1] The aggregate payroll share fell from approximately 57% in 2001 to 53% by 2019, while the gross job reallocation rate (JRR)—the sum of job creation and destruction rates measured from the Census Business Dynamics Statistics—declined from 31% to 22% over the same period.

these two trends may have distinct underlying drivers.

This paper provides a distinct, *within-firm* mechanism rooted in the labor market that bridges these two literatures. Our contribution is not to displace the product-market-power or capital-deepening channels, but to identify a complementary force that they do not directly address. We build the argument in three steps.

First, we document a novel micro-level fact that we term the "rise of inaction." Using establishment-level data from the BLS Business Employment Dynamics (BED), we show that the aggregate decline in the job reallocation rate is disproportionately driven by a sharp increase in the fraction of establishments reporting *exactly zero* net employment change in a given quarter. This "no-change" fraction rose from approximately 44% in the late 1990s to over 52% by 2019. The timing is critical: this rise begins around 2000, precisely when the declines in both the labor share and job reallocation accelerate. This fact is difficult to reconcile with standard workhorse models of firm dynamics that assume smooth convex (*e.g.*, quadratic) adjustment costs, which predict continuous adjustment and zero mass at the no-change point. The existence and growth of this inaction mass is strongly suggestive of kinked adjustment costs.

Second, we show that the macro-level co-movement between the labor share and job reallocation mirrors the micro-level inaction trend. Using a state-by-industry panel built from Census Business Dynamics Statistics (BDS) job-flow data and BEA payroll-share data over 2001–2019, we show that states and industries experiencing larger declines in their job reallocation rate also experienced larger declines in their labor share, and this relationship operates through both firm entry/exit and continuing-firm adjustments, with entry and exit playing the dominant role. The relationship is robust to the inclusion of state-industry and year fixed effects and is not explained away by controls for market concentration, unionization, and manufacturing exposure. A multi-dimensional heterogeneity analysis shows that the JRR–labor-share nexus is most precisely estimated in services, strengthens with routine task intensity, is concentrated in the pre-Great-Recession period, and is larger in lower-concentration cells. A new bridge analysis then connects the micro and macro evidence more directly: the same state-industry cells with larger increases in BED inaction also experienced larger declines in BDS reallocation and payroll shares. This bridge is present for both continuing-firm and birth/death margins, though it is quantitatively larger for firm turnover.

Third, guided by the micro-level evidence, we build a dynamic general equilibrium model in the spirit of Hopenhayn and Rogerson (1993). The model's crucial departure is the incorporation of a kinked linear labor adjustment cost $\gamma|\Delta n|$, where γ is the per-worker cost of changing employment, which endogenously creates a band of inaction. Critically, we calibrate the baseline (2001) economy with a *positive* adjustment cost, so that the benchmark already features a realistic inaction region. Because the model is annual while the BED no-change fact is quarterly, we use the quarterly BED no-change rate as an empirical proxy for latent inaction and target the share of firms in the stationary distribution that lie inside a positive-width (S, s) band. The counterfactual then raises γ to match the corresponding 2019 BED target, and the resulting declines in the labor share and job reallocation rate emerge as *untargeted model predictions*.

The model's distinctive value lies in its ability to connect the rise of inaction to the decline in aggregate reallocation and the labor share within a single quantitative framework. Superstar-firm and capital-deepening explanations can speak to the labor share decline, but they are largely silent on inaction dynamics. Quantitatively, targeting the observed rise in establishment inaction, the calibrated model accounts for approximately 45 percent of the observed 3.9 percentage-point decline in the U.S. labor share and 39 percent of the decline in the job reallocation rate—both as untargeted predictions. We interpret this as evidence that rising labor adjustment frictions can explain a meaningful share of the joint decline, not that they are the only force at work.

Our paper relates to several strands of the literature beyond the core labor-share and business-dynamism literatures discussed above. The "declining responsiveness" finding of Decker, Haltiwanger, Jarmin and Miranda (2020)—that the post-2000 decline in dynamism reflects weaker firm-level employment responses to productivity shocks, not smaller shocks—is the direct structural-estimation analog of our rising-adjustment-cost hypothesis. Our "rise of inaction" finding provides a specific, granular manifestation of this pattern at the establishment level. Our modeling approach builds on the foundational literature on adjustment costs surveyed by Hamermesh and Pfann (1996), and draws on the insight of Bloom (2009) that non-smooth adjustment costs—whether fixed or kinked linear—generate bands of inaction. Our contribution is to show that specifically *linear* per-unit costs, rather than fixed costs, best match the core inaction and large-adjustment facts we document. Lee and Mukoyama (2018) and Cooper, Haltiwanger and Willis (2015) show that non-smooth labor adjustment cost models are essential for explaining plant-level employment dynamics, while Cooper, Haltiwanger and Willis (2024) argues that rising labor adjustment costs are a leading explanation for the decline in establishment responsiveness over time. More broadly, our paper suggests that rising labor adjustment frictions can connect declining responsiveness at the micro level to aggregate factor-share movements, thereby complementing markup-based and capital-deepening explanations rather than displacing them. On the institutional side, we document that rising regulatory compliance burdens (Trebbi and Zhang, 2022) and employer-provided health insurance costs have grown substantially over our sample period, contributing to the plausibility of rising adjustment frictions. On the labor-market side, our work complements the growing literature on labor market power (Berger, Herkenhoff and Mongey, 2022, Yeh, Macaluso and Hershbein, 2022, Mertens and Schoefer, 2024, Gouin-Bonenfant, 2022), which documents rising monopsony power and between-firm productivity dispersion as potential sources of the very frictions we model. Finally, Hubmer and Restrepo (2026) and Smith, Yagan, Zidar and Zwick (2022) highlight the importance of firm heterogeneity and organizational form in understanding the labor share, while Bergholt, Furlanetto and Maffei-Faccioli (2022) provide new time-series evidence on the decline.

The remainder of the paper is organized as follows. Section 2 describes the data sources and documents the aggregate co-movement between the labor share and job reallocation. Section 3 presents the formal empirical analysis, including task-based heterogeneity and temporal stability. Section 4 documents the rise of establishment-level inaction and assembles empirical evidence on the sources of rising adjustment costs, building the micro-level foundation for the model. Section 5 details the theoretical model and its calibration. Section 6 presents the quantitative results. Section 7 concludes.

2 Data

We construct a state-by-industry panel that combines data on labor income shares, job reallocation, and market concentration from several administrative and survey sources. Our main sample covers the period 2001–2019, chosen to begin with the first year of reliable NAICS-based data and to end before the COVID-19 pandemic, which introduced extreme and transitory disruptions to employment dynamics. We discuss each data source in turn; additional detail on variable construction is provided in the Data Appendix (Appendix A).

2.1 Labor Income Share

Our measure of the labor income share is the *payroll share*, following the methodology of Elsby et al. (2013) and Karabarbounis (2024). For each state-industry-year cell, we define:

$$\lambda_{ist} = \frac{\text{Compensation of Employees}_{ist}}{\text{Gross Value Added}_{ist}}$$

where i indexes industry, s indexes state, and t indexes year. The numerator is total compensation of employees (wages, salaries, and supplements), and the denominator is gross domestic product by state (value added). Both are sourced from the Bureau of Economic Analysis (BEA) Regional Economic Accounts, specifically the SAGDP4 (compensation) and SAGDP2 (GDP) tables. These data are available annually at the state-by-industry level for broad NAICS-based sectors from 1997 to 2024. We map these to 11 private-sector industry groups that correspond to the industry classifications available in the Census Business Dynamics Statistics. By construction, this payroll share excludes proprietors' income, which avoids the well-known ambiguity of allocating self-employment income between labor and capital (Elsby et al., 2013, Grossman and

Oberfield, 2022).

2.2 Job Reallocation

We construct measures of labor market fluidity from the Business Dynamics Statistics (BDS), produced by the Census Bureau from the Longitudinal Business Database (LBD). The BDS is an annual dataset that tracks firm-level job flows—job creation from expanding and entering firms, and job destruction from contracting and exiting firms—at the state-by-sector level from 1978 to 2023.

Following the standard definitions in Davis, Haltiwanger and Schuh (1998), we define job creation (JC_t) as the sum of all employment gains at expanding and entering firms and job destruction (JD_t) as the sum of all employment losses at contracting and exiting firms. These gross flows are normalized by average employment to yield the Job Creation Rate (JCR_t) and Job Destruction Rate (JDR_t). Our primary measure of labor market fluidity is the Job Reallocation Rate:

$$JRR_{ist} = JCR_{ist} + JDR_{ist}$$

which captures the total "churn" or dynamism in the labor market. The BDS reports data at the state-by-sector level using 2-digit NAICS codes, which we aggregate to 11 private-sector industry groups matching the BEA classification (see Appendix A.2 for the sector mapping). For the establishment-level size-of-change analysis (Section 4.1), we use the BLS Business Employment Dynamics (BED), which uniquely provides data on the size of establishment-level employment changes. BED results are reported as a robustness check in Appendix B.4.

2.3 Market Concentration

To control for the product-market-power channel, we construct an employment-based Herfindahl-Hirschman Index (HHI) at the state-by-industry-year level using the Statistics of U.S. Businesses (SUSB) from the Census Bureau. The SUSB reports the number of firms (N^k_{ist}) and total employment (E^k_{ist}) by detailed firm-size class for each state-industry cell. Because the SUSB groups firms into size bins rather than reporting individual firm sizes, we approximate the firm-level HHI by assuming equal-sized firms within each bin:

$$HHI_{ist} = \sum_k \frac{1}{N^k_{ist}} \left(\frac{E^k_{ist}}{E_{ist}} \right)^2$$

where the division by N^k_{ist} converts each bin's squared employment share into the sum of squared shares of the individual firms in that bin (under the equal-size-within-bin assumption). Higher

HHI values indicate greater employment concentration. The SUSB data are available from 1998 to 2022.

2.4 Additional Controls: Unionization and Manufacturing Share

To control for alternative explanations of the labor share decline, we construct two additional state-level variables. First, we measure *private-sector union membership rates* from the database of Hirsch and Macpherson (2003), which provides state-by-year union density estimates derived from the Current Population Survey. Declining unionization has been widely cited as a contributing factor to the falling labor share (Stansbury and Summers, 2020), and we use union rates to ensure that our JRR–labor-share finding is not driven by this channel. Second, we compute the *manufacturing employment share* at the state level using the SUSB data. Because the secular decline in manufacturing is associated with both lower labor shares (Elsby et al., 2013) and lower job reallocation, controlling for the structural transformation channel strengthens our identification. Both variables vary at the state-year level and are merged into our panel accordingly.

2.5 Establishment-Level Size-of-Change Data

A key contribution of this paper exploits BLS data on establishments classified by the *size* of their quarterly employment change. The BED program reports, for the aggregate private sector, the number of establishments and associated employment flows broken down by the magnitude of their employment change: 1–4 jobs, 5–19 jobs, 20 or more jobs, and crucially, *zero change*. These data are available from 1992 through 2025 and are not seasonally adjusted. We compute the inaction rate as the fraction of all establishments reporting zero net employment change in a given quarter. Details on the construction of the size-of-change series are provided in Appendix A.

2.6 Analysis Panel

We merge these data sources into a state-by-industry-year panel. The unit of observation is a state (s) × industry (i) × year (t) cell, where industries correspond to 11 private-sector groups: Construction; Manufacturing; Wholesale Trade; Retail Trade; Transportation and Warehousing; Information; Finance, Insurance, and Real Estate; Professional and Business Services; Education and Health Services; Leisure and Hospitality; and Other Services. The panel also includes state-level union membership rates and manufacturing employment shares as additional controls. After restricting to cells with non-missing observations for the payroll share, job reallocation

Table 1: Summary Statistics

Panel A: Full Sample (2001–2019)						
Variable	Mean	Std.Dev.	Min	Median	Max	Obs.
Payroll share (%)	56.63	16.85	12.21	59.28	90.93	10631
Job reallocation rate (%)	27.47	7.71	10.39	26.55	90.07	10631
Job creation rate (%)	13.97	4.20	4.34	13.43	41.16	10631
Job destruction rate (%)	13.50	4.62	4.34	12.84	62.49	10631
HHI (employment)	0.0019	0.0027	0.0001	0.0011	0.0625	10631
Union membership (%)	6.59	3.46	0.81	5.82	16.91	10631
Mfg. employment share	0.11	0.05	0.00	0.11	0.26	10631

Panel B: Means by Year					
Variable	2001	2005	2010	2015	2019
Payroll share (%)	60.22	56.72	55.41	55.52	55.36
Job reallocation rate (%)	31.59	28.67	26.75	24.21	22.74
Job creation rate (%)	15.99	14.86	11.84	13.14	12.15
Job destruction rate (%)	15.60	13.80	14.91	11.06	10.59
HHI (employment)	0.0012	0.0020	0.0020	0.0019	0.0020
Union membership (%)	8.22	7.01	6.34	6.07	5.87
Mfg. employment share	0.15	0.12	0.10	0.10	0.10

Panel C: Long-Differences (2001–2019)			
Variable	Mean Δ	Std.Dev.	Obs.
Payroll share (p.p.)	-4.81	7.31	551
Job reallocation rate (p.p.)	-8.91	6.15	551
Job creation rate (p.p.)	-3.84	3.35	551
Job destruction rate (p.p.)	-5.07	3.87	551
HHI (employment)	0.0007	0.0012	551
Union membership (p.p.)	-2.35	1.86	551
Mfg. employment share	-0.05	0.02	551

Notes: The table reports summary statistics at the state-industry level. The payroll share is the ratio of compensation of employees to gross value added (%). Job reallocation, creation, and destruction rates are from the Census BDS (%). The HHI is a grouped-data employment concentration proxy constructed from Census SUSB firm counts and employment by size class. The union membership rate covers the private sector and is drawn from unionstats.com (Hirsch & Macpherson). The manufacturing employment share is computed from SUSB data.

rate, and HHI, our main estimation sample contains 10,631 state-industry-year observations spanning 561 state-industry cells over 2001–2019.

Table 1 presents the summary statistics in three panels. Panel A reports the pooled distribution: the average payroll share is 56.6% with substantial cross-sectional variation (standard deviation of 16.9%). The mean JRR is 27.5%, and the mean employment HHI is 0.0019. Private-sector union membership averages 6.6% across states and years, while manufacturing accounts

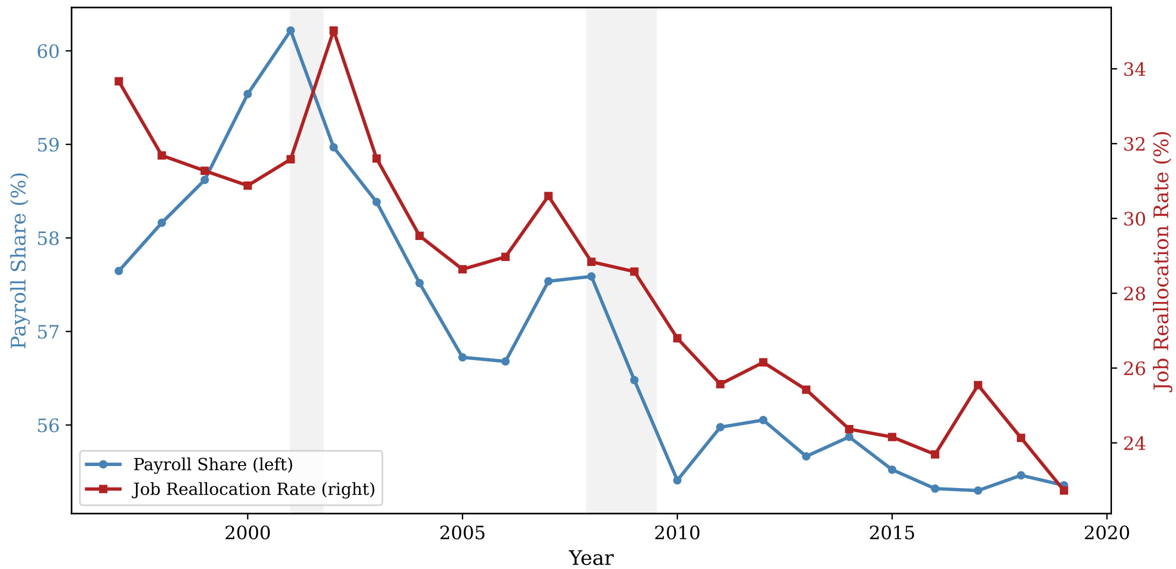

Figure 1: U.S. Payroll Share and Job Reallocation Rate, 1997–2019

Notes: This figure plots the average payroll share (left axis) and job reallocation rate (right axis) across state-by-industry cells. The payroll share is the ratio of compensation of employees to gross value added, constructed from BEA Regional Economic Accounts. The job reallocation rate is the sum of job creation and destruction rates from the Census Business Dynamics Statistics. Shaded regions indicate NBER recession dates.
Source: Bureau of Economic Analysis, Regional Economic Accounts (SAGDP); Census Bureau, Business Dynamics Statistics.

for 11% of state employment on average. Panel B documents the secular trends: between 2001 and 2019, the mean payroll share fell from 60.2% to 55.4%, the JRR declined from 31.6% to 22.7%, and union membership fell from 8.2% to 5.9%. The grouped-data HHI rises modestly over the sample but remains small in level terms, underscoring that it should be interpreted as an imperfect concentration proxy rather than a direct markup measure. Panel C reports the distribution of long-differences: the average state-industry cell experienced a 4.8 percentage-point decline in the payroll share and an 8.9 percentage-point decline in the JRR.

Figure 1 plots these trends visually. The payroll share and JRR exhibit a pronounced joint decline beginning around 2000–2001, tracking each other closely throughout the 2001–2019 sample period. The decline in job reallocation is pervasive across all 11 industry groups, with no exceptions.[2]

[2]Appendix Figure B.4 plots the JRR for each industry with fitted trend lines.

3 Empirical Analysis

The aggregate co-movement between the payroll share and the JRR documented in Section 2 suggests a systematic link between declining business dynamism and the falling labor share. We now test this hypothesis formally using panel fixed-effects regressions that exploit *within*-cell variation over time, controlling for permanent cross-cell differences and common aggregate shocks.

Our empirical strategy proceeds in four steps. First, we estimate the core JRR–labor-share relationship using panel fixed effects, progressively adding controls that correspond to the leading alternative explanations for the labor share decline (Section 3.1). Second, we examine *where* this relationship holds through heterogeneity analysis along four dimensions—industry type, time period, market structure, and task content (Section 3.2)—and then document the temporal stability of the relationship across sub-periods (Section 3.3). Third, we probe the mechanism through alternative specifications—excess JRR, lagged JRR, creation-destruction decomposition, and margin decompositions (Section 3.4). Finally, we synthesize the evidence and discuss identification limitations (Section 3.5).

3.1 Panel Analysis

Our primary specification is a panel fixed-effects model using all annual observations:

$$\ln \lambda_{ist} = \alpha_{is} + \gamma_t + \beta \ln JRR_{ist} + \mathbf{X}'_{ist}\boldsymbol{\delta} + \epsilon_{ist}$$

where α_{is} are state-industry fixed effects that absorb all permanent differences across cells (industry composition, geography, institutional environment), γ_t are year fixed effects that absorb common macro shocks (business cycles, policy changes, aggregate technology trends), and $\mathbf{X}_{ist}$ is a vector of time-varying controls. The coefficient β is thus identified solely from *within-cell* variation over time: does a given state-industry cell's labor share fall in years when its reallocation rate declines, beyond what can be explained by aggregate trends and the leading competing channels?

Table 2 reports the results for our baseline sample (2001–2019) and the pre-Great-Recession sub-period (2001–2007). We build up the specification progressively, adding controls that correspond to the leading alternative explanations for the labor share decline. The goal is to determine whether the JRR–labor-share relationship survives once these competing channels are accounted for.

The baseline specification (column 1) yields an elasticity of 0.044: within a given state-industry cell, a 10% decline in the JRR is associated with a 0.4% decline in the payroll share. For

Table 2: Panel Regression Results

	log(Payroll share)					
	(1)	(2)	(3)	(4)	(5)	(6)
Panel: 2001–2019						
Job reallocation rate	0.044***	0.045***		0.045***	0.045***	0.045***
	(0.008)	(0.008)		(0.008)	(0.008)	(0.008)
HHI (concentration)		0.012	0.008	0.013	0.011	0.012
		(0.013)	(0.013)	(0.013)	(0.013)	(0.013)
Union membership				0.003*		0.003*
				(0.001)		(0.001)
Mfg. share					-0.324	-0.305
					(0.250)	(0.249)
Year FE	O	O	O	O	O	O
State×Industry FE	O	O	O	O	O	O
Observations	10631	10631	10631	10631	10631	10631
Panel: 2001–2007						
Job reallocation rate	0.033***	0.032***		0.032***	0.032***	0.032***
	(0.011)	(0.011)		(0.011)	(0.011)	(0.011)
HHI (concentration)		-0.023	-0.026*	-0.023	-0.022	-0.023
		(0.015)	(0.015)	(0.015)	(0.015)	(0.015)
Union membership				-0.000		-0.000
				(0.002)		(0.002)
Mfg. share					-0.384	-0.383
					(0.248)	(0.248)
Year FE	O	O	O	O	O	O
State×Industry FE	O	O	O	O	O	O
Observations	3904	3904	3904	3904	3904	3904

Notes: The dependent variable is the log payroll share (compensation of employees divided by gross value added), and all regressors enter in logs. Column (1) includes only the JRR; column (2) adds the employment HHI; column (3) includes only HHI; column (4) adds the state-level private-sector union membership rate; column (5) adds the state-level manufacturing employment share; and column (6) includes all controls simultaneously. All specifications include state×industry and year fixed effects. Standard errors are clustered at the state×industry level and reported in parentheses. ***, **, * denote significance at the 1%, 5%, and 10% levels, respectively.

the pre-recession panel (2001–2007), this elasticity is somewhat smaller (0.033)—reflecting the shorter sample window but still highly significant.

The most prominent alternative explanation for the labor share decline centers on rising product market power. Autor et al. (2020) argue that technological change and globalization have increasingly favored a small number of highly productive "superstar" firms that capture disproportionate market shares; because these dominant firms tend to be less labor-intensive, the reallocation of economic activity toward them mechanically reduces the aggregate labor share. Relatedly, De Loecker et al. (2020) document a broad rise in firm-level markups since the 1980s, implying that firms increasingly price above marginal cost and retain a larger share of revenue as profits at the expense of both labor and capital. If this same rise in concentration simultaneously reduces competitive entry and exit—thereby lowering measured job reallocation—then our JRR coefficient could be spuriously capturing a market-power effect rather than a labor-friction effect. To test this, we add log employment HHI as a grouped-data proxy for local market concentration (column 2). The JRR coefficient is unchanged at 0.045, and HHI itself enters insignificantly. When we exclude JRR and retain HHI alone (column 3), concentration cannot explain within-cell labor share variation on its own. We therefore conclude that the JRR–labor-share relationship is not explained away by our concentration proxy. This finding is consistent with Albrecht and Decker (2026), who show that there is no systematic negative correlation between rising markups and declining business dynamism across industries—the sectors with the largest markup increases are not those with the largest dynamism declines.

A second alternative emphasizes declining worker bargaining power. Stansbury and Summers (2020) argue that the erosion of unions and labor-market institutions has weakened workers' ability to claim rents, contributing to the labor share decline independently of changes in technology or market structure. Declining unionization could also confound our estimates if rigid union wage structures simultaneously discourage labor reallocation: as unions weaken, both the labor share and reallocation barriers fall together, producing a spurious positive correlation. Adding state-level private-sector union membership rates (column 4) has no effect on the JRR coefficient, indicating that our friction channel operates independently of the bargaining-power channel.

A third explanation highlights the structural transformation of the U.S. economy away from manufacturing—a sector with historically above-average labor shares—toward services (Elsby et al., 2013). Although our state-industry fixed effects absorb permanent between-industry composition differences, deindustrialization could still confound our estimates through state-level equilibrium channels: states losing manufacturing may experience wage spillovers, demand shifts, or labor-market tightening that affect both the JRR and the labor share within non-manufacturing industries. Controlling for state-level manufacturing employment share (column 5) leaves the JRR coefficient unchanged, confirming that our results are not driven by these deindustrialization spillovers.

Finally, column 6 includes all three controls simultaneously alongside JRR. The JRR coefficient remains 0.045 and statistically significant, while none of the alternative channels—concentration, unionization, or manufacturing exposure—enters significantly. The main takeaway from Table 2

Table 3: Heterogeneity Analysis: Panel Regressions

	By Industry		By Period		By Concentration		Task
	Goods	Services	2001–2007	2008–2019	Low HHI	High HHI	Interact.
log(JRR)	0.062	0.045***	0.032***	0.017*	0.072***	0.033***	0.040***
	(0.042)	(0.007)	(0.011)	(0.009)	(0.012)	(0.009)	(0.008)
log(HHI)	0.053*	-0.015	-0.023	-0.009	-0.053***	0.029	0.014
	(0.030)	(0.014)	(0.015)	(0.016)	(0.016)	(0.021)	(0.013)
log(JRR) × RTI_z							0.022**
							(0.009)
Year FE	O	O	O	O	O	O	O
State×Ind FE	O	O	O	O	O	O	O
Observations	1927	8704	3904	6727	5244	5212	10631

Notes: The dependent variable is the log payroll share, and all regressors enter in logs. Columns (1)–(2) split the sample into goods-producing and service industries; columns (3)–(4) split by the pre- and post-Great-Recession subperiods; columns (5)–(6) split by below- and above-median initial employment HHI; and column (7) interacts log JRR with the standardized routine task intensity index (RTI_z) following Autor and Dorn (2013). All specifications include state×industry and year fixed effects and control for log HHI. Standard errors are clustered at the state×industry level and reported in parentheses. ***, **, * denote significance at the 1%, 5%, and 10% levels, respectively.

is therefore straightforward: the JRR–labor-share relationship is not absorbed by the leading alternative mechanisms we can measure in this setting.

3.2 Heterogeneity

Having established the baseline positive relationship between JRR and the labor share, we now ask: *where* does this relationship hold, and where does it break down? The answers to these questions are critical for distinguishing a labor-friction interpretation from competing explanations. If the JRR–labor-share relationship were driven by a confounding aggregate trend, it would appear uniformly across all subsamples. If, instead, it reflects an economically meaningful labor-friction channel, it should be strongest in settings where such frictions are likely to matter most.

We estimate separate panel regressions by industry type, time period, initial market concentration, and task content. Table 3 reports the results.

By industry type. The JRR–labor-share relationship is precisely estimated and highly significant in service industries (column 2, coefficient 0.045, $p < 0.01$), which account for the vast majority of U.S. employment. The goods-producing coefficient is larger in magnitude (0.062) but statistically insignificant, reflecting the limited cross-sectional variation available from only two broad goods-producing industry groups (construction and manufacturing). The key finding is that the relationship is most precisely estimated in services—the part of the economy where the labor share decline has the greatest aggregate consequence.

By time period. The relationship is stronger before the Great Recession (2001–2007, coefficient 0.032, $p < 0.01$) than afterward (2008–2019, coefficient 0.017, $p < 0.10$), but remains positive in both sub-periods. We show in Section 3.3 that this difference is influenced by the immediate post-recession recovery years and that the relationship remains positive across rolling windows.

By initial concentration. In lower-concentration cells (below-median HHI), the JRR coefficient is large (0.072) and highly significant; in higher-concentration cells, it is smaller (0.033) but still significant. This gradient is consistent with a labor-friction channel operating more strongly where the grouped-data concentration proxy is least likely to dominate the interpretation. Put differently, the JRR–labor-share relationship is larger in lower-concentration environments, though we do not interpret this proxy as a definitive measure of product-market power.

By task content. The adjustment-cost theory predicts that the JRR–labor-share relationship should be stronger in industries where employment adjustment occurs primarily on the *head-count* margin—i.e., where production depends on standardized, easily quantified labor inputs—and weaker where firms adjust through hours, task composition, or worker quality. To test this, we interact log JRR with a standardized Routine Task Intensity (RTI) index following Autor and Dorn (2013), constructed from occupational composition data in the O*NET task importance database (see Appendix A.7 for details). Column (7) confirms the prediction: the JRR$\times$RTI interaction is positive and significant (0.022, $p < 0.05$), indicating that moving from the least to the most routine-intensive industry substantially amplifies the JRR–labor-share elasticity. This result is more consistent with a labor-friction interpretation than with a purely spurious aggregate trend.

3.3 Temporal Stability

The heterogeneity analysis in Table 3 reveals a modest difference in the JRR–labor-share elasticity between the pre-Great-Recession (0.032) and post-recession (0.017) periods. We now investigate whether this difference reflects a genuine structural change or is an artifact of particular sub-period conditions.

Figure 2 presents rolling 7-year panel regressions plotted by end year. The JRR coefficient remains positive throughout the sample, though it varies in magnitude over time. The coefficient dips in windows that include the Great Recession—when common macro shocks dominate

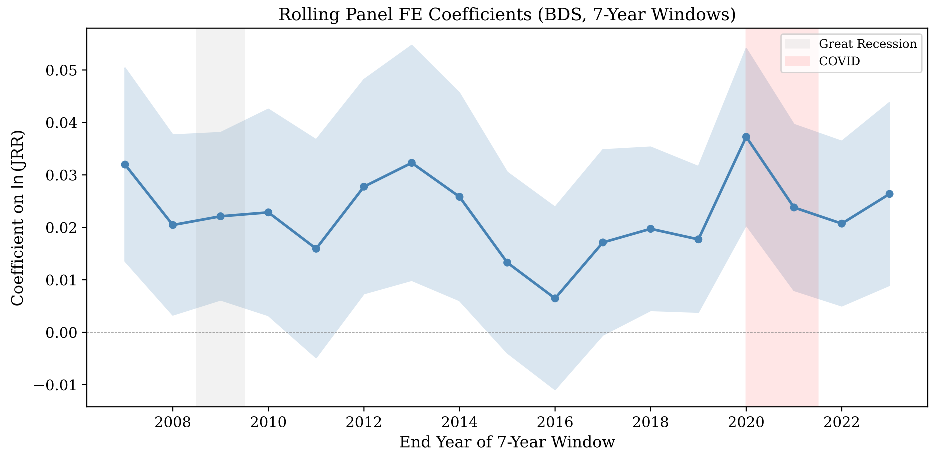

Figure 2: Rolling Panel FE Coefficients (7-Year Windows)

Notes: Each point is the coefficient on ln(JRR) from a panel regression of ln(payroll share) on ln(JRR), with state-industry and year FE, estimated over a 7-year window ending at the indicated year. Shaded bands are 90% confidence intervals (standard errors clustered at the state×industry level). Gray and red shaded regions mark the Great Recession and COVID periods, respectively. Job reallocation from Census BDS.

Table 4: JRR–Labor Share Relationship by Sub-Period

Period	Years	β(JRR)	SE	N
Full pre-COVID	2001–2019	0.045***	(0.008)	10631
Pre-GR	2001–2007	0.032***	(0.011)	3904
Post-GR	2010–2019	0.017*	(0.010)	5605
Post-GR (excl. recovery)	2012–2019	0.024***	(0.009)	4483
Full incl. COVID	2001–2022	0.042***	(0.008)	12308

Notes: Each row reports the coefficient from a panel fixed-effects regression of the log payroll share on the log JRR, with state-industry and year fixed effects. Standard errors are clustered at the state-industry level and reported in parentheses. ***, **, * denote significance at the 1%, 5%, and 10% levels, respectively.

within-cell variation—and later recovers. The 90% confidence band remains above zero for most of the sample period. The COVID-era spike (windows ending 2020–2021) is consistent with the relationship re-emerging when labor markets experienced exceptionally high churn, before settling back toward its pre-COVID range.

Table 4 confirms this pattern in a sub-period decomposition. The post-recession coefficient (2010–2019) is smaller than the pre-recession coefficient (0.017 vs. 0.032), but this difference is largely driven by the immediate recovery years: excluding 2010–2011, the 2012–2019 coefficient is 0.024 and highly significant ($p < 0.01$). The full-sample coefficient including the COVID period (2001–2022) remains positive and precisely estimated at 0.042.

The modest pre- vs. post-recession difference is consistent with the timing of institutional

Table 5: Robustness and Alternative Specifications

Specification	Coefficient	(s.e.)	N
Excess JRR	0.031***	(0.006)	10631
Lagged JRR ($t-1$)	0.030***	(0.008)	10070
JC rate (separate regression)	0.012*	(0.006)	10631
JD rate (separate regression)	0.033***	(0.005)	
Intensive margin (continuers)	0.019**	(0.007)	10631
Extensive margin (births+deaths)	0.027***	(0.004)	10631

Notes: The dependent variable is the log payroll share. All specifications use 2001–2019 BDS data with state×industry and year fixed effects. The "excess JRR" removes the year-level mean from the JRR before entering it as a regressor. The "separate regression" rows report coefficients from a single regression with log JC and log JD entered simultaneously. The intensive margin is job reallocation from continuing-firm expansions and contractions; the extensive margin is job creation from firm births plus job destruction from firm deaths. Sample sizes vary across rows because each specification uses the available observations for its specific reallocation measure. Standard errors are clustered at the state×industry level and reported in parentheses. ***, **, * denote significance at the 1%, 5%, and 10% levels, respectively.

friction growth. Employer healthcare costs—the single largest identified component of per-worker adjustment costs (Section 4.4)—grew at approximately 8% per year in real terms between 2001 and 2005 (KFF), before decelerating to 3–5% annual growth after 2010 as the Affordable Care Act and market forces moderated premium inflation. The inaction rate itself rose most steeply before 2007 (Figure 3(a)), with minimal further increase thereafter. If the bulk of the friction increase was front-loaded in the early 2000s, the pre-recession period would naturally exhibit a somewhat stronger relationship—not because the mechanism weakened, but because the economy had already transitioned to a higher-friction regime where marginal increases in γ have smaller effects on both the JRR and the labor share.

3.4 Robustness and Alternative Specifications

Having established the core relationship and investigated its heterogeneity, we now subject the results to alternative specifications that probe the mechanism more deeply.

Table 5 reports the key checks. The *excess JRR* specification—which removes the aggregate (year-level) component of JRR variation—yields a significant coefficient of 0.031. The idiosyncratic, cell-specific component of JRR variation is linked to the labor share, consistent with a micro-level friction channel rather than a common macro trend driving the correlation.

The *lagged JRR* specification (0.030, $p < 0.01$) helps address reverse causality—the possibility that labor share movements drive reallocation rather than vice versa. If declining labor shares caused firms to restructure, the contemporaneous JRR would be endogenous, but the lagged JRR would be predetermined. The significant lagged coefficient is consistent with, though it does not establish, the direction of causation emphasized by our model.

Decomposing the JRR into its job creation and job destruction components reveals an asymmetry: the job destruction rate (0.033, $p < 0.01$) accounts for nearly all of the baseline relationship, while the job creation rate (0.012, $p < 0.10$) contributes modestly. This asymmetry is consistent with the adjustment-cost mechanism operating primarily through the destruction margin—firing costs, severance obligations, and regulatory compliance impose larger fixed costs on separations than on hiring, making the reallocation-labor share link predominantly a story about barriers to workforce downsizing.

The *entry/exit decomposition* is equally informative. The BDS reports job flows from firm births and deaths separately from continuing firms, allowing us to decompose the JRR into an entry/exit component (births + deaths) and a continuing-firm component (expansions + contractions of surviving firms). Both components are significant: the entry/exit margin (0.027, $p < 0.01$) and the continuing-firm margin (0.019, $p < 0.05$). That both contribute is natural at the firm level, where "continuing firm" adjustments may involve restructuring across establishments—but the larger entry/exit coefficient indicates that firm turnover is the primary channel through which reallocation dynamics are linked to the labor share.

Finally, we verify that our results are not driven by the inclusion of employer-provided supplements (such as health insurance) in the compensation numerator. Replacing the compensation share with a *wages-only share*—wages and salaries divided by GDP, excluding all supplements—yields a panel coefficient that is essentially unchanged (Appendix Table B.6 in Appendix B.6), confirming that the mechanism operates through the wage channel, not through mechanical variation in benefits costs.

3.5 Discussion

The empirical results tell a coherent story. The JRR–labor-share relationship passes a demanding set of diagnostic tests that help distinguish a labor-friction interpretation from the leading alternatives.

First, it survives stringent fixed effects and controls. State-industry and year FE absorb permanent cross-cell differences and common aggregate shocks. The JRR coefficient is not absorbed by adding HHI, unionization, and manufacturing controls—the primary confounders implied by the superstar-firm, declining-bargaining-power, and deindustrialization hypotheses. Second, the relationship is larger in lower-concentration markets than in higher-concentration markets, consistent with a labor-friction channel operating alongside rather than being subsumed by market-power forces. Third, the relationship is most clearly estimated in services and strengthens significantly with routine task intensity, consistent with the friction channel operating more strongly where headcount adjustment is the primary margin. Fourth, the JC/JD decomposition reveals that job destruction drives the relationship, consistent with firing costs being an especially important non-smooth friction. Fifth, the entry/exit decomposition shows that firm turnover is the dominant channel, though continuing-firm adjustments also contribute—and replacing total compensation with wages-only leaves the panel coefficient unchanged, ruling out a mechanical supplements-driven correlation.

These panel results establish *that* the JRR is robustly linked to the labor share. They do not, however, reveal *what has caused the JRR to decline*. If rising adjustment frictions are responsible, they should leave a specific micro-level fingerprint: a growing mass of establishments that forgo employment adjustment altogether. We investigate this prediction directly in the next section, turning to establishment-level data that allow us to observe the adjustment decision at the individual-unit level and to bridge that micro margin back to the state-industry panel.

We acknowledge an important limitation: these regressions document robust conditional correlations, not causal effects. The heterogeneity, decomposition, and bridge results sharpen the interpretation of those correlations, but they do not by themselves identify exogenous variation in adjustment frictions. We therefore use the structural model in Section 5 to quantify a disciplined mechanism, not to convert the reduced-form regressions into causal estimates.

4 The Rise of Inaction and Sources of Rising Adjustment Costs

The preceding section established a robust link between the JRR and the labor share. We now ask a different question: what has caused the JRR itself to decline? For this analysis, we turn to the BLS BED, which provides establishment-level size-of-change data that are not available in the BDS.[3] This micro-level anatomy of the JRR decline raises two immediate questions. First, is there direct evidence that more establishments are forgoing adjustment altogether? Second, can we identify specific economic forces that would cause adjustment to become costlier? We address each in turn.

[3]The BDS reports aggregate job flows by firm demographic category (births, deaths, expansions, contractions) but does not report the distribution of employment changes across individual units. The BED, by contrast, tabulates the number of establishments in each size-of-change category (e.g., 1–4 jobs, 5–19 jobs, 20+ jobs), allowing us to observe the adjustment decision at the establishment level. Because approximately 95% of firms in the BDS are single-establishment enterprises, the firm-level JRR used in our panel regressions closely tracks establishment-level reallocation; Appendix B.4 confirms that the panel relationship is robust when using the BED's establishment-level JRR instead.

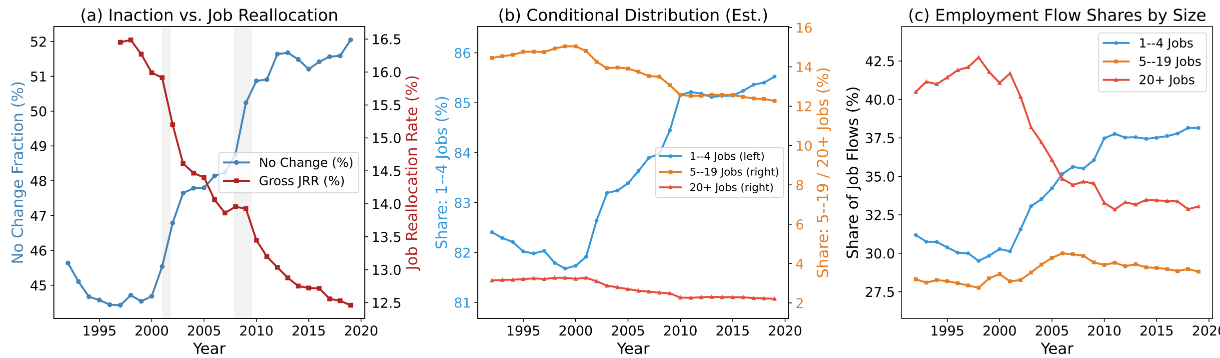

Figure 3: The Rise of Inaction and the Shift in the Size-of-Change Distribution, 1992–2019

Notes: Panel (a) overlays the inaction rate—the fraction of establishments reporting zero net employment change (left axis)—with the average gross job reallocation rate (right axis). Panel (b) shows the conditional distribution of changing establishments by size of change; the small-adjustment share (1–4 jobs, left axis) has risen while the medium (5–19) and large (20+) shares (right axis) have declined. Panel (c) shows the share of total employment flows (job gains plus losses) accounted for by each size category. All series are annual averages of quarterly data. Shaded areas indicate NBER recession dates.
Source: Bureau of Labor Statistics, Business Employment Dynamics.

4.1 The Rise of Inaction

The aggregate JRR can decline through two channels: establishments may adjust less *frequently*—a rising share of units reporting zero employment change—or they may adjust by smaller amounts when they do adjust. We refer to the first as the *inaction margin* and the second as the *size-of-change margin*. Note that these are conceptually distinct from the entry/exit decomposition in Section 3.4, which partitions the JRR by firm demographic status (births and deaths versus continuing firms); here, we examine the adjustment decision at the individual-establishment level. A detailed examination of the BED data reveals that the inaction margin is overwhelmingly the primary force behind the JRR decline.

The inaction margin: more establishments doing nothing. Figure 3(a) plots the inaction rate—the fraction of private-sector establishments reporting zero net employment change—against the aggregate job reallocation rate, revealing a striking mirror-image pattern. The inaction rate was approximately 44% in the late 1990s and has risen steadily, reaching 52% by 2019. The timing of this increase—beginning around 2000 and accelerating through the 2000s—coincides precisely with the acceleration of both the labor share and JRR declines documented in Figure 1. As the fraction of inactive establishments rises, the JRR falls; this negative co-movement is near-perfect in the pre-COVID period.

The size-of-change margin: smaller adjustments among changers. Panels (b) and (c) of Figure 3 examine the conditional distributions among establishments that *do* adjust. Panel (b) shows the distribution of changing establishments by size of change: among firms that adjust, the share making small adjustments (1–4 jobs) has risen from 82% to 86% (left axis), while the share making medium (5–19 jobs) and large adjustments (20 or more jobs) has fallen (right axis). Panel (c) shows the corresponding employment flow shares: the fraction of total job flows accounted for by small-change establishments (1–4 jobs) has risen from approximately 30% to 38%, while the share from large-change establishments has declined.

4.2 Bridging Establishment Inaction to the BDS Panel

The national BED size-of-change data in Figure 3 establish the aggregate rise of establishment-level inaction, but our main panel regressions are estimated on state-industry cells using firm-level BDS job flows. To connect these two pieces of evidence more directly, we construct a public BED-based state-industry inaction proxy from the BED all-items establishment rate file. Specifically, for each state-sector-quarter, we compute the fraction of establishments with no employment change as 100 – gain rate – loss rate, and then average these quarterly rates to the annual frequency. At the national level, this rate-based proxy tracks the size-of-change no-change series extremely closely (correlation 0.99 over 1992–2024), giving confidence that it captures the same underlying margin of adjustment.

Figure 4 shows the basic bridge relationship: state-industry cells with larger increases in establishment-level inaction experienced larger declines in firm-level job reallocation. Table 6 formalizes this connection in both panel fixed-effects regressions and long-difference regressions.

The bridge results sharpen the interpretation of the main panel evidence in three ways. First, the BED inaction proxy is strongly negatively associated with the BDS JRR within the same state-industry cells: in panel FE regressions, a 10% increase in establishment inaction is associated with a 4.5% decline in total reallocation. Second, this relationship is present for both continuing-firm and birth/death margins, but it is quantitatively larger for the extensive margin. In the panel FE specification, the corresponding elasticities are -0.335 for continuing-firm reallocation and -0.696 for births plus deaths. This pattern is important for interpretation: establishment no-change in BED does not map one-for-one into the continuing-firm component of BDS job flows. Rather, the two datasets appear to capture related manifestations of a broader slowdown in employment adjustment, one that affects both within-firm changes and firm turnover. Third, the bridge extends to the labor share itself. Cells with larger inaction increases also experienced larger payroll-share declines, and in a horse-race specification both the BDS JRR and BED inaction

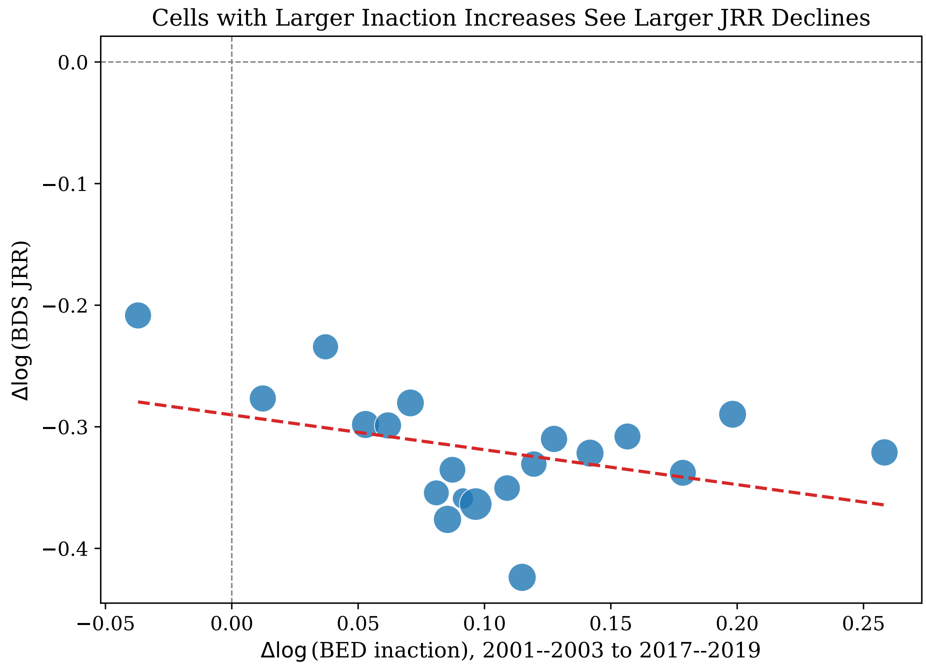

Figure 4: State-Industry Changes in BED Inaction and BDS Reallocation

Notes: Each point is a GDP-weighted bin of state-industry cells. The x-axis is the change in log BED inaction between 2001–2003 and 2017–2019; the y-axis is the corresponding change in log BDS gross job reallocation. The dashed line is the weighted linear fit across all state-industry cells.
Source: Bureau of Labor Statistics, Business Employment Dynamics; Census Bureau, Business Dynamics Statistics; Bureau of Economic Analysis.

remain significant. The long-difference results point in the same direction, although there the continuing-firm margin is imprecisely estimated while the extensive-margin and payroll-share bridges remain strong.

The bridge is not uniform across the sample. Appendix Table B.5 shows that the continuing-firm bridge is strongest in lower-concentration cells and in the post-2008 period, while the bridge from inaction to payroll shares is strongest in services and especially pronounced before the Great Recession. We therefore interpret the BED and BDS evidence as complementary rather than mechanically identical: the BED inaction series disciplines the micro margin of adjustment most directly, while the BDS panel shows where broader reallocation declines coincide with lower labor shares.

4.3 Implications for Modeling

Taken together, these facts reveal that since 2000, establishments have adjusted their employment less frequently and, when they do adjust, they do so by smaller amounts. This specific pattern has sharp implications for the class of models that can credibly explain the data.

Table 6: Bridging BED Inaction to BDS Reallocation and Payroll Shares

	(1)	(2)	(3)	(4)	(5)
Panel A: State-industry panel fixed effects, 2001–2019					
Dependent variable	log JRR	log intensive JRR	log extensive JRR	log payroll share	log payroll share
log(BED inaction)	-0.449***	-0.335***	-0.696***	-0.146***	-0.125***
	(0.052)	(0.044)	(0.088)	(0.040)	(0.040)
log(BDS JRR)					0.047***
					(0.007)
Entity FE	O	O	O	O	O
Year FE	O	O	O	O	O
N	9473	9473	9473	9473	9473
R^2 (within)	0.139	0.098	0.125	0.063	0.105
Panel B: State-industry long differences, 2001–2003 to 2017–2019					
Dependent variable	Δ log JRR	Δ log intensive JRR	Δ log extensive JRR	Δ log payroll share	Δ log payroll share
Δ log(BED inaction)	-0.287***	-0.087	-0.749***	-0.267***	-0.248***
	(0.094)	(0.091)	(0.170)	(0.094)	(0.094)
Δ log(BDS JRR)					0.065
					(0.064)
N	499	499	499	499	499
R^2	0.033	0.003	0.067	0.027	0.030

Notes: The BED inaction measure is constructed from establishment gain and loss rates in the BED all-items file as 100 – gain rate – loss rate. Panel A uses state-industry and year fixed effects with standard errors clustered by state-industry cell. Panel B uses state-industry long differences between 2001–2003 and 2017–2019, weighted by initial-period cell GDP shares, with HC1 robust standard errors. The intensive JRR is the continuing-firm margin from BDS; the extensive JRR is the births-plus-deaths margin. $^{***} p < 0.01$, $^{**} p < 0.05$, $^{*} p < 0.10$.

Ruling out convex costs. The existence of a large and *growing* mass of establishments at the no-change point is fundamentally inconsistent with models that assume convex (*e.g.*, quadratic) adjustment costs. In such models, the marginal cost of the first unit of adjustment is zero, meaning firms find it optimal to adjust continuously in response to every shock. These models predict zero mass at the no-change point. The large and growing inaction mass points directly to *non-smooth* adjustment costs—whether fixed costs or kinked linear per-unit costs—which create a band of inaction. Within this band, firms find it optimal to forgo adjustment because the expected gain is smaller than the cost.

Ruling out fixed costs. The decline in large adjustments (Figure 3(b)–(c)) further disciplines the choice among non-smooth cost structures. A pure fixed cost would encourage lumpier, less frequent but larger adjustments once the threshold is crossed. A *linear* cost, by contrast, makes very large adjustments progressively more expensive, consistent with the observed shift toward smaller adjustments. The data thus favor a linear per-unit adjustment cost.

We formalize these arguments through model counterfactuals in Section 6.2 (Table 12). Holding the calibrated benchmark fixed, we introduce each alternative cost structure—pure fixed, pure quadratic, and linear—each calibrated to match the same 2019 JRR target. The calibrated model suggests that the linear specification is the only one that remains broadly consistent with the rise of inaction and with a decline in conditional large adjustments, although it does not reproduce the increase in the conditional small-adjustment share. We view that miss as a limitation of the one-parameter counterfactual, not as evidence in favor of returning to a zero-friction benchmark.

The inaction evidence thus pins down both the *type* of friction (kinked, specifically linear) and the *direction* of change (increasing over time). But what concrete economic forces could be driving this increase?

4.4 Empirical Support for Rising Adjustment Costs

The preceding subsection established that the micro-level evidence points toward rising kinked linear per-unit adjustment costs as a plausible explanation for declining job reallocation. We now assemble independent empirical evidence—from the structural estimation literature, from institutional trends, and from our own data—supporting the premise that per-worker adjustment costs have indeed risen secularly in the United States.

Direct structural evidence. Decker et al. (2020) provide especially useful evidence. Using Census Longitudinal Business Database microdata, they decompose the post-2000 decline in business dynamism into two components: declining dispersion of productivity shocks ("less to reallocate") versus declining responsiveness to given shocks ("harder to reallocate"). They find that the decline is overwhelmingly driven by the latter: within-industry dispersion of total factor productivity has actually *risen* since 2000, while the marginal responsiveness of establishment-level employment growth to productivity shocks has fallen to roughly half its 1990s level. This pattern is a direct empirical analog of rising per-worker adjustment costs: conditional on receiving the same productivity draw, firms are behaving as though adjustment has become more costly.

Institutional sources of rising friction. Several specific institutional channels have grown in magnitude and could contribute to rising adjustment costs.

Regulatory compliance costs. Trebbi and Zhang (2022) estimate that the average U.S. firm spends between 1.3 and 3.3 percent of its total wage bill on regulatory compliance, and that these costs have grown by approximately 1 percent per year in real terms from 2002 to 2014. Compliance costs are largely labor-intensive and are disproportionately incurred when firms

Table 7: Healthcare Costs, Job Reallocation, and Labor Share

	(1)	(2)	(3)	(4)
Dependent variables:	log(LS) Baseline	log(LS) +HC Cost	log(LS) HC Only	log(JRR) Reduced Form
log(JRR)	0.044***	0.042***		
	(0.008)	(0.007)		
log(Premium)		-0.128***	-0.153***	-0.590***
		(0.011)	(0.011)	(0.009)
State-Ind FE	Yes	Yes	Yes	Yes
Year FE	Yes	No	No	No
Observations	10631	10631	10631	10631
R^2 (within)	0.050	0.150	0.144	0.410

Notes: Column (1) includes state-industry and year fixed effects. Columns (2)–(4) include only state-industry fixed effects; year fixed effects are omitted because the premium variable is constructed from a national time trend interacted with a state cost index (see Appendix A.8). The premium measure is the real (2019$) employer contribution to family health insurance, constructed as the KFF national trend multiplied by the MEPS-IC state cost index. Column (4) regresses the log JRR on the log premium to show the reduced-form association between healthcare costs and reallocation. Standard errors are clustered at the state-industry level and reported in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

change the scale or composition of their workforce, making them functionally equivalent to an adjustment cost.

Employer-provided health insurance. Employer health insurance costs per worker have risen sharply in real terms—the average employer contribution to a family premium grew from approximately $7,400 to $14,300 (in 2019 dollars) between 2001 and 2019 (see Appendix A.8). Because these benefits are largely per-worker (as opposed to per-hour), they raise the cost of adding a worker and can create a disincentive to adjust the headcount margin, particularly for moderate employment changes.[4] Exploiting state-level variation in healthcare cost indices from the MEPS Insurance Component, we find a strong negative association between employer premium levels and the JRR (Table 7): conditional on entity fixed effects, a 10% increase in real employer premiums is associated with a 5.9% decline in the JRR ($p < 0.01$). Because the premium measure is identified from a national trend interacted with cross-state cost differences, this result should be interpreted as institutional consistency rather than as causal identification; omitting year fixed effects leaves open the possibility that other aggregate trends co-move with premiums and reallocation. As we discuss in Section 6.2, a back-of-envelope comparison nonetheless indicates that the real increase in per-worker healthcare costs over 2001–2019 is of

[4]Employer supplements are included in our compensation numerator (SAGDP4), so rising healthcare costs mechanically *raise* it. Our channel operates through the denominator: higher per-worker costs widen the inaction band, reduce allocative efficiency, and lower output Y. The model's labor share (wN/Y) excludes supplements entirely, so the model-implied decline is conservative. Appendix B.6 confirms this: a wages-only share (excluding all supplements) yields essentially unchanged panel coefficients.

the same order of magnitude as the model-implied adjustment cost, suggesting that healthcare is a plausible component of the aggregate friction.

Evidence from our data. Our own micro-level findings are consistent with a secular increase in the effective adjustment cost. The rise in the inaction rate (Section 4.1) documents a growing mass of establishments reporting exactly zero employment change—a pattern that is the hallmark of non-smooth adjustment costs in the structural estimation literature (Cooper et al., 2015, Lee and Mukoyama, 2018). Moreover, the size-of-change data show that, conditional on adjusting, establishments are making smaller changes. Both patterns—more inaction and smaller adjustments—are consistent with an increase in labor adjustment costs over time, as in Cooper et al. (2024). The new bridge analysis in Section 4.2 shows that this establishment-level inaction margin is not merely an aggregate time-series fact: the same state-industry cells with larger increases in BED inaction also experienced larger declines in BDS reallocation and payroll shares. The rolling-window panel regressions (Figure 2) reinforce this interpretation: the JRR–labor-share relationship is strongest in the early 2000s, when the rise of inaction was steepest, consistent with adjustment costs being the binding friction during this period.

Taken together, the evidence from this section and the preceding empirical analysis forms a coherent case: (i) the labor share co-moves with the JRR in panel data, with firm entry and exit as the dominant BDS margin; (ii) at the establishment level, a growing mass of units are forgoing adjustment altogether, driving the aggregate JRR decline; (iii) the new bridge analysis shows that the same state-industry cells with larger inaction increases also experienced larger declines in BDS reallocation and payroll shares, even though the BED and BDS margins are not mechanically identical; (iv) the pattern of rising inaction and smaller conditional adjustments is the signature of kinked, specifically linear, adjustment costs; and (v) concrete institutional frictions—healthcare costs and regulatory compliance—have risen in magnitude over precisely the period when inaction increased. We now build a structural model disciplined by these facts to quantify how much of the labor share decline can be attributed to rising adjustment frictions.

5 Model

The micro-level evidence in Section 4 places a sharp restriction on the class of models that can credibly explain the data. The growing mass of inactive establishments is inconsistent with convex adjustment cost models, which predict continuous adjustment, and points directly to non-smooth costs. The decline in large conditional adjustments further disciplines the functional form: linear per-unit costs fit this margin much better than fixed costs. We therefore build a dynamic general equilibrium model incorporating this specific friction.

We extend the standard firm dynamics framework of Hopenhayn and Rogerson (1993) by incorporating a kinked linear labor adjustment cost. This feature endogenously creates a band of inaction where firms optimally choose not to adjust employment, matching the empirical pattern. We use this framework to test whether an increase in this friction can simultaneously generate both the observed decline in job reallocation and the fall in the labor share, while being transparent about which intensive-margin moments the one-parameter counterfactual does and does not match.

5.1 Environment

The economy comprises a continuum of heterogeneous firms (plants) and a representative household. Firms use labor as their sole input to produce a homogeneous good. Firm-level dynamics are governed by idiosyncratic productivity shocks and endogenous entry and exit decisions.

Incumbent Firms. An incumbent firm begins period t with state (s_{t-1}, n_{t-1}), where s_{t-1} is its previous-period productivity and n_{t-1} is its previous employment. Its value is $W(s_{t-1}, n_{t-1})$.

Exit/Stay Decision. The firm first draws a stochastic exit value x_t from distribution $\xi(\cdot)$ and then decides whether to exit. If it exits, it collects x_t but must pay a dismissal cost $h(0, n_{t-1}) = \gamma n_{t-1}$ to release its workforce. If it stays, it observes its new productivity s_t and proceeds to the adjustment decision. The integrated value is:

$$W(s_{t-1}, n_{t-1}) = \int \max\big(x_t - h(0, n_{t-1}),\ \mathbb{E}[V^c(s_t, n_{t-1})|s_{t-1}]\big)\, d\xi(x_t)$$

Adjust/Inaction Decision. A firm that stays, now knowing its state (s_t, n_{t-1}), must decide whether to adjust its labor force or do nothing:

$$V^c(s_t, n_{t-1}) = \max\Big\langle \underbrace{V^a(s_t, n_{t-1})}_{\text{Adjust}},\ \underbrace{V^n(s_t, n_{t-1})}_{\text{Do Nothing}} \Big\rangle$$

This explicit choice is the core mechanism. If the productivity shock s_t is not large enough to warrant paying the adjustment cost, the firm optimally chooses inaction (V^n), setting its net employment change to zero. This maps directly to the rising mass of "no-change" establishments in Figure 3.

If the firm does not adjust:

$$V^n(s_t, n_{t-1}) = f(s_t, n_{t-1}) - w_t n_{t-1} + \beta W(s_t, n_{t-1})$$

If the firm adjusts to a new labor level n_t:

$$V^a(s_t, n_{t-1}) = \max_{n_t} \{ f(s_t, n_t) - w_t n_t - h(n_t, n_{t-1}) + \beta W(s_t, n_t) \}$$

where the continuation value W already integrates over the exit draw (as defined above), and the expectation over next-period productivity s_{t+1} is embedded recursively through the fixed-point structure of W.

The adjustment cost $h(n_t, n_{t-1}) = \gamma |n_t - n_{t-1}|$ is a kinked linear function—the key theoretical instrument. As motivated by the size-of-change evidence (Figure 3), a linear cost is preferred over a fixed cost because it makes very large adjustments progressively more expensive, consistent with the observed decline in the intensive margin of adjustment.

Potential Entrants. In any period, a mass of potential entrants can enter by paying a sunk entry cost c_e. Upon entry, a firm starts with zero prior employment and draws its initial productivity from distribution $\nu(s_t)$. The value of entry is:

$$V^e = \int V^c(s_t, 0)\, d\nu(s_t)$$

Free entry ensures $V^e = c_e$, endogenously determining the mass of entrants.

Households. A representative household owns all firms, supplies labor, and consumes. Following Hopenhayn and Rogerson (1993), utility is linear in consumption ($u(C_t) = C_t$), so the discount rate for firms is β. The household's labor supply is determined by:

$$v'(N_t) = w_t$$

where $v(\cdot)$ is the disutility of labor.

5.2 Stationary Equilibrium

A stationary equilibrium is characterized by a wage w^*, a mass of entrants M^*, and a distribution of incumbents $\mu^*(s, n)$ such that: (i) firms optimize given w^*; (ii) free entry holds; and (iii) the labor market clears:

$$L^s(w^*) = \int \phi(s, n; w^*)\, d\mu^*(s, n) + M^* \int \phi(s, 0; w^*)\, d\nu(s)$$

The free-entry condition pins down w^*; the labor-supply curve determines L^*; and the market-clearing condition determines M^*.

Table 8: Model Calibration

Parameter	Value	Description
Panel A: Externally Set		
β	0.94	Discount factor
ρ	0.934	Productivity persistence
Panel B: Internally Calibrated (Benchmark)		
θ	0.623	Labor elasticity of output
σ_ϵ	0.233	Shock volatility
a	0.032	Productivity drift
c_e	646.0	Sunk entry cost
x_0	0.771	Exit shock mass at zero
γ_0	0.132	Benchmark adjustment cost
Panel C: Counterfactual		
γ_1	0.192	Counterfactual adjustment cost

Notes: The model period is one year. Productivity persistence $\rho = 0.934$ is set externally, consistent with annual plant-level TFP estimates in Cooper and Haltiwanger (2006). The six benchmark parameters (θ, σ_ϵ, a, c_e, x_0, γ_0) are jointly calibrated via multi-start Nelder-Mead to minimize a weighted loss over six 2001-based moments: labor share (57.1%), JRR (31.1%), entry rate (11.0%), average establishment size (18.0), average entrant size (5.6), and the share of firms inside a positive-width (S, s) inaction band (45.5%, from the BED quarterly no-change rate), yielding an exactly identified system. The counterfactual γ_1 is then chosen so that the model's inaction band share matches the 2019 BED no-change rate (52.0%). All other parameters are held fixed between the benchmark and counterfactual; the resulting changes in the labor share and job reallocation rate are untargeted model predictions.

5.3 Calibration

We calibrate the model in two stages. In the first stage, we calibrate the *benchmark* economy to match the U.S. economy in 2001, the initial year of our empirical panel. Crucially, the benchmark has a *positive* adjustment cost ($\gamma > 0$), so that the baseline economy already features a realistic inaction region. We target six moments, including a new inaction moment: the share of firms in the stationary distribution whose state (s, n) falls inside a positive-width (S, s) band. Because the model period is annual while the BED no-change fact is quarterly, we use the 2001 BED quarterly no-change rate (45.5%) as an empirical proxy for the prevalence of latent inaction rather than as a literal annual no-adjustment probability. In the second stage, we increase γ to match the corresponding 2019 BED quarterly no-change rate (52.0%). The resulting changes in the labor share and job reallocation rate emerge as untargeted model predictions, providing a disciplined test of the adjustment-cost mechanism.

Production is Cobb-Douglas: $f(s, n) = s\, n^{\theta}$, where θ governs the labor elasticity of output. In a frictionless economy, θ equals the labor share. The idiosyncratic productivity process follows a log-AR(1): $\log(s_t) = a + \rho \log(s_{t-1}) + \epsilon_t$, with $\epsilon_t \sim N(0, \sigma_\epsilon^2)$, discretized on a 50-point grid using the Tauchen method. The labor choice set is discretized on an 80-point grid with geometric spacing to capture the heavy right tail of the firm size distribution. The exit value distribution $\xi(\cdot)$ follows a uniform distribution on $[0, \bar{x}]$, with a point mass x_0 at zero that governs exit rates; x_0 close to zero implies more exit value heterogeneity and lower average exit rates.

One period in the model is one year. The discount factor $\beta = 0.94$ is standard. Productivity persistence $\rho = 0.934$ is set externally, consistent with the annual plant-level TFP estimates in Cooper and Haltiwanger (2006). Six parameters (θ, σ_ϵ, a, c_e, x_0, γ_0) are jointly calibrated to match six 2001-based moments: the labor share (57.1%), the gross job reallocation rate (31.1%), the establishment entry rate (11.0%), average establishment size (18.0 employees, from SUSB), average entrant size (5.6 employees), and the share of firms inside the (S, s) inaction band (45.5%, from the BED quarterly no-change rate). Targets are drawn from BEA (labor share), Census BDS (JRR, entry rate, establishment size, entrant size), and BLS BED (inaction share). For the counterfactual, the 2019 BED no-change rate of 52.0% is the sole target; the observed 2019 labor share of 53.2% and JRR of 22.4% are not targeted—they are the outcomes we wish to explain. We use a multi-start Nelder-Mead algorithm with five initial parameter vectors, evaluating each candidate at the general equilibrium wage w^* that satisfies free entry ($V^e = c_e$) via bisection, and selecting the combination that minimizes a weighted loss function over the six targeted moments. Table 8 reports the calibrated parameter values.

Several features of the calibrated benchmark merit comment. With persistence set externally at $\rho = 0.934$ following Cooper and Haltiwanger (2006), the calibrated shock volatility ($\sigma_\epsilon = 0.233$) is pinned down by the need to match the gross job reallocation rate in the BDS data (31.1% in 2001): reproducing that level of churn requires sufficiently volatile idiosyncratic productivity shocks. The benchmark adjustment cost ($\gamma_0 = 0.132$) ensures that the baseline economy already features a realistic inaction region, with 44.9% of firms lying inside the stationary inaction band—closely matching the 2001 BED quarterly no-change target. The exit shock mass at zero ($x_0 = 0.771$) implies limited heterogeneity in exit values, generating realistic exit rates and average firm sizes. The entry cost ($c_e = 646$, in units of output, the numeraire) reflects the substantial sunk costs—licensing, regulatory compliance, physical plant, organizational setup—that new establishments must incur.

Panel A of Table 9 shows that the benchmark ($\gamma_0 = 0.132$) closely matches all six targeted moments. The model labor share is 56.9% (data: 57.1%), the JRR is 30.8% (data: 31.1%), the entry rate is 11.2% (data: 11.0%), average establishment size is 17.0 (data: 18.0), average entrant size is 5.4 (data: 5.6), and the inaction band share is 44.9% (data: 45.5%). All targets are matched within 6%. Panel B reports untargeted moments. The mean (S, s) band width is 18 employees, indicating that the benchmark economy already features substantial inaction regions. The model's

Table 9: Model Fit: Targeted and Untargeted Moments

Moment	Model ($\gamma_0 = 0.132$)	Data (2001)
Panel A: Targeted Moments		
Labor share	0.569	0.571
Job reallocation rate	0.308	0.311
Entry rate	0.112	0.110
Avg. establishment size	17.0	18.0
Avg. entrant size	5.4	5.6
Inaction band share	0.449	0.455
Panel B: Untargeted Moments		
Mean (S, s) band width (empl.)	18	—
Cond. share: small (1–4 jobs)	53.3	82.4
Cond. share: large (20+ jobs)	13.9	3.1

Notes: Panel A reports the six moments targeted in the benchmark calibration ($\gamma_0 = 0.132$). All targets are matched within 6%. Panel B reports untargeted moments. The mean (S, s) band width measures the average range of employment levels (in number of employees) over which firms optimally choose inaction, weighted by the stationary distribution. The inaction band share is the fraction of incumbent firms in the stationary distribution whose state (s, n) falls inside a positive-width inaction band, matched to the 2001 BED quarterly no-change rate (45.5%), interpreted as a proxy for latent inaction in the annual model. The model's conditional size-of-change distribution differs from the data in levels, reflecting the coarser employment grid and absence of establishment-level heterogeneity. Labor share is from BEA; JRR, entry rate, and establishment size are from Census BDS; inaction and size-of-change data are from BLS BED.

conditional size-of-change distribution differs from the data in levels—the model predicts a smaller share of 1–4 job adjustments (53.3% vs. 82.4%) and a larger share of 20+ job adjustments (13.9% vs. 3.1%)—reflecting the coarser employment grid and absence of establishment-level heterogeneity beyond productivity. We therefore treat the model's strongest quantitative discipline as coming from the inaction target and the aggregate counterfactual moments, while viewing the size-of-change evidence as qualitative guidance about the relevant class of frictions. In particular, we do not revert to a zero-friction benchmark simply because it can improve the sign of one conditional-share response: the positive-γ_0 benchmark is preferred because the 2001 economy already exhibits substantial inaction.

6 Quantitative Results

We now use the calibrated model to quantify the contribution of rising adjustment frictions to the labor share decline. The central exercise compares two stationary equilibria. The benchmark economy ($\gamma_0 = 0.132$) is disciplined by 2001 moments, including the prevalence of inaction in the stationary distribution. The counterfactual then raises the adjustment cost just enough to match the 2019 BED quarterly no-change rate through that same stationary inaction object. Because the model is annual while the BED moment is quarterly, the targeted object should be read as a proxy for the prevalence of latent inaction rather than as a literal annual no-adjustment probability. Crucially, both the labor share and the job reallocation rate are *untargeted* model predictions.

Holding all other structural parameters fixed, we increase γ from 0.132 to 0.192 so that the model's inaction-band share rises from 44.9% to 52.0%, matching the 2019 BED target. Table 10 compares the model-implied changes with the data. With this single change, the labor share falls from 56.9% to 55.1%, a decline of 1.8 percentage points—about 45% of the observed 3.9 pp decline. The job reallocation rate falls from 30.8% to 27.4%, a decline of 3.4 pp—about 39% of the observed 8.6 pp decline. The quantitative message is therefore that a friction increase large enough to match the rise in inaction can account for a substantial, though far from complete, portion of both aggregate declines.

The equilibrium wage falls from 0.405 to 0.386 (a 4.7% decline), as firms' demand for labor at the margin is reduced by the prospect of higher future adjustment costs. The entry rate is essentially unchanged in the model (compared with a decline from 11.0% to 9.6% in the BDS data), because the counterfactual deliberately varies only γ. The model's limited entry response is therefore not a failure of the exercise so much as a reflection of scope: other forces behind declining business formation—rising startup costs, tighter credit conditions, and housing-market frictions (Davis and Haltiwanger, 2024)—operate through channels outside the model. Average establishment size rises from 17.0 to 17.6 employees as incumbents tolerate larger deviations from their frictionless targets inside the wider inaction region.

The value of the exercise is that a counterfactual targeted solely to the rise in inaction also produces sizeable declines in both the labor share and job reallocation. The model-implied labor share decline of 1.8 percentage points accounts for approximately 45 percent of the observed 3.9 percentage-point decline (from 57.1% to 53.2%) between 2001 and 2019. The model-implied JRR decline of 3.4 percentage points accounts for approximately 39 percent of the observed 8.6 percentage-point decline (from 31.1% to 22.4%). We view that magnitude as economically meaningful: it suggests that rising labor adjustment frictions can be an important part of the story even though they do not exhaust the full decline in either aggregate series. The same one-parameter counterfactual is less successful on the conditional composition of adjustments: in the model, the conditional small-change share falls slightly rather than rising as in the BED data.

Table 11 decomposes the model-implied labor-share decline into an exact Shapley wage channel and a residual allocation/productivity channel. The main quantitative force is the

Table 10: Counterfactual Analysis: Rising Adjustment Costs

	Model		Data	
Moment	γ: 0.132 → 0.192	Δ (pp)	2001 → 2019	Δ (pp)
Targeted moment				
Inaction band share	0.449 → 0.520	+7.1	0.455 → 0.520	+6.5
Untargeted predictions				
Labor share	0.569 → 0.551	−1.8	0.571 → 0.532	−3.9
Job reallocation rate	0.308 → 0.274	−3.4	0.311 → 0.224	−8.6
Cond. share: small (1–4)	53.3 → 52.1	−1.2	82.4 → 85.9	+3.5
Cond. share: large (20+)	13.9 → 13.9	0.0	3.1 → 2.2	−0.9
Entry rate	0.112 → 0.112	0.0	0.110 → 0.096	−1.4
Additional model outcomes				
Equilibrium wage (w^*)	0.405 → 0.386	−4.7%		
Avg. establishment size	17.0 → 17.6	+0.6		
Mean (S, s) band width (empl.)	18 → 26	+8		

Notes: The counterfactual holds all structural parameters fixed and increases γ from 0.132 to 0.192 to match the 2019 BED quarterly no-change rate (52.0%) through the model's stationary inaction-band share, which is the only targeted counterfactual moment. The labor share and job reallocation rate are untargeted model predictions. The model-implied labor share decline of 1.8 pp accounts for approximately 45% of the observed 3.9 pp decline. The model-implied JRR decline of 3.4 pp accounts for approximately 39% of the observed 8.6 pp decline. The mean (S, s) band width widens from 18 to 26 employees, confirming the expansion of inaction regions. Data: labor share from BEA; JRR and entry rate from Census BDS; inaction and size-of-change from BLS BED.

equilibrium wage decline: holding the counterfactual allocation fixed, the lower wage reduces the labor share by 2.68 percentage points. Changes in the stationary distribution and policy functions partially offset this effect, contributing +0.89 percentage points. Between the benchmark and counterfactual, the wage falls by 4.8%, labor input rises by 3.8%, and output rises by 2.2%. Output per worker declines by 1.6%, so the net allocation/productivity effect does not amplify the labor-share decline; instead it cushions part of the wage effect. Firm mass increases only modestly (+0.3%), confirming that entry and mass effects are quantitatively secondary relative to the wage channel.

Beyond the aggregate moments, the model delivers a transparent prediction about the inaction region itself. The mean (S, s) band width widens from 18 to 26 employees, indicating that the range of employment levels over which firms optimally forgo adjustment expands substantially as γ rises. This widening of the inaction band is the model counterpart of the empirical rise of inaction documented in Figure 3(a). The mechanism is intuitive: a higher

Table 11: Model Channel Decomposition of the Labor-Share Decline

	Benchmark	Counterfactual
Labor share	0.569	0.551
Wage	0.405	0.386
Labor input	151.78	157.55
Output	108.13	110.48
Firm mass	8.94	8.96
Exact Shapley decomposition (percentage points)		
Total labor-share change	-1.79	
Wage channel	-2.68	
Allocation / productivity channel	0.89	
Exact log decomposition		
Δ log wage	-0.048	
Δ log labor	0.037	
-Δ log output	-0.021	
Δ log output per worker	-0.016	
Δ log firm mass	0.003	

Notes: The table decomposes the model-implied labor-share decline between the benchmark ($\gamma_0 = 0.132$) and the counterfactual ($\gamma_1 = 0.192$) using the discrete-time annual model. The wage and allocation/productivity contributions are an exact two-component Shapley decomposition in levels: the wage channel holds the counterfactual allocation fixed and isolates the effect of the equilibrium wage decline, while the allocation/productivity channel captures the change in labor and output generated by the altered stationary distribution and policy functions. The lower panel reports the exact log-identity $\Delta \log LS = \Delta \log w + \Delta \log L - \Delta \log Y$ and supporting diagnostics. Firm mass is reported as a memo item; because both labor and output scale with firm mass, its effect on the labor share is indirect rather than mechanical.

per-unit adjustment cost makes it optimal for firms to tolerate larger deviations from their frictionless employment target before incurring the cost of adjustment.

This intensive-margin miss is economically interpretable. In a model that starts from $\gamma_0 = 0$, introducing a positive linear cost removes many marginal adjusters from the set of firms that change employment, and those marginal adjusters tend to be small changers; depending on the calibration, the conditional small-change share can therefore rise mechanically. In our preferred benchmark, however, the 2001 economy already has a positive adjustment cost and a well-developed inaction band. Raising γ further then pushes the remaining marginal adjusters at the edge of that existing band into inaction, so the small-adjustment margin is itself "eaten" by the expanding band. The resulting fall in the conditional small-change share is therefore a selection effect around a positive-friction benchmark, not a reason to abandon that benchmark. We nonetheless treat it as evidence that a one-parameter increase in linear adjustment costs is better suited to explaining the rise of inaction and the decline in large changes than the full conditional distribution of adjustment sizes.

6.1 The Mechanism in Detail

To understand why rising adjustment costs lower the labor share, it is useful to distinguish three conceptual channels, while keeping in mind that the quantitative decomposition in Table 11 indicates that the equilibrium wage channel is the dominant force in the current calibration.

Channel 1: Entry and firm-mass effects. The adjustment cost $h(n_t, n_{t-1}) = \gamma|n_t - n_{t-1}|$ is a pure resource cost—it does not flow to workers as compensation or to firm owners as profit. In principle, higher adjustment costs reduce the option value of entry and can alter the equilibrium mass of firms. In our calibration, however, this margin is quantitatively small: the entry rate falls only modestly and firm mass changes little. Its main importance is indirect, through general equilibrium effects on wages and the composition of active firms, rather than through a large standalone contribution to the labor-share decline.[5]

Channel 2: Allocative inefficiency (the "inaction wedge"). When $\gamma > 0$, firms that receive moderate productivity shocks find it optimal not to adjust employment. These firms operate at suboptimal labor levels: some have "too many" workers (negative productivity shocks that are not large enough to warrant paying γ per dismissed worker) and others have "too few" (positive shocks not large enough to justify the hiring cost). This inaction wedge changes both labor input and output. In our calibration, the net allocation/productivity component partially offsets the wage decline documented in Table 11. Thus, the inaction wedge remains economically important for matching the reallocation and inaction moments, but it is not the main quantitative source of the labor-share decline. The mean (S, s) band width widens from 18 to 26 employees between the benchmark and counterfactual, confirming the expansion of the inaction region as adjustment costs rise.

Channel 3: General equilibrium wage adjustment. The channels above operate at the firm level. In general equilibrium, the prospect of future adjustment costs reduces firms' expected returns to employment, lowering the demand for labor at every wage. The equilibrium wage w^* falls—from 0.405 to 0.386 in our calibration (a 4.7% decline). This wage decline is the general-equilibrium amplification of the underlying adjustment friction, and Table 11 shows that it is the

[5] The model's labor share is wN/Y, where $Y = \sum_i f(s_i, n_i)$ is gross output (before adjustment costs). Adjustment costs are pure resource destruction that does not appear in measured gross output. The labor share decline therefore operates through equilibrium changes in wages and the joint evolution of labor and output, rather than through a mechanical accounting change in Y.

dominant quantitative driver of the labor-share decline. Intuitively, once firms internalize that future employment changes are more costly, labor becomes less attractive at the margin across the entire stationary distribution, putting downward pressure on the equilibrium wage.

All three channels operate simultaneously, and their combined effect produces the 1.8 percentage-point decline in the labor share. Quantitatively, the wage channel dominates, while the allocation/productivity component partially offsets it and entry/mass effects remain modest. The mechanism nonetheless remains empirically grounded: it relies only on the premise that per-worker adjustment costs have risen, a claim supported by the evidence assembled in Section 4.4.

To confirm that the mechanism operates smoothly across the parameter space, we compute a sequence of stationary equilibria for a range of γ values, holding all other parameters at their benchmark values. Figure 5 plots the model-implied labor share, JRR, and inaction band share as functions of γ. The labor share and JRR both decline monotonically in γ, with the response approximately linear over the empirically relevant range, while the inaction band share rises monotonically. The labor share response exhibits diminishing sensitivity for large γ (the function is convex, $f'' > 0$, as the decline flattens), as the labor share approaches a floor determined by θ and the degree of misallocation. Average establishment size rises monotonically with γ, consistent with the secular increase in average establishment size observed in U.S. data.

The diminishing sensitivity of the labor share response is consistent with the temporal stability documented in Section 3.3. The modest difference between pre- and post-recession coefficients is exactly what this convex response function predicts: once the economy has transitioned to a higher-friction regime, further increases in γ produce smaller marginal effects on both the JRR and the labor share. The task-heterogeneity finding in Section 3.2—the significant JRR×RTI interaction—confirms that the mechanism is strongest in routine-task-intensive industries, where headcount adjustment costs bind most tightly.

6.2 Discussion

The model's distinctive value lies in its ability to connect three otherwise disconnected empirical phenomena—the decline in the labor share, the decline in JRR, and the rise of inaction—within a single quantitative exercise. The counterfactual targets only the 2019 inaction rate; the labor share and job reallocation declines emerge as model predictions. Product-market-power explanations can rationalize the labor share decline but have no direct prediction for the rise of inaction, and Albrecht and Decker (2026) show that rising markups do not even predict declining dynamism at the industry level; capital-deepening explanations are similarly silent about inaction dynamics. Quantitatively, the adjustment-cost mechanism accounts for approximately 45%

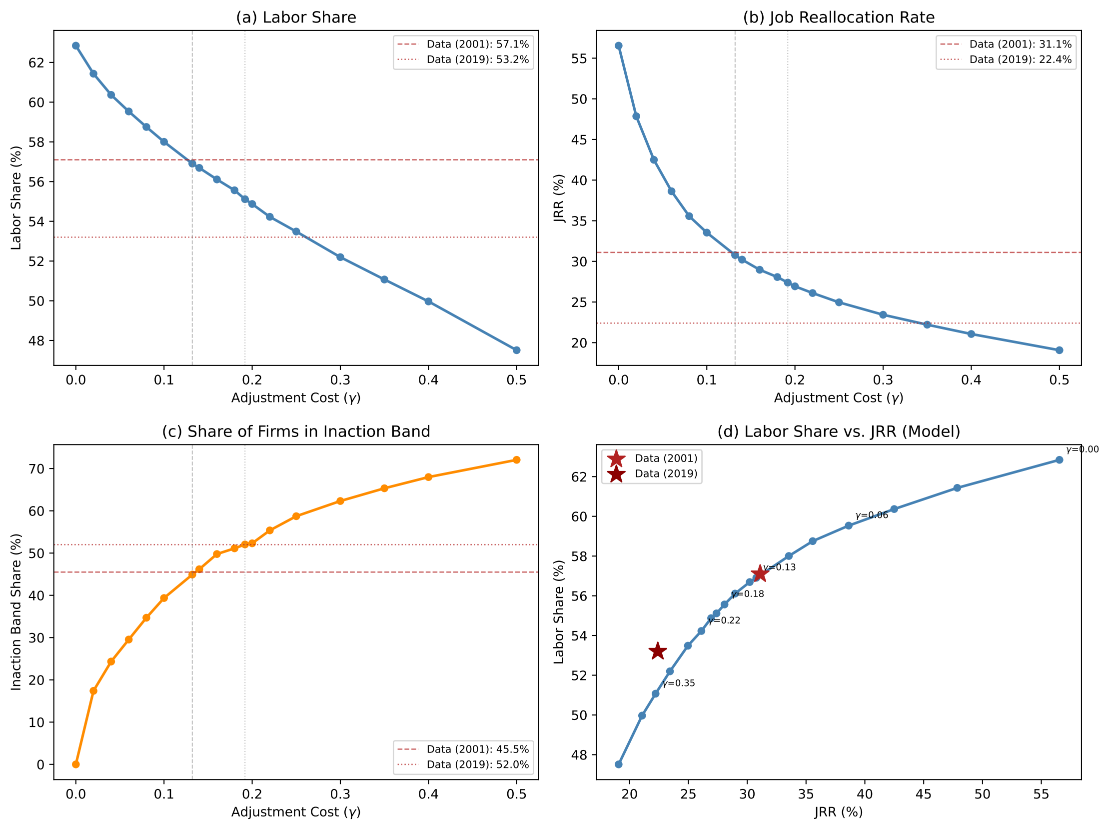

Figure 5: Comparative Statics: Model Outcomes as Functions of γ

Notes: Each point represents a stationary general equilibrium of the calibrated model at a given value of the linear adjustment cost γ. Panels (a)–(c) plot the labor share, job reallocation rate, and inaction band share as functions of γ; panel (d) traces the implied labor share–JRR relationship. Dashed horizontal lines indicate the 2001 and 2019 data values. Vertical dashed lines in panels (a)–(c) mark the benchmark $\gamma_0 = 0.132$ and the counterfactual $\gamma_1 = 0.192$. Stars in panel (d) mark the data values.

of the observed labor share decline and 39% of the observed JRR decline. The model features a single friction operating through one production factor (labor, not capital), abstracts from hours-margin adjustments, and is calibrated to match only one counterfactual moment (the 2019 BED no-change rate). We therefore read the result as evidence that rising labor adjustment frictions can explain a meaningful share of both aggregate declines, while leaving substantial room for rising market power (Autor et al., 2020, De Loecker et al., 2020), capital deepening (Karabarbounis and Neiman, 2014), globalization, and compositional shifts. Our finding identifies a complementary mechanism rather than a complete account.

A natural concern is whether the choice of linear per-unit adjustment costs is driving our results, or whether a fixed cost or quadratic cost could generate similar predictions. We address this through formal counterfactuals: holding all benchmark parameters fixed (including $\gamma_0 = 0.132$), we add each alternative cost structure on top and calibrate its magnitude to match the same JRR decline as the linear counterfactual. Table 12 reports the results. The mean (S, s) band width provides a sharp diagnostic. Adding a fixed cost ($F = 3.50$) on top of the baseline linear cost generates a much wider mean inaction band (48 employees, compared to 26 for the linear counterfactual), meaning firms accumulate large deviations before adjusting—and when they do, they make large discrete jumps to the target. This catastrophically fails on the size-of-change distribution: conditional on adjusting, the share of small (1–4 job) changes collapses to zero while the share of large (20+) changes rises to 72.3%. Adding a quadratic cost ($\phi = 0.0035$) fails differently: because the marginal cost of the first unit of adjustment is zero, the quadratic component does not create a well-defined inaction region.[6] By contrast, increasing the linear per-unit cost generates a moderate, well-defined expansion of the inaction band and moves the conditional large-adjustment share in the right direction. Its main remaining miss is that the conditional small-adjustment share falls slightly rather than rises. That miss is informative: under a positive-γ_0 benchmark, additional increases in γ remove the marginal small adjusters that remain at the edge of the existing inaction band. A zero-friction benchmark can reverse that sign, but we do not treat it as the preferred baseline because the 2001 economy already exhibited substantial inaction. For the purposes of this paper, the positive-γ_0 benchmark is therefore the more disciplined specification, and the linear-cost model's main quantitative strength is on inaction, large adjustments, and the aggregate labor-share and JRR responses.

The model's adjustment cost parameter γ is a reduced-form summary of many underlying forces. The institutional evidence in Section 4.4—rising regulatory compliance costs and expanding employer health insurance burdens—suggests that multiple specific mechanisms contribute to the aggregate increase in γ. A back-of-envelope calculation suggests that rising employer healthcare costs alone can account for a meaningful share of the implied increase in γ.[7]

7 Conclusion

This paper documents three connected facts about the U.S. economy since 2000: a decline in job reallocation, a rise in establishment-level inaction, and a positive association between payroll shares and reallocation across state-industry cells. We show that these patterns are

[6] The reported mean band width of approximately 5,000 employees under the quadratic specification spans the entire employment grid, indicating the absence of a meaningful inaction region. In continuous space, quadratic costs would eliminate inaction entirely; the residual model inaction is an artifact of the discrete employment grid.

[7] Real employer health premium contributions rose by roughly \$6,900 per family between 2001 and 2019 (KFF). Adjusting for coverage rates yields a per-worker increase of approximately \$2,760. The implied increase in γ is $\Delta\gamma = 0.192 - 0.132 = 0.060$. Since γ is in output units and $w^* \approx 0.386$, the implied per-worker increase in adjustment cost is $(\Delta\gamma / w^*) \times \$53{,}490 \approx \$8{,}312$, so healthcare accounts for roughly $2{,}760/8{,}312 \approx 33\%$ of the implied increase. This is an upper bound, as it equates a recurring annual cost with a one-time adjustment cost. To the extent that firms internalize the present value of future healthcare obligations upon hiring, the recurring annual cost functions as an effective adjustment cost at the margin of employment changes.

Table 12: Alternative Adjustment Cost Structures

	Benchmark ($\gamma_0 = 0.132$)	Linear ($\gamma_1 = 0.192$)	Fixed ($F = 3.50$)	Quadratic ($\phi = 0.0035$)
Panel A: Targeted moment				
Job reallocation rate	0.308	0.273	0.274	0.273
Panel B: Untargeted moments				
Labor share	0.569	0.551	0.567	0.550
		(−1.8 pp)	(−0.2 pp)	(−1.9 pp)
Cond. share: 1–4 jobs (%)	53.8	51.9	0.0	55.6
		(-2.0)	(-53.8)	(+1.8)
Cond. share: 20+ jobs (%)	15.0	13.6	72.3	11.2
		(**−1.4**)	(+57.3)	(−3.8)
Entry rate	0.112	0.112	0.111	0.112
Avg. establishment size	17.0	17.6	20.0	16.1
Mean (S, s) band width (empl.)	18	26	48	≈5,000
Panel C: Match with data				
Well-defined (S, s) band		✓	✓	×
Cond. small (1–4) rises		×	×	✓
Cond. large (20+) falls		✓	×	✓

Notes: Each counterfactual holds all structural parameters at their benchmark-calibrated values ($\gamma_0 = 0.132$) and introduces additional adjustment costs, calibrated to match the linear counterfactual's JRR (27.3%). All other rows are untargeted. The mean (S, s) band width measures the average range of employment levels over which firms optimally choose inaction, weighted by the stationary distribution. This table uses a finer employment grid ($N_n = 200$) than the baseline calibration ($N_n = 80$) to resolve the size-of-change bins more accurately; aggregate moments are virtually unchanged, while conditional shares may differ slightly from Table 10. Changes from benchmark are reported in parentheses; bold indicates the correct direction of change relative to 2001–2019 data trends.

consistent with rising labor adjustment frictions. The main panel relationship survives controls for concentration, unionization, and manufacturing exposure; the new bridge analysis shows that the same state-industry cells with larger increases in BED inaction also experienced larger declines in BDS reallocation and payroll shares; and a calibrated firm dynamics model with linear adjustment costs can account for approximately 45 percent of the observed labor share decline and 39 percent of the JRR decline when the counterfactual is disciplined only by the rise in inaction.

Our findings complement rather than compete with the leading explanations for the labor share decline. Rising labor adjustment frictions provide one way to connect declining responsiveness at the micro level to aggregate factor-share movements, and may therefore help explain why lower dynamism and a lower labor share have emerged together. In that sense, the paper points to a complementary macro mechanism rather than a substitute for markup-based, capital-deepening, globalization, or compositional accounts. The "declining responsiveness" finding of Decker et al. (2020)—that firm-level employment responses to productivity shocks have weakened—is the direct structural analog of our rising γ, while institutional trends in regulatory compliance costs and employer healthcare burdens provide concrete reasons why effective per-worker adjustment costs may have risen. The quantitative exercise does not claim that rising frictions are the entire explanation; its contribution is to show that a calibration disciplined by inaction can explain a meaningful share of both aggregate declines.

Several avenues for future research remain. Stronger instruments—perhaps from state-level policy reforms—are needed to establish a causal link from specific institutional frictions to the JRR–labor-share nexus. Extending the model to incorporate capital-labor substitution would allow a more complete assessment of whether capital deepening and rising adjustment costs are complementary or substitutable forces, and could amplify the quantitative contribution beyond the current 45%. Finally, the task-heterogeneity results suggest that distinguishing between headcount and hours adjustment margins would sharpen predictions for heterogeneous industries.

References

Acemoglu, Daron and Pascual Restrepo, "Tasks, Automation, and the Rise in US Wage Inequality," *Econometrica*, 2022, *90* (5), 1973–2016.

Akcigit, Ufuk and Sina T. Ates, "What Happened to US Business Dynamism?," *Journal of Political Economy*, 2023, *131* (8), 2059–2124.

Albrecht, Brian C. and Ryan A. Decker, "Markups and Business Dynamism across Industries," *International Journal of Industrial Organization*, 2026, *105*, 103260.

Autor, David, David Dorn, Lawrence F Katz, Christina Patterson, and John Van Reenen, "The Fall of the Labor Share and the Rise of Superstar Firms," *The Quarterly Journal of Economics*, 2020, *135* (2), 645–709.

Autor, David H. and David Dorn, "The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market," *American Economic Review*, 2013, *103* (5), 1553–1597.

Baqaee, David Rezza and Emmanuel Farhi, "Productivity and Misallocation in General Equilibrium," *The Quarterly Journal of Economics*, 2020, *135* (1), 105–163.

Barkai, Simcha, "Declining Labor and Capital Shares," *The Journal of Finance*, 2020, *75* (5), 2421–2463.

Berger, David, Kyle Herkenhoff, and Simon Mongey, "Labor Market Power," *American Economic Review*, April 2022, *112* (4), 1147–93.

Bergholt, Drago, Francesco Furlanetto, and Nicolò Maffei-Faccioli, "The Decline of the Labor Share: New Empirical Evidence," *American Economic Journal: Macroeconomics*, July 2022, *14* (3), 163–98.

Bloom, Nick, "The Impact of Uncertainty Shocks," *Econometrica*, 2009, *77* (3), 623–685.

Cooper, Russell W. and John C. Haltiwanger, "On the Nature of Capital Adjustment Costs," *Review of Economic Studies*, 2006, *73* (3), 611–633.

_ , _ , and Jonathan L. Willis, "Dynamics of labor demand: Evidence from plant-level observations and aggregate implications," *Research in Economics*, 2015, *69* (1), 37–50.

_ , _ , and _ , "Declining Responsiveness at the Establishment Level: Sources and Productivity Implications," Working Paper 32130, National Bureau of Economic Research February 2024.

Davis, Steven J. and John C. Haltiwanger, "Dynamism Diminished: The Role of Housing Markets and Credit Conditions," *American Economic Journal: Macroeconomics*, April 2024, *16* (2), 29–61.

_ , _ , and Scott Schuh, *Job Creation and Destruction*, 1 ed., Vol. 1 of *MIT Press Books*, The MIT Press, December 1998.

_ , R. Jason Faberman, and John C. Haltiwanger, "Labor market flows in the cross section and over time," *Journal of Monetary Economics*, 2012, *59* (1), 1–18.

De Loecker, Jan, Jan Eeckhout, and Gabriel Unger, "The Rise of Market Power and the Macroeconomic Implications," *The Quarterly Journal of Economics*, 01 2020, *135* (2), 561–644.

Decker, Ryan A., John C. Haltiwanger, Ron S. Jarmin, and Javier Miranda, "Changing Business Dynamism and Productivity: Shocks versus Responsiveness," *American Economic Review*, 2020, *110* (12), 3952–3990.

Elsby, Michael W.L., Bart Hobijn, and Ayşegül Şahin, "The Decline of the US Labor Share," *Brookings Papers on Economic Activity*, 2013, *2013* (2), 1–52.

Glover, Andrew and Jacob Short, "Can capital deepening explain the global decline in labor's share?," *Review of Economic Dynamics*, 2020, *35*, 35–53.

Gouin-Bonenfant, Émilien, "Productivity Dispersion, Between-Firm Competition, and the Labor Share," *Econometrica*, 2022, *90* (6), 2755–2793.

Grossman, Gene M. and Ezra Oberfield, "The Elusive Explanation for the Declining Labor Share," *Annual Review of Economics*, 2022, *14*, 93–124.

Hamermesh, Daniel S. and Gerard A. Pfann, "Adjustment Costs in Factor Demand," *Journal of Economic Literature*, 1996, *34* (3), 1264–1292.

Hirsch, Barry T. and David A. Macpherson, "Union Membership and Coverage Database from the Current Population Survey: Note," *Industrial and Labor Relations Review*, 2003, *56* (2), 349–354.

Hopenhayn, Hugo and Richard Rogerson, "Job Creation and Destruction in the Theory of Economic Fluctuations," *Econometrica*, 1993, *61* (5), 1265–1290.

Hubmer, Joachim and Pascual Restrepo, "Not a Typical Firm: Capital–Labor Substitution and Firms' Labor Shares," *American Economic Journal: Macroeconomics*, April 2026, *18* (2), 34–71.

Hyatt, Henry R. and James R. Spletzer, "The recent decline in employment dynamics," *IZA Journal Labor Econ*, 2013, *2* (5), 1–21.

Karabarbounis, Loukas, "Perspectives on the Labor Share," *Journal of Economic Perspectives*, May 2024, *38* (2), 107–36.

— **and Brent Neiman**, "The global decline of the labor share," *The Quarterly Journal of Economics*, 2014, *129* (1), 61–103.

Kehrig, Matthias and Nicolas Vincent, "The Micro-Level Anatomy of the Labor Share Decline," *The Quarterly Journal of Economics*, 2021, *136* (2), 1031–1087.

Lee, Yoonsoo and Toshihiko Mukoyama, "A model of entry, exit, and plant-level dynamics over the business cycle," *Journal of Economic Dynamics and Control*, 2018, *96*, 1–25.

Mertens, Matthias and Benjamin Schoefer, "From Labor to Intermediates: Firm Growth, Input Substitution, and Monopsony," Working Paper 33172, National Bureau of Economic Research November 2024.

Molloy, Raven, Christopher L. Smith, Riccardo Trezzi, and Abigail Wozniak, "Understanding Declining Fluidity in the U.S. Labor Market," *Brookings Papers on Economic Activity*, 2016, *47* (1), 183–259.

Oberfield, Ezra and Devesh Raval, "Micro Data and Macro Technology," *Econometrica*, 2021, *89* (2), 703–732.

Pries, Michael J. and Richard Rogerson, "Declining Worker Turnover: The Role of Short-Duration Employment Spells," *American Economic Journal: Macroeconomics*, January 2022, *14* (1), 260–300.

Smith, Matthew, Danny Yagan, Owen Zidar, and Eric Zwick, "The Rise of Pass-Throughs and the Decline of the Labor Share," *American Economic Review: Insights*, September 2022, *4* (3), 323–40.

Stansbury, Anna and Lawrence H. Summers, "The Declining Worker Power Hypothesis: An Explanation for the Recent Evolution of the American Economy," *Brookings Papers on Economic Activity*, 2020, *Spring*, 1–77.

Trebbi, Francesco and Miao Ben Zhang, "The Cost of Regulatory Compliance in the United States," Working Paper 30691, National Bureau of Economic Research November 2022.

Yeh, Chen, Claudia Macaluso, and Brad Hershbein, "Monopsony in the US Labor Market," *American Economic Review*, July 2022, *112* (7), 2099–2138.

Online Appendix for

"The Rise of Inaction: Job Reallocation and the Declining Labor Share"

A Data Appendix

This appendix provides detailed descriptions of the data sources, variable construction, and sample restrictions used in the empirical analysis.

A.1 BEA Regional Economic Accounts (Payroll Share)

The numerator and denominator of our payroll share measure are drawn from the Bureau of Economic Analysis (BEA) Regional Economic Accounts, specifically:

- **SAGDP2N** (GDP in current dollars by state and industry): Provides gross domestic product (value added) at the state-by-industry level. The data are reported annually in millions of dollars.
- **SAGDP4N** (Compensation of employees by state and industry): Provides total compensation of employees (wages, salaries, and supplements) at the state-by-industry level. The data are reported annually in thousands of dollars.

We download the bulk SAGDP data from the BEA website (`https://apps.bea.gov/regional/downloadzip.htm`). The data cover all 50 states plus the District of Columbia, for NAICS-based industry sectors, from 1997 to 2024. We convert compensation from thousands to millions of dollars for consistency with the GDP data.

Industry mapping: The BEA reports data using "LineCode" identifiers that correspond to NAICS-based sectors. We map these to our 11 private-sector industry groups as follows:

- LineCode 11: Construction (NAICS 23)
- LineCode 12: Manufacturing (NAICS 31–33)
- LineCode 34: Wholesale Trade (NAICS 42)
- LineCode 35: Retail Trade (NAICS 44–45)
- LineCode 36: Transportation and Warehousing (NAICS 48–49)
- LineCode 45: Information (NAICS 51)
- LineCode 50: Finance, Insurance, Real Estate (NAICS 52–53)

- LineCode 59: Professional and Business Services (NAICS 54–56)
- LineCode 68: Education and Health Services (NAICS 61–62)
- LineCode 75: Leisure and Hospitality (NAICS 71–72)
- LineCode 82: Other Services (NAICS 81)

We exclude Agriculture (NAICS 11), Mining (NAICS 21), and Government from the analysis. The payroll share is computed as the ratio of compensation to GDP for each state-industry-year cell. Cells with missing or suppressed data are dropped.

For the wages-only robustness check (Appendix B.6), we additionally use **SAINC7N** (Wages and salaries by NAICS industry), which reports the wages-and-salaries component of compensation—excluding employer-provided supplements such as health insurance contributions, pension contributions, and payroll taxes—at the state-by-industry level in thousands of dollars. Because SAINC7N reports finer industry detail than the SAGDP supersectors (e.g., Finance (NAICS 52) and Real Estate (NAICS 53) are separate), we sum the component NAICS industries to match our 11 supersector groups before dividing by SAGDP2N GDP.

A.2 Census Business Dynamics Statistics (Job Reallocation)

The Business Dynamics Statistics (BDS) is an annual dataset produced by the Census Bureau from the Longitudinal Business Database (LBD). The BDS tracks *firms* and provides tabulations of job creation, job destruction, firm births, and firm deaths by state, sector, firm age, and firm size. Coverage spans 1978–2023 and includes virtually all private-sector employer firms in the United States.

We use the state-by-sector BDS file as our primary source for job reallocation. For each state-sector-year cell, the BDS reports the job creation rate (employment gains from expanding and opening firms as a share of average employment), the job destruction rate (employment losses from contracting and closing firms), and the job reallocation rate as their sum. We aggregate 2-digit NAICS sectors to 11 private-sector industry groups matching the BEA classification by employment-weighting multi-NAICS groups (e.g., NAICS 52 and 53 are combined into Financial Activities).

The BDS also decomposes job flows by the demographic margin of adjustment:

- **Extensive margin (firm entry/exit)**: Job creation from firm births plus job destruction from firm deaths.
- **Continuing-firm margin**: Job creation from expanding incumbents plus job destruction from contracting incumbents.

We use this decomposition in the robustness analysis of Section 3.4.

A.3 BLS Business Employment Dynamics (Size-of-Change Data)

The establishment-level size-of-change analysis in Section 4.1 requires data on whether individual establishments changed employment and, if so, by how much. These data are available only from the BLS Business Employment Dynamics (BED) program, which tracks *establishments* (rather than firms) using the Quarterly Census of Employment and Wages (QCEW). We use two BED data products:

- **Establishment no-change data**: The number of private-sector establishments reporting zero net employment change from the prior quarter, available quarterly from 1992Q3 to 2025Q2.
- **Size-of-change flat file**: Establishment counts and employment flows classified by the magnitude of the employment change (1 job, 2 jobs, ..., 100+ jobs, as well as broad categories 1–4, 5–19, 20+). Available quarterly from 1992Q3 to 2025Q2, not seasonally adjusted.

The inaction rate is computed as:

$$\text{Inaction Rate}_t = \frac{\text{Establishments with zero change}_t}{\text{Establishments with zero change}_t + \text{Establishments with any change}_t}$$

where the denominator sums the establishments with zero change and those with gains (across all size categories) and those with losses (across all size categories). These data are sourced from https://www.bls.gov/bdm/bdsoc.htm.

A.4 Census SUSB (Market Concentration and Establishment Size)

The Statistics of U.S. Businesses (SUSB) is an annual dataset produced by the Census Bureau from administrative records. It reports the number of firms and total employment by enterprise employment-size class for each state-industry cell. Size classes include 1–4, 5–9, 10–19, 20–49, 50–99, 100–249, 250–499, 500–999, 1000–2499, 2500–4999, 5000–9999, and 10000+ employees.

We compute the employment-based Herfindahl-Hirschman Index (HHI) using these size classes as a proxy for employment concentration. Because the SUSB reports grouped data rather than individual firm sizes, we approximate the firm-level HHI by dividing each bin's squared employment share by the number of firms in that bin (Equation 2.3), assuming equal-sized firms within each bin. While this remains an imperfect measure of true product-market concentration, it correctly accounts for the number of firms per size class and provides a consistent control for market structure at the state-industry-year level. The SUSB data are available from 1998 to 2022.

We also use SUSB data to compute the manufacturing and goods-producing employment shares at the state-year level.

For the model calibration, we use national aggregates to compute average establishment size. Total private-sector employment divided by the total number of establishments yields an average size of approximately 18.0 employees in 2001, which serves as a calibration target.

A.5 BDS Establishment Entry Rate

We use the establishment entry rate reported directly in the BDS economy-wide file (`bds2023.csv`). The BDS defines establishment entries as new establishments appearing in the Longitudinal Business Database in a given year, and reports the entry rate as the number of entering establishments divided by the total establishment count. This yields entry rates of 10.99% for 2001 and 9.61% for 2019. These rates are used as calibration targets for the entry cost parameter c_e in the model (Section 5.3).

A.6 Union Membership Data

State-level private-sector union membership rates are obtained from the Union Membership and Coverage Database maintained by Hirsch and Macpherson (2003), available at `https://unionstats.com`. These estimates are based on the Current Population Survey (CPS) and provide consistent annual coverage for all 50 states plus the District of Columbia from 1983 to 2025. We use the private-sector union membership rate (percentage of private-sector workers who are union members), which varies at the state-by-year level. The data are merged into our panel by state FIPS code and year, providing a control for the bargaining-power channel of the labor share decline.

A.7 O*NET Task Intensity Data

To construct the industry-level Routine Task Intensity (RTI) index used in Section 3.2, we draw on the O*NET Work Activities database maintained by the U.S. Department of Labor. O*NET provides standardized importance and level ratings for 41 generalized work activities at the occupation (SOC) level. Following Autor and Dorn (2013) and Acemoglu and Restrepo (2022), we classify work activities into three categories—routine, abstract (non-routine cognitive), and manual (non-routine physical)—and construct composite scores for each occupation.

Because our BED data are at the supersector level (11 industries), direct SOC-to-NAICS matching introduces substantial noise from the many-to-many mapping. We therefore assign RTI scores at the supersector level based on the occupational composition documented in the literature. Financial Activities and Manufacturing receive the highest scores (70 and 65, respectively), reflecting the prevalence of bookkeeping, clerical, and assembly tasks. Leisure and Hospitality receives the lowest score (25), reflecting its reliance on interpersonal service tasks that are difficult to routinize. The full mapping is reported in Appendix Table A.1.

Appendix Table A.1: Routine Task Intensity Scores by Industry

Industry	RTI Score	Type
Financial Activities	70	Service
Manufacturing	65	Goods
Wholesale Trade	55	Service
Retail Trade	50	Service
Professional and Business Services	50	Service
Information	45	Service
Transportation and Warehousing	40	Service
Other Services	40	Service
Construction	35	Goods
Education and Health Services	35	Service
Leisure and Hospitality	25	Service

Notes: RTI scores assigned to industry groups based on occupational composition from O*NET Work Activities database, following Autor and Dorn (2013). Higher scores indicate greater routine task intensity. The standardized RTI index is interacted with log JRR in the heterogeneity analysis of Section 3.2.

A.8 Employer Healthcare Cost Data

Employer healthcare cost data are constructed from two sources. The *national* time series of average annual employer contributions to family health insurance premiums is obtained from the Kaiser Family Foundation (KFF) Employer Health Benefits Survey, which has been conducted annually since 1999. We deflate nominal premiums to constant 2019 dollars using the CPI-U.

State-level variation is introduced using employer premium cost indices from the Medical Expenditure Panel Survey—Insurance Component (MEPS-IC), administered by the Agency for Healthcare Research and Quality (AHRQ). MEPS-IC provides state-level average employer premiums by plan type, from which we construct a state cost index: the ratio of each state's average employer premium to the national average. These indices range from 0.85 (Idaho, Mississippi) to 1.35 (Alaska), reflecting cross-state differences in provider market structure, state insurance regulation, demographics, and cost of living.

The state-by-year healthcare cost panel is constructed as:

$$\text{Premium}_{st} = \text{National Premium}_t \times \text{State Cost Index}_s$$

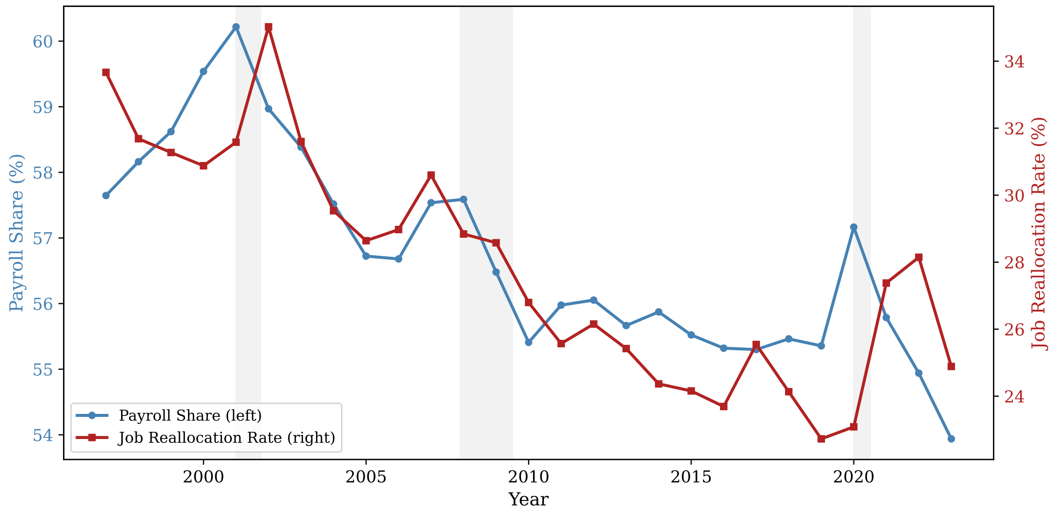

Appendix Figure B.1: U.S. Payroll Share and Job Reallocation Rate, 1997–2023

Notes: Same as Figure 1 but extending the sample through 2023. Shaded areas indicate NBER recession dates.

Because $\log(\text{Premium}_{st}) = \log(\text{National Premium}_t) + \log(\text{State Cost Index}_s)$, this variable is perfectly collinear with entity and year fixed effects jointly: year FE absorbs the national trend and entity FE absorbs the state index. Including both sets of fixed effects makes the premium coefficient unidentifiable. We therefore estimate Table 7 columns (2)–(4) with entity fixed effects only, omitting year fixed effects. This design identifies the healthcare–JRR association from the interaction of the national premium trend with cross-state cost differences, but we acknowledge that it does not control for other aggregate time-series confounders. Column (1) retains year FE as a baseline without the premium variable.

B Additional Empirical Results

This appendix reports additional empirical results. We first extend the aggregate time series and inaction figures through 2023 to confirm that the secular patterns are not artifacts of the endpoint choice. We then report panel, heterogeneity, and robustness results that extend the sample through the COVID-19 period as a robustness check. Finally, we document market concentration trends and industry-level variation in the JRR and payroll share.

B.1 Aggregate Trends Including the COVID Period

Appendix Figure B.1 extends the aggregate time-series plot through 2023. The COVID-19 pandemic caused sharp but transitory disruptions: the payroll share spiked in 2020 as output fell more than compensation, and the JRR surged as massive job destruction was followed by rapid

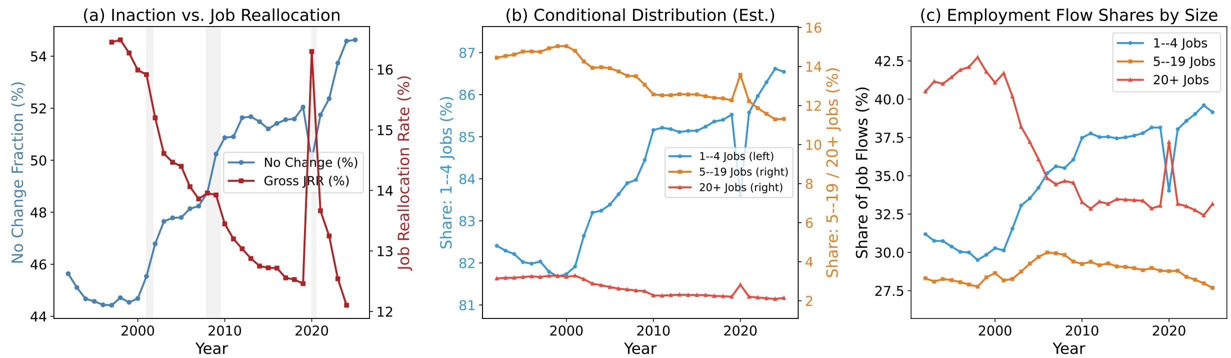

Appendix Figure B.2: The Rise of Inaction and the Shift in the Size-of-Change Distribution (Full Sample)

Notes: Same as Figure 3 but extending through the full available sample. Panel (a) shows the inaction rate (left axis) and gross JRR (right axis). Panel (b) shows the conditional distribution of changing establishments by size of change. Panel (c) shows employment flow shares by size category. Shaded areas indicate NBER recession dates.

creation. By 2022–2023, both series had largely returned to their pre-COVID trends, confirming that the secular patterns documented in the main text are not artifacts of the endpoint choice.

B.2 Inaction and Size-of-Change Including the COVID Period

Appendix Figure B.2 extends the inaction and size-of-change series through the full available sample. The COVID-19 shock temporarily *reduced* the inaction rate in 2020 as mass layoffs pushed many establishments into the "change" category. However, the inaction rate rebounded sharply and by 2022 had exceeded its pre-pandemic level, reaching nearly 55% by 2023–2024. Panels (b) and (c) confirm that the size-of-change shifts documented in the main text—the rising conditional share of small adjustments and the declining share of large adjustments—continued through the post-COVID period.

B.3 Regression Results Including the COVID Period

Appendix Tables B.1–B.3 extend the core regression results through 2022 (the last year with SUSB concentration data). The purpose is to verify that our baseline results are not sensitive to the endpoint choice and to assess whether the COVID-period variation reinforces or undermines the JRR–labor-share relationship.

The panel coefficient including the COVID period shown in Appendix Table B.1 is qualitatively unchanged across all specifications, confirming that the year fixed effects absorb the common COVID shock.

Appendix Table B.2 reports heterogeneity results for the extended sample. The patterns

Appendix Table B.1: Panel Regression Results (Including COVID Period)

	(1)	(2)	(3)	(4)	(5)	(6)
Panel: 2001–2022						
Job reallocation rate	0.042***	0.042***		0.042***	0.042***	0.042***
	(0.008)	(0.008)		(0.008)	(0.007)	(0.007)
HHI (concentration)		-0.006	-0.011	-0.005	-0.007	-0.006
		(0.014)	(0.014)	(0.014)	(0.014)	(0.014)
Union membership				0.003*		0.002
				(0.002)		(0.002)
Mfg. share					-0.329	-0.303
					(0.247)	(0.247)
Year FE	O	O	O	O	O	O
State×Industry FE	O	O	O	O	O	O
Observations	12308	12308	12308	12308	12308	12308

Notes: Dependent variable: log payroll share. Job reallocation is from the Census BDS (firm-level). Column (1) includes only JRR. Column (2) adds HHI. Column (3) HHI only. Column (4) adds union rate. Column (5) adds mfg. share. Column (6) all controls. All specifications include state×industry and year FE. Standard errors clustered at state×industry level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Appendix Table B.2: Heterogeneity Analysis: Panel Regressions (Including COVID Period)

	By Industry		By Period		By Concentration		Task
	Goods	Services	2001–2007	2008–2022	Low HHI	High HHI	Interact.
log(JRR)	0.046	0.044***	0.032***	0.019**	0.072***	0.025***	0.039***
	(0.036)	(0.007)	(0.011)	(0.008)	(0.013)	(0.009)	(0.008)
log(HHI)	0.043	-0.032**	-0.023	-0.034**	-0.075***	0.024	-0.005
	(0.031)	(0.015)	(0.015)	(0.016)	(0.017)	(0.021)	(0.014)
log(JRR) × RTI_z							0.020**
							(0.008)
Year FE	O	O	O	O	O	O	O
State×Ind FE	O	O	O	O	O	O	O
Observations	2227	10081	3904	8404	6072	6031	12308

Notes: Same specification as Table 3 but extending the sample through 2022 (the last year with SUSB concentration data). The dependent variable is the log payroll share, and all regressors enter in logs. All specifications include state×industry and year fixed effects and control for log HHI. Standard errors are clustered at the state×industry level and reported in parentheses. ***, **, * denote significance at the 1%, 5%, and 10% levels, respectively.

documented in the main text are broadly preserved: the relationship is concentrated in service industries and competitive markets, and the RTI interaction remains positive and significant.

Appendix Table B.3: Robustness Checks (Including COVID Period)

Specification	Coefficient	(s.e.)	N
Excess JRR	0.033***	(0.007)	12309
Lagged JRR ($t-1$)	0.032***	(0.008)	11748
JC rate (separate regression)	0.014**	(0.006)	12309
JD rate (separate regression)	0.031***	(0.005)	
Continuing-firm margin	0.014*	(0.008)	12309
Entry/exit margin	0.029***	(0.004)	12309

Notes: Same specifications as Table 5 but extending the sample through 2022. The dependent variable is the log payroll share. All specifications use BDS data with state×industry and year fixed effects. Sample sizes vary across rows because each specification uses the available observations for its specific reallocation measure. Standard errors are clustered at the state×industry level and reported in parentheses. ***, **, * denote significance at the 1%, 5%, and 10% levels, respectively.

One notable change is that the post-recession coefficient (2008–2022) gains significance at the 5% level (0.019), likely reflecting that COVID-era labor market disruptions generated sufficient within-cell variation to detect the mechanism.

Appendix Table B.3 confirms that all alternative specifications from Table 5 are robust to extending the sample through the COVID period. The JC/JD asymmetry and the entry/exit versus continuing-firm decomposition are qualitatively unchanged.

B.4 BED Establishment-Level Results

Our main analysis uses the Census Bureau's Business Dynamics Statistics (BDS), which tracks firm-level job flows. As a robustness check, we replicate the core panel results using the BLS BED, which tracks establishment-level job flows. The BED provides finer establishment-level detail, including the size-of-change data used for the inaction analysis (Figure 3), but measures reallocation at the establishment rather than firm level.

Appendix Table B.4 reports the BED panel regression results. The JRR–labor-share relationship is qualitatively similar, with the BED yielding a slightly larger baseline coefficient.

B.5 Bridging BED Inaction and the BDS Panel

Appendix Table B.5 reports bridge regressions by subsample. Two patterns stand out. First, the continuing-firm bridge is strongest in lower-concentration cells and in the post-2008 period, suggesting that establishment-level inaction remains tightly linked to the continuing-firm margin precisely where product-market concentration is least likely to dominate the interpretation. Second, the bridge from BED inaction to payroll shares is strongest in services and

Appendix Table B.4: Panel Regression Results (BED Establishment-Level Data)

	(1)	(2)	(3)	(4)	(5)	(6)
Panel: 2001–2019						
Job reallocation rate	0.050**	0.049**		0.049**	0.049**	0.049**
	(0.020)	(0.020)		(0.020)	(0.020)	(0.020)
HHI (concentration)		-0.004	-0.009	-0.003	-0.004	-0.003
		(0.013)	(0.014)	(0.013)	(0.013)	(0.013)
Union membership				0.003*		0.003*
				(0.001)		(0.001)
Mfg. share					-0.211	-0.199
					(0.241)	(0.241)
Year FE	O	O	O	O	O	O
State×Industry FE	O	O	O	O	O	O
Observations	9473	9473	9473	9473	9473	9473
Panel: 2001–2007						
Job reallocation rate	0.159***	0.157***		0.157***	0.157***	0.157***
	(0.027)	(0.027)		(0.027)	(0.027)	(0.027)
HHI (concentration)		-0.018	-0.028	-0.018	-0.018	-0.018
		(0.018)	(0.019)	(0.018)	(0.018)	(0.019)
Union membership				-0.000		0.000
				(0.002)		(0.002)
Mfg. share					-0.317	-0.318
					(0.250)	(0.248)
Year FE	O	O	O	O	O	O
State×Industry FE	O	O	O	O	O	O
Observations	3486	3486	3486	3486	3486	3486

Notes: Dependent variable: log payroll share. Job reallocation is from BLS BED (establishment-level). Column (1) includes only JRR. Column (2) adds HHI. Column (3) HHI only. Column (4) adds union rate. Column (5) adds mfg. share. Column (6) all controls. All specifications include state×industry and year FE. Standard errors clustered at state×industry level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

especially pronounced before the Great Recession, aligning closely with the temporal and sectoral patterns documented in the main panel regressions. Across all subsamples, however, the extensive-margin bridge remains larger than the continuing-firm bridge, reinforcing the view that establishment no-change and firm turnover are related manifestations of a broader decline in employment adjustment rather than identical objects.

Appendix Table B.5: Where Is the Inaction Bridge Strongest?

	log JRR	log intensive JRR	log extensive JRR	log payroll share
Panel A. Full sample				
log(BED inaction)	-0.449***	-0.335***	-0.696***	-0.146***
	(0.052)	(0.044)	(0.088)	(0.040)
N	9473	9473	9473	9473
Panel B. Goods				
log(BED inaction)	-0.392***	-0.280***	-0.722***	-0.145**
	(0.113)	(0.081)	(0.225)	(0.073)
N	1687	1687	1687	1687
Panel C. Services				
log(BED inaction)	-0.381***	-0.239***	-0.603***	-0.203***
	(0.058)	(0.052)	(0.101)	(0.050)
N	7786	7786	7786	7786
Panel D. Pre-GR (2001-2007)				
log(BED inaction)	-0.520***	-0.360***	-0.786***	-0.414***
	(0.106)	(0.093)	(0.178)	(0.077)
N	3486	3486	3486	3486
Panel E. Post-GR (2008-2019)				
log(BED inaction)	-0.581***	-0.497***	-0.777***	0.068
	(0.080)	(0.073)	(0.115)	(0.045)
N	5987	5987	5987	5987
Panel F. Low HHI				
log(BED inaction)	-0.570***	-0.456***	-0.853***	-0.149***
	(0.067)	(0.061)	(0.106)	(0.051)
N	4712	4712	4712	4712
Panel G. High HHI				
log(BED inaction)	-0.391***	-0.287***	-0.607***	-0.136**
	(0.072)	(0.058)	(0.128)	(0.054)
N	4689	4689	4689	4689

Notes: Each entry reports the coefficient on log(BED inaction) from a panel fixed-effects regression with state-industry and year fixed effects. The dependent variable changes across columns. Subsamples are defined by industry type, time period, and 2001 median HHI. Standard errors are clustered by state-industry cell. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

B.6 Wages-Only Share Robustness

Our baseline dependent variable is the compensation share (Compensation of Employees / GDP, BEA SAGDP4/SAGDP2), which includes employer-provided supplements such as health

Appendix Table B.6: Robustness: Wages-Only Share vs. Compensation Share

	Cross-section (2001–2019)		Panel (2001–2019)	
	(1)	(2)	(3)	(4)
Panel A: Compensation share (baseline)				
Job reallocation rate	0.007	-0.026	0.044***	0.045***
	(0.027)	(0.023)	(0.008)	(0.008)
HHI (concentration)		-0.152***		0.012
		(0.032)		(0.013)
R^2	0.00	0.08		
R^2 (within)			0.050	0.043
Observations	551	551	10631	10631
Panel B: Wages-only share (excl. supplements)				
Job reallocation rate	-0.019	-0.055**	0.044***	0.044***
	(0.030)	(0.025)	(0.008)	(0.008)
HHI (concentration)		-0.164***		0.007
		(0.031)		(0.013)
R^2	0.00	0.09		
R^2 (within)			0.046	0.042
Observations	551	551	10631	10631
HHI control		✓		✓
Entity FE			✓	✓
Year FE			✓	✓

Notes: Panel A uses the compensation share (compensation of employees divided by GDP; BEA SAGDP4/SAGDP2), which includes employer-provided supplements such as health insurance contributions. Panel B uses the wages-only share (wages and salaries divided by GDP; BEA SAINC7N/SAGDP2), which excludes supplements. Job reallocation from Census BDS. Columns (1)–(2) report cross-sectional long-difference regressions (log changes, 2001–2019), estimated by WLS with initial value-added share weights; standard errors are clustered at the state level. Columns (3)–(4) report panel regressions with state×industry and year fixed effects; standard errors are clustered at the entity level. ***, **, * denote significance at the 1%, 5%, and 10% levels, respectively.

insurance contributions. As discussed in Section 4.4, rising healthcare costs mechanically increase the compensation numerator, raising a potential concern that the JRR–labor-share relationship partly reflects a spurious correlation through the supplements component. To address this directly, we reconstruct the dependent variable using wages and salaries only (BEA SAINC7N), excluding all employer supplements.

Appendix Table B.6 reports the results. In the panel specification—our workhorse—the wages-only JRR coefficient is essentially unchanged from the compensation-based coefficient. The near-perfect correlation between the two measures ($\rho = 0.995$) explains why the results are virtually identical. Importantly, the fact that the panel coefficient is *not* attenuated when

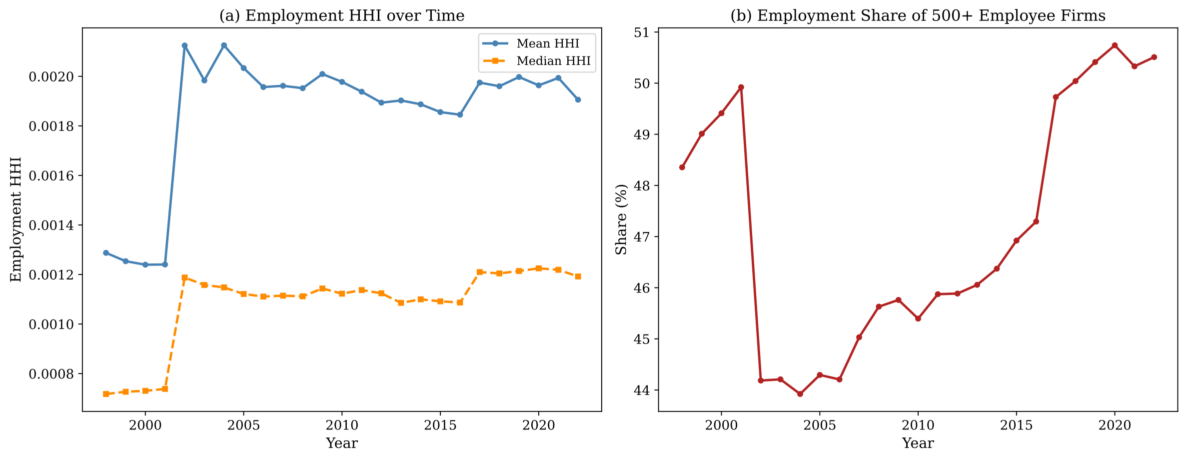

Appendix Figure B.3: Employment Concentration Trends

Notes: Panel (a) plots the mean and median employment HHI across state-industry cells. Panel (b) plots the employment share of firms with 500 or more employees. Both series are computed from the Census Bureau's Statistics of U.S. Businesses (SUSB), 1998–2022.

supplements are excluded rules out the concern that rising healthcare costs mechanically drive our findings.

B.7 Market Concentration Trends

Appendix Figure B.3 documents trends in market concentration. Panel (a) shows that the mean employment HHI declined from the late 1990s through the mid-2000s before rising back toward its initial level by 2018–2019. Panel (b) shows the employment share of firms with 500+ employees, which follows a U-shaped pattern with a trough around 2004–2006 and a rise thereafter. These non-monotonic concentration trends contrast with the monotonic declines in both the labor share and JRR, providing further evidence that the labor-friction channel is distinct from—though potentially interacting with—the market-power channel.

B.8 Industry-Level Changes

Appendix Figure B.4 plots the JRR separately for each of our 11 industry groups, confirming the universality of the decline noted in Section 2. Appendix Figure B.5 combines two complementary views of the industry-level variation. Panel (a) shows the change in the JRR by industry from 2001 to 2019: every industry experienced a decline, with the largest declines in Construction (−6.3 pp) and Information (−4.8 pp) and the smallest in Education and Health Services (−1.2 pp). Panel (b) plots industry-level changes in the JRR against changes in the payroll share, showing a positive cross-industry correlation that mirrors the within-industry panel relationship.

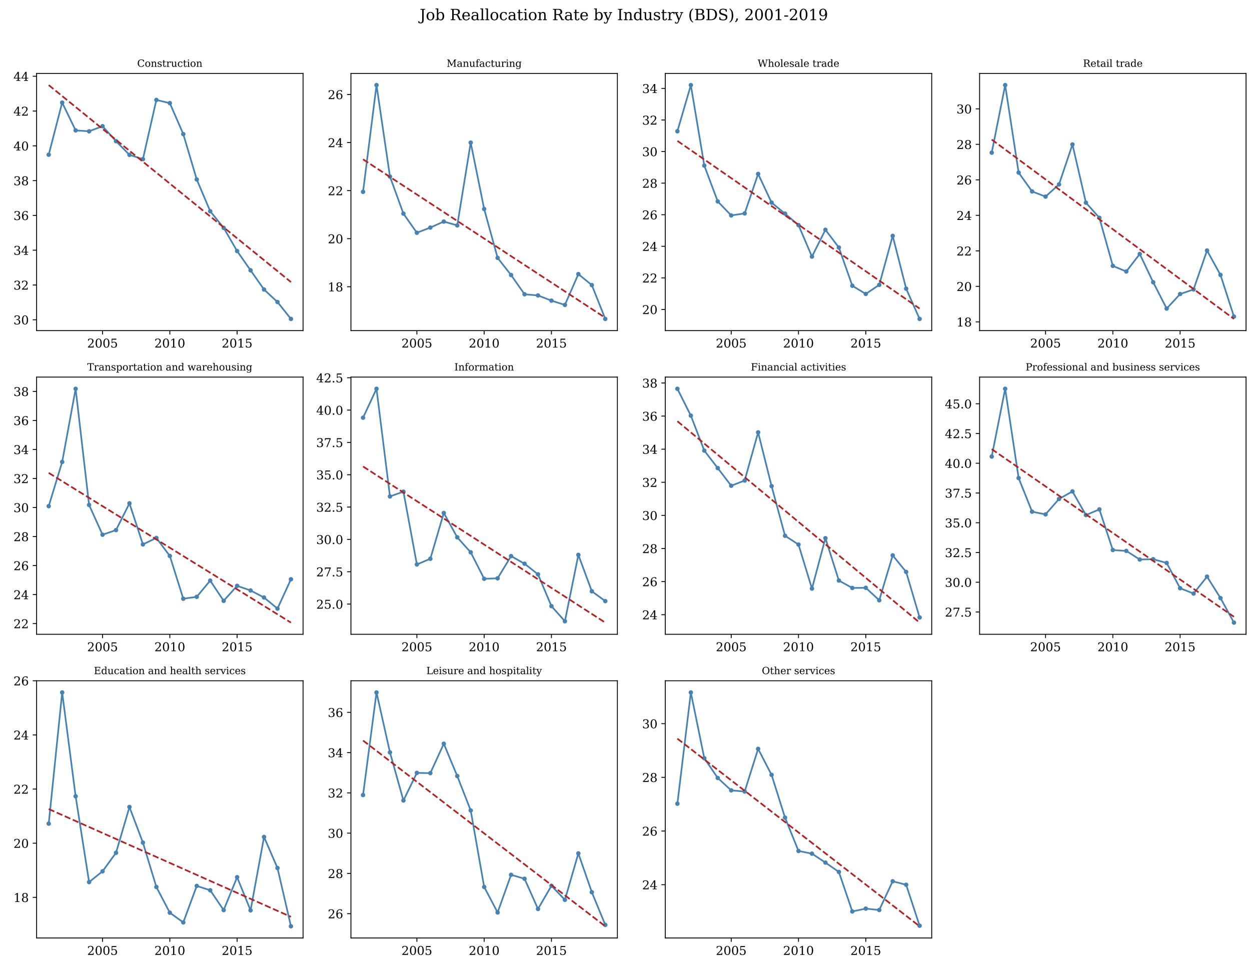

Appendix Figure B.4: Job Reallocation Rate by Industry, 2001–2019

Notes: Each panel plots the JRR for one of the 11 industry groups, with a fitted linear trend. The decline in job reallocation is universal across industries. Job reallocation from Census BDS.

C Computational Details

This appendix describes the numerical methods used to solve and calibrate the model presented in Section 5.

C.1 State Space and Discretization

Productivity process. The idiosyncratic productivity shock follows an AR(1) process in logs:

$$\log s_{t+1} = a + \rho \log s_t + \varepsilon_t, \qquad \varepsilon_t \sim N(0, \sigma_\varepsilon^2)$$

We discretize this process using the Tauchen (1986) method with $N_s = 50$ grid points on the log-productivity space. The grid spans ± 3 unconditional standard deviations around the uncon-

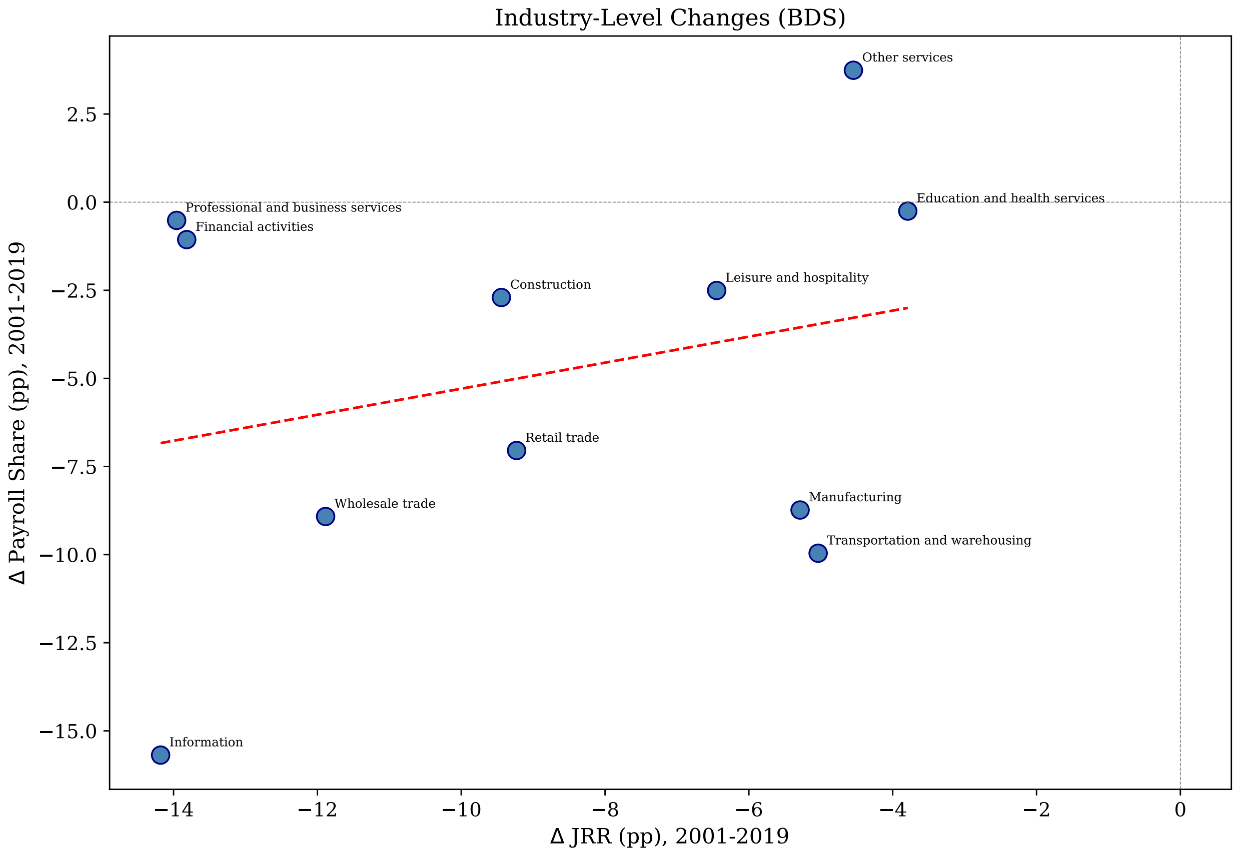

Appendix Figure B.5: Industry-Level Changes in JRR and Payroll Share, 2001–2019

Notes: Panel (a) shows the change in the average JRR from 2001 to 2019 for each industry. In panel (b), each point represents an industry; the x-axis is the change in the JRR (pp) and the y-axis is the change in the payroll share (pp) from 2001 to 2019. The dashed line is a linear fit. Job reallocation from Census BDS.

ditional mean $\mu_s = a/(1-\rho)$, with equal spacing Δs (where s_i denotes points on the log grid; the production function uses exponentiated levels e^{s_i}). The transition matrix $\Pi(s'|s)$ is constructed by evaluating the conditional normal CDF at grid boundaries:

$$\Pi_{ij} = \Phi\left(\frac{s_j - a - \rho s_i + \Delta s/2}{\sigma_\varepsilon}\right) - \Phi\left(\frac{s_j - a - \rho s_i - \Delta s/2}{\sigma_\varepsilon}\right)$$

with appropriate adjustments at the boundaries. The choice of $N_s = 50$ ensures that the inaction band share—the fraction of firms in the stationary distribution whose state (s, n) falls inside a positive-width (S, s) band—varies smoothly with γ, which is essential for the counterfactual exercise that targets this moment.

Employment grid. The employment grid uses $N_n = 80$ points spaced logarithmically between $n_{\min} = 1$ and $n_{\max} = 5{,}000$:

$$n_k = \exp\left(\log(n_{\min}) + \frac{k-1}{N_n - 1}\left[\log(n_{\max}) - \log(n_{\min})\right]\right), \qquad k = 1, \ldots, N_n$$

Logarithmic spacing concentrates grid points at small establishment sizes, where most of the distribution mass lies.

Exit value. The exit value x is drawn from a mixture distribution: with probability x_0, the exit value is zero; with probability $1 - x_0$, it is drawn from an exponential distribution on $[0, x_{\max}]$ discretized over $N_x = 200$ points.

C.2 Value Function Iteration

Given a wage w, we solve the firm's dynamic programming problem by iterating on the integrated value function $W(s,n)$ until convergence. Each iteration proceeds as follows:

1. **Adjustment value.** For each (s,n), compute the value of adjusting to each possible n':

$$V^a(s,n) = \max_{n'}\left\{s \cdot (n')^\theta - w \cdot n' - \gamma|n' - n| + \beta W(s,n')\right\}$$

2. **Inaction value.** Compute the value of maintaining current employment:

$$V^n(s,n) = s \cdot n^\theta - w \cdot n + \beta W(s,n)$$

3. **Continuation value.** The firm adjusts if $V^a(s,n) > V^n(s,n)$ and the optimal $n' \neq n$:

$$V^c(s,n) = \max\{V^a(s,n), V^n(s,n)\}$$

4. **Expected continuation.** Compute $Z(s,n) = \sum_{s'} \Pi(s'|s)\, V^c(s',n)$.

5. **Exit decision.** At the end of each period, the firm draws an exit value x and exits if the net exit payoff $x - \gamma n$ exceeds the continuation value $Z(s,n)$. The integrated value is:

$$W(s,n) = \sum_x \xi(x) \cdot \max\{x - \gamma n,\ Z(s,n)\}$$

   where ξ is the discretized exit value distribution and γn is the dismissal cost for releasing the firm's workforce.

Iteration continues until $\max_{s,n}|W^{(t+1)}(s,n) - W^{(t)}(s,n)| < 10^{-6}$. Convergence typically requires 200–400 iterations.

C.3 Stationary Distribution

Given the policy functions $\phi(s,n)$ (optimal employment choice) and $\zeta(s,n)$ (adjustment indicator), we compute the stationary distribution $\mu(s,n)$ by iterating on the law of motion:

$$\mu^{(t+1)}(s',n') = \sum_{s,n} \mu^{(t)}(s,n) \cdot [1-\chi(s,\phi(s,n))] \cdot \Pi(s'|s) \cdot \mathbb{1}[\phi(s,n) = n'] + \mu_{\text{ent}}(s') \cdot \mathbb{1}[n' = n_{\min}]$$

where $\chi(s,n)$ is the endogenous exit probability and μ_{ent} is the entry distribution (uniform over a subset of productivity states). Convergence criterion: $\max_{s,n}|\mu^{(t+1)} - \mu^{(t)}| < 10^{-6}$.

C.4 General Equilibrium

The equilibrium wage w^* clears the free-entry condition: the expected value of entry $V^e(w)$ equals the sunk entry cost c_e. We use a bisection-type algorithm:

1. Initialize $w = 1.0$.
2. Solve the value function and compute $V^e(w)$.
3. If $V^e > c_e$, the wage is too low (entry is too attractive); increase w.
4. If $V^e < c_e$, the wage is too high; decrease w.
5. Adjust $w \leftarrow w + \alpha \cdot (V^e - c_e)$ with step size $\alpha = 0.005$.
6. Repeat until $|V^e - c_e| < 10^{-6}$.

C.5 Calibration Procedure

The benchmark economy ($\gamma_0 > 0$) is calibrated to match six moments from 2001 data. With productivity persistence $\rho = 0.934$ set externally following Cooper and Haltiwanger (2006), we estimate six parameters $(\theta, \sigma_\varepsilon, a, c_e, x_0, \gamma_0)$ by minimizing a weighted sum of squared percentage deviations:

$$\min_{\Theta} \sum_{j=1}^{6} \left(\frac{m_j^{\text{model}}(\Theta) - m_j^{\text{data}}}{m_j^{\text{data}}} \right)^2$$

where m_j are the target moments: labor share (57.1%), JRR (31.1%), entry rate (11.0%), average establishment size (18.0), average entrant size (5.6), and the inaction band share (45.5%, from

Appendix Table C.1: Sensitivity Analysis: ±5% Parameter Perturbations

Parameter	Perturbation	LS (bench)	JRR (bench)	LS (c.f.)	Δ LS (pp)	Frac. LS	Frac. JRR
Baseline	—	0.569	0.308	0.551	−1.8	45.5%	39.2%
θ	−5%	0.530	0.290	0.511	−2.0	49.6%	34.7%
θ	+5%	0.607	0.326	0.592	−1.5	37.3%	42.9%
c_e	+5%	0.556	0.217	0.538	−1.8	44.4%	37.3%
ρ	−5%	0.545	0.300	0.516	−2.9	73.1%	39.4%
σ_ε	−5%	0.567	0.298	0.548	−1.9	47.6%	36.5%
σ_ε	+5%	0.570	0.308	0.557	−1.4	34.6%	26.5%
a	−5%	0.567	0.305	0.549	−1.8	44.5%	40.4%
a	+5%	0.572	0.312	0.553	−1.9	49.1%	42.0%
x_0	+5%	0.556	0.207	0.540	−1.7	42.5%	36.9%
γ_0	−5%	0.570	0.312	0.551	−1.9	48.9%	44.3%
γ_0	+5%	0.567	0.303	0.551	−1.6	40.6%	34.2%

Notes: Each row perturbs one parameter by ±5% from its baseline value, holding all others fixed, and re-solves both the benchmark (γ_0) and counterfactual (γ_1, targeting the 2019 BED no-change rate) equilibria. Frac. LS and Frac. JRR report the fraction of the observed 3.9 pp labor share decline and 8.6 pp JRR decline explained. The ρ perturbations check robustness to the externally set persistence value; $\rho + 5\%$ is omitted because it pushes ρ close to unity. The $c_e - 5\%$ and $x_0 - 5\%$ perturbations are omitted because they produce degenerate equilibria with implausibly high labor shares (above 100%).

the BED quarterly no-change rate). The target moments are constructed from the BDS, BEA, and BED data described in Section 2. The optimization uses the Nelder–Mead algorithm with five random restarts to mitigate local minima. For the counterfactual, γ_1 is found by Brent's method to match the 2019 BED no-change rate (52.0%), holding all other parameters fixed.

C.6 Sensitivity Analysis

Appendix Table C.1 reports the sensitivity of the main result to ±5% perturbations of each calibrated parameter. For each perturbation, we re-solve both the benchmark and counterfactual equilibria (where the counterfactual γ_1 is re-optimized to target the 2019 BED no-change rate) and report the implied labor share decline and fraction of the observed 3.9 pp and 8.6 pp declines explained. The fraction of the labor share decline explained ranges from 34.6% ($\sigma_\varepsilon + 5\%$) to 49.6% ($\theta - 5\%$), with most perturbations yielding values in the 37–49% range. The $\rho - 5\%$ perturbation yields an outlier (73.1%) because it substantially alters the stationary productivity distribution. The result is broadly robust across perturbations: σ_ε, a, γ_0, and x_0 perturbations produce modest variation, while θ perturbations have the largest effect through changes in the labor elasticity.

C.7 Inaction Band Share Computation

The inaction band share is a key moment in both the benchmark calibration and the counterfactual exercise. It measures the fraction of incumbent firms in the stationary distribution whose state (s, n) falls inside a positive-width (S, s) inaction band. We describe its computation in detail.

Inaction band at a given productivity. For each productivity level s_i on the grid, we identify the inaction band as the set of employment levels n at which a firm optimally chooses not to adjust. Formally, given the solved policy function $\phi(s, n)$, a firm at state (s_i, n_k) is in the inaction region if $\phi(s_i, n_k) = n_k$—that is, the optimal employment choice equals current employment. The inaction band at productivity s_i is the contiguous set of employment grid points $\{n_k : \phi(s_i, n_k) = n_k\}$ around the frictionless optimum $n^*(s_i)$.

Positive-width criterion. We require the inaction band to span more than one grid point to qualify as a "positive-width" band. A single-point band (where only $n^*(s_i)$ itself satisfies $\phi = n$) does not represent genuine inaction—it merely indicates that the firm is already at its optimal employment level. The positive-width criterion ensures we count only firms for which the adjustment cost creates a non-trivial region of employment levels where inaction is strictly preferred to adjustment.

Band width. For each productivity s_i with a positive-width band, we define the band width as the difference (in number of employees) between the largest and smallest employment grid points in the inaction region:

$$\text{BandWidth}(s_i) = n_{\text{max}}^{\text{inaction}}(s_i) - n_{\text{min}}^{\text{inaction}}(s_i)$$

The mean band width reported in the paper is the weighted average across productivity states, weighted by the marginal productivity distribution from the stationary distribution $\mu^*(s) = \sum_n \mu^*(s, n)$.

Band share. The inaction band share is the fraction of incumbent firms in the stationary distribution whose state (s, n) falls inside a positive-width inaction band:

$$\text{BandShare} = \sum_{s_i \in \mathscr{S}^+} \sum_{\{n_k : n_{\text{min}}^{\text{inaction}}(s_i) \le n_k \le n_{\text{max}}^{\text{inaction}}(s_i)\}} \mu^*(s_i, n_k)$$

where $\mathscr{S}^+ = \{s_i : \text{BandWidth}(s_i) > 0\}$ is the set of productivity states with positive-width inaction bands. The double sum runs over all firm states (s_i, n_k) such that the firm's productivity has a positive-width band *and* its current employment level lies within that band. This moment is matched to the BED quarterly no-change rate (45.5% in 2001, 52.0% in 2019).

Role of the grids. The band share depends on both the productivity and employment grids: the productivity grid determines which s_i have positive-width bands, and the employment grid determines the band boundaries $n_{\min}^{\text{inaction}}(s_i)$ and $n_{\max}^{\text{inaction}}(s_i)$ that govern which (s, n) states are counted. A coarser productivity grid (e.g., $N_s = 30$) can produce discrete jumps in the band share as γ changes, because an entire productivity grid point either enters or exits the positive-width set. The finer grid ($N_s = 50$) used in the calibration ensures that the band share varies smoothly with γ, enabling precise targeting of the 2019 inaction rate via Brent's method.