Discussion Paper No. 1321

ISSN (Print) 0473-453X
ISSN (Online) 2435-0982

HETEROGENEOUS FACTOR-AUGMENTING PRODUCTIVITY

Wei Qiao
Lianming Zhu

September 2026

The Institute of Social and Economic Research
The University of Osaka
6-1 Mihogaoka, Ibaraki, Osaka 567-0047, Japan

Heterogeneous Factor-Augmenting Productivity[*]

Wei Qiao
Jinan University

Lianming Zhu
University of Osaka

September 10, 2026

Abstract

This paper studies how heterogeneous factor-augmenting productivity shapes factor allocation, aggregate labor share, and misallocation. In Chinese manufacturing, more labor-productive firms have lower labor shares, higher capital intensity, and lower employment. We first prove that, under empirically motivated restrictions on factor-market frictions, this joint pattern requires heterogeneity in relative factor-augmenting productivity. We then develop an approach that simultaneously identifies firm-level capital- and labor-augmenting productivities, labor markdowns, and marginal capital expenditures. As a validation exercise, we exploit China's value-added tax reform that lowered equipment costs and find that the calibrated ratio of capital-augmenting to labor-augmenting productivity rises. Counterfactuals show that this ratio is central to the employment-productivity relationship and the aggregate labor share. Accounting for this heterogeneity shifts the inferred composition of misallocation toward capital and reverses the comovement of the marginal revenue products of capital and labor.

Keywords: Factor-augmenting productivity; labor share; misallocation.

[*]Qiao: weighqiao@gmail.com. Zhu: lianmingzhu@iser.osaka-u.ac.jp. We are grateful to Michael Song and Xiaodong Zhu for valuable comments and discussions during the early stages of this project.

1 Introduction

The aggregate consequences of firm heterogeneity depend on why firms choose different mixes of capital and labor. Dispersion in firms' input mixes may arise either from factor market distortions that alter relative factor costs or from heterogeneity in factor-augmenting productivity. This paper develops a theory establishing when heterogeneity in relative factor productivity is required to rationalize observed patterns in the joint allocation of capital and labor across firms. We then build a quantitative framework that separates firm-level capital- and labor-augmenting productivity from frictions in labor and capital markets, and use it to study the aggregate labor share and factor misallocation.

We motivate the analysis with firm-level evidence from Chinese manufacturing. Within narrowly defined markets, firms with higher value-added per worker have lower labor shares, higher capital-labor ratios, and lower employment, even though they produce more value-added and use more capital. These patterns are robust to alternative constructions of value-added, alternative labor productivity measures that address mechanical relationships with the outcomes, and controls for systematic differences across ownership types. The joint pattern of higher capital intensity and lower employment among more labor-productive firms provides the key empirical restriction for distinguishing heterogeneity in relative factor productivity from differences in relative factor costs.

We first show that the joint pattern above cannot be explained by scalar productivity heterogeneity or factor market distortions. Scalar productivity scales the marginal products of capital and labor proportionally and thus cannot generate cross-firm variation in relative factor demand. We impose two empirically motivated restrictions on labor- and capital-market frictions: marginal labor expenditure is weakly increasing in employment, and capital access does not worsen with firm size. These restrictions imply that smaller employers face a weakly lower marginal cost of labor relative to capital and would therefore be less capital intensive if relative factor productivity were constant. The observed rise in capital intensity therefore requires relative factor productivity to vary across firms. Given that capital and labor are gross substitutes in Chinese manufacturing (Meng et al., 2025), capital-augmenting productivity must rise relative to labor-augmenting productivity along the labor-productivity distribution.

Second, we provide a method for calibrating firm-level factor-augmenting productivities. The key challenge is that the labor share jointly reflects the output elasticity of labor and the labor markdown. Separating relative factor-augmenting productivity from labor market power requires additional structure: Hicks-neutral technology, perfectly competitive factor markets, or an equilibrium model of factor market. We employ a Burdett-Mortensen wage-posting framework to infer labor markdowns from the joint distribution of wages and employment. Conditional on the inferred markdowns, the labor share identifies the output elasticity of labor. The capital-labor ratio and value added per worker then identify capital- and labor-augmenting productivity and hence relative factor-augmenting productivity. The capital first-order condition in turn recovers firm-specific marginal capital expenditure.

We validate the calibrated objects using information not targeted in the calibration. We exploit China's value-added tax (VAT) reform, which lowered the tax-adjusted cost of equipment investment for eligible firms in the Northeast. Lower equipment costs can facilitate the adoption of newer vintages, increasing the productive effectiveness of capital without directly affecting labor. Consistent with this prediction, the VAT reform substantially increases calibrated capital-augmenting productivity, while leaving labor-augmenting productivity largely unchanged. This factor-specific response provides external validation of our separation of the two productivity components. We also examine calibrated marginal capital expenditure. It is lower for state-owned and foreign-invested firms, consistent with their comparatively better access to credit. Its decline with firm size parallels the decline in observed borrowing costs and supports the restriction on capital access imposed in our theoretical analysis.

Our third contribution quantifies how heterogeneity in factor-augmenting productivity affects the allocation of factors, the aggregate labor share, and the assessment of misallocation. Cross-sectionally, this heterogeneity is central to why more labor productive firms are smaller employers: compressing dispersion in relative factor-augmenting productivity by 10 percent reverses the employment-productivity correlation, making more labor productive firms larger employers. This counterfactual highlights the quantitative importance of relative factor-augmenting productivity for the joint pattern that motivates the paper.

Heterogeneity in factor-augmenting productivity also shapes the aggregate labor share. The calibrated model reproduces much of the observed decline over the sample period, reflecting both lower average firm labor shares and a greater concentration of value-added in firms with lower labor shares. Counterfactuals that reduce cross-firm dispersion in relative factor-augmenting productivity increase the aggregate labor share. This result is driven by a weaker pattern in which firms with lower labor shares account for a larger share of value-added. Reversing the observed changes in relative factor augmentation deepens the labor share decline.

Firm-level heterogeneity in factor-augmenting productivity also changes the assessment of factor misallocation. In Hsieh-Klenow accounting, firms within an industry share the same output elasticities with respect to capital and labor. Our calibrated factor-augmenting productivities instead imply firm-specific output elasticities and capital-labor mixes. Incorporating this heterogeneity substantially increases the potential output gain from capital reallocation and reduces the gain from labor reallocation relative to the Hsieh-Klenow benchmark. It also reverses the correlation between the marginal revenue products of capital and labor from positive to negative, indicating a greater role for variation in relative marginal returns associated with firms' capital-labor mixes. Thus, factor-augmenting productivity changes not only the magnitude of measured misallocation but also its composition, affecting whether the implied reallocation operates primarily through firm scale or firms' capital-labor mixes.

Related Literature. Our paper contributes to the macroeconomic literature on factor shares and heterogeneity in factor intensity. Karabarbounis and Neiman (2014) emphasize declining investment-good prices and substitution toward capital as a force behind the decline in the labor share, while Oberfield and Raval (2021) attribute an important role to the bias of technical change in the decline of the U.S. manufacturing labor share. Kehrig and Vincent (2021) emphasize the reallocation of value added toward low labor share plants, and Gouin-Bonenfant (2022) shows how productivity heterogeneity and between-firm wage competition shape the aggregate labor share. More recently, Kaymak and Schott (2023) and Hubmer and Restrepo (2026) study how changes in the cost of capital interact with heterogeneous factor intensity and technology choice to affect aggregate factor shares. We complement this literature by distinguishing heterogeneity in relative factor-augmenting productivity from differences in relative factor costs. We show that relative factor-augmenting productivity is quantitatively important for both the cross-sectional allocation of capital and labor, and the evolution of the aggregate labor share.

We also build on the literature on resource misallocation. Standard misallocation accounting interprets differences in factor use across firms through distortions to labor and capital allocation, typically allowing firms to differ in scalar productivity (Restuccia and Rogerson, 2008; Hsieh and Klenow, 2009). A large macroeconomic literature studies how financial frictions distort capital allocation and aggregate productivity (Song et al., 2011; Midrigan and Xu, 2014; Moll, 2014; Bai et al., 2026), while Brandt et al. (2013) quantify distortions in the allocation of labor and capital across provinces and sectors in China. Closest to our analysis, David and Venkateswaran (2019) use the comovement of the average revenue products of capital and labor to distinguish forces operating through firm scale from those affecting the capital-labor mix, while their quantitative analysis centers on capital misallocation. Asturias and Rossbach (2023) show that allowing for group-specific production technologies can materially alter inferred marginal revenue products and the gains from eliminating factor misallocation. We deepen this perspective by introducing firm-specific capital- and labor-augmenting productivity as a structural source of input-mix heterogeneity and separating it from differences in relative factor costs.

A growing body of evidence shows that firms have wage-setting power and that firm-specific rents are only partly passed through to workers (Card et al., 2018; Lamadon et al., 2022; Berger et al., 2022). A related literature estimates firm-level labor markdowns from production data (Yeh et al., 2022). In contrast, we recover firm-specific labor markdowns from equilibrium wage and employment patterns, jointly with heterogeneous factor-augmenting productivity.

Finally, a related empirical literature measures input-specific productivity and shows that firms differ not only in scalar productivity but also in the efficiency with which they use particular inputs (Doraszelski and Jaumandreu, 2018; Zhang, 2019; Rubens et al., 2026; Demirer, 2026). A parallel macroeconomic literature studies capital-labor substitution jointly with biased or factor-augmenting technical change in aggregate CES frameworks (Antràs, 2004; León-Ledesma et al., 2010). We connect these strands of literature by recovering firm-level capital- and labor-augmenting productivities, jointly with the relative factor costs that shape input allocation, and embedding these elements in an equilibrium framework to study their implications for the aggregate labor share and measured misallocation.

The remainder of the paper is organized as follows. Section 2 documents the motivating facts. Section 3 describes the theoretical framework and demonstrates why the motivating facts require heterogeneity in factor-augmenting productivities. Section 4 builds the quantitative framework. Section 5 presents the calibrated firm-level objects, the VAT reform validation exercise, and their role in factor allocation. Section 6 studies the aggregate labor share. Section 7 examines the implications for misallocation. Section 8 concludes.

2 Motivating Evidence

2.1 Data

The analysis uses the Annual Survey of Industrial Firms, conducted by China's National Bureau of Statistics, for the period 1998–2007. The survey covers all state-owned industrial enterprises and all non-state industrial enterprises with annual sales above 5 million RMB, and accounts for roughly 90 percent of total industrial output. It reports detailed information on firm characteristics and accounting variables, including industry affiliation, geographic location, sales, intermediate inputs, employment, fixed assets, and wage bills.

We define labor compensation as wage bills plus supplementary benefits such as bonuses and unemployment insurance. We construct value added in two ways. Our baseline measure follows the production approach and defines value added as gross output minus intermediate inputs plus value-added tax. As a robustness check, we also construct value-added using the income approach, following Qian and Zhu (2013). We construct real capital using the perpetual inventory method following Brandt et al. (2012). Appendix A provides details on the construction of the capital stock and the alternative value-added measure. The labor share is measured as labor compensation divided by value added, the capital-labor ratio as real capital divided by employment, and labor productivity as value added per employee.[1]

2.2 Factor Allocation across Labor Productivity Distribution

Figure 1 presents binned scatterplots of the labor share, the log capital-labor ratio, and log employment against firms' percentiles in the labor productivity distribution. We define markets by two-digit industry and province. We residualize log value added per employee and each outcome with respect to market and year fixed effects. Within each year, firms are grouped into 100 percentile bins according to residualized log labor

[1]We restrict the sample to observations with positive employment, sales, value added, and labor compensation, and trim the top and bottom 1 percent of constructed real capital in each year.

productivity, and each panel plots the mean residualized outcome within each bin, recentered at the corresponding year mean.

Figure 1 documents three related patterns. As firms move up the labor productivity distribution, the labor share declines, the capital-labor ratio rises, and employment falls. More labor productive firms are therefore more capital intensive but employ fewer workers, and they have lower labor shares. This joint pattern motivates our analysis of the allocation of capital and labor across firms.

Figure 1: Factor Allocation across Labor Productivity Distribution

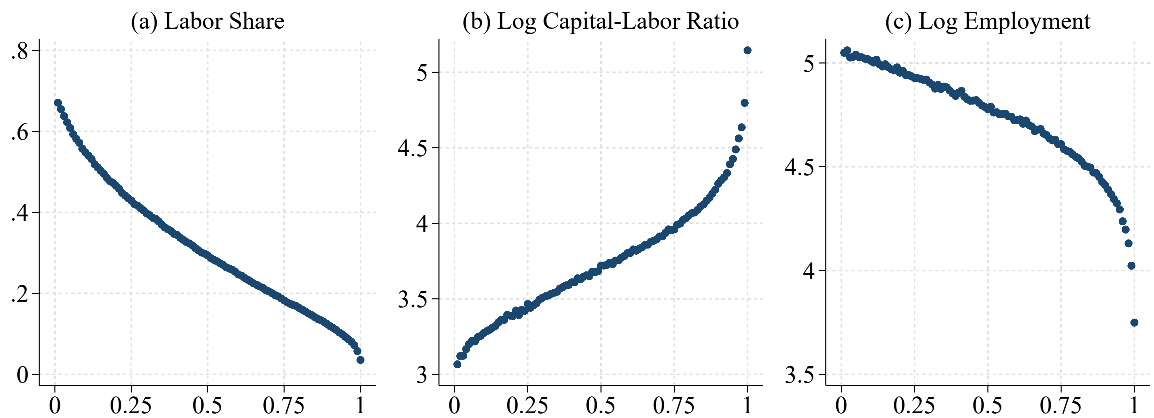

Notes: Each panel presents a binned scatterplot in which each point represents the mean outcome within one of 100 equal-frequency bins of residualized log labor productivity. Labor productivity and the outcomes are residualized with respect to market and year fixed effects. The residualized outcomes are recentered at year-specific means. Panel (a) reports the labor share, Panel (b) the log capital-labor ratio, and Panel (c) log employment.

Table 1 summarizes these patterns using percentile gaps. Within each year, we sort firms by residualized log labor productivity and compute the difference in each residualized outcome between firms at the 80th and 20th percentiles of the productivity distribution. Panel A uses current-year labor productivity. Moving from the 20th to the 80th percentile is associated with lower labor share, higher capital-labor ratio, and lower log employment, corroborating the patterns in Figure 1.

A concern with the contemporaneous sorting is that value added and employment enter both labor productivity and some of the outcomes, potentially generating mechanical correlations. Panel B addresses this concern using an odd-even measure of labor productivity. Even-year observations are sorted by firms' average residualized labor productivity in odd years, while odd-year observations are sorted by their average productivity in even years. The current observation therefore does not enter the pro-

Table 1: Factor Allocation across Labor Productivity Percentiles

	Labor share	Log capital-labor ratio	Log employment
	(1)	(2)	(3)
Panel A. Current-year labor productivity			
80–20 gap	–0.302***	0.666***	–0.405***
	(0.006)	(0.025)	(0.033)
Panel B. Odd-even labor productivity measure			
80–20 gap	–0.198***	0.594***	–0.321***
	(0.006)	(0.022)	(0.029)

Notes: Each estimate reports the difference in the outcome listed in each column between the 80th and 20th percentile groups of the relevant labor productivity measure. All variables are residualized with respect to market and year fixed effects. Panel A uses current-year residualized log labor productivity. Panel B uses an odd-even measure: even-year observations are sorted by firms' average residualized log labor productivity in odd years, and vice versa for odd-year observations. Estimates are averages of annual gaps over 1998–2007. Standard errors are reported in parentheses. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

ductivity measure used for sorting. The three gradients remain large and statistically significant under this alternative measure.

Appendix Figures A1–A8 and Table A2 provide additional robustness checks using the income-based accounting measures, lagged and leave-one-out measures of labor productivity, three- and five-year differences, ownership controls, and alternative percentile gaps. Across these exercises, higher labor productivity remains associated with lower labor shares, higher capital-labor ratios, and lower employment. Additional anatomy in Appendix Figures A9–A10 shows that more labor productive firms produce more value-added, use more capital, and pay higher average wages.

Taken together, the data show that more labor productive firms combine higher capital intensity with lower employment and lower labor shares, and that these relationships are robust across alternative specifications. The next section develops a theoretical framework that characterizes the variation in relative factor-augmenting productivity required to reconcile this joint empirical pattern.

3 Model

This section explains why heterogeneous relative factor-augmenting productivity is needed to account for the empirical patterns outlined in the previous section. We consider a two-factor CES environment with general labor and capital market frictions. We impose two empirically motivated regularity conditions: marginal labor expenditure rises with employment, and larger firms do not face higher marginal capital expenditure. We show that factor market frictions satisfying these conditions cannot alone generate the observed patterns. Relative factor augmentation must vary sufficiently across firms. We then characterize the magnitude and direction of the required variation.

3.1 Environment

Consider a continuum of firms indexed by i. Firm i produces output Y_i using capital K_i and labor L_i. Define

$$k_i \equiv \frac{K_i}{L_i}, \qquad y_i \equiv \frac{Y_i}{L_i},$$

where k_i is the capital-labor ratio, and y_i is labor productivity.[2]

Within narrowly defined markets, firms with higher labor productivity have lower employment, higher capital intensity, and lower labor shares. We express these facts as an ordered comparison along the labor productivity distribution. Let h and ℓ denote two representative firms in the same market, with h having higher labor productivity:

$$y_h > y_\ell, \qquad L_h < L_\ell, \qquad k_h > k_\ell.$$

[2] Firm-level output quantities and prices are not observed in the data, and labor productivity is measured as value-added per worker. This distinction does not affect the relative factor-augmenting productivity that is central to our analysis. As shown below, a multiplicative output price scales the two productivity components proportionally and therefore cancels from their ratio. Output prices only matter for interpreting the levels of these components. In the benchmark quantification we calibrate the model separately by market and year, and assume a common output price across firms within a market-year, with the common price absorbed into the corresponding productivity scale. We can further relax the common-price assumption by introducing differentiated-product demand. Under CES demand, the model can be written as a revenue production function, with the demand elasticity governing the mapping between revenue and physical output (Klette and Griliches, 1996; De Loecker, 2011).

Firm i has the CES production technology

$$Y_i = \left[\frac{1}{2}(A_{Ki}K_i)^{\rho} + \frac{1}{2}(A_{Li}L_i)^{\rho}\right]^{1/\rho}, \qquad A_{Ki}, A_{Li} > 0. \tag{1}$$

Here A_{Ki} is capital-augmenting productivity and A_{Li} is labor-augmenting productivity. We normalize Hicks-neutral productivity to one and the constant CES weights to one half. For $\rho \neq 0$, both normalizations can be absorbed into the two factor-augmenting productivities. They leave cross-firm variation in relative factor augmentation A_{Ki}/A_{Li} unchanged.[3]

The elasticity of substitution between capital and labor is $\sigma \equiv 1/(1-\rho) > 0$. The cases $\sigma > 1$ and $\sigma < 1$ correspond to gross substitutability and gross complementarity, respectively. Cobb-Douglas is obtained as the limit $\rho \to 0$. Using $k_i = K_i/L_i$, labor productivity is

$$y_i = A_{Ki}\left[\frac{1}{2}k_i^{\rho} + \frac{1}{2}\left(\frac{A_{Ki}}{A_{Li}}\right)^{-\rho}\right]^{1/\rho}. \tag{2}$$

The output elasticity with respect to capital is

$$\mu_i \equiv \frac{\partial \log Y_i}{\partial \log K_i} = \frac{k_i^{\rho}}{k_i^{\rho} + (A_{Ki}/A_{Li})^{-\rho}}. \tag{3}$$

Under constant returns to scale, the output elasticity with respect to labor is $1 - \mu_i$.

3.2 Factor Markets and Firm Decisions

Firms may face frictions in both factor markets. On the labor side, let $w_i = W_i(L_i)$ denote the inverse labor supply curve faced by firm i. The firm pays a common wage to all its workers, so total labor expenditure is $W_i(L_i)L_i$. Marginal labor expenditure is

$$m_i \equiv \frac{d[W_i(L_i)L_i]}{dL_i} = W_i(L_i) + L_i W_i'(L_i). \tag{4}$$

In the competitive benchmark with perfectly elastic labor supply, $m_i = W_i(L_i) = w_i$.

On the capital side, let $R_i(K_i)$ denote the inverse capital supply schedule faced by firm i. This schedule may reflect borrowing costs, collateral constraints, default premia,

[3]The normalization affects the level of A_{Ki}/A_{Li}, but not its cross-firm variation. Any alternative common CES weights rescale this ratio by the same constant for all firms.

subsidies, or other features of capital markets. Total capital expenditure is $R_i(K_i)K_i$, and marginal capital expenditure is

$$n_i \equiv \frac{d[R_i(K_i)K_i]}{dK_i} = R_i(K_i) + K_i R_i'(K_i). \tag{5}$$

In the competitive benchmark with a constant rental rate r_i, $n_i = R_i(K_i) = r_i$.

Our theoretical result relies on the following empirically grounded regularity conditions.

Assumption 1 (Factor market friction regularity)**.** *For any two firms i and j in the same market:*

(i) Marginal labor expenditure is weakly increasing with employment:

$$L_i > L_j \quad \Rightarrow \quad m_i \geq m_j.$$

(ii) Capital access is size-neutral or size-favoring. Thus, greater employment is associated with weakly lower marginal capital expenditure:

$$L_i > L_j \quad \Rightarrow \quad n_i \leq n_j.$$

The labor-side condition nests both competitive and imperfect labor markets. Under perfect competition, firms in the same labor market face a common wage, so marginal labor expenditure is constant across firms. Under imperfect competition, standard models generally imply that marginal labor expenditure rises with employment. Appendix C establishes the corresponding ordering for the employment component generated by worker flows in a Burdett-Mortensen wage-posting model and derives the analogous property in a BHM-style model of finite-firm competition.[4]

The capital side condition is consistent with size-dependent financial frictions.[5]

[4] Azar and Marinescu (2024) distinguish three broad sources of labor market power: oligopsony, job differentiation, and search frictions. The model of Burdett and Mortensen (1998) represents the search-friction mechanism. The BHM framework of Berger et al. (2022) combines differentiated jobs with strategic interactions among a finite number of firms and therefore provides a benchmark for both job differentiation and oligopsony. Together, the two models cover the three broad mechanisms relevant for our assumption.

[5] Consistent with this condition, Panel (a) of Figure 2 shows that observed borrowing costs are negatively correlated with firm size in the data. In the quantitative analysis, Panel (b) shows that calibrated marginal capital expenditure is also negatively correlated with firm size.

Larger firms may have more collateral, greater retained earnings, stronger repayment histories, and more established banking relationships. These advantages can relax borrowing constraints and lower the marginal cost of capital. This is the benchmark mechanism in Bai et al. (2026): productive firms grow, accumulate internal funds, and gradually relax their borrowing constraints. The condition still allows capital costs to be heterogeneous and capital supply schedules to slope upward. It requires only that larger firms do not face higher marginal capital expenditure along the equilibrium allocation profile.

At an interior optimum, the firm equates each factor's marginal revenue product to its marginal expenditure:

$$MRPK_i = n_i, \qquad MRPL_i = m_i.$$

Under constant returns to scale, these first-order conditions imply

$$n_i k_i = \mu_i y_i, \tag{6}$$

$$m_i = (1 - \mu_i) y_i. \tag{7}$$

Taking the ratio of the factor first-order conditions and using the CES technology gives

$$k_i = \left(\frac{m_i}{n_i}\right)^{\sigma} \left(\frac{A_{Ki}}{A_{Li}}\right)^{\sigma-1}. \tag{8}$$

Equation (8) is the central relative factor-demand condition. Capital intensity is determined by relative marginal factor expenditure, m_i/n_i, and relative factor augmentation, A_{Ki}/A_{Li}.

3.3 Relative Factor-Augmenting Productivity

Proposition 1 shows that the motivating facts require relative factor-augmenting productivity to vary systematically across firms.

Proposition 1 (Relative factor augmentation along the labor-productivity profile)**.** *Suppose production is CES and Assumption 1 holds. Consider two representative points h and ℓ on the within-market allocation profile such that*

$$y_h > y_\ell, \qquad L_h < L_\ell, \qquad k_h > k_\ell.$$

Relative factor augmentation must satisfy

$$\begin{cases} \dfrac{A_{Kh}/A_{Lh}}{A_{K\ell}/A_{L\ell}} \geq \left(\dfrac{k_h}{k_\ell}\right)^{1/(\sigma-1)} > 1, & \sigma > 1, \\ 0 < \dfrac{A_{Kh}/A_{Lh}}{A_{K\ell}/A_{L\ell}} \leq \left(\dfrac{k_h}{k_\ell}\right)^{1/(\sigma-1)} < 1, & \sigma < 1. \end{cases}$$

If $\sigma = 1$, no configuration of relative factor augmentation can generate the observed empirical pattern.

Proof. Taking logs of equation (8) and subtracting the equation for ℓ from that for h gives

$$\begin{aligned}(\sigma - 1)&\left[\log\left(\frac{A_{Kh}}{A_{Lh}}\right) - \log\left(\frac{A_{K\ell}}{A_{L\ell}}\right)\right] \\ &= \log\left(\frac{k_h}{k_\ell}\right) - \sigma\left[\log\left(\frac{m_h}{n_h}\right) - \log\left(\frac{m_\ell}{n_\ell}\right)\right]. \end{aligned} \tag{9}$$

Because $L_h < L_\ell$, Assumption 1 implies

$$m_h \leq m_\ell, \qquad n_h \geq n_\ell,$$

and hence

$$\frac{m_h}{n_h} \leq \frac{m_\ell}{n_\ell}.$$

Together with $k_h > k_\ell$, this result implies that the right-hand side of equation (9) is at least $\log(k_h/k_\ell) > 0$. Dividing by $\sigma - 1$ and exponentiating gives the stated bounds for $\sigma > 1$ and $\sigma < 1$.

When $\sigma = 1$, equation (8) implies

$$\frac{k_h}{k_\ell} = \frac{m_h/n_h}{m_\ell/n_\ell} \leq 1,$$

which contradicts $k_h > k_\ell$. □

The proposition has a simple economic interpretation. Firms with higher labor productivity are smaller but more capital intensive. Under Assumption 1, smaller firms face weakly lower marginal labor expenditure and weakly higher marginal capital expenditure. Their relative factor cost, m_i/n_i, is therefore weakly lower. If relative factor augmentation were constant across firms, equation (8) would imply that these firms are weakly less capital intensive. The data show the opposite. Relative factor augmentation must therefore vary sufficiently to offset the factor-cost channel. The ratio A_{Ki}/A_{Li} must be higher among high-labor-productivity firms when $\sigma > 1$ and lower when $\sigma < 1$. Recent evidence from Chinese manufacturing supports the former case. In particular, Meng et al. (2025) estimate capital-labor substitution elasticities well above one at both the firm and aggregate manufacturing levels. We therefore use $\sigma > 1$ as our benchmark in the quantitative analysis.

The capital-side restriction in Assumption 1 is sufficient but not necessary. The proposition continues to hold whenever relative marginal factor expenditure is weakly increasing with employment:

$$L_i > L_j \quad \Rightarrow \quad \frac{m_i}{n_i} \geq \frac{m_j}{n_j}.$$

Thus, marginal capital expenditure may rise with firm size, provided that it does not rise enough to reverse the relationship between m_i/n_i and employment. A constant rental rate and size-favoring capital access are both special cases.

The constant-returns specification is not essential for the relative factor-demand result. Consider the homothetic extension

$$Y_i = \left[\frac{1}{2}(A_{Ki}K_i)^\rho + \frac{1}{2}(A_{Li}L_i)^\rho\right]^{\nu/\rho}, \qquad \nu > 0,$$

where ν governs returns to scale and $\nu = 1$ is the benchmark. The parameter ν affects the marginal products of capital and labor through the same multiplicative term, which cancels from their ratio. Equation (8), Proposition 1, and its proof are therefore unchanged. Non-constant returns affect the level mapping from productivity to output and factor shares, but not the necessary condition on relative factor augmentation.

The result also clarifies why factor-augmenting productivities matter for the observed positive relationship between labor shares and employment. Define the labor markdown ratio as $\tau_i \equiv m_i/w_i$. Equation (7) implies

$$LS_i \propto \frac{1-\mu_i}{\tau_i}. \tag{10}$$

Under our benchmark case with $\sigma > 1$, firms with higher labor productivity must have higher A_{Ki}/A_{Li}. Together with their higher capital intensity, this raises μ_i and reduces the output elasticity of labor. Factor-augmenting productivity therefore helps explain why the high-labor-productivity firms in the data have both lower employment and lower labor shares.

4 Quantification

More labor productive firms are smaller, more capital intensive, and have lower labor shares. The previous section shows that rationalizing these patterns requires substantial heterogeneity in relative factor augmentation. This section quantifies that heterogeneity using firm-level data. Wages, employment, labor shares, capital intensity, and labor productivity are observed. The labor market structure uses wages and employment to recover a firm's labor markdown. The CES production structure then uses the recovered markdowns and the remaining observables to recover output elasticity of labor and factor-augmenting productivities.

4.1 Labor Market Structure

We focus on steady-state equilibrium in each local labor market and year. For exposition, we suppress the market and time indices whenever there is no confusion.

The theoretical result allows for general frictions in both factor markets. Quantification requires additional structure because labor markdowns are not directly observed. We therefore model the labor market as a Burdett-Mortensen wage-posting environment.[6] We treat firm-specific marginal capital expenditure n_i as a reduced-form capital market wedge. We discuss this asymmetric treatment of the two factor markets in Section 4.3.

Firms post wages, and all workers are employed and search on the job. Workers receive outside offers and move when an offer pays more than their current wage. Employment relationships end exogenously at rate $\delta^L > 0$, and affected workers are immediately reallocated. The model therefore abstracts from unemployment. Worker flows determine the rank-dependent component of employment along the wage ladder,

[6] We use the Burdett-Mortensen structure as our benchmark because it allows labor markdowns to decline with firm employment. Consistent with this feature, previous studies recover firm-level wage markdowns from production function estimation and find that they are negatively correlated with firm employment (Lu et al., 2019; Pham, 2023). The markdowns recovered from our calibration exhibit the same pattern. In differentiated-oligopsony models, larger firms instead typically face less elastic labor supply and have larger markdowns, as in Berger et al. (2022).

and firms account for this response when choosing their wages.[7]

We initially take the distribution of $MRPL$ at firms' optima as given. Firms with higher $MRPL$ post higher wages in equilibrium. We call this alignment between the $MRPL$ ordering and the wage ordering the rank-preservation property. We index firms by their $MRPL$ rank $q \in [0,1]$, which also indexes their position on the equilibrium wage ladder. Workers receive outside offers at rate $\lambda > 0$. This rate is estimated separately and may therefore differ across market-years. We also allow for a firm-specific employment shifter ψ_i.[8] The firm treats ψ_i as given when posting its wage, even though the realized shifter may be statistically associated with the firm's equilibrium wage rank. The employment shifter captures amenities, location, hiring capacity, and other factors that affect employment conditional on rank.

Let $w(q)$, $L(q)$, and $m(q)$ denote the equilibrium wage, the Burdett-Mortensen employment component, and marginal labor expenditure at rank q. At the firm's optimum, marginal labor expenditure equals $MRPL$, so $m(q)$ also denotes the equilibrium $MRPL$ profile. For firm i with rank q_i,

$$w_i = w(q_i), \qquad m_i = m(q_i), \qquad L_i = \psi_i L(q_i). \tag{11}$$

Steady-state worker-flow balance implies

$$L(q) = \frac{B}{\left[\delta^L + (1-q)\lambda\right]^2}, \tag{12}$$

where $B > 0$ is a constant and determines the scale of baseline employment. Differentiating equation (12) gives

$$g(q) \equiv \frac{d \log L(q)}{dq} = \frac{2\lambda}{\delta^L + (1-q)\lambda}. \tag{13}$$

The wage-posting first-order condition determines the wage profile:

$$\frac{dw(q)}{dq} = g(q)[m(q) - w(q)]. \tag{14}$$

[7]Appendix C derives the worker-flow equations and rank representation used below.

[8]Assumption 1 and Proposition 1 provide a sharp theoretical benchmark rather than restrictions imposed in the calibration. The quantitative model distinguishes the employment profile $L(q_i)$ generated by Burdett-Mortensen worker flows from observed firm employment $L_i = \psi_i L(q_i)$. The employment shifter accommodates additional determinants of firm employment and measurement error, while marginal capital expenditure is recovered without imposing the size ordering in Assumption 1.

The corresponding labor markdown is

$$\tau(q) \equiv \frac{m(q)}{w(q)}. \tag{15}$$

As a simplifying assumption, we impose a minimum admissible wage equal to the lowest firm's marginal revenue product of labor, so that $w(0) = m(0)$. The lowest-ranked firm therefore has no labor markdown. This boundary condition anchors the level of the markdown profile while allowing its shape to vary with rank.

We refer to the job-arrival rate λ and the firm-specific employment shifters $\{\psi_i\}$ as the *labor market primitives*. Conditional on the distribution of m_i, λ determines the wage and markdown profiles as well as the Burdett–Mortensen employment component, while ψ_i maps that component into observed firm employment.

$$\{m_i, \lambda, \psi_i\} \longrightarrow \{w_i, L_i, \tau_i\}.$$

The distribution of m_i determines firms' latent ranks.

4.2 Mapping the Primitives to the Data

The previous subsection characterized the labor market equilibrium while taking the distribution of marginal labor expenditures m_i as given. We now determine this distribution from the CES technology and firms' factor choices. This closes the model.

Under the normalizations in Section 3, firm i is characterized by A_{Ki}, A_{Li}, and n_i. We refer to these objects as the *production primitives*. We treat n_i as a reduced-form primitive because the labor market structure is sufficient for recovering the factor-augmenting productivities. Section 4.3 provides the rationale for this choice.

Given the production primitives and ρ, equation (6) and the CES technology determine the output elasticity of capital:

$$\mu_i = \left[\frac{n_i}{A_{Ki}(1/2)^{1/\rho}}\right]^{\rho/(\rho-1)}. \tag{16}$$

Conditional on μ_i, capital intensity and labor productivity are

$$k_i = \frac{A_{Li}}{A_{Ki}} \left(\frac{\mu_i}{1-\mu_i}\right)^{1/\rho}, \tag{17}$$

and

$$y_i = A_{Ki} \left[\frac{(1/2)(A_{Li}/A_{Ki})^{\rho}}{1-\mu_i} \right]^{1/\rho}. \tag{18}$$

Equation (7) then determines m_i.

Collecting all firms, the cross-sectional distribution of m_i determines each firm's latent rank q_i. Given λ and ψ_i, the labor market equilibrium determines w_i, L_i, and τ_i. The labor share follows from equation (10). We refer to the labor market primitives and production primitives jointly as the *model primitives*. We collect the relationships from these model primitives to the observables in the equilibrium mapping

$$\mathcal{E} : \{A_{Ki}, A_{Li}, n_i, \lambda, \psi_i\} \longmapsto \{w_i, L_i, LS_i, k_i, y_i\}.$$

The mapping is defined jointly across firms within a market-year because each firm's labor market outcomes depend on its rank in the distribution of marginal labor expenditures.

4.3 Calibration

We use the production normalization in Section 3. We set $\sigma = 4.36$, the aggregate capital-labor substitution elasticity estimated for Chinese manufacturing by Meng et al. (2025), and set $\rho = 1 - 1/\sigma$. We set the annual job-separation rate to $\delta^L = 0.08$ and measure the job-arrival rate in annual units. The benchmark job-arrival rate λ varies across market-years.

The following proposition solves the pointwise factor-augmenting productivities conditional on labor markdown.

Proposition 2 (Recovery of factor-augmenting productivities)**.** *Suppose $\rho \neq 0$ is known and the labor markdown τ_i has been recovered. If $0 < \tau_i LS_i < 1$, then LS_i, k_i, and y_i uniquely determine the output elasticity of labor, the two factor-augmenting productivities, and marginal capital expenditure. In particular,*

$$1 - \mu_i = \tau_i LS_i, \tag{19}$$

$$\frac{A_{Li}}{A_{Ki}} = k_i \left(\frac{1-\mu_i}{\mu_i} \right)^{1/\rho}, \tag{20}$$

and

$$A_{Ki} = \frac{y_i}{\left[\frac{1}{2}k_i^{\rho} + \frac{1}{2}(A_{Li}/A_{Ki})^{\rho}\right]^{1/\rho}}, \quad A_{Li} = \left(\frac{A_{Li}}{A_{Ki}}\right) A_{Ki}. \tag{21}$$

Marginal capital expenditure is

$$n_i = \frac{\mu_i y_i}{k_i}.$$

Proof. Equation (10) gives (19). Solving equation (3) for A_{Li}/A_{Ki} gives (20). Substituting the recovered ratio into equation (2) gives (21). The level of A_{Li} then follows from the recovered ratio. Equation (6) gives $n_i = \mu_i y_i / k_i$. □

The proposition explains why labor market structure is essential for the quantification. The wage w_i is observed, but the labor markdown $\tau_i = m_i/w_i$ is not. Equation (10) shows that the labor share alone identifies only a combination of the output elasticity of labor and the markdown. Production function approaches that infer markdowns from factor shares face the same problem when factor-augmenting productivity is allowed to vary across firms. The Burdett-Mortensen equilibrium provides the additional restrictions needed to recover markdowns from wages and employment.

Recovering the factor-augmenting productivities requires an equilibrium model for at least one factor market. Our benchmark models the labor market, after which the labor share identifies μ_i. Alternatively, an explicit capital market structure could recover n_i, and equation (6) would then identify μ_i. We model the labor market because wages and employment are measured relatively directly in the firm data. Firm-level rental-rate proxies are noisier and more frequently missing. A structural capital market model could explain the cross-firm variation in n_i, but it is not needed to recover factor-augmenting productivity conditional on observed capital intensity.

We next describe how the marginal distributions of wages and employment reveal labor markdowns. Let d and t index markets and years. We impose the following assumption on the employment shifter ψ_{idt}:

Assumption 2 (Employment shifters). *For a fixed finite $\bar{p}$, the employment shifter satisfies*

$$\log \psi_{idt} = \phi_i + \xi_{dt} + \sum_{p=1}^{\bar{p}} \beta_{\psi,p} q_{idt}^{p} + \upsilon_{idt}, \tag{22}$$

We require υ_{idt} to have mean zero conditional on firm identity and its entire rank and market-year histories.

Here ϕ_i captures persistent employment differences associated with location or hiring capacity; ξ_{dt} captures employment conditions common to a market-year; the polynomial $\sum_{p=1}^{\bar{p}} \beta_{\psi,p} q_{idt}^p$ captures the systematic association between the shifter and firm rank; and υ_{idt} captures transitory differences and measurement error.[9]

Combining Assumption 2 with equation (12) gives

$$\log L_{idt} = \underbrace{\log B_{dt} - 2\log\left[\delta^L + (1-q_{idt})\lambda_{dt}\right]}_{\equiv \log L_{dt}(q_{idt})} + \phi_i + \xi_{dt} + \sum_{p=1}^{\bar{p}} \beta_{\psi,p} q_{idt}^p + \upsilon_{idt}. \tag{23}$$

Throughout the calibration, w_{idt}, L_{idt}, LS_{idt}, k_{idt}, and y_{idt} denote observed firm-level variables, and q_{idt} is the firm's observed wage rank. Let us define

$$r_{idt} = \begin{pmatrix} \log w_{idt} - \log w_{dt}(q_{idt}) \\ \log L_{idt} - \left[\log L_{dt}(q_{idt}) + \phi_i + \xi_{dt} + \sum_{p=1}^{\bar{p}} \beta_{\psi,p} q_{idt}^p\right] \\ LS_{idt} - \dfrac{(1-\mu_{idt}) w_{dt}(q_{idt})}{m_{dt}(q_{idt})} \\ \log y_{idt} - \log\left[\dfrac{m_{dt}(q_{idt})}{1-\mu_{idt}}\right] \end{pmatrix}. \tag{24}$$

For a market-year containing N_{dt} firms, we impose the normalization condition

$$\frac{1}{N_{dt}} \sum_{i\in(d,t)} \log \psi_{idt} = 0. \tag{25}$$

Let Θ collect the firm and market-year effects ϕ_i and ξ_{dt}, polynomial coefficients $\{\beta_{\psi,p}\}$, arrival rates λ_{dt}, employment scales B_{dt}, and $MRPL$ profiles $m(q)$. Let Ω_{dt} standardize the four residuals in (24). For numerical stability, we introduce a penalty term $\mathcal{I}_{dt}$ for candidates near the economic boundary such that $\mu_{idt} \approx 0$. The calibration solves the following minimum-distance problem:

$$\widehat{\Theta} = \arg\min_{\Theta} \sum_{d,t} \left\{ \sum_{i\in(d,t)} r'_{idt} \Omega_{dt} r_{idt} + N_{dt}\mathcal{I}_{dt} \right\}. \tag{26}$$

[9] Even if Assumption 2 does not hold, our main objects remain identified when additional information on job-to-job transition rates identifies λ_{dt}. The wage-posting equations then recover labor markdowns, and Proposition 2 recovers factor-augmenting productivities.

The minimization is subject to $\lambda_{dt} > 0$, $B_{dt} > 0$, a positive, weakly increasing $m(q)$, $0 < \mu_{idt} < 1$, the equilibrium equations, the boundary condition $w_{dt}(0) = m_{dt}(0)$, and normalization (25). Appendix D describes the implementation details. Proposition 3 summarizes the identification result.

Proposition 3 (Identification of markdowns and productivities)**.** *Under Assumption 2, the (non-degenerate) joint panel distribution of wages, employment, labor shares, capital intensity, and labor productivity uniquely identifies firm-level labor markdowns and factor-augmenting productivities.*

Proof. By rank preservation, the observed wage ordering coincides with the $MRPL$ ordering. q_{idt} is therefore observed from the wage distribution. We assume that, within each market, firms' wage ranks in each year have continuous variation even after conditioning on their ranks in all other years. Consider two sets of primitives satisfying the assumption 2 and generating the same joint population distribution of observables. Taking within-firm differences between years t and s removes firm effects. Conditional on the firm's wage-rank history, differentiating $\bar{p}+1$ times with respect to its rank in year t, holding all other ranks fixed, eliminates market-year effects and the rank polynomial:

$$\frac{\partial^{\bar{p}+1}}{\partial q_{idt}^{\bar{p}+1}}\mathbb{E}\left[\log\left(\frac{L_{idt}}{L_{ids}}\right)\,\middle|\,\{q_{idu}\}_u, d\right] = 2\bar{p}!\left[\frac{\lambda_{dt}}{\delta^L + (1-q_{idt})\lambda_{dt}}\right]^{\bar{p}+1}. \tag{27}$$

For every $q \in (0,1)$, the expression in brackets is strictly increasing in λ_{dt}. The two sets of model primitives must therefore imply the same λ_{dt} and, by equation (13), the same employment gradient $g_{dt}(q)$ along the wage ladder. Rearranging the wage-posting condition (14) gives

$$m_{dt}(q) = w_{dt}(q) + \frac{w'_{dt}(q)}{g_{dt}(q)}.$$

The common $g_{dt}(q)$ therefore implies the same $MRPL$ schedule $m_{dt}(q)$. Equation (15) then gives the same labor markdown:

$$\tau_{dt}(q) = \frac{m_{dt}(q)}{w_{dt}(q)}.$$

Conditional on this markdown, Proposition 2 provides a pointwise one-to-one mapping from observed labor shares, capital intensity, and labor productivity into A_{Li} and

A_{Ki}. Labor markdowns and factor-augmenting productivities are therefore uniquely identified by the joint population distribution of observables. $\square$

In the Burdett-Mortensen equilibrium, a worker employed by a firm at rank q receives a better offer at rate $(1-q)\lambda_{dt}$. Firms higher on the wage ladder therefore lose fewer workers to competing firms and accumulate more employment. The arrival rate λ_{dt} determines how rapidly this retention advantage increases with rank, generating curvature in the log employment profile. Consider the case $\bar{p}=1$. The employment shifter ψ_{idt} contributes only a linear relationship between log employment and wage rank. After removing firm and market-year effects, this linear component can affect the slope of the employment profile but not its curvature. The curvature of $\log L_{idt}$ with respect to q_{idt} identifies $\widehat{\lambda}_{dt}$, while the remaining linear relationship identifies $\widehat{\beta}_{\psi,1}$. The normalization (25) fixes the employment scale $\widehat{B}_{dt}$. Given $\widehat{\lambda}_{dt}$, the wage-posting condition reveals the model-implied wage schedule $\widehat{w}_{dt}(q)$, the firm's surplus from hiring an additional worker, $\widehat{m}_{dt}(q)-\widehat{w}_{dt}(q)$, and hence its labor markdown $\widehat{\tau}_{idt}$.

4.4 Model Fit and Validation

This section evaluates the benchmark calibration along two dimensions. We first examine its in-sample fit by comparing selected marginal and joint features of firm outcomes in the data and the benchmark model. We then evaluate the economic interpretation of recovered marginal capital expenditure, n_i, using firms' borrowing costs and ownership, neither of which enters the calibration. The first exercise assesses whether the calibrated model reproduces the firm-level allocation, while the second provides nontargeted evidence on the recovered capital market object.

Model Fit. Table 2 assesses how well the benchmark model reproduces selected features of the marginal and joint distributions of firm outcomes. Because wages, employment, labor shares, and labor productivity enter the minimum-distance criterion, while capital intensity enters the recovery of the production primitives, the table summarizes the in-sample fit of the calibration.

Panel A reports means and standard deviations. The model closely reproduces the distribution of labor productivity: the model-implied mean and standard deviation are 3.89 and 0.95, compared with 3.91 and 0.96 in the data. Employment and capital intensity are matched exactly by construction. The model also closely reproduces the means and standard deviations of labor shares and wages.

Table 2: Model Fit

	Data		Model	
	(1)		(2)	
Panel A. Means and standard deviations				
	Mean	SD	Mean	SD
$\log y$	3.91	0.96	3.89	0.95
$\log L$	4.72	1.06	4.72	1.06
$\log(K/L)$	3.82	1.14	3.82	1.14
LS	0.30	0.21	0.30	0.20
$\log w$	2.43	0.62	2.41	0.54
Panel B. Correlations				
$\text{Corr}(LS, \log y)$	–0.73		–0.80	
$\text{Corr}(\log(K/L), \log y)$	0.42		0.41	
$\text{Corr}(\log L, \log y)$	–0.20		–0.20	
$\text{Corr}(\log w, \log y)$	0.45		0.38	
$\text{Corr}(\log L, \log(K/L))$	–0.05		–0.05	
$\text{Corr}(\log L, LS)$	0.21		0.22	
$\text{Corr}(\log L, \log w)$	0.02		0.04	
$\text{Corr}(\log(K/L), LS)$	–0.27		–0.29	
$\text{Corr}(\log(K/L), \log w)$	0.26		0.26	
$\text{Corr}(LS, \log w)$	0.13		0.08	

Notes: This table reports statistics from the data and from the benchmark model. Data statistics are computed using Chinese manufacturing firm data for 1998–2007. Panel A characterizes the marginal distribution of each variable using its mean and standard deviation. Panel B characterizes the joint distribution by reporting pairwise correlations among variables after removing market-year fixed effects. The variable y denotes value-added per worker, L employment, K real capital, LS the labor share, and w average labor compensation per worker.

Panel B reports pairwise correlations among variables within market-year cells, with all variables residualized using market-year fixed effects. The model reproduces the negative relationships of labor productivity with employment and the labor share and its positive relationships with capital intensity and wages. It also closely matches the relationships of labor shares and wages with employment and capital intensity. The employment-wage correlation is slightly stronger in the model than in the data. This difference may partly reflect the model's abstraction from worker heterogeneity and the sorting of workers across firms, which generate additional variation in firm-level average wages through workforce composition.[10] Overall, the benchmark model captures the principal relationships governing the joint allocation of inputs and outcomes across firms.

Validation of Marginal Capital Expenditure. The calibration yields firm-specific marginal capital expenditure, n_i. We assess its economic interpretation using two sets of moments that are not targeted in the calibration: the size gradient in n_i and differences in n_i across ownership types.

Figure 2 examines the relationship between capital costs and firm size. Panel (a) documents an empirical pattern in the data: borrowing costs, measured as interest payments divided by total liabilities, decline with firm size. Panel (b) shows that calibrated marginal capital expenditure also declines with firm size, providing a nontargeted validation of the firm-specific capital cost measure.

As a complementary validation, Table 3 relates recovered marginal capital expenditure to firm ownership. We regress n_i on indicators for state-owned and foreign-invested enterprises. Columns (1) and (2) include market and year fixed effects, while columns (3) and (4) include market-year fixed effects. State-owned enterprises have substantially lower recovered marginal capital expenditure than private firms. Foreign-invested enterprises also have lower marginal capital expenditure, although the difference is smaller. These ownership gradients are consistent with more favorable access to capital among state-owned and foreign-invested firms.

5 Calibrated Firm-Level Objects

5.1 Factor-Augmenting Productivity

Figure 3 plots the evolution of capital- and labor-augmenting productivity over 1998–2007, normalized to zero in 1998. Both rise substantially over the sample period. The two series track each other closely through 2000. Thereafter, capital-augmenting productivity grows more rapidly, shifting relative factor-augmenting productivity toward capital.

[10] Gouin-Bonenfant (2022) similarly argues that a model with homogeneous workers cannot account for the full dispersion of observed wages because part of earnings variation reflects worker-specific

Figure 2: Borrowing Costs, Marginal Capital Expenditure, and Firm Size

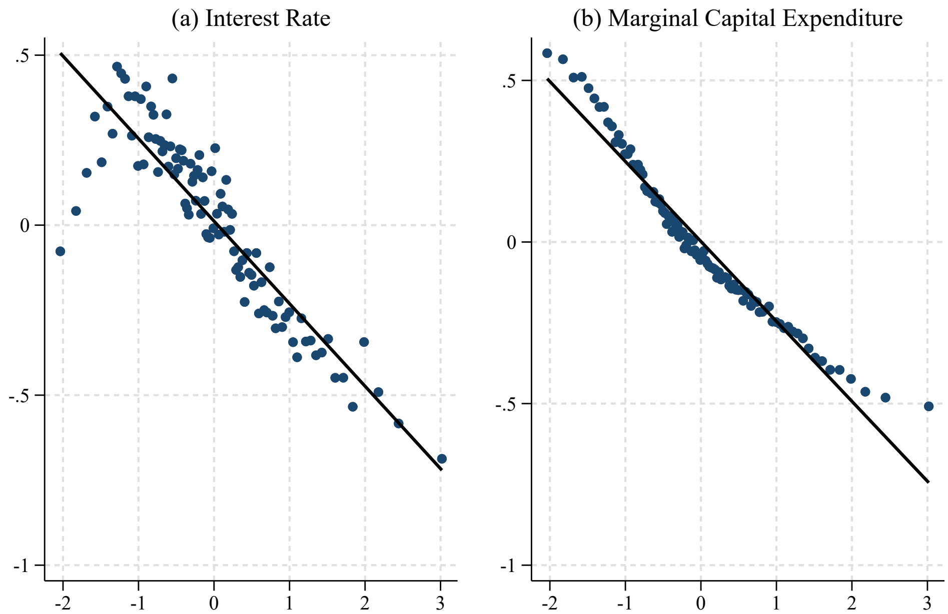

Notes: This figure plots borrowing costs and calibrated marginal capital expenditure against firm size. Panel (a) measures borrowing costs as interest payments divided by total liabilities, and Panel (b) reports calibrated marginal capital expenditure. Firm size is measured by log employment. Firm size and both outcomes are residualized with respect to market-year fixed effects. For each year, firms are sorted into 100 equal-frequency bins based on firm size. Each point reports the 1998–2007 average of the within-bin mean firm size and corresponding mean outcome. The solid lines report linear fits across the bin averages.

Table A1 summarizes the cross-sectional distribution of relative factor-augmenting productivity and its two components. Relative factor augmentation exhibits substantial firm-level heterogeneity. Across the sample, the average annual interquartile range of $\log(A_K/A_L)$ is 2.38. The dispersion is also similar across industries, with interquartile ranges between 2.07 and 2.59. Dispersion is substantially greater in capital-augmenting productivity, with an average annual interquartile range of 2.04 for $\log A_K$ and 0.67 for $\log A_L$.

5.2 VAT Reform and Factor-Augmenting Productivity

We next use a policy-induced reduction in the cost of capital to assess the economic interpretation of the recovered A_K and A_L. The reform directly altered incentives for

heterogeneity.

Table 3: Ownership and Marginal Capital Expenditure

	Marginal capital expenditure, n_i			
	(1)	(2)	(3)	(4)
SOE dummy	–0.641***	–0.434***	–0.638***	–0.440***
	(0.017)	(0.016)	(0.015)	(0.015)
FIE dummy	–0.339***	–0.340***	–0.338***	–0.338***
	(0.024)	(0.020)	(0.025)	(0.020)
Firm controls	No	Yes	No	Yes
Market fixed effects	Yes	Yes	No	No
Year fixed effects	Yes	Yes	No	No
Market-year fixed effects	No	No	Yes	Yes
Observations	1,446,386	1,446,386	1,446,386	1,446,386

Notes: This table reports coefficients from firm-level regressions of calibrated marginal capital expenditure, n_i, on indicators for state-owned enterprises (SOE) and foreign-invested enterprises (FIE). Columns (1) and (2) include market and year fixed effects, whereas columns (3) and (4) include market-year fixed effects. Firm controls include exporter status, size indicators, and age-category indicators. Standard errors clustered at the market level are reported in parentheses. ***, **, and * denote significance at the 1, 5, and 10 percent levels, respectively.

equipment investment, providing external variation with which to examine whether the two recovered productivity components respond differently to the reform.

Under China's production-based value-added tax system, firms could deduct input VAT on intermediate inputs but not on purchases of fixed assets, thereby raising the effective acquisition cost of machinery and equipment. In 2003, the government announced a pilot transition to a consumption-based VAT system as part of its strategy to revitalize the Northeast industrial base. Beginning in July 2004, eligible firms in the Northeast could deduct input VAT on qualifying fixed-asset purchases.[11] Liu and Lu (2015) show that the pilot raised fixed-asset investment among eligible firms. Because the reform was confined to the Northeast, it provides a geographically concentrated reduction in the effective price of equipment investment for firms across a broad range of manufacturing industries.[12] We exploit this policy-induced capital-cost

[11]The pilot covered equipment manufacturing, petroleum processing and chemicals, metallurgy, shipbuilding, automobile manufacturing, and agricultural-product processing, with coking, electrolytic aluminum, and small-scale steel production explicitly excluded.

[12]The pilot was extended to 26 prefectures in six central provinces in 2007, and the consumption-based VAT system was adopted nationwide in 2009. We restrict the sample to 1998–2006, before these subsequent extensions.

Figure 3: Evolution of Aggregate Factor-Augmenting Productivity

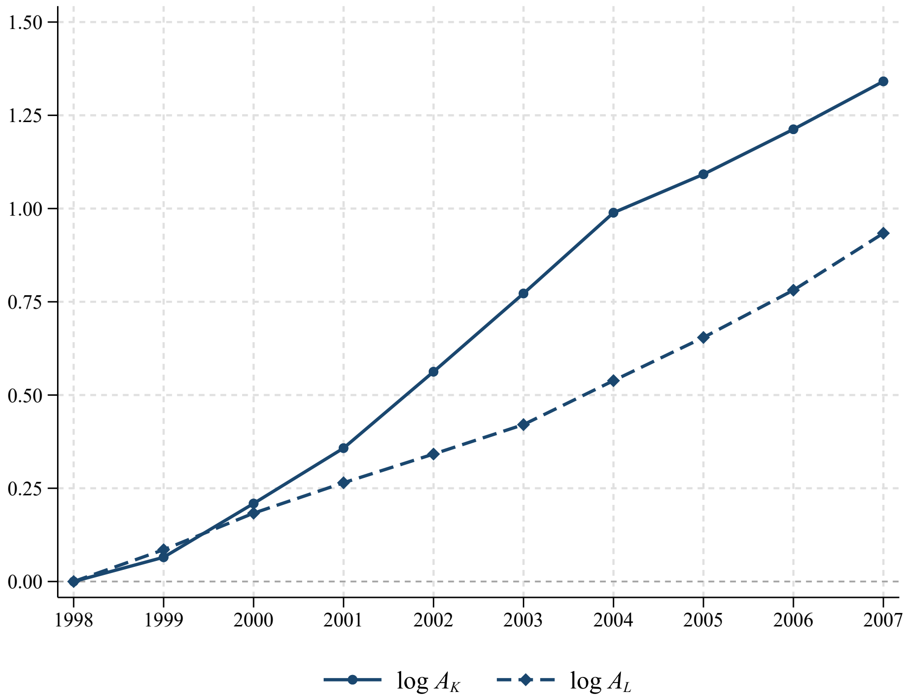

Notes: The figure plots the annual mean of firm-level capital- and labor-augmenting productivity, $\log A_K$ and $\log A_L$, over 1998–2007. The means are unweighted across firms, and both series are normalized to zero in 1998.

shock to examine whether the calibrated productivity components respond in the factor-specific direction implied by the reform. If lower equipment costs facilitate the adoption of newer vintages, the reform would raise capital-augmenting productivity relative to labor-augmenting productivity, with little direct effect on labor-augmenting productivity.

We consider the following event-study specification:

$$y_{idt} = \gamma_d + \gamma_{st} + \sum_{\tau \neq 2002} \beta_\tau D_d \mathbf{1}\{t = \tau\} + \sum_{\tau \neq 2002} \mathbf{X}_d' \Gamma_\tau \mathbf{1}\{t = \tau\} + \varepsilon_{idt},$$

where i, d, and t index firms, markets, and years, respectively. The outcome y_{idt} is $\log(A_{K,idt}/A_{L,idt})$, $\log A_{L,idt}$, or $\log A_{K,idt}$. The treatment indicator D_d equals one if market d is exposed to the Northeast VAT pilot, and zero otherwise. The specification includes market fixed effects, γ_d, and industry-year fixed effects, γ_{st}, which absorb nationwide changes specific to each industry. The vector $\mathbf{X}_d$ contains the initial-year values of SOE share, foreign ownership share, log employment, and export share in market d, each interacted with year indicators. We normalize the treatment coefficients to zero in 2002, the year preceding the policy announcement. Standard errors are clustered at the market level.

Figure 4 reports the dynamic effects of the VAT reform. Prior to the policy announcement, there is no evidence of a differential trend in relative factor-augmenting productivity. The estimate in 2003 remains small, suggesting a limited response to the announcement before the reform was implemented. Following implementation, relative factor-augmenting productivity rises steadily, with the estimated effect reaching approximately 40 percent by 2006. This increase is driven primarily by capital-augmenting productivity, while labor-augmenting productivity remains largely unchanged.

Figure 4: Dynamic Effects of VAT Reform on Factor-Augmenting Productivity

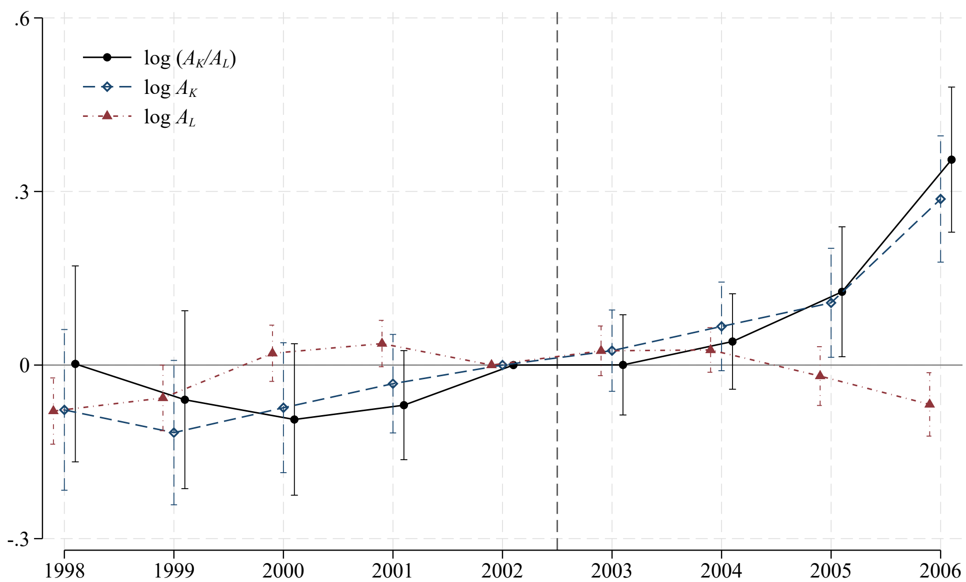

Notes: This figure plots event-study estimates of the effects of the VAT reform on calibrated A_K, A_L, and relative factor-augmenting productivity. Treatment equals one for markets exposed to the Northeast VAT pilot. Coefficients are normalized to zero in 2002, the year preceding the policy announcement. All regressions include market and industry-year fixed effects, together with interactions between year indicators and the initial-year values of SOE share, foreign ownership share, log employment, and export share. Standard errors are clustered at the market level. Vertical bars show 95 percent confidence intervals. The dashed vertical line marks the 2003 policy announcement.

One interpretation of this response is capital-embodied upgrading. By lowering the effective acquisition cost of machinery and equipment, the reform may facilitate the adoption of newer vintages. New equipment may embody improved technologies or provide higher-quality capital services, raising the productive effectiveness of capital beyond the increase in its recorded quantity. This interpretation is consistent with the investment-specific technical change literature, which distinguishes improvements embodied in new capital goods from Hicks-neutral productivity growth (Greenwood et al., 1997).

5.3 Labor Markdowns

Table 4 summarizes the calibrated labor markdowns and examines their relationship with firm characteristics. Panel A shows substantial labor market power in the data. The unweighted mean and median markdowns are 1.37 and 1.32, respectively, while the corresponding employment-weighted values are 1.34 and 1.30. The median markdown implies that workers receive about 76 percent of their marginal revenue product.

Panel B relates log markdowns to firm size and other firm characteristics. The first specification controls for industry, province, and year fixed effects. The second specification additionally includes indicators for state-owned enterprises and foreign-invested enterprises, exporter status, and firm age. Labor markdowns decline significantly with firm size in both specifications. An interquartile increase in log employment is associated with about a 0.9 percent decline in markdowns, equivalent to roughly 5 percent of the interquartile range of log markdowns. The negative size gradient is consistent with labor markdowns estimated using the production function approach (Lu et al., 2019; Pham, 2023). State-owned, foreign-invested, and exporting firms also exhibit lower markdowns.

5.4 Role of Factor-Augmenting Heterogeneity

The calibration recovers substantial heterogeneity in relative factor-augmenting productivity across firms within markets. We conduct counterfactual analysis to quantify the role of this heterogeneity in shaping the joint allocation of capital and labor across firms.

Our main counterfactual compresses relative factor-augmenting productivity, A_{Ki}/A_{Li}, by 10 percent toward its market-year mean, while holding each firm's overall level of factor-augmenting productivity fixed. This exercise isolates the role of relative factor-augmenting productivity emphasized by Proposition 1. We also consider separate 10 percent compressions of capital- and labor-augmenting productivity toward their

Table 4: Labor Markdowns and Firm Characteristics

Panel A. Summary statistics			
	Mean	Median	SD
Markdown (unweighted)	1.37	1.32	0.27
Markdown (employment-weighted)	1.34	1.30	0.25

Panel B. Regressions on firm characteristics		
Log employment	–0.009***	–0.008***
	(0.001)	(0.000)
SOE dummy		–0.017***
		(0.002)
FIE dummy		–0.005***
		(0.002)
Export dummy		–0.006***
		(0.001)
Observations	1,446,389	1,446,389

Notes: Panel A reports the mean, median, and standard deviation of calibrated labor markdowns in levels. Each statistic is first computed separately by year and then averaged across 1998–2007. The first row is unweighted, while the second row weights firms by employment within each year. Panel B reports coefficients from firm-level regressions of log labor markdowns on firm characteristics. The first specification includes log employment and industry, province, and year fixed effects. The second specification additionally includes indicators for state-owned enterprises (SOE) and foreign-invested enterprises (FIE), exporter status, and age-category indicators. Standard errors clustered at the market level are reported in parentheses. ***, **, and * denote significance at the 1, 5, and 10 percent levels, respectively.

market-year means. In each counterfactual, we hold the remaining primitives at their benchmark values and re-solve the equilibrium using the mapping $\mathcal{E}$ established in section 4.2. Appendix E.1 provides the detailed construction of these counterfactuals.

Figure 5 presents the counterfactual allocation profiles for a 10 percent compression. Compressing relative factor augmentation flattens the labor share profile, steepens the capital-labor ratio profile, and reverses the negative employment-productivity relationship. Because the calibration does not impose Assumption 1 and permits observed employment to depart from the worker-flow profile through ψ_i, these results show that the mechanism highlighted by Proposition 1 remains quantitatively important under the more flexible employment specification. Compressing capital-augmenting productivity generates qualitatively similar changes in all three profiles, although this counterfactual also changes the overall level of factor-augmenting productivity.

By contrast, compressing labor-augmenting productivity leaves the labor share and capital-labor ratio profiles close to the benchmark while attenuating the decline in employment.

Figure 5: Compressing Factor-Augmenting Heterogeneity

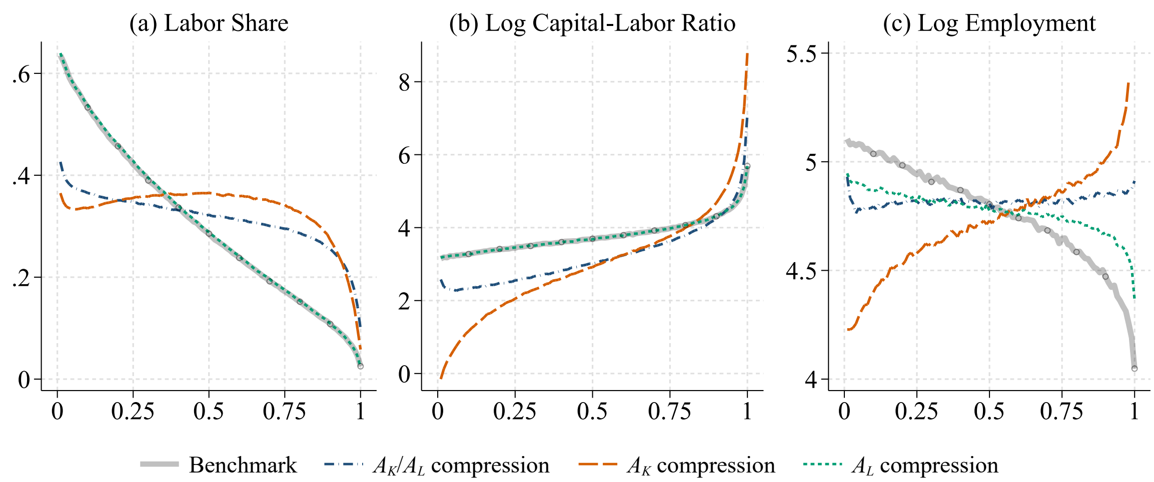

Notes: Each panel presents binned profiles of the indicated outcome against residualized log labor productivity. Labor productivity and the outcomes are residualized with respect to market and year fixed effects. The residualized outcomes are recentered at year-specific means. The benchmark reports the equilibrium outcomes implied by the calibrated model. Each counterfactual compresses heterogeneity in the indicated measure by 10 percent toward its market-year mean and re-solves the equilibrium holding the remaining primitives at their benchmark values. The A_K/A_L counterfactual holds each firm's overall level of factor-augmenting productivity fixed, while the A_K and A_L counterfactuals compress the corresponding productivity component separately. Panel (a) reports the labor share, Panel (b) the log capital-labor ratio, and Panel (c) log employment.

6 Factor-Augmenting Productivity and Labor Share

Heterogeneous factor-augmenting productivity can affect the aggregate labor share through both firm labor shares and the allocation of value added across firms. In our model, firm labor shares depend on output elasticities and wage markdowns, while the aggregate labor share also depends on how value added is distributed across firms. Understanding aggregate labor share movements therefore requires distinguishing between how firm labor shares are determined and how they are aggregated across firms. We first use accounting and structural decompositions to separate these margins. We then use counterfactuals to examine how factor-augmenting productivity shapes the aggregate labor share.

6.1 Anatomy of the Aggregate Labor Share

We first use an accounting decomposition to separate the aggregate labor share into the labor share of the average firm and value-added reallocation across firms. This decomposition shows how firm heterogeneity enters aggregation, but does not identify the economic forces underlying differences in firm labor shares. We therefore complement it with a structural decomposition that uses the model to distinguish the output elasticity component from the markdown adjustment.

Let $\omega_{it} = Y_{it}/\sum_j Y_{jt}$ denote firm i's value-added share, and let N_t denote the number of firms. The aggregate labor share can be written as

$$LS_t = \sum_i \omega_{it} LS_{it} = \underbrace{N_t^{-1} \sum_i LS_{it}}_{\overline{LS}_t} + \underbrace{\sum_i \left(\omega_{it} - N_t^{-1}\right) LS_{it}}_{LS_t^{\text{realloc}}}, \tag{28}$$

where $\overline{LS}_t$ is the unweighted mean firm labor share. The second term, LS_t^{realloc}, is the reallocation component. It captures the relationship between firms' value-added shares and labor shares and is negative when firms accounting for more value-added have lower labor shares.

To identify the economic forces underlying firm labor shares, we use the firm's labor first-order condition. Aggregating equation (10) gives

$$LS_t = \sum_i \omega_{it} \frac{1-\mu_{it}}{\tau_{it}} = \underbrace{\sum_i \omega_{it}(1-\mu_{it})}_{LS_t^{\text{elas}}} + \underbrace{\sum_i \omega_{it}(1-\mu_{it}) \left(\frac{1}{\tau_{it}} - 1\right)}_{LS_t^{\text{adj}}}. \tag{29}$$

The first term, LS_t^{elas}, is the value-added weighted output elasticity with respect to labor. The second term, LS_t^{adj}, is the markdown adjustment and captures the departure of the aggregate labor share from the output elasticity component arising from wage markdowns.

The accounting decomposition shows how firm heterogeneity enters aggregation but does not identify the economic forces underlying differences in firm labor shares. Because both structural components are weighted by value-added, changes in either can reflect changes for the average firm or reallocation across firms. Thus, neither decomposition subsumes the other, and they are different groupings of the same underlying components. Formally, let $e_{it} = 1 - \mu_{it}$ denote the output elasticity with respect to labor and

$a_{it} = e_{it}(\tau_{it}^{-1} - 1)$ denote the markdown adjustment, so that $LS_{it} = e_{it} + a_{it}$. Define their unweighted means as $\bar{e}_t = N_t^{-1} \sum_i e_{it}$ and $\bar{a}_t = N_t^{-1} \sum_i a_{it}$. For $x \in \{e, a\}$, define the corresponding reallocation component as $R_t^x = \sum_i \left(\omega_{it} - N_t^{-1}\right) x_{it}$. It follows that

$$LS_t = \underbrace{(\bar{e}_t + \bar{a}_t)}_{\overline{LS}_t} + \underbrace{(R_t^e + R_t^a)}_{LS_t^{\text{realloc}}} = \underbrace{(\bar{e}_t + R_t^e)}_{LS_t^{\text{elas}}} + \underbrace{(\bar{a}_t + R_t^a)}_{LS_t^{\text{adj}}}.$$

The accounting decomposition therefore groups the four underlying terms along the average-firm and reallocation margins, whereas the structural decomposition regroups them into the output elasticity and markdown adjustment components. The former shows how firm heterogeneity enters aggregation, while the latter provides a model-based interpretation of the same aggregate labor share.

This distinction is particularly relevant in our framework. Under the Cobb-Douglas technology, the output elasticity with respect to labor is fixed across firms, implying $R_t^e = 0$ and that heterogeneity in firm labor shares reflects variation in the markdown adjustment. Under our CES technology,

$$1 - \mu_{it} = \frac{(A_{Lit} L_{it})^\rho}{(A_{Kit} K_{it})^\rho + (A_{Lit} L_{it})^\rho}.$$

The output elasticity with respect to labor therefore varies with relative factor-augmenting productivity and the capital-labor ratio. Heterogeneous factor-augmenting productivity can consequently affect all four underlying terms through changes in output elasticities, equilibrium markdowns, and the allocation of value-added across firms. The accounting decomposition shows whether these effects appear along the average firm or reallocation margins, while the structural decomposition shows whether they appear in the output elasticity or markdown adjustment components.

The two decompositions also play different empirical roles. The accounting decomposition can be constructed directly from the data, whereas the structural decomposition relies on the model recovered output elasticities and markdowns. The model's ability to reproduce the accounting margins therefore provides observable discipline on its structural interpretation of aggregate labor share movements.

Panel A of Table 5 applies the accounting decomposition directly to the data. The aggregate labor share falls by 5.5 percentage points from 1998 to 2007. The mean firm labor share falls by 3.3 percentage points, while the reallocation component becomes 2.2 percentage points more negative. This reallocation margin is closely related to the mechanism emphasized by Kehrig and Vincent (2021), in which value added shifts toward firms with lower labor shares. The decline therefore reflects both a lower labor share for the average firm and a larger share of value-added accruing to firms with lower labor shares.

Table 5: Anatomy of Aggregate Labor Share

	1998	2007	Change
	(1)	(2)	(3)
Panel A. Data			
Aggregate labor share	24.90	19.39	−5.50
Accounting decomposition			
Mean labor share	32.48	29.16	−3.33
Reallocation	−7.59	−9.76	−2.17
Panel B. Calibrated model			
Aggregate labor share	23.01	18.62	−4.38
Accounting decomposition			
Mean labor share	31.93	29.03	−2.90
Reallocation	−8.92	−10.41	−1.49
Structural decomposition			
Output elasticity	29.85	24.49	−5.36
Markdown adjustment	−6.84	−5.86	0.98

Notes: This table reports the anatomy of the aggregate labor share in 1998 and 2007. Panel A decomposes the aggregate labor share in the data into the mean firm labor share and reallocation according to equation (28). Panel B reports the corresponding decomposition in the calibrated model and structurally decomposes the aggregate labor share into output elasticity and markdown adjustment according to equation (29). Labor share levels are reported in percent, and changes from 1998 to 2007 are reported in percentage points.

Panel B shows that the calibrated model captures both the aggregate labor share decline and its two accounting components. These components are not directly targeted in the calibration. The model closely matches the decline in the mean firm labor share and reproduces a sizable contribution from value-added reallocation, providing a useful non-targeted validation of the model. The structural decomposition shows that the output elasticity component declines by 5.4 percentage points, while the markdown adjustment rises by 1 percentage point, partially offsetting the decline. Thus, the model-generated decline is driven primarily by the output elasticity component rather than the markdown adjustment.

These results complement Gouin-Bonenfant (2022), who emphasizes the interaction between productivity heterogeneity and wage-setting frictions in shaping the aggregate labor share. In our framework, factor-augmenting productivity affects firms' output elasticities and endogenous capital intensity.

6.2 Counterfactual Analysis of the Aggregate Labor Share

We conduct two sets of counterfactual experiments. The first compresses cross-sectional heterogeneity in relative factor-augmenting productivity, A_K, and A_L by 10 percent toward their market-year means, using the counterfactuals described in Section 5.4 and detailed in Appendix E.1. The second rolls back 10 percent of the 1998–2007 changes in these productivity profiles toward their 1998 profiles. Appendix E.2 provides the detailed construction of the over-time counterfactuals.

Panel A of Table 6 reports the effects of compressing cross-sectional heterogeneity in factor-augmenting productivity. Compressing heterogeneity in relative factor-augmenting productivity raises the aggregate labor share by 6.4 percentage points on average. In the accounting decomposition, the mean firm labor share rises by 1.2 percentage points, while the reallocation component rises by 5.2 percentage points. The structural decomposition provides a complementary interpretation: the output elasticity component rises by 11.5 percentage points, partly offset by a 5.1 percentage-point decline in the markdown adjustment. The accounting decomposition shows that most of the counterfactual increase in the aggregate labor share comes from value-added reallocation, while the structural decomposition attributes most of the increase to the output elasticity component. Compressing capital- or labor-augmenting productivity separately lowers the aggregate labor share by 4.2 and 2.5 percentage points, respectively.

Panel B reports the effects on the 1998–2007 change in the aggregate labor share. Entries are differences between the counterfactual and benchmark changes, so a negative entry indicates a larger decline under the counterfactual. Rolling back 10 percent of the growth in relative factor-augmenting productivity increases the decline by 11.8 percentage points. In the accounting decomposition, most of this effect operates through value-added reallocation, while the mean firm labor share makes a smaller contribution. The structural decomposition attributes most of the effect to the output elasticity component, with the markdown adjustment partially offsetting the decline. Thus, as in the cross-sectional exercise, value-added reallocation is the main accounting margin,

Table 6: Counterfactual Analysis of Aggregate Labor Share

	A_K/A_L	A_K	A_L
	(1)	(2)	(3)
Panel A. Compressing cross-sectional heterogeneity			
Aggregate labor share	6.39	−4.16	−2.51
Accounting decomposition			
Mean labor share	1.23	2.91	0.29
Reallocation	5.16	−7.07	−2.79
Structural decomposition			
Output elasticity	11.53	−3.33	−3.84
Markdown adjustment	−5.14	−0.83	1.34
Panel B. Rolling back factor-augmenting productivity growth			
Aggregate labor share	−11.76	−7.02	−1.81
Accounting decomposition			
Mean labor share	−2.88	5.98	0.12
Reallocation	−8.88	−13.01	−1.93
Structural decomposition			
Output elasticity	−12.91	−6.00	−2.88
Markdown adjustment	1.15	−1.03	1.07

Notes: Panel A reports the effects of a 10 percent compression of cross-sectional heterogeneity in relative factor-augmenting productivity, capital-augmenting productivity, and labor-augmenting productivity toward their market-year means, averaged over 1998–2007. Panel B reports the effects of rolling back 10 percent of the 1998–2007 changes in relative factor augmentation, capital-augmenting productivity, and labor-augmenting productivity toward their corresponding 1998 profiles. In column (1), the relative factor augmentation counterfactual holds each firm's overall level of factor-augmenting productivity fixed, while columns (2) and (3) vary capital- and labor-augmenting productivity separately. In each counterfactual, the equilibrium is re-solved holding the remaining primitives at their benchmark values. Mean labor share and reallocation correspond to the decomposition in equation (28). Output elasticity and markdown adjustment correspond to the decomposition in equation (29). Entries in Panel A are percentage-point differences relative to the benchmark. Entries in Panel B are percentage-point differences between the counterfactual and benchmark changes from 1998 to 2007.

while the output elasticity component is the main structural margin. Rolling back capital- and labor-augmenting productivity growth separately increases the decline by 7 and 1.8 percentage points, respectively.

Taken together, the results distinguish the effects of cross-sectional heterogeneity from those of changes in the productivity profile over time. Cross-sectionally, heterogeneity in relative factor augmentation lowers the aggregate labor share by strengthening the negative relationship between firms' value-added shares and labor shares. Over time, however, the observed evolution of relative factor augmentation cushions the aggregate labor share decline: moving the 2007 profile toward its 1998 counterpart makes the decline substantially larger.

7 Factor-Augmenting Productivity and Misallocation

This section studies how heterogeneous factor-augmenting productivity changes the assessment of factor misallocation. We first introduce the aggregation framework used to measure potential output gains from factor reallocation under the two production specifications. We then focus on two questions. How does factor augmentation change the gains from reallocating capital and labor? Which source is quantitatively more important for measured misallocation: differences in firm scale or differences in firms' capital-labor mixes?

7.1 Framework

For welfare evaluation, we take the calibrated production functions as given and aggregate firm outputs following Hsieh and Klenow (2009). Aggregate output, denoted by Q, is Cobb-Douglas across industries:

$$Q = \prod_s Q_s^{\omega_s}, \qquad \sum_s \omega_s = 1,$$

where s indexes industry, Q_s is industrial output, and ω_s is the industry's share in aggregate value added. Within each industry, firm varieties are combined using the CES aggregator,

$$Q_s = \left[\sum_{i \in s} Y_i^{\frac{\varepsilon^D - 1}{\varepsilon^D}} \right]^{\frac{\varepsilon^D}{\varepsilon^D - 1}}, \tag{30}$$

where $\varepsilon^D > 1$ is the elasticity of substitution across firm varieties.

We compare our benchmark with the Cobb-Douglas specification in Hsieh and Klenow (2009), $Y_i = A_i K_i^{\alpha} L_i^{1-\alpha}$, where A_i represents firm-specific scalar productivity, and output elasticities are common across firms. By contrast, in our benchmark, the output elasticity with respect to capital, μ_i, varies across firms.

Given the CES aggregator in equation (30), the marginal revenue products relevant for factor allocation are proportional to

$$MRPK_i \propto Y_i^{-1/\varepsilon^D} \frac{\partial Y_i}{\partial K_i}, \qquad MRPL_i \propto Y_i^{-1/\varepsilon^D} \frac{\partial Y_i}{\partial L_i}. \tag{31}$$

An efficient allocation equalizes the marginal revenue product of each factor across firms within a market, subject to fixed market-level factor supplies.[13]

For the quantitative exercises, we set the elasticity of substitution across firm varieties to $\varepsilon^D = 3$, as in Hsieh and Klenow (2009). For the Hsieh-Klenow specification, we set the output elasticity with respect to capital to $\alpha = 0.5$.

7.2 Quantitative Implications for Misallocation

We now quantify how the production specification changes both the magnitude and the composition of measured factor misallocation.

Potential Gains from Factor Reallocation. Panel A of Table 7 shows sharply different potential gains from factor reallocation across the two production specifications. Under the Hsieh-Klenow specification, eliminating capital misallocation raises aggregate output by 46 percent, compared with 283 percent in our benchmark. The pattern is reversed for labor. Eliminating labor misallocation raises aggregate output by 24 percent under Hsieh-Klenow and by 13 percent in our benchmark. Thus, relative to Hsieh-Klenow, our benchmark implies a substantially larger potential output gain from capital reallocation and a smaller gain from labor reallocation.[14]

Two features help account for the larger gain from reallocating capital under our benchmark: wider cross-firm dispersion in $MRPK$ and lower curvature of the $MRPK$ schedule. As shown in Panel B, the standard deviation of log $MRPK$ is 1.65 under our benchmark, compared with 1.05 under the Cobb-Douglas specification. This difference follows directly from the production technology. Using the definition of μ_i, equation (31)

[13] In this section, $MRPK$ and $MRPL$ incorporate the firm-specific demand adjustment implied by the aggregate output aggregator.

[14] The qualitative comparison is robust to alternative elasticities of substitution across firm varieties and to excluding observations with extreme inputs or marginal revenue products.

Table 7: Misallocation and Distribution of Marginal Revenue Products

	Hsieh-Klenow	Benchmark
	(1)	(2)
Panel A. Potential gains from reallocation		
Capital reallocation	46	283
Labor reallocation	24	13
Panel B. Dispersion and curvature of marginal revenue products		
SD($\log MRPK$)	1.05	1.65
SD($\log MRPL$)	0.69	0.44
Median κ^K	0.67	0.30
Median κ^L	0.67	0.27
Panel C. Joint variation in marginal revenue products		
Corr($\log MRPK, \log MRPL$)	0.32	−0.07
Var[$\log(MRPK\, MRPL)$]	2.05	2.90
Var[$\log(MRPK/MRPL)$]	1.11	3.11

Notes: Panel A reports the potential increase in aggregate output, in percent, from reallocating capital and labor separately across firms within each market-year to equalize their marginal revenue products. Panel B reports the dispersion and curvature of marginal revenue products. Panel C reports the correlation between the marginal revenue products of capital and labor and the variances of their log product and log ratio.

implies

$$MRPK_i \propto \mu_i \frac{Y_i}{K_i} Y_i^{-1/\varepsilon^D}.$$

Under the Cobb-Douglas specification, $\mu_i = \alpha$ is identical across firms, whereas μ_i varies across firms in our benchmark. Because the two specifications are evaluated at the same observed Y_i and K_i, the larger dispersion in $MRPK$ arises from cross-firm variation in $\log \mu_i$. This variation in output elasticities with respect to capital is generated by heterogeneous factor-augmenting productivity, which substantially widens $MRPK$ dispersion in our benchmark.

The larger gain from reallocating capital is also driven by a flatter $MRPK$ schedule. We measure its curvature with

$$\kappa_i^K \equiv -\frac{\partial \log MRPK_i}{\partial \log K_i} = (1-\rho)(1-\mu_i) + \frac{\mu_i}{\varepsilon^D},$$

holding the firm's labor and productivity fixed. The first term captures the curvature of the physical marginal product of capital, while the second captures the demand adjustment associated with higher firm output. In the Cobb-Douglas limit, $\rho \to 0$ and $\mu_i \to \alpha = 0.5$. With $\varepsilon^D = 3$, this implies $\kappa_i^K = 0.67$ for every firm. Under our benchmark, $\rho = 0.77$ and μ_i varies across firms, yielding a median curvature of 0.3. For a given initial gap in $MRPK$, lower curvature causes the gap to close more slowly as capital is reallocated. The marginal gain from reallocation therefore remains positive over a larger transfer of capital, contributing to a larger output gain.

The smaller gain from reallocating labor primarily reflects lower cross-firm dispersion in $MRPL$ under our benchmark. The standard deviation of $\log MRPL$ falls from 0.69 under the Cobb-Douglas specification to 0.44 under our benchmark. Analogously to the capital margin,

$$MRPL_i \propto (1-\mu_i)\frac{Y_i}{L_i} Y_i^{-1/\varepsilon^D}.$$

Firm-specific output elasticities with respect to labor offset variation in the remaining revenue-product term. Firms with high labor productivity tend to have lower output elasticities with respect to labor, compressing the dispersion in $MRPL$. The $MRPL$ schedule is also flatter under our benchmark, with median curvature falling from 0.67 to 0.27.[15] As on the capital margin, lower curvature would, for a given initial gap, increase the potential gain from reallocation. In the labor case, however, the smaller initial dispersion in $MRPL$ dominates this curvature effect, resulting in the smaller output gain from reallocating labor.

Scale versus Input-Mix Misallocation. The reallocation exercises quantify the potential output gains but do not identify the underlying sources of misallocation. Measured misallocation may reflect differences in firm scale associated with scalar productivity or differences in firms' capital-labor mixes. The comovement of $MRPK$ and $MRPL$ reveals which dimension is quantitatively more important. If a firm exhibits both a high $MRPK$ and a high $MRPL$, it is employing too little capital and labor compared with the efficient allocation and should increase its overall scale of operation. If it instead has high $MRPK$ but low $MRPL$, it should increase capital relative to labor and adjust its input mix.

Panel C shows that the two production specifications imply sharply different directions of factor reallocation. The correlation between $\log MRPK$ and $\log MRPL$

[15] Analogously to the capital margin, the curvature of the $MRPL$ schedule under CES technology is $\kappa_i^L \equiv -\frac{\partial \log MRPL_i}{\partial \log L_i} = (1-\rho)\mu_i + \frac{1-\mu_i}{\varepsilon^D}$.

is 0.32 under the Hsieh-Klenow specification but -0.07 under our benchmark. Under the Hsieh-Klenow specification, firms with relatively high returns to capital also tend to have relatively high returns to labor, indicating that capital and labor are reallocated toward the same firms. Under our benchmark, firms with relatively high returns to capital instead tend to have relatively low returns to labor, indicating a larger role for reallocations that change firms' capital-labor mixes.

To formalize this distinction, the covariance between the two marginal revenue products can be decomposed as

$$\mathrm{Cov}\left(\log MRPK_i, \log MRPL_i\right) = \frac{1}{4}\left\{\mathrm{Var}\left[\log(MRPK_i\, MRPL_i)\right] - \mathrm{Var}\left[\log\left(\frac{MRPK_i}{MRPL_i}\right)\right]\right\}.$$

The first term captures common variation in the marginal revenue products of capital and labor, which moves the returns to both factors in the same direction and points to differences in firm scale. The second captures variation in relative marginal revenue products, which moves the return to one factor relative to the other and points to differences in the capital-labor mix. Under the Hsieh-Klenow specification, the variance of the common component is 2.05, compared with 1.11 for the relative component. Under our benchmark, the common component rises modestly to 2.90, while the relative component increases to 3.11. The reversal in the correlation therefore reflects the larger role of variation in relative marginal revenue products under our benchmark.

The larger relative-variation term under our benchmark arises from heterogeneity in relative factor augmentation. Scalar productivity affects the marginal revenue products of capital and labor proportionally and therefore drops out of their ratio. Under our benchmark,

$$\log\frac{MRPL_i}{MRPK_i} = \rho\log\frac{A_{Li}}{A_{Ki}} + (1-\rho)\log\frac{K_i}{L_i}.$$

Under the Cobb-Douglas specification, by contrast, firms share the same output elasticities with respect to capital and labor, and relative marginal revenue products vary only with the capital-labor ratio. Moreover, because $1-\rho < 1$, the CES technology attenuates the contribution of variation in the capital-labor ratio relative to Cobb-Douglas. The larger dispersion in relative marginal revenue products under our benchmark therefore reflects the additional variation associated with relative factor-augmenting productivity and its covariance with firms' capital-labor ratios.

Relative factor augmentation changes the firm-specific efficient capital-labor mix against which the observed allocation is evaluated. Accounting for this heterogeneity therefore changes which firms appear short of capital relative to labor and, consequently, the composition of inferred misallocation. These findings are related to David and Venkateswaran (2019), who use the comovement of the average revenue products of capital and labor to assess the role of heterogeneous production elasticities. Our recovered factor-augmenting productivities provide a firm-specific technological source of heterogeneity in these elasticities and connect it to firms' capital-labor mixes. Our analysis is also related to Asturias and Rossbach (2023), who show that allowing for group-specific production technologies can materially alter inferred marginal revenue products and the gains from eliminating factor misallocation. In our framework, heterogeneity in factor elasticities arises from firm-specific capital- and labor-augmenting productivity, which changes both the magnitude and the composition of inferred factor misallocation.

8 Conclusion

Differences in how intensively firms use capital relative to labor reflect both heterogeneity in firms' input mixes and differences in the factor costs they face. This paper shows that the two can be separated, and that the separation matters. In Chinese manufacturing, more labor-productive firms are more capital intensive yet employ fewer workers and have lower labor shares. This joint pattern requires systematic heterogeneity in relative factor-augmenting productivity under empirically motivated conditions on factor markets. We recover firm-specific capital- and labor-augmenting productivity jointly with labor markdowns and the relative factor costs that firms face. The value-added tax reform that lowered the cost of equipment raises capital-augmenting productivity while leaving labor-augmenting productivity largely unchanged, providing external validation of the factor-specific interpretation of the recovered productivities.

The recovered heterogeneity has broad implications for factor allocation, the aggregate labor share, and measured misallocation. In our counterfactual, compressing relative factor-augmenting productivity reverses the negative employment-productivity relationship, showing that this heterogeneity is central to why more labor-productive firms are smaller employers. Over time, growth in relative factor augmentation partly offsets the decline in the aggregate labor share, largely through the equilibrium response of firms' output elasticities. Accounting for heterogeneous factor-augmenting productivity also changes the assessment of misallocation. It raises the inferred importance of capital relative to labor and reverses the sign of the comovement between the marginal revenue products of capital and labor.

The results point to a broader agenda on the origins of heterogeneous factor-augmenting productivity. Related theories of directed and factor-augmenting technical change provide a natural starting point (Acemoglu, 2002, 2003). Introducing heterogeneous firms into these frameworks could endogenize firm-level differences in factor-augmenting productivity and deepen our understanding of their aggregate consequences.

References

Acemoglu, Daron, "Directed Technical Change," *Review of Economic Studies*, 2002, *69* (4), 781–809.

__ , "Labor- and Capital-Augmenting Technical Change," *Journal of the European Economic Association*, 2003, *1* (1), 1–37.

Antràs, Pol, "Is the U.S. Aggregate Production Function Cobb-Douglas? New Estimates of the Elasticity of Substitution," *Contributions in Macroeconomics*, 2004, *4* (1).

Asturias, Jose and Jack Rossbach, "Grouped Variation in Factor Shares: An Application to Misallocation," *International Economic Review*, 2023, *64* (1), 325–360.

Azar, José and Ioana Marinescu, "Monopsony Power in the Labor Market," in Christian Dustmann and Thomas Lemieux, eds., *Handbook of Labor Economics*, Vol. 5, Elsevier, 2024, pp. 761–827.

Bai, Yan, Dan Lu, Xu Tian, and Yajie Wang, "When Do Financial Frictions Matter for Misallocation?," NBER Working Paper 34930, National Bureau of Economic Research February 2026.

Berger, David, Kyle Herkenhoff, and Simon Mongey, "Labor Market Power," *American Economic Review*, 2022, *112* (4), 1147–1193.

Brandt, Loren, Johannes Van Biesebroeck, and Yifan Zhang, "Creative Accounting or Creative Destruction? Firm-Level Productivity Growth in Chinese Manufacturing," *Journal of Development Economics*, 2012, *97* (2), 339–351.

__ , Trevor Tombe, and Xiaodong Zhu, "Factor Market Distortions Across Time, Space, and Sectors in China," *Review of Economic Dynamics*, 2013, *16* (1), 39–58.

Burdett, Kenneth and Dale T. Mortensen, "Wage Differentials, Employer Size, and Unemployment," *International Economic Review*, 1998, *39* (2), 257–273.

Card, David, Ana Rute Cardoso, Jörg Heining, and Patrick Kline, "Firms and Labor Market Inequality: Evidence and Some Theory," *Journal of Labor Economics*, 2018, *36* (S1), S13–S70.

David, Joel M. and Venky Venkateswaran, "The Sources of Capital Misallocation," *American Economic Review*, 2019, *109* (7), 2531–2567.

De Loecker, Jan, "Product Differentiation, Multiproduct Firms, and Estimating the Impact of Trade Liberalization on Productivity," *Econometrica*, 2011, *79* (5), 1407–1451.

Demirer, Mert, "Production Function Estimation with Factor-Augmenting Technology: An Application to Markups," *Econometrica*, 2026.

Doraszelski, Ulrich and Jordi Jaumandreu, "Measuring the Bias of Technological Change," *Journal of Political Economy*, 2018, *126* (3), 1027–1084.

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

Greenwood, Jeremy, Zvi Hercowitz, and Per Krusell, "Long-Run Implications of Investment-Specific Technological Change," *American Economic Review*, 1997, *87* (3), 342–362.

Hsieh, Chang-Tai and Peter J. Klenow, "Misallocation and Manufacturing TFP in China and India," *Quarterly Journal of Economics*, 2009, *124* (4), 1403–1448.

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.

Karabarbounis, Loukas and Brent Neiman, "The Global Decline of the Labor Share," *Quarterly Journal of Economics*, 2014, *129* (1), 61–103.

Kaymak, Barış and Immo Schott, "Corporate Tax Cuts and the Decline of the Manufacturing Labor Share," *Econometrica*, 2023, *91* (6), 2371–2408.

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

Klette, Tor Jakob and Zvi Griliches, "The Inconsistency of Common Scale Estimators When Output Prices Are Unobserved and Endogenous," *Journal of Applied Econometrics*, 1996, *11* (4), 343–361.

Lamadon, Thibaut, Magne Mogstad, and Bradley Setzler, "Imperfect Competition, Compensating Differentials, and Rent Sharing in the U.S. Labor Market," *American Economic Review*, 2022, *112* (1), 169–212.

León-Ledesma, Miguel A, Peter McAdam, and Alpo Willman, "Identifying the Elasticity of Substitution with Biased Technical Change," *American Economic Review*, 2010, *100* (4), 1330–1357.

Liu, Qing and Yi Lu, "Firm Investment and Exporting: Evidence from China's Value-Added Tax Reform," *Journal of International Economics*, 2015, *97* (2), 392–403.

Lu, Yi, Yoichi Sugita, and Lianming Zhu, "Wage and Markdowns and FDI Liberalization," *Discussion paper series HIAS-E-83, Hitotsubashi University*, 2019.

Meng, Lingsheng, Yifan Zhang, and Yunbin Zhang, "Estimating Capital-Labor Substitution in China: Evidence from Firm-Level Data," *Working Paper*, 2025.

Midrigan, Virgiliu and Daniel Yi Xu, "Finance and Misallocation: Evidence from Plant-Level Data," *American Economic Review*, 2014, *104* (2), 422–458.

Moll, Benjamin, "Productivity Losses from Financial Frictions: Can Self-Financing Undo Capital Misallocation?," *American Economic Review*, 2014, *104* (10), 3186–3221.

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

Pham, Hoang, "Trade Reform, Oligopsony, and Labor Market Distortion: Theory and Evidence," *Journal of International Economics*, 2023, *144*, 103787.

Qian, Zhenjie and Xiaodong Zhu, "Is China's Labor Share Really Low? Evidence from an International Comparison of Manufacturing," *Journal of World Economy*, 2013, (10), 27–53.

Restuccia, Diego and Richard Rogerson, "Policy Distortions and Aggregate Productivity with Heterogeneous Plants," *Review of Economic Dynamics*, 2008, *11* (4), 707–720.

Rubens, Michael, Yingjie Wu, and Mingzhi Xu, "Exploiting or Augmenting Labor?," *American Economic Review: Insights*, 2026, *8* (1), 72–89.

Song, Zheng, Kjetil Storesletten, and Fabrizio Zilibotti, "Growing Like China," *American Economic Review*, 2011, *101* (1), 196–233.

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

Zhang, Hongsong, "Non-Neutral Technology, Firm Heterogeneity, and Labor Demand," *Journal of Development Economics*, 2019, *140*, 145–168.

Online Appendix

Appendix A Data Appendix

A.1 Construction of Capital Stock

We construct firm-level capital stock using the perpetual inventory method following Brandt et al. (2012). Let BK_{it} denote the gross book value of fixed assets for firm i in year t. We use the change in the book value of fixed assets, $BK_{it} - BK_{i,t-1}$, as a measure of nominal investment, and deflate it by the national capital-goods price index P_t^K. We construct real capital recursively as

$$K_{it} = (1 - \delta)K_{i,t-1} + \frac{BK_{it} - BK_{i,t-1}}{P_t^K},$$

where $\delta = 0.05$ is the annual depreciation rate.

To initialize the capital stock, we use firms' reported gross book values of fixed assets. For firms founded after the beginning of the sample period, we use the book value reported in the founding year as the initial book value. For firms founded before the sample period, we impute the initial book value by projecting the first observed book value backward to the founding year. Specifically, if firm i was founded in year t_0 but first appears in the data in year t_1, we impute the founding-year gross book value as $BK_{i,t_0} = BK_{i,t_1}/(1 + g_i)^{t_1 - t_0}$, where g_i is the firm's geometric average growth rate of book fixed assets over the years in which the firm is observed after t_1. Starting from this imputed value, we apply the perpetual inventory equation to obtain firm-level real capital in each subsequent year.

A.2 Construction of Income Approach Value-Added

We construct value added using the income approach, following the accounting framework in Qian and Zhu (2013). For firm i in year t, income approach value added is defined as

$$VA_{it}^I = DEP_{it} + LC_{it}^I + OS_{it}^I + NT_{it}^I,$$

where DEP_{it} is current depreciation, LC^I_{it} is labor compensation, OS^I_{it} is operating surplus, and NT^I_{it} is net production taxes.

We construct labor compensation from wage bills, welfare expenses, labor and unemployment insurance, pension and medical insurance, housing related payments, and the labor compensation component of selling, administrative, and financial expenses. Pension and medical insurance expenses and housing related payments are reported only from 2004 onward. For earlier years, we impute these components using industry-province level ratios computed from years in which the components are observed.

We allocate the relevant portion of selling, administrative, and financial expenses to labor compensation, operating surplus, and net production taxes using industry level allocation ratios from the National Bureau of Statistics. Let X_{it} denote the sum of selling, administrative, and financial expenses, and let $\omega^{LC}_{j(i)t}$, $\omega^{OS}_{j(i)t}$, and $\omega^{NT}_{j(i)t}$ denote the corresponding allocation ratios for industry $j(i)$. The labor compensation adjustment is $\omega^{LC}_{j(i)t}X_{it}$.

Operating surplus is constructed from profit related accounting items. Specifically, we start from total profits, subtract administrative and financial expenses, add subsidy income, and then add the operating surplus component of selling, administrative, and financial expenses, $\omega^{OS}_{j(i)t}X_{it}$. Net production taxes are constructed as business taxes and surcharges plus taxes reported under administrative expenses plus value-added tax payable, net of subsidy income. We then add the production tax component of selling, administrative, and financial expenses, $\omega^{NT}_{j(i)t}X_{it}$.

The labor share based on income approach value added is

$$LS^I_{it} = \frac{LC^I_{it}}{VA^I_{it}}.$$

Appendix B Additional Motivating Evidence

B.1 Robustness of Factor Allocation Patterns

We assess the robustness of the allocation profiles in Section 2 along several dimensions. Figure A1 reconstructs value-added and labor compensation using the income approach. Figures A2–A5 address potential mechanical overlap between current labor productivity and the outcomes. Figure A2 uses the odd-even measure described in Section 2.2. Figures A3–A5 provide further evidence using productivity measures that exclude the current-year observation. Figure A3 uses one-year lagged labor productivity, Figure A4 uses average labor productivity over the previous two years, and Figure A5 uses leave-one-out labor productivity, defined for each firm-year as the firm's average residualized log labor productivity over all other available years. Figures A6–A7 examine three-year and five-year differences in labor productivity and the corresponding differences in the outcomes. Figure A8 additionally controls for ownership type using indicators for state-owned, foreign-invested, and private firms.

Table A2 summarizes these robustness exercises using percentile gaps. Panels A–G correspond to the alternative value added measure, lagged and leave-one-out productivity measures, three- and five-year differences, and ownership controls described above. Panels H and I report 75–25 and 90–10 percentile gaps. Across these specifications, the three factor-allocation gradients remain robust.

B.2 Anatomy of the Labor Productivity Distribution

Figure A9 provides additional anatomy of the labor productivity distribution by plotting log value-added, log capital, log average wages, and log labor compensation against firms' labor productivity percentiles. Firms with higher labor productivity produce more value-added, use more capital, and pay higher average wages. Labor compensation, however, changes much less and declines near the top of the distribution. The increase in value added is therefore accompanied by substantial capital deepening and limited labor absorption.

Figure A10 repeats the exercise using the odd-even labor productivity measure. The profiles closely resemble those in Figure A9: value-added, capital, and average wages rise across the distribution, while labor compensation changes much less and declines near the top.

Appendix C Labor Market Derivations

This appendix serves two purposes. It shows why marginal labor expenditure increases with employment in two standard models of labor market power, and it derives the Burdett-Mortensen equilibrium used in the quantification. We first consider differentiated employers and finite-firm competition, and then turn to the wage-posting model.

Suppose firm i faces inverse labor supply $w_i = W_i(L_i)$. Its marginal labor expenditure is

$$m_i = \frac{d[W_i(L_i)L_i]}{dL_i} = W_i(L_i) + L_i W_i'(L_i). \tag{C1}$$

C.1 BHM-Style Finite-Firm Competition

Consider the following representation of the labor-supply system in Berger et al. (2022):

$$L_i = C w_i^{\eta} \left(\sum_{h \in \mathcal{J}} w_h^{1+\eta} \right)^{\frac{\theta-\eta}{1+\eta}}, \tag{C2}$$

where $\eta > \theta > 0$. The parameter η governs substitution across firms within a labor market, while θ governs labor supply to the market.

Define

$$s_i = \frac{w_i^{1+\eta}}{\sum_{h \in \mathcal{J}} w_h^{1+\eta}} \in (0, 1).$$

The labor-supply elasticity faced by firm i is

$$\varepsilon_i \equiv \frac{d \log L_i}{d \log w_i} = \eta(1 - s_i) + \theta s_i. \tag{C3}$$

Because $\eta > \theta > 0$, labor supply is increasing in the wage. Moreover,

$$\frac{d\varepsilon_i}{d \log w_i} = (\theta - \eta)(1 + \eta)s_i(1 - s_i) \leq 0. \tag{C4}$$

Using equation (C3), marginal labor expenditure can be written as

$$m_i = w_i \left(1 + \frac{1}{\varepsilon_i} \right). \tag{C5}$$

Differentiating gives

$$\frac{d \log m_i}{d \log w_i} = 1 - \frac{1}{\varepsilon_i(1+\varepsilon_i)} \frac{d\varepsilon_i}{d \log w_i} > 0. \tag{C6}$$

Since $d \log L_i / d \log w_i = \varepsilon_i > 0$, equations (C3) and (C6) imply

$$\frac{dm_i}{dL_i} > 0.$$

Thus, the BHM model implies that marginal labor expenditure increases with employment.

C.2 Burdett-Mortensen Wage Posting

We first derive rank preservation and the relationship between marginal labor expenditure and employment in the baseline Burdett-Mortensen model. We then introduce the employment shifters used in the quantification and show why they leave the wage-posting first-order condition unchanged. Throughout, $L(q)$ denotes baseline employment at rank q.

All workers are employed and search on the job. Workers receive outside offers at rate $\lambda > 0$. Employment relationships end at rate $\delta^L > 0$, and affected workers are immediately reassigned to a new employer. Both rates are constant within a market-year. Outside offers and reassignment destinations are drawn uniformly across firms.

C.2.1 Baseline Employment and Rank Preservation

Normalize the mass of firms to one. Initially, let $q \in [0, 1]$ denote a firm's wage rank. We show below that this rank also coincides with its rank in the $MRPL$ distribution. Workers accept outside offers from higher-wage firms.

The total mass of workers is $\int_0^1 L(s)\, ds$. At rank q, firms hire workers reassigned after separations and workers moving from lower-ranked firms. They lose workers through separations and offers from higher-ranked firms. Steady-state flow balance per unit of firm rank therefore requires

$$\delta^L \int_0^1 L(s)\, ds + \lambda \int_0^q L(s)\, ds = \left[\delta^L + (1-q)\lambda\right] L(q). \tag{C7}$$

Differentiating with respect to q gives

$$\left[\delta^L + (1-q)\lambda\right] L'(q) = 2\lambda L(q).$$

Solving for the employment profile yields

$$L(q) = \frac{B}{\left[\delta^L + (1-q)\lambda\right]^2}, \qquad B > 0, \tag{C8}$$

where B determines the scale of baseline employment. Hence,

$$g(q) \equiv \frac{d \log L(q)}{dq} = \frac{2\lambda}{\delta^L + (1-q)\lambda} > 0. \tag{C9}$$

Baseline employment is therefore strictly increasing in wage rank.

Consider a firm with production-side $MRPL$ m_i. If it chooses wage rank q, its payoff is

$$\Pi_i(q) = [m_i - w(q)]\, L(q). \tag{C10}$$

This objective satisfies

$$\frac{\partial^2 \Pi_i(q)}{\partial m_i\, \partial q} = L'(q) > 0.$$

The gain from choosing a higher wage rank is therefore greater for firms with higher $MRPL$. Firms with higher $MRPL$ choose weakly higher wage ranks and have weakly higher baseline employment. The wage-rank representation therefore also orders firms by $MRPL$.

At an interior optimum, differentiating (C10) with respect to q gives

$$-w'(q)L(q) + [m_i - w(q)]\, L'(q) = 0.$$

Writing $m(q)$ for equilibrium $MRPL$ at rank q, the wage profile therefore satisfies

$$w'(q) = g(q)\,[m(q) - w(q)]\,. \tag{C11}$$

Equivalently,

$$m(q) = w(q) + \frac{w'(q)}{g(q)}. \tag{C12}$$

The right-hand side is marginal labor expenditure. Indeed,

$$w(q) + L(q)\frac{dw(q)}{dL(q)} = w(q) + \frac{w'(q)}{g(q)} = m(q).$$

Rank preservation and the increase in $L(q)$ imply

$$L(q_i) > L(q_j) \quad \Rightarrow \quad m(q_i) \geq m(q_j).$$

Thus, marginal labor expenditure is weakly increasing with employment in the baseline model, supporting the labor-side condition in Assumption 1.

C.2.2 Employment Shifters in the Quantitative Model

For quantification, we extend the baseline employment schedule with a positive firm-specific shifter ψ_i:

$$L_i = \psi_i L(q_i). \tag{C13}$$

We take this shifted employment as the firm's labor supply when posting its wage, with ψ_i held fixed. It may be statistically associated with the firm's equilibrium rank.

The firm's wage-posting payoff becomes

$$\widetilde{\Pi}_i(q) = \psi_i \left[m_i - w(q)\right] L(q) = \psi_i \Pi_i(q). \tag{C14}$$

Because $\psi_i > 0$ multiplies the entire payoff and is fixed under a wage deviation, it does not change the firm's wage posting decision. Rank preservation therefore continues to hold. To see explicitly, differentiate (C14) while holding ψ_i fixed:

$$\frac{d\widetilde{\Pi}_i(q)}{dq} = \psi_i \left\{-w'(q)L(q) + \left[m_i - w(q)\right] L'(q)\right\} = 0.$$

Dividing by $\psi_i L(q) > 0$ gives

$$w'(q) = g(q)\left[m(q) - w(q)\right],$$

which is exactly the wage-posting condition (14) used in the main text.

The same cancellation applies to marginal labor expenditure. Holding ψ_i fixed,

$dL_i/dq = \psi_i L'(q) = g(q)L_i$, and hence

$$w(q) + L_i \frac{dw(q)}{dL_i} = w(q) + \frac{w'(q)}{g(q)} = m(q).$$

The labor markdown remains

$$\tau(q) = \frac{m(q)}{w(q)}. \tag{C15}$$

We impose the boundary condition

$$w(0) = m(0), \qquad \tau(0) = 1. \tag{C16}$$

Conditional on the distribution of m_i, the employment shifters therefore leave the wage and markdown profiles unchanged.

Appendix D Calibration Details

This appendix describes the numerical algorithm used to implement the calibration in problem (26). The calibration determines employment shifters and the labor market parameters in two steps. According to the rank preservation property, we first fix q_{idt} at the firm's wage rank in market d and year t. In the first step, given an initial guess on $\{\lambda_{dt}\}$, variation across firms and years identifies $\beta_{\psi,1}$ and the fixed-effect components of the employment equation (23). In the second step, using the resulting firm and rank component of $\log \psi_{idt}$, the labor market equilibrium in each market and year determines B_{dt}, $m_{dt}(q)$, and an update for $\{\lambda_{dt}\}$ by minimizing its contribution to problem (26). This two step approach reflects the economic structure of the problem: employment shifters are disciplined by firm-level variation over time, whereas worker flows and equilibrium schedules are disciplined by cross-sectional variation within each market and year. Because these two sets of objects depend on one another, we alternate between them until the changes in $\beta_{\psi,1}$, and λ_{dt} satisfy a fixed point condition.

We normalize the two CES weights to one half. We set $\sigma = 4.36$, following Meng et al. (2025), and use $\rho = 1 - 1/\sigma$. The annual separation rate is $\delta^L = 0.08$. The job-arrival rate λ_{dt} is constant within each market-year but is estimated separately across market-years. In our benchmark calibration, we consider the case that $\bar{p} = 1$. The calibration is implemented on annual data from 1998–2007. For calibration, we retain market-years containing at least 35 firms.

D.1 Wage Rank and Employment Shifters

For a market-year (d, t) containing N_{dt} firms, the calibration fixes each firm's structural rank to its observed wage rank within the market-year:

$$q_{idt} = \frac{\text{rank}_{dt}(w_{idt}) - \frac{1}{2}}{N_{dt}}. \tag{D1}$$

Conditional on a candidate λ_{dt}, equation (23) can be written as

$$\log L_{idt} + 2\log\left[\delta^L + (1 - q_{idt})\lambda_{dt}\right] = \log B_{dt} + \xi_{dt} + \phi_i + \beta_{\psi,1} q_{idt} + \upsilon_{idt}. \tag{D2}$$

We estimate this equation by least squares. The regression uses observation weights $s^{-2}_{\log L,dt}$, consistent with the standardization of the employment residual in the calibration objective. The terms $\log B_{dt}$ and ξ_{dt} are absorbed by market-year fixed effect. We normalize the predicted firm and rank components to have zero mean within each market and year:

$$\widehat{\log \psi}^{P}_{idt} = \widehat{\phi}_i + \widehat{\beta}_{\psi,1} q_{idt} - \frac{1}{N_{dt}} \sum_{h \in (d,t)} \left(\widehat{\phi}_h + \widehat{\beta}_{\psi,1} q_{hdt} \right). \tag{D3}$$

The employment scale B_{dt} is determined separately in the next step.

D.2 Labor Market Equilibrium

Given $\widehat{\log \psi}^{P}_{idt}$, the calibration minimizes the problem (26) over λ_{dt} and 12 nodes of $\log MRPL$, which approximates the $MRPL$ distribution. B_{dt} and model-implied labor productivity are determined conditionally on λ_{dt} and $\log MRPL$ schedule. Let u_ℓ, $\ell = 1, \ldots, 12$, denote 12 equally spaced nodes on firm ranks in $[0, 1]$. We parameterize the $MRPL$ values on these nodes as

$$\log m_{dt}(u_\ell) = \theta_1 + \sum_{h=2}^{\ell} \log\left(1 + \exp(\theta_h)\right). \tag{D4}$$

This specification guarantees that $m(q)$ is weakly increasing in q. We linearly interpolate $\log m_{dt}(q)$ between nodes and restrict the search for λ_{dt} within $10^{-4} \leq \lambda_{dt} \leq 20$.

Given λ_{dt} and $m_{dt}(q)$, equations (13) and (14), along with $w_{dt}(0) = m_{dt}(0)$, determine the wage profile $w_{dt}(q)$. In the numerical implementation, we solve the ODE (14) over grids

$$q_0 = 0, \qquad q_j = \frac{j - 1/2}{N_{dt}}, \quad j = 1, \ldots, N_{dt}. \tag{D5}$$

Let $\bar{m}_j$ and $\bar{g}_j$ denote the averages of $m_{dt}(q)$ and $g_{dt}(q)$, respectively, over the endpoints of $[q_{j-1}, q_j]$. We advance the wage equation according to

$$w_{dt}(q_j) = \min\left\{ m_{dt}(q_j), \bar{m}_j + \left[w_{dt}(q_{j-1}) - \bar{m}_j \right] \exp\left[-\bar{g}_j (q_j - q_{j-1}) \right] \right\}.$$

Define

$$D_{dt}(q) = \delta^L + (1 - q)\lambda_{dt}.$$

The Burdett-Mortensen employment profile is

$$L_{dt}^{BM}(q) = \frac{B_{dt}}{D_{dt}(q)^2}. \tag{D6}$$

For a candidate value of λ_{dt}, normalization (25) determines the employment scale:

$$\log B_{dt} = \frac{1}{N_{dt}} \sum_{i \in (d,t)} \left[\log L_{idt} - \widehat{\log \psi}_{idt}^{P} + 2 \log D_{dt}(q_{idt}) \right]. \tag{D7}$$

Let $s_{\log w,dt}$, $s_{\log L,dt}$, $s_{LS,dt}$, and $s_{\log y,dt}$ denote the within market-year standard deviations of the four observables. The weighting matrix in problem (26) is

$$\Omega_{dt} = \operatorname{diag}\left(s_{\log w,dt}^{-2}, s_{\log L,dt}^{-2}, s_{LS,dt}^{-2}, s_{\log y,dt}^{-2} \right). \tag{D8}$$

If the within market-year standard deviation is degenerate, the corresponding scale is set to one. No separate wage-employment rank moment enters the objective because q_{idt} is fixed at the observed wage rank.

Let z_{idt} denote log model-implied labor productivity. Conditional on the wage and $MRPL$ schedules, the calibration fits this object separately for each observation:

$$\widehat{z}_{idt} = \arg\min_{z} \left\{ \frac{(\log y_{idt} - z)^2}{s_{\log y,dt}^2} + \frac{\left[LS_{idt} - \exp\big(\log w_{dt}(q_{idt}) - z\big) \right]^2}{s_{LS,dt}^2} \right\} \tag{D9}$$

subject to

$$z \geq \log m_{dt}(q_{idt}) - \log(1 - \varepsilon_\mu), \qquad \varepsilon_\mu = 10^{-6}. \tag{D10}$$

The model-implied labor share is

$$LS_{idt}^{M} = \exp\big(\log w_{dt}(q_{idt}) - \widehat{z}_{idt}\big).$$

The associated output elasticity with respect to capital is

$$\widehat{\mu}_{idt} = 1 - \exp\left(\log m_{dt}(q_{idt}) - \widehat{z}_{idt}\right). \tag{D11}$$

Thus the lower bound imposes $\widehat{\mu}_{idt} \geq \varepsilon_\mu$, ensuring a strictly positive output elasticity with respect to capital.

To numerically stabilize solutions near the lower economic boundary, define $b_\mu(\mu) =$

$[1+\exp((\mu-\varepsilon_b)/h_b)]^{-1}$, where $\varepsilon_b > 0$ and $h_b > 0$ are small fixed constants determining the penalty threshold and smoothing width, respectively. The four residuals in the market-year problem are those collected in r_{idt} in equation (24), evaluated after substituting B_{dt} and $\widehat{z}_{idt}$ as above. The employment residual is $\log L_{idt} - \log L^{BM}_{dt}(q_{idt}) - \widehat{\log\psi}^{P}_{idt}$. The market-year objective is therefore

$$J_{dt} = \frac{1}{N_{dt}} \sum_{i\in(d,t)} r'_{idt}\Omega_{dt} r_{idt} + 100 \left[\frac{1}{N_{dt}} \sum_{i\in(d,t)} b_\mu(\widehat{\mu}_{idt}) \right]^2 . \tag{D12}$$

The second term is denoted as $\mathcal{I}_{dt}$ in (26). The boundary penalty enters only the objective J_{dt}, evaluated at the labor productivity values determined by the solution to (D9). It permits individual observations near the boundary while discouraging solutions in which they are pervasive.

D.3 Global Iteration

Let (j) index global iterations. We initialize $\lambda^{(0)}_{dt} = 0.16$ in every market-year and use equation (D2) to construct $\widehat{\log\psi}^{P,(0)}_{idt}$.

At iteration $j \geq 1$, each market-year takes $\widehat{\log\psi}^{P,(j-1)}_{idt}$ as given and solves (D12). After solving all market-years for 1998–2007, we use the resulting $\widehat{\lambda}^{(j)}_{dt}$ to re-estimate (D2) and construct $\widehat{\log\psi}^{P,(j)}_{idt}$ from (D3). Unless the fixed point has converged, this component enters iteration $j+1$.

We stop the iteration when the root-mean-squared change in $\widehat{\log\psi}^{P}_{idt}$, the absolute change in $\widehat{\beta}_{\psi,1}$, and the 95th percentile and maximum of $|\Delta \log \widehat{\lambda}_{dt}|$ fall below their respective convergence tolerances.

D.4 Labor Markdowns and Production Primitives

Observed k_{idt} enters the production-side inversion after the labor markdown and model labor productivity have been recovered.

The calibrated labor markdown and total employment shifter are

$$\widehat{\tau}_{idt} = \frac{\widehat{m}_{dt}(q_{idt})}{\widehat{w}_{dt}(q_{idt})}, \qquad \log\widehat{\psi}_{idt} = \log L_{idt} - \log \widehat{L}^{BM}_{dt}(q_{idt}). \tag{D13}$$

The employment shifter includes the panel component and the remaining employment residual, so $\widehat{L}_{dt}^{BM}(q_{idt})\widehat{\psi}_{idt}$ reproduces observed employment by construction. The employment residual in (D12) instead uses $\widehat{\log\psi}_{idt}^{P}$.

Equation (D11) implies the output elasticity with respect to labor:

$$1 - \widehat{\mu}_{idt} = \frac{\widehat{m}_{dt}(q_{idt})}{\exp(\widehat{z}_{idt})}. \tag{D14}$$

Applying Proposition 2 to the calibrated labor share and labor productivity, together with observed k_{idt}, yields

$$\frac{\widehat{A}_{L,idt}}{\widehat{A}_{K,idt}} = k_{idt}\left(\frac{1-\widehat{\mu}_{idt}}{\widehat{\mu}_{idt}}\right)^{1/\rho}, \tag{D15}$$

and

$$\widehat{A}_{K,idt} = \frac{\exp(\widehat{z}_{idt})}{\left[\frac{1}{2}k_{idt}^{\rho} + \frac{1}{2}\left(\widehat{A}_{L,idt}/\widehat{A}_{K,idt}\right)^{\rho}\right]^{1/\rho}}, \qquad \widehat{A}_{L,idt} = \left(\frac{\widehat{A}_{L,idt}}{\widehat{A}_{K,idt}}\right)\widehat{A}_{K,idt}. \tag{D16}$$

Finally, marginal capital expenditure is

$$\widehat{n}_{idt} = \frac{\widehat{\mu}_{idt}\exp(\widehat{z}_{idt})}{k_{idt}}. \tag{D17}$$

These firm-level objects form the benchmark panel used in the quantitative exercises.

Appendix E Construction of Counterfactuals

E.1 Cross-Sectional Counterfactuals

Our main counterfactual compresses cross-sectional heterogeneity in relative factor-augmenting productivity. Let $b_i \equiv \log(A_{Ki}/A_{Li})$. We move b_i toward its market-year mean, $\bar{b}_{dt}$, while holding the overall level of factor-augmenting productivity fixed:

$$\log A_{Ki}(\chi) = \log A_{Ki} - \frac{\chi}{2}\big(b_i - \bar{b}_{dt}\big), \quad \log A_{Li}(\chi) = \log A_{Li} + \frac{\chi}{2}\big(b_i - \bar{b}_{dt}\big). \tag{E1}$$

Here, $\bar{b}_{dt}$ is the mean of b_i among firms in market d and year t. The two components move by equal and opposite amounts, leaving $\log A_{Ki} + \log A_{Li}$ unchanged for every firm while compressing relative factor-augmenting productivity toward its market-year mean.[16] This exercise isolates the role of relative factor augmentation emphasized by Proposition 1.

We also examine the two factor-augmenting productivity components separately. For productivity component $A \in \{A_K, A_L\}$, we compress $\log A_i$ toward its market-year mean according to

$$\log A_i(\chi) = (1-\chi)\log A_i + \chi\,\overline{\log A}_{dt}, \tag{E2}$$

where $\overline{\log A}_{dt}$ is the mean of $\log A$ among firms in market d and year t. Throughout these exercises, we set $\chi = 0.1$, corresponding to a 10 percent compression toward the market-year mean. Unlike the relative factor-augmenting productivity exercise, compressing either A_K or A_L changes both the firm's relative factor augmentation and the overall scale of factor-augmenting productivity. In all counterfactuals, the marginal capital expenditure n_i, the offer-arrival rate λ_d, and the employment shifter ψ_i are held at their benchmark values, and we re-solve the equilibrium using the mapping $\mathcal{E}$ defined in section 4.2.

[16] This construction holds fixed the Hicks-neutral component of factor-augmenting productivity. In the equal-share Cobb-Douglas limit, $Y_i = (A_{Ki}K_i)^{1/2}(A_{Li}L_i)^{1/2}$, productivity depends on A_{Ki} and A_{Li} only through their geometric mean. Since $\log A_{Ki} + \log A_{Li}$ remains unchanged, the production function is also unaffected in this case. The counterfactual therefore isolates the consequences of relative factor-augmenting productivity, which has no independent role under the alternative Cobb-Douglas specification.

E.2 Over-Time Counterfactuals

We construct the over-time counterfactuals by moving the 2007 factor-augmenting productivity profiles toward their corresponding 1998 profiles. For a firm i observed in 2007, let $q(i)$ denote its 2007 labor productivity rank within market d, and define

$$b_{d,q(i)}^{1998} \equiv \log\left(\frac{A_{K,d,q(i)}^{1998}}{A_{L,d,q(i)}^{1998}}\right),$$

where the 1998 productivity components are evaluated at the same rank in the same market. The rollback counterfactual for relative factor-augmenting productivity is

$$\log \widetilde{A}_{Ki}(\chi) = \log A_{Ki} - \frac{\chi}{2}\left(b_i - b_{d,q(i)}^{1998}\right), \quad \log \widetilde{A}_{Li}(\chi) = \log A_{Li} + \frac{\chi}{2}\left(b_i - b_{d,q(i)}^{1998}\right). \tag{E3}$$

This construction moves relative factor-augmenting productivity toward its 1998 profile while leaving $\log A_{Ki} + \log A_{Li}$ unchanged for every firm.

We also roll back the capital- and labor-augmenting productivities separately. For $A \in \{A_K, A_L\}$, let $A_{d,q(i)}^{1998}$ denote the value of the 1998 productivity profile at the same rank in market d. The counterfactual assigns

$$\log \widetilde{A}_i(\chi) = (1-\chi)\log A_i + \chi \log A_{d,q(i)}^{1998}. \tag{E4}$$

We set $\chi = 0.1$, corresponding to a 10 percent rollback toward the 1998 profile. In each counterfactual, we re-solve the 2007 equilibrium holding all other primitives at their benchmark values.

Appendix F Appendix Figures

Figure A1: Factor Allocation across Labor Productivity Distribution: Income-Approach Value Added

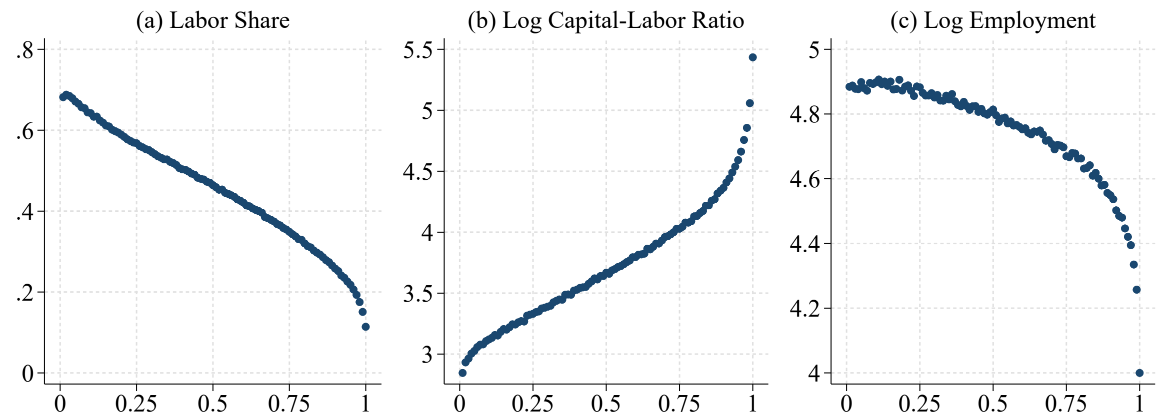

Notes: This figure uses value added and labor compensation constructed using the income approach. Each panel presents a binned scatterplot in which each point represents the mean outcome within one of 100 equal-frequency bins of residualized log labor productivity. Labor productivity and the outcomes are residualized with respect to market and year fixed effects. The residualized outcomes are recentered at year-specific means. Panel (a) reports the labor share, Panel (b) the log capital-labor ratio, and Panel (c) log employment.

Figure A2: Factor Allocation across Odd-Even Labor Productivity Distribution

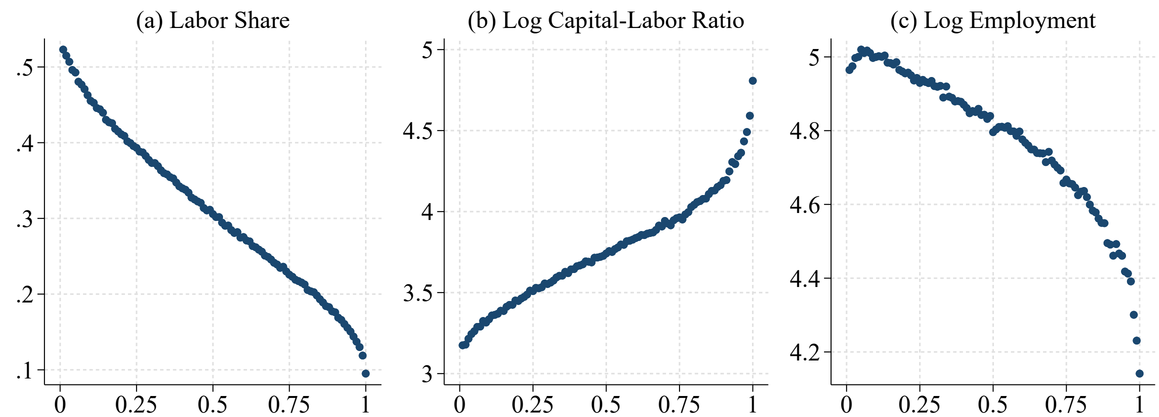

Notes: Each panel presents a binned scatterplot in which each point represents the mean outcome within one of 100 equal-frequency bins of the odd-even labor productivity measure. Even-year observations are sorted by firms' average residualized log labor productivity in odd years, and vice versa for odd-year observations. Labor productivity and the outcomes are residualized with respect to market and year fixed effects. The residualized outcomes are recentered at year-specific means. Panel (a) reports the labor share, Panel (b) the log capital-labor ratio, and Panel (c) log employment.

Figure A3: Factor Allocation across One-Year Lagged Labor Productivity Distribution

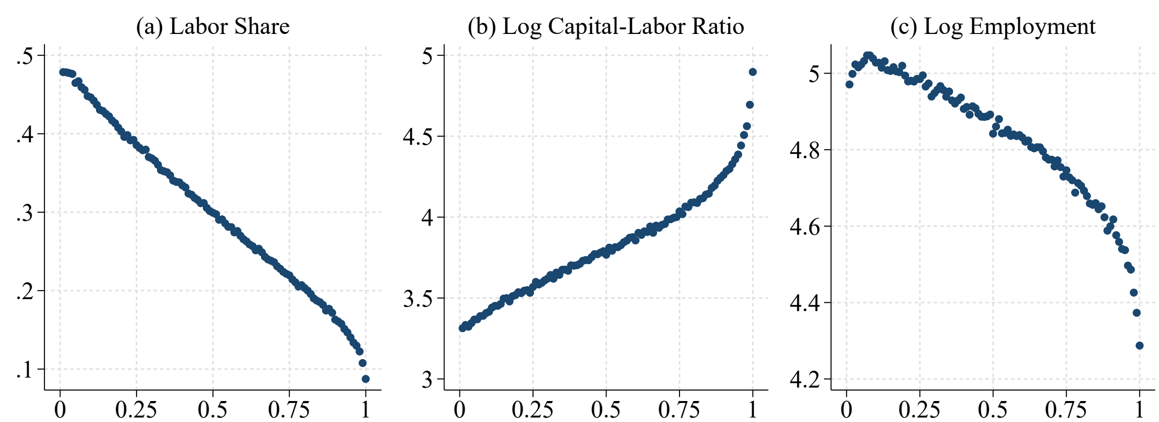

Notes: Each panel presents a binned scatterplot in which each point represents the mean outcome within one of 100 equal-frequency bins of one-year lagged residualized log labor productivity. Labor productivity and the outcomes are residualized with respect to market and year fixed effects. The residualized outcomes are recentered at year-specific means. Panel (a) reports the labor share, Panel (b) the log capital-labor ratio, and Panel (c) log employment.

Figure A4: Factor Allocation across Distribution of Two-Year Average Labor Productivity

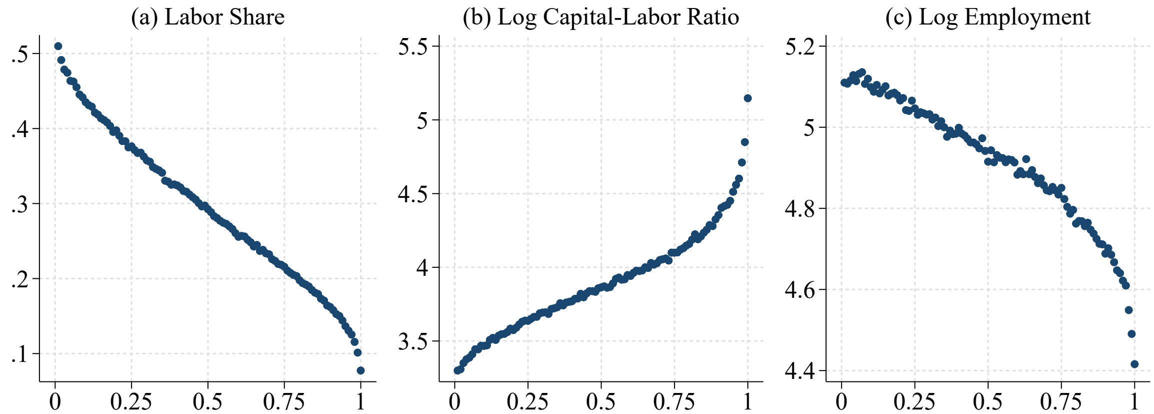

Notes: Each panel presents a binned scatterplot in which each point represents the mean outcome within one of 100 equal-frequency bins of firms' average residualized log labor productivity over the previous two years. Labor productivity and the outcomes are residualized with respect to market and year fixed effects. The residualized outcomes are recentered at year-specific means. Panel (a) reports the labor share, Panel (b) the log capital-labor ratio, and Panel (c) log employment.

Figure A5: Factor Allocation across Leave-One-Out Labor Productivity Distribution

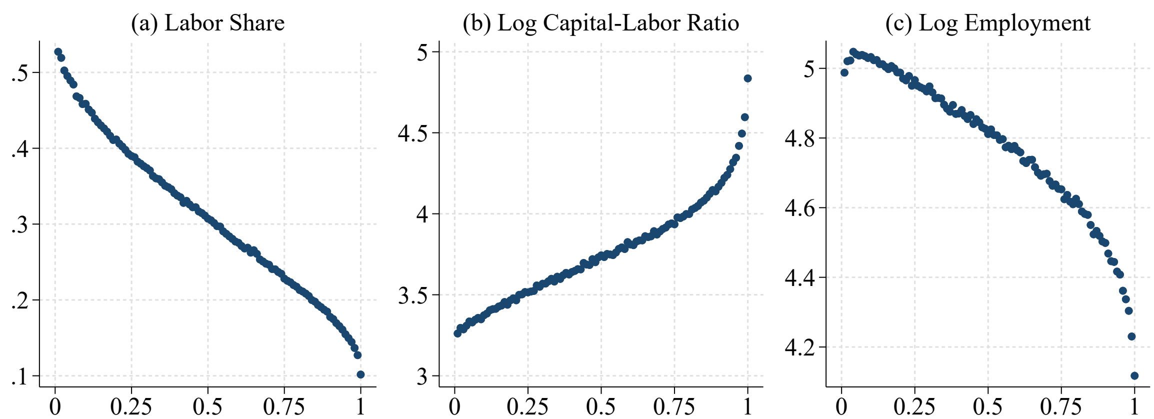

Notes: Each panel presents a binned scatterplot in which each point represents the mean outcome within one of 100 equal-frequency bins of firms' leave-one-out average residualized log labor productivity, excluding the current firm-year observation. Labor productivity and the outcomes are residualized with respect to market and year fixed effects. The residualized outcomes are recentered at year-specific means. Panel (a) reports the labor share, Panel (b) the log capital-labor ratio, and Panel (c) log employment.

Figure A6: Factor Allocation across Three-Year Changes in Labor Productivity

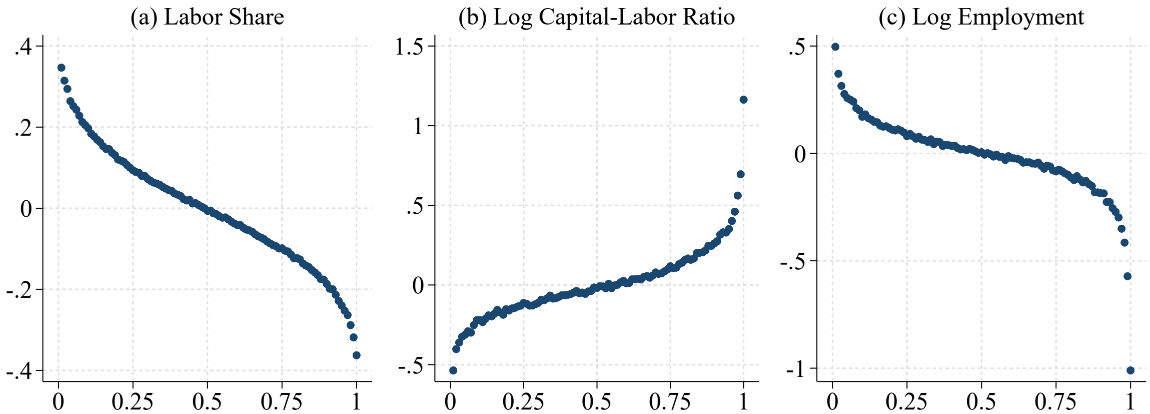

Notes: Each panel presents a binned scatterplot in which each point represents the mean three-year change in the outcome within one of 100 equal-frequency bins of three-year changes in residualized log labor productivity. Labor productivity and the outcomes are residualized with respect to market and year fixed effects. The residualized outcomes are recentered at year-specific means. Panel (a) reports the three-year change in the labor share, Panel (b) the three-year change in the log capital-labor ratio, and Panel (c) the three-year change in log employment.

Figure A7: Factor Allocation across Five-Year Changes in Labor Productivity

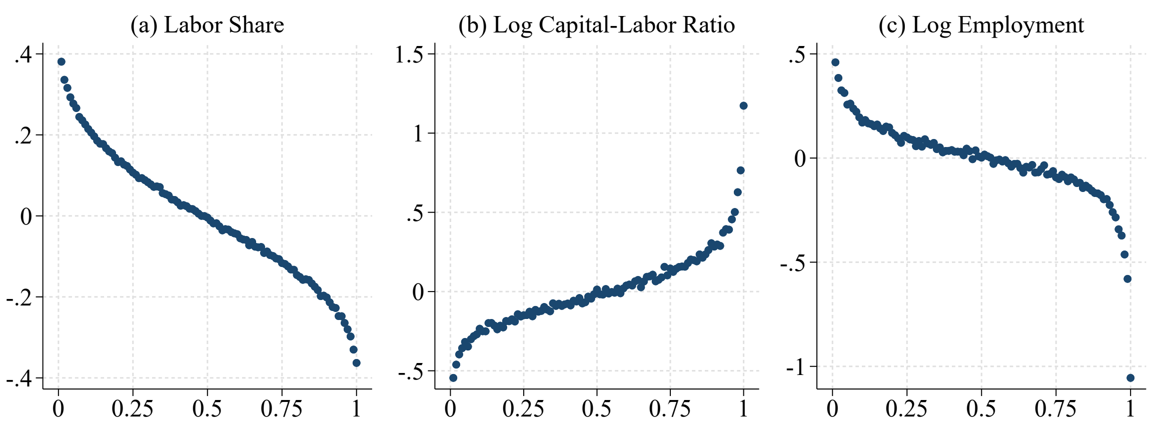

Notes: Each panel presents a binned scatterplot in which each point represents the mean five-year change in the outcome within one of 100 equal-frequency bins of five-year changes in residualized log labor productivity. Labor productivity and the outcomes are residualized with respect to market and year fixed effects. The residualized outcomes are recentered at year-specific means. Panel (a) reports the five-year change in the labor share, Panel (b) the five-year change in the log capital-labor ratio, and Panel (c) the five-year change in log employment.

Figure A8: Factor Allocation across Labor Productivity Distribution: Controlling for Firm Ownership

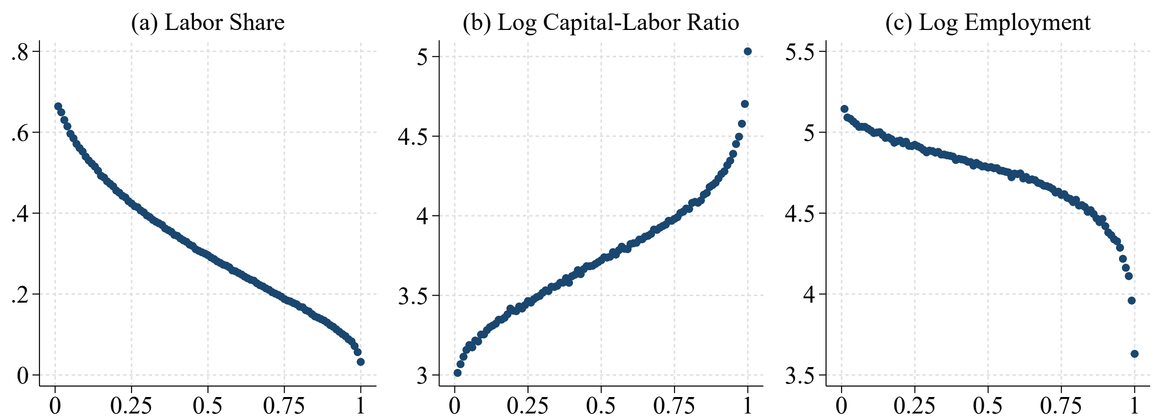

Notes: Each panel presents a binned scatterplot in which each point represents the mean outcome within one of 100 equal-frequency bins of residualized log labor productivity. Labor productivity and the outcomes are residualized with respect to market fixed effects, ownership dummies, and year fixed effects. The residualized outcomes are recentered at year-specific means. Panel (a) reports the labor share, Panel (b) the log capital-labor ratio, and Panel (c) log employment.

Figure A9: Firm Scale and Compensation across Labor Productivity Distribution

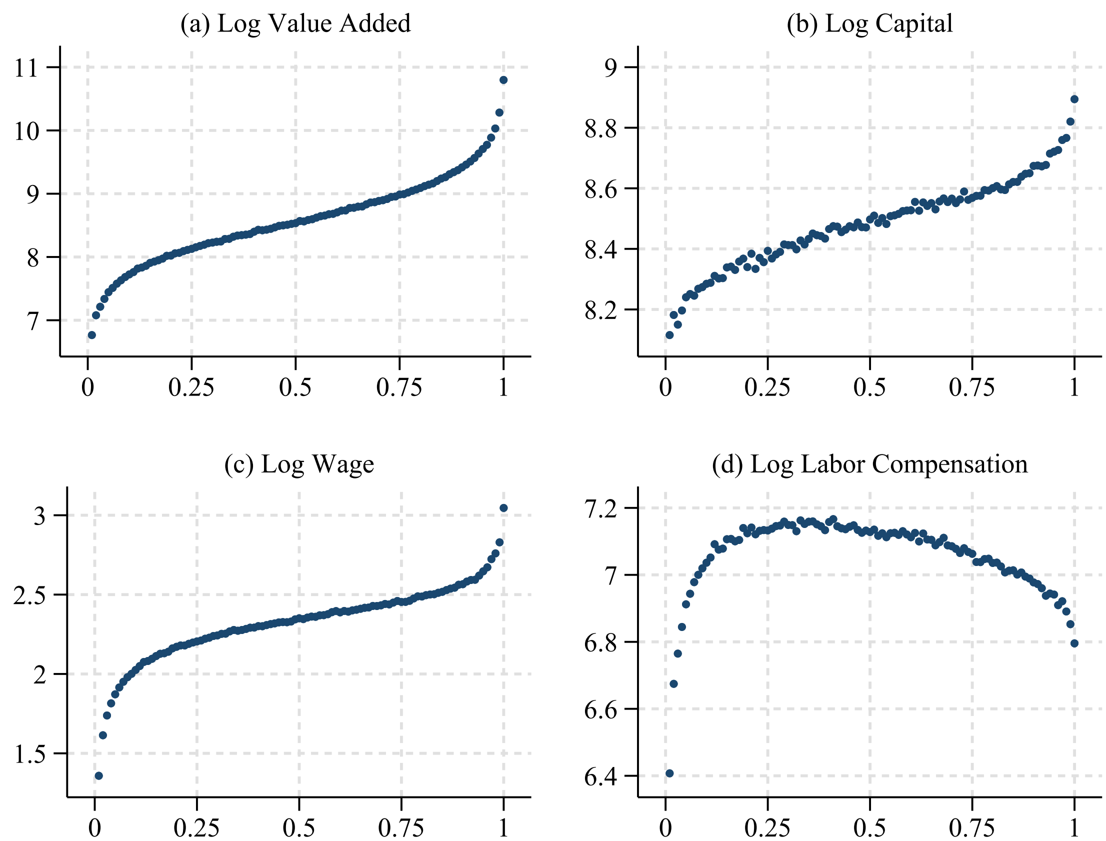

Notes: Each panel presents a binned scatterplot in which each point represents the mean outcome within one of 100 equal-frequency bins of residualized log labor productivity. Labor productivity and the outcomes are residualized with respect to market and year fixed effects. The residualized outcomes are recentered at year-specific means. Panel (a) reports log value-added, Panel (b) log capital, Panel (c) log wage, and Panel (d) log labor compensation.

Figure A10: Firm Scale and Compensation across Odd-Even Labor Productivity Distribution

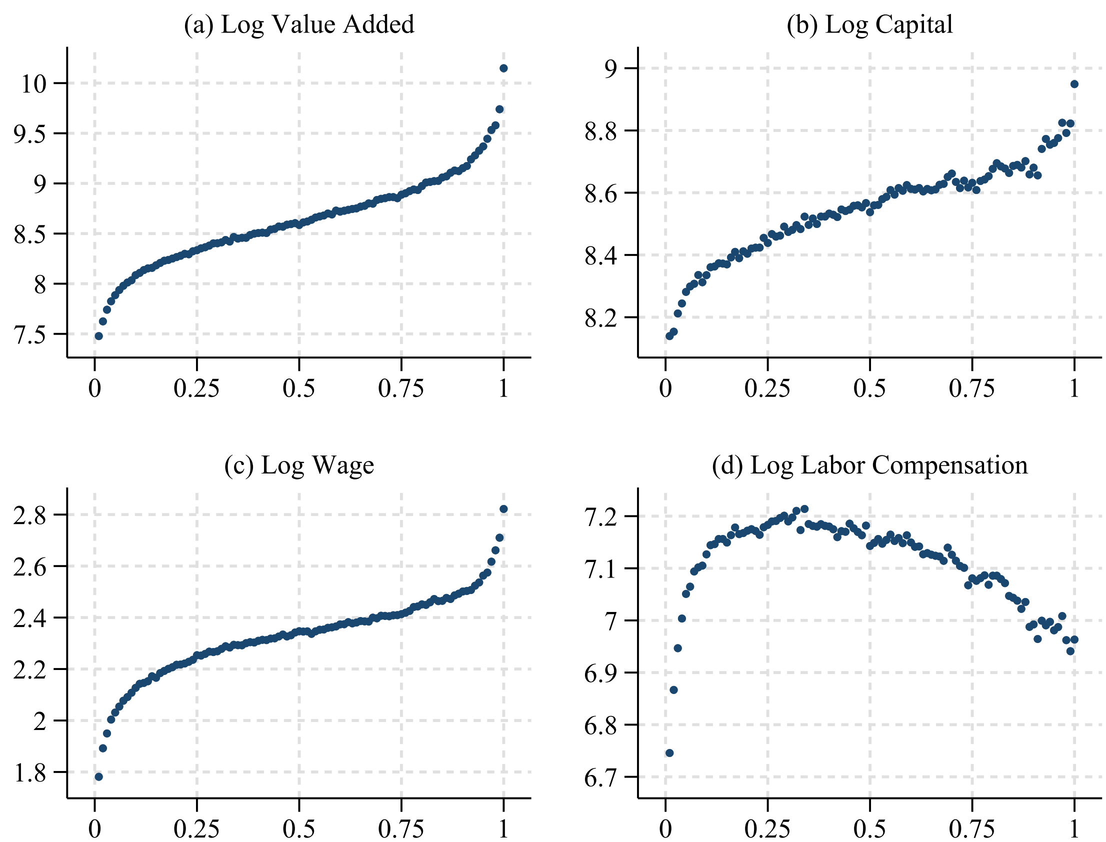

Notes: Each panel presents a binned scatterplot in which each point represents the mean outcome within one of 100 equal-frequency bins of the odd-even labor productivity measure. Even-year observations are sorted by firms' average residualized log labor productivity in odd years, and vice versa for odd-year observations. Labor productivity and the outcomes are residualized with respect to market and year fixed effects. The residualized outcomes are recentered at year-specific means. Panel (a) reports log value-added, Panel (b) log capital, Panel (c) log wage, and Panel (d) log labor compensation.

Appendix G Appendix Tables

Table A1: Summary Statistics of Factor-Augmenting Productivity

Industry	$\log(A_K/A_L)$		$\log A_K$		$\log A_L$	
	Mean	IQR	Mean	IQR	Mean	IQR
	(1)	(2)	(3)	(4)	(5)	(6)
Food processing	–2.74	2.47	0.21	2.09	2.95	0.76
Food manufacturing	–3.41	2.59	–0.34	2.20	3.08	0.78
Beverages	–3.36	2.37	–0.33	1.99	3.03	0.75
Textiles	–3.35	2.27	–0.29	1.94	3.07	0.62
Apparel and footwear	–3.08	2.33	0.01	2.05	3.09	0.51
Leather and related products	–3.04	2.33	0.05	2.01	3.09	0.55
Wood and related products	–2.95	2.41	0.09	2.05	3.04	0.69
Furniture	–3.29	2.56	–0.07	2.15	3.23	0.67
Paper products	–3.31	2.17	–0.19	1.84	3.12	0.65
Printing and media reproduction	–4.11	2.33	–0.85	2.01	3.26	0.64
Cultural, educational, and sports goods	–3.40	2.41	–0.22	2.14	3.17	0.54
Petroleum processing and coking	–2.89	2.17	0.21	1.80	3.10	0.69
Chemicals	–3.09	2.32	0.07	1.97	3.16	0.71
Pharmaceuticals	–3.41	2.29	–0.20	1.94	3.21	0.74
Chemical fibers	–3.44	2.07	–0.21	1.72	3.23	0.52
Rubber products	–3.39	2.36	–0.20	1.97	3.19	0.66
Plastic products	–3.35	2.26	–0.15	1.89	3.19	0.63
Nonmetallic minerals	–3.43	2.33	–0.43	2.01	3.00	0.69
Ferrous metals	–2.87	2.22	0.24	1.89	3.11	0.65
Nonferrous metals	–2.83	2.40	0.33	2.03	3.16	0.63
Metal products	–3.22	2.28	–0.01	1.94	3.20	0.60
General machinery	–3.33	2.28	–0.11	1.94	3.22	0.62
Special machinery	–3.42	2.43	–0.18	2.10	3.23	0.66
Transport equipment	–3.50	2.31	–0.23	1.98	3.27	0.60
Electrical machinery	–3.15	2.26	0.09	1.92	3.24	0.63
Electronics and communication equipment	–3.51	2.43	–0.15	2.16	3.37	0.60
Instruments and office machinery	–3.60	2.45	–0.18	2.22	3.43	0.56
Other manufacturing	–3.07	2.47	0.01	2.14	3.08	0.61
Whole sample	–3.26	2.38	–0.12	2.04	3.14	0.67

Notes: This table reports summary statistics for calibrated firm-level relative factor-augmenting productivity, $\log(A_K/A_L)$, capital-augmenting productivity, $\log A_K$, and labor-augmenting productivity, $\log A_L$, by industry. For each industry, the mean and interquartile range are calculated separately in each year and then averaged over 1998–2007. The interquartile range is the difference between the 75th and 25th percentiles.

Table A2: Robustness of Factor Allocation across Labor Productivity Percentiles

	Labor share	Log capital-labor ratio	Log employment
	(1)	(2)	(3)
Panel A. Income-approach value-added			
80–20 gap	–0.269***	0.869***	–0.222***
	(0.002)	(0.022)	(0.017)
Panel B. One-year lagged labor productivity			
80–20 gap	–0.200***	0.556***	–0.288***
	(0.005)	(0.024)	(0.028)
Panel C. Two-year average labor productivity			
80–20 gap	–0.200***	0.584***	–0.304***
	(0.004)	(0.026)	(0.034)
Panel D. Leave-one-out labor productivity			
80–20 gap	–0.198***	0.521***	–0.362***
	(0.005)	(0.020)	(0.018)
Panel E. Three-year differences in labor productivity			
80–20 gap	–0.243***	0.318***	–0.223***
	(0.006)	(0.017)	(0.007)
Panel F. Five-year differences in labor productivity			
80–20 gap	–0.279***	0.344***	–0.214***
	(0.004)	(0.034)	(0.024)
Panel G. Controlling for firm ownership			
80–20 gap	–0.287***	0.639***	–0.366***
	(0.003)	(0.023)	(0.013)
Panel H. 75–25 percentile gap			
75–25 gap	–0.246***	0.493***	–0.318***
	(0.005)	(0.022)	(0.030)
Panel I. 90–10 percentile gap			
90–10 gap	–0.429***	0.988***	–0.599***
	(0.007)	(0.025)	(0.030)

Notes: In Panels A–G, each estimate reports the difference in the outcome listed in each column between the 80th and 20th percentile groups of the relevant labor productivity measure. Panel A constructs value added and labor compensation using the income approach. Panels B–D use one-year lagged, two-year average, and leave-one-out labor productivity, respectively. Panels E and F sort firms by three-year and five-year differences in labor productivity and report the corresponding gaps in differences in each outcome. Panel G additionally controls for firm ownership by residualizing labor productivity and the outcomes with respect to market fixed effects, ownership dummies, and year fixed effects. Panels H and I report differences between the 75th and 25th percentile groups and between the 90th and 10th percentile groups of residualized log labor productivity, respectively.