Misallocation versus Inequality

Chang-Tai Hsieh Peter J. Klenow

Emmanuella Kyei Manu Emma Rockall *

Thursday 26th March, 2026
Click here for the latest version

Abstract

A stream of research contends that input misallocation takes a toll on aggregate productivity. If this is true, then why does it persist? One answer could be that policies to mitigate misallocation exacerbate inequality. We explore this possibility using data on Indian manufacturing firms and households in a model with both heterogeneous firms and workers. We focus on size-dependent misallocation given the ample evidence for size-dependent policies in India. We find that abandoning such policies would boost aggregate manufacturing productivity by 1-4%, but would also raise the skill premium and overall inequality because larger firms employ more skilled workers and earn more profits for their owners. We estimate that size-based reforms could lower consumption-equivalent welfare on the order of 9%. This suggests removing misallocation may be unpopular because it worsens consumption inequality.

*Hsieh: University of Chicago Booth School of Business and the NBER; Klenow: Stanford University and the NBER; Kyei Manu: Stanford University; Rockall: IMF. We are grateful to seminar participants at the Minneapolis Fed, Columbia University, and Notre Dame BIG for helpful discussions. The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the IMF.

1 Introduction

In Indian manufacturing, larger firms appear to be substantially more productive than smaller ones. Within a sector, firms with 500 or more employees have revenue productivity (value added per unit of inputs) roughly 45 percent above that of firms with 5 or fewer employees. This steep size-productivity gradient is consistent with substantial input misallocation: too many inputs are crowded into small, relatively unproductive firms and too few flow to large, highly productive firms. A growing literature attributes such misallocation in part to size-dependent policies (such as regulations, taxes, and subsidies) that favor small firms and estimates the resulting drag on aggregate productivity to be large.[1]

This raises a natural question: if correcting size-based misallocation raises average incomes, why do such policies persist? One possible answer is that the policies are politically popular. We argue that this is plausible because the same features of the economy that make size-dependent policies costly in aggregate, i.e., large firms are more productive and employ more-skilled workers, also mean that dismantling those policies exacerbates consumption inequality. Workers and small-firm owners who lose their policy-granted advantages are numerous. If most agents would lose even as average output rises, size-dependent reforms could face political headwinds.

We investigate this tradeoff using data on Indian manufacturing firms and households. Our data combine the Annual Survey of Industries (ASI), which covers formal manufacturing firms, and the Survey of Unorganized Manufacturing (SUM), which covers informal firms, for the years 2000, 2005, and 2012, for which both surveys are available. Together they are representative of all manufacturing firms in India. We supplement this with the National Sample Survey (NSS), a representative labor force survey that records consumption, income, hours worked, and employer size.

[1] Examples include Guner, Ventura and Xu (2008), Cai and Pandey (2013), Restuccia and Rogerson (2017), and Guner and Ruggieri (2022).

Three facts from the Indian data motivate our analysis. First, within a sector, revenue productivity is strongly increasing in firm size, consistent with larger firms facing bigger distortions. Second, worker wages are also increasing in firm size, so that workers at the largest firms earn roughly 60 percent more than those at the smallest. Third, the skill mix shifts sharply with size: the share of workers with secondary education or above rises from about 30 percent among firms with 1-5 employees to over 60 percent among firms with more than 20 employees. Wages are higher at larger firms in substantial part because larger firms employ more skilled workers, not merely because those firms share rents.[2]

To quantify the tradeoff, we build a general equilibrium model with heterogeneous firms and heterogeneous workers. Firms differ in factor-neutral productivity, skill-biased productivity, and capital intensity, and face revenue wedges that stand in for size-dependent policies such as low taxation of informal firms. There is a fixed aggregate supply of skilled workers, unskilled workers, and a fixed number of firms. Firms are monopolistically competitive and choose skilled labor, unskilled labor, and capital to maximize profits. It is a small open economy with a perfectly elastic supply of capital at the world interest rate. Workers are hand-to-mouth and supply labor inelastically. Profits flow to absentee firm owners, each of whom owns one firm. Both workers and owners are assumed to have CRRA preferences over consumption.

In this environment, the skill premium depends not just on the supply of skilled and unskilled workers, but on the allocation across skill-biased firm productivities. When wedges are correlated with skill-biased productivity, in particular favoring small firms who are skill-light, removing those wedges expands employment in skill-intensive large firms and contracts it in skill-light small firms. This raises the demand for skilled labor relative to unskilled labor, increasing the skill premium and worsening wage inequality.

[2]We confirm this by showing that controlling for education substantially reduces the firm-size wage gradient in the NSS data.

We use the model to infer firm-level productivities and revenue wedges in the Indian data in order to exactly match each firm's observed value added, wage bill, employment, and physical capital. We then conduct a counterfactual exercise removing the component of revenue wedges attributable to firm size, holding all other firm characteristics fixed. This neutralizes the size-dependent aspect of misallocation while preserving the remaining dispersion in wedges.

Our main findings are as follows. Removing size-dependent wedges raises aggregate manufacturing output by 1–4 percent across the three years. Aggregate profits rise by a similar amount, and high-skill wages increase by roughly 7 percent. Low-skill wages *fall* by about 7 percent, so the skill premium rises by around 14 percentage points. The impact on profits is highly heterogeneous: profits rise for large-firm owners by 15–40 percent and fall for small-firm owners by 15–20 percent. When we map these income changes through to consumption we find that approximately 70–90 percent of all agents, workers and owners alike, experience a consumption loss. The welfare calculation makes the aggregate cost of reform transparent. Under standard preferences, consumption-equivalent welfare for the economy as a whole falls by roughly 9 percent: a modest level gain of about 2.5 percent is swamped by an inequality loss of about 11 percent. Intuitively, the losses are increasing in the substitutability of output across firms and curvature in individual utility over consumption, and decreasing in the elasticity of substitution of skilled vs. unskilled labor within firms. Our results suggest that size-dependent misallocation may be a political equilibrium.

After a brief discussion of the related literature, the rest of the paper is organized as follows. Section 2 documents the motivating facts on the size gradient in revenue productivity, wages, and skill intensity. Section 3 presents the model. Section 4 describes how we measure model objects using the data. Section 5 reports the counterfactual results, including aggregate output, the skill premium, profits, and welfare. Section 6 concludes.

Related Literature. The misallocation literature, following Restuccia and Rogerson (2008) and Hsieh and Klenow (2009), has documented that dispersion in revenue productivity across firms is large in developing countries and potentially translates into substantial TFP losses (Baqaee and Farhi, 2020; Bento and Restuccia, 2017; Bau and Matray, 2023). Size-dependent policies have been identified as a leading source of such misallocation in India (Alfaro and Chari, 2014; Martin, Nataraj and Harrison, 2017; Bertrand, Hsieh and Tsivanidis, 2021) and other developing countries, or even Europe (Garicano, Lelarge and Van Reenen, 2016). We ask what share of misallocation is size-driven, and provide an estimate of the welfare cost of removing it when inequality is accounted for.

A smaller literature has studied the interaction between misallocation and inequality. Jovanovic (2014), Boar and Midrigan (2024), and Guner and Ruggieri (2022) show that misallocation and inequality can be jointly determined. Atkin, Faber and Gonzalez-Navarro (2018) and Atkin, Bernadac, Donaldson, Garg and Huneeus (2025) document distributional consequences of market integration in developing countries. We contribute by showing that the inequality channel is large enough to reverse the welfare sign of a reform that raises aggregate output, and by tracing the mechanism through the labor market: the skill premium rises because large firms are skill-intensive and reform shifts activity toward them.

The literature on size-dependent policies — such as Guner et al. (2008), Alfaro and Chari (2014), Garcia-Santana and Pijoan-Mas (2014), Garicano et al. (2016), Martin et al. (2017), Dabla-Norris, Jaramillo, Lima and Sollaci (2018), Ramaswamy (2021), and Bertrand et al. (2021) — has focused primarily on aggregate productivity effects. We complement this work by quantifying the potential distributional consequences of size-dependent reforms in Indian manufacturing and showing that they may dominate the aggregate gains in welfare terms.

2 Motivating Facts

This section documents three empirical patterns in Indian manufacturing that motivate our model. First, within manufacturing sectors, revenue productivity is increasing in firm size, suggesting that inputs may be misallocated toward smaller and less productive firms. Second, wages and household consumption are increasing in firm size, implying that a reallocation toward larger firms may have important distributional consequences. Third, larger firms employ substantially more educated workers, suggesting their higher average wages stem at least in part from skill intensity. This opens the door to the possibility that size-dependent reforms may increase not only aggregate output but also the skill premium and wage inequality in Indian manufacturing.

2.1 Data sources, construction and cleaning

Our analysis draws on three datasets covering Indian manufacturing firms and households for the years 2000, 2005, and 2012: the Annual Survey of Industries (ASI), the Survey of Unorganized Manufacturing (SUM), and the National Sample Survey (NSS).[3]

Firm data. We combine the ASI and SUM, both conducted by India's National Sample Survey Organization. The ASI is a census of formally registered manufacturing firms with 100 or more employees, supplemented by a random sample of smaller registered firms.[4] The SUM covers the remaining universe of informal manufacturing enterprises, including self-employed proprietors and firms too small to appear in the ASI. Together, the ASI and

[3] The ASI, SUM, and NSS are conducted over fiscal years (April-March). We combine the closest available rounds of each survey and refer to them by the approximate calendar years 2000, 2005, and 2012.

[4] Under India's Factories Act of 1948, firms with more than 20 workers (or 10 or more workers if the firm uses electricity) are required to be formally registered. More exactly, the ASI surveys a firm's combined establishments in a given industry-state; multi-establishment firms may appear as separate observations in different industry-state cells.

SUM are representative of all manufacturing firms in India. In our analysis years, the combined dataset contains approximately 30,000 ASI firms and 75,000 SUM firms per year. Sectors are defined at the four-digit level of the National Industrial Classification (NIC), yielding between 100 and 150 manufacturing sectors depending on the year.

From both surveys we observe firm revenue (more exactly, value added), the book value of fixed capital, the number of workers, and total labor compensation. We pool the ASI and SUM using the official sampling weights provided with each survey. We drop observations with missing or zero reported capital and sectors with fewer than 10 observations in any year. For firms with unpaid family workers, we adjust the wage bill and labor input assuming they work half-time at half the wage of paid workers in the same firm. Where firm-level wages are missing, we impute using the weighted average wage in the same year, sector, survey type, and size category. We trim the distribution of average wages per worker at the 5th and 95th percentiles within each survey type and year, and separately trim the inferred revenue wedges at the 5th and 95th percentiles within each sector and year. Appendix B.1 provides full details, and we report robustness to alternative trimming thresholds in Section 5.4.

Household data. We supplement the firm data with the Employment and Unemployment rounds of India's National Sample Survey (NSS), a nationally representative household survey. The NSS records each individual's education, usual activity status, weekly time use, labor income, household consumption expenditure, and, importantly for our purposes, the number of workers at their place of employment. We restrict the NSS sample to individuals employed in manufacturing for whom employer size is reported. We classify individuals as owners (self-employed or employers) or workers (all other employed individuals, including unpaid family workers). Consumption is measured as annualized monthly household consumer expenditure per capita. We use

these data to document the empirical gradients in wages, consumption, and education across firm sizes in the remainder of this section, and to map the model's income predictions into consumption-equivalent welfare in Section 5.

2.2 Revenue productivity is increasing in firm size

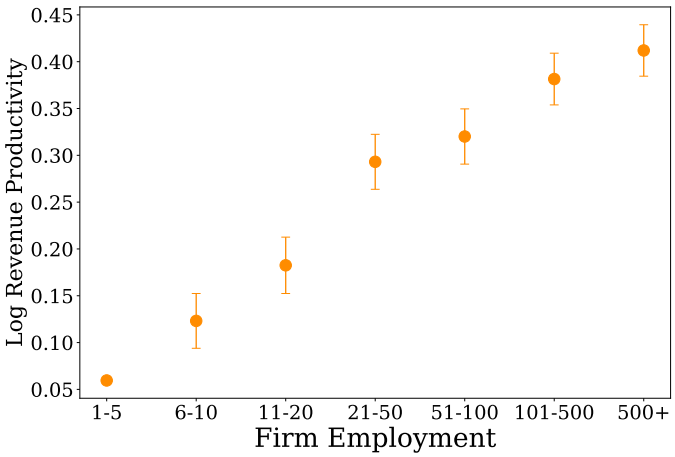
Using the firm data, we construct a measure of revenue productivity proportional to value added per unit of total labor and capital expenditure.[5] Intuitively, this measure is higher for firms that generate more value added from a given amount of labor and capital expenditure. Figure 1 shows that, within a sector, revenue productivity is sharply increasing in firm employment. Firms with 500 or more employees have revenue productivity roughly 45 percent higher than firms with 5 or fewer employees. This pattern is consistent with substantial input misallocation: too many resources are allocated to small, relatively unproductive firms and too few to the large, highly productive firms where they would generate more output.

Not all dispersion in revenue productivity necessarily reflects policy-driven misallocation. Some may arise from measurement error, adjustment costs, or heterogeneous markups that do not correspond to distortions a policymaker could readily target. Misallocation that projects onto firm size, however, is plausibly linked to extensive size-dependent policies in Indian manufacturing. Small firms in India benefit from lower corporate income tax rates, and exemptions from excise and goods-and-services taxes, while the vast informal sector operates largely outside the tax system.[6] Small firms also face

[5] Specifically, revenue productivity is defined as firm value added, scaled by total spending on labor and capital, and adjusted by the standard monopolistic-competition markup. This measure of revenue productivity is consistent with the model we introduce in Section 3 which is a model of monopolistically competitive firms that vary in their capital intensity. We set the elasticity of substitution between firms, $\sigma = 3$.

[6] See PRS Legislative Research (2019) for information on Taxation Laws Ordinance 2019. See Ramaswamy (2016), Velayudhan (2018), Government of India (2006) for summaries of the SSI excise exemption thresholds and their implications for firm behavior. Only about one-fifth of India's 63 million enterprises are registered in the GST network, and over 90 percent of informal workers earn below the income tax exemption threshold (Prasad, 2024; Mehrotra and Parida, 2019; PTI, 2022).

Figure 1: Log revenue productivity by firm size



Notes: The figure shows log revenue productivity by firm size, where revenue productivity is measured as firm value added relative to total expenditure on labor and capital, adjusted by the monopolistic-competition term $(\sigma - 1)/\sigma$ with $\sigma = 3$. Firm size is measured by the number of workers in the firm. We use data from the India ASI + SUM for 2012. The dots are coefficients from revenue-weighted regressions with sector fixed effects. The bars are 95% confidence intervals.

substantially lower labor costs: firms below 20 workers are exempt from Employees' Provident Fund contributions and, in most states, from Employees' State Insurance, creating a 15-16 percent wage cost advantage relative to larger firms.[7] In addition, the government has historically provided direct benefits to small and medium enterprises through subsidized financing, preferential access to government procurement, and a product reservation policy that reserved over 1,000 products for exclusive production by small-scale industry (Garcia-Santana and Pijoan-Mas, 2014; Martin et al.,

[7]See Employees' Provident Fund Organisation (n.d.a,n); Employees' State Insurance Corporation (n.d.); Employees' Provident Fund Organisation (2025) for official information on the statutory coverage thresholds, employer contribution rates, and registration and compliance requirements under the Employees' Provident Fund and Employees' State Insurance schemes. The Industrial Disputes Act (1947) further raises firing costs for firms with more than 100 workers, who must obtain government authorization for retrenchments (Government of India, 1947; Bertrand et al., 2021).

2017).[8] The estimated productivity costs of these policies are large: studies of product reservation (Martin et al., 2017), preferential tax treatment (Bento and Restuccia, 2017), and size-based labor regulations (Bertrand et al., 2021) have documented substantial gains from reform.

The fact that misallocation along the size dimension is closely connected to specific, identifiable policies makes the firm size dimension a natural object for reform. A policymaker who neutralized distortions that favor small firms could, in principle, capture a substantial share of the efficiency gains from reallocation without needing to identify the full set of firm-specific distortions. The central question for the remainder of the paper is whether such apparently efficiency-enhancing reforms may exacerbate inequality (and hence be politically unattractive).

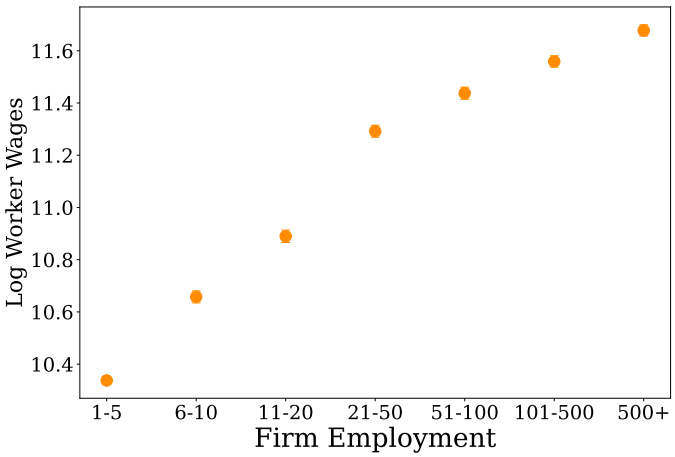
2.3 Worker wages are increasing in firm size

Figure 2 shows that average wages per worker are also strongly increasing in firm size. Within a sector, workers at firms with 500 or more employees earn roughly 60 percent more than workers at firms with 5 or fewer employees. This suggests that a size-neutralizing reform would benefit firms that pay the highest wages. Reallocation toward larger firms may therefore have first-order distributional consequences: even if aggregate output rises, the gains will depend on which workers are employed at the firms that expand and which are employed at the firms that contract.

This wage gradient along the firm size dimension translates into meaningful differences in living standards. Using manufacturing workers and firm owners from the NSS, we regress individual-level outcomes on firm-size bin indicators with controls for gender, age, state, and sector by year fixed

[8] See the National Small Industries Corporation's official materials for information on its mandate to promote MSMEs and its support services in marketing, technology, finance, raw-material assistance, buyer-seller linkages, and government procurement facilitation (National Small Industries Corporation, n.d., 2024)

Figure 2: Log worker wages by firm size



Notes: The figure shows log of average wage per worker within a firm by the firm size, measured as the number of workers in the firm. We use data from the India ASI + SUM for 2012. The regression is revenue-weighted and includes sector fixed effects. The bars are the 95% confidence intervals.

effects.[9] Workers at the largest firms have 62 percent higher labor income relative to those at the smallest firms, and this maps into 38 percent higher household consumption per capita. The pattern is similar for firm owners: consumption is approximately 40 percent higher for owners of the largest firms compared to the smallest. Thus income differences across firm sizes are not fully smoothed through savings or transfers. These consumption differences are not offset by differences in leisure: leisure hours vary by less than 1 percent across firm size bins for both workers and owners.[10]

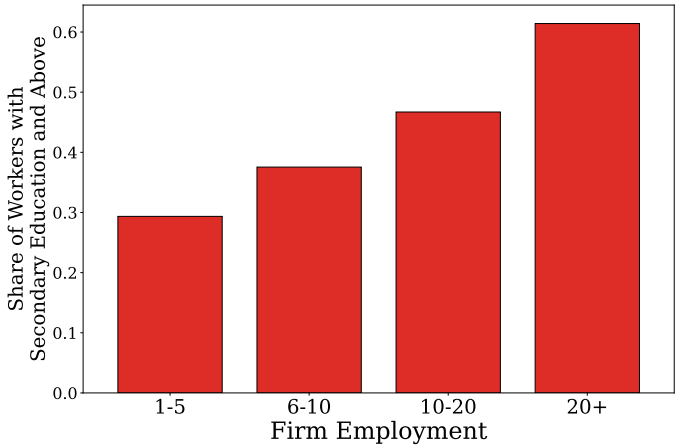
2.4 Education and firm size

Figure 3 shows that the education composition of a firm's workforce rises sharply with its number of employees. The share of workers with secondary education or above rises from about 30 percent among firms with 1–5 employees to over 60 percent among firms with more than 20 employees. This

[9] See Appendix A.1 for the full regression specification and results.
[10] See Appendix A.1 for details on the construction of leisure and the leisure-size gradient.

pattern has potential implications for the distributional consequences of reform: if skilled and unskilled workers are imperfect substitutes in production, then reallocating labor demand toward larger, more skill-intensive firms may increase the relative demand for skilled labor, raising the skill premium and widening wage inequality.

Figure 3: Share of workers with secondary and above education by firm size



Notes: The figure shows the share of workers that have secondary school education and above by the size of the firm that they are employed in, as encoded in the India NSS from 2012.

The education gradient can help explain the wage gradient documented in Section 2.3. Higher wages at larger firms could reflect positive sorting of skilled workers into large firms and/or rent sharing by large firms that pay above-market wages to their workers. To distinguish between these, we add controls for individual education level to the wage-size regressions. Including education attenuates the firm-size gradient in both labor income and consumption, and once we additionally interact education with firm-size bins the residual size gradient is small within education groups.[11] This suggests that the wage-size relationship is primarily driven by the composition of education across firm sizes rather than by rent sharing, and that the relevant margin for understanding the distributional effects of reform is the skill premium rather than firm-specific rents.

[11] See Appendix A.1 for the full regression results with education controls and interactions.

Taken together, the three facts documented in this section point to a tension at the heart of size-dependent policies. The productivity gradient (Section 2.2) implies efficiency gains from reallocation toward larger firms, but the wage and consumption gradients (Section 2.3) imply negative distributional consequences from such reforms, driven by the skill-intensity gradient we document. The popularity and optimality of these reforms will therefore depend on the strength of these two channels, which we quantify in the rest of this paper.

3 Modeling Framework

The empirical patterns documented in Section 2 motivate a general equilibrium model designed to quantify the tradeoff between efficiency gains from reallocation and its adverse distributional consequences. We build on Hsieh and Klenow (2009) by incorporating heterogeneous firms with factor-neutral productivity and firm-specific revenue wedges, and combine this with a Katz and Murphy (1992)-style framework of heterogeneous labor in production, so that firms also differ in skill-biased productivity.

3.1 Economic environment

Production. The economy consists of S sectors. Final output Y is a Cobb-Douglas aggregator of sectoral output, Y_s:

$$Y = \prod_{s=1}^{S} Y_s^{\theta_s} \tag{1}$$

where θ_s is the Cobb-Douglas share of sector s in final output. Sectoral output is in turn a CES aggregator of output from M_s firms in sector s:

$$Y_s = \left(\sum_{i=1}^{M_s} Y_{si}^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}} \tag{2}$$

where $\sigma > 1$ is the elasticity of substitution between firms in a given sector. Firms produce with a Cobb-Douglas production function using capital K_{si} and labor composite Z_{si}:

$$Y_{si} = A_{si} K_{si}^{\alpha_{si}} Z_{si}^{1-\alpha_{si}}. \tag{3}$$

A_{si} is firm i's factor-neutral productivity, and firms also vary in their capital intensity α_{si}.[12] The firm's labor composite is a CES aggregator of low-skill labor, L_{si}, and high-skill labor, H_{si}:

$$Z_{si} = \left(\left(A_{si}^H H_{si} \right)^{\frac{\eta-1}{\eta}} + \left(L_{si} \right)^{\frac{\eta-1}{\eta}} \right)^{\frac{\eta}{\eta-1}}, \tag{4}$$

where $\eta > 1$ is the elasticity of substitution between skill types, and A_{si}^H is the firm's skill-biased productivity.

Preferences. The economy has two types of agents: owners and workers. Both types of agents are hand-to-mouth and have CRRA preferences over final consumption with risk aversion coefficient γ. Each firm has one absentee owner i who consumes the profits from the firm they own. There is a fixed mass of high skill workers, $\overline{H}$, and low skill workers, $\overline{L}$. Workers are endowed with, and inelastically supply, one unit of labor. They spend their wage income on the final consumption good, C.

[12] In the measurement exercise, this assumption allows us to match the heterogeneity of firms' capital shares without relying on modeling mix distortions as in Hsieh and Klenow (2009).

3.2 Market structure and optimizing behavior

We normalize the price of final output, $P = 1$. Sectoral expenditures are

$$P_s Y_s = \theta_s Y \tag{5}$$

where $P_s = \left(\sum_{i=1}^{M_s} P_{si}^{1-\sigma} \right)^{\frac{1}{1-\sigma}}$ is the price of the sectoral good given firm prices P_{si}.

Firms face heterogeneous size distortions, $\tau_{Y_{si}}$, which capture in reduced form the size-dependent policies discussed in Section 2.2, along with other distortions that differentially constrain or expand firm size beyond what would be implied by fundamental productivity differences. The monopolistically competitive firms choose H_{si}, L_{si}, and K_{si} to maximize profits

$$\pi_{si} = (1 - \tau_{Y_{si}}) P_{si} Y_{si} - w_H H_{si} - w_L L_{si} - R K_{si}. \tag{6}$$

Firms hire high and low skill labor in a competitive market and pay wages w_H and w_L. We assume a small open economy, with foreign-owned capital that is supplied perfectly elastically at rental rate R.[13] Consumer demand for each firm's output is given by

$$C_{si} = P_{si}^{-\sigma} P_s^{\sigma} Y_s. \tag{7}$$

The firm's first-order conditions are

$$L_{si}: \quad \frac{w_L}{1 - \tau_{Y_{si}}} = \frac{\sigma - 1}{\sigma} (1 - \alpha_{si}) P_{si} A_{si} \left(\frac{Z_{si}}{K_{si}} \right)^{\alpha_{si} - 1} \left(\frac{Z_{si}}{L_{si}} \right)^{\frac{1}{\eta}} \tag{8}$$

$$H_{si}: \quad \frac{w_H}{1 - \tau_{Y_{si}}} = \frac{\sigma - 1}{\sigma} (1 - \alpha_{si}) P_{si} A_{si} A_{si}^H \left(\frac{Z_{si}}{K_{si}} \right)^{\alpha_{si} - 1} \left(\frac{Z_{si}}{A_{si}^H H_{si}} \right)^{\frac{1}{\eta}} \tag{9}$$

$$K_{si}: \quad \frac{R}{1 - \tau_{Y_{si}}} = \frac{\sigma - 1}{\sigma} \alpha_{si} P_{si} A_{si} \left(\frac{K_{si}}{Z_{si}} \right)^{\alpha_{si} - 1} \tag{10}$$

[13] In Section 5.4 we gauge the robustness of our results to instead assuming capital is domestically-owned and inelastically supplied.

These optimality conditions state that firms choose labor and capital so that each input's marginal revenue product equals its price. Under monopolistic competition, marginal revenue products are scaled down by the markup $(\sigma - 1)/\sigma$. A higher size distortion $\tau_{Y_{si}}$ further lowers the revenue return to production, shrinking the firm's demand for all inputs. Differences in A_{si} shift overall input demand, whereas differences in A_{si}^H tilt demand toward high-skill labor by raising its relative productivity within the labor composite. The optimal skill ratio is given by

$$\frac{H_{si}}{L_{si}} = \left(\frac{w_L}{w_H}\right)^{\eta} \left(A_{si}^H\right)^{\eta-1}. \tag{11}$$

This is the analogue, across heterogeneous firms at a point in time, of the Katz and Murphy (1992) relationship over time. The optimal skill ratio is increasing in A_{si}^H: firms with higher skill-biased productivity optimally employ relatively more skilled workers, with the magnitude of these differences determined by the elasticity of substitution between skill types, η.

3.3 Competitive equilibrium and allocative efficiency

Equilibrium definition. Given R, the equilibrium consists of labor allocations $\{H_{si}, L_{si}\}_{s,i}$, capital allocations $\{K_{si}\}_{s,i}$, and prices w_H, w_L, $\{P_{si}\}_{s,i}$ such that

1. Firms maximize profits given in (6)
2. Consumer demand is given by (7)
3. Sectoral expenditures are given by (5)
4. Labor markets clear: $\sum_{i,s} H_{si} = \overline{H}$ and $\sum_{i,s} L_{si} = \overline{L}$

Efficient allocation. Consider a planner who takes firm-level fundamentals $\{A_{si}, A_{si}^H, \alpha_{si}\}$ as given and chooses the allocation of inputs across firms to maximize final output subject to the aggregate feasibility constraints for skilled

and unskilled labor, $\overline{H}$ and $\overline{L}$. Following Hsieh and Klenow (2009), the planner's allocation coincides with the decentralized equilibrium if and only if revenue distortions are equalized across firms, that is $\tau_{Y_{si}} = \tau_{Y_{sj}}$ for all i, j.[14] Dispersion in revenue wedges therefore distorts the allocation of inputs across firms and reduces aggregate output relative to the efficient benchmark.

Because the model does not yield a closed-form expression for the output loss from dispersion in revenue distortions, we illustrate the relationship numerically in a simplified environment with one sector and no capital. Specifically, we assume a joint lognormal distribution for A_{si}, A_{si}^H, and $(1 - \tau_{Y_{si}})$, and plot allocative efficiency, defined as aggregate output in the decentralized equilibrium ("DE") relative to aggregate output from this planner problem ("SP"), Y^{DE}/Y^{SP}, as a function of the variance of $(1 - \tau_{Y_{si}})$. This is shown in Panel (a) of Figure 4.

In Figure 4a, log allocative efficiency is decreasing linearly in the variance of $(1-\tau_{Y_{si}})$ with a slope of $-\frac{1}{2}\sigma$, where σ denotes the elasticity of substitution across firms. Thus, greater dispersion in firm-level wedges leads to a larger aggregate efficiency loss, and the magnitude of that loss is increasing in σ along with the variance of wedges. When firm output is more substitutable, the wedges distort input allocations more severely and generate a larger misallocation cost.

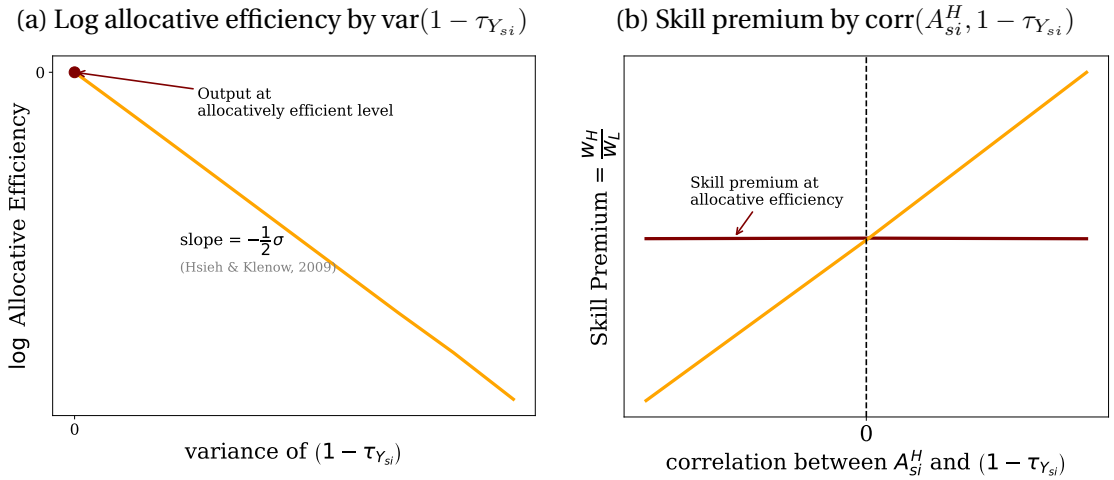
3.4 The tradeoff between misallocation and inequality

Having characterized how revenue wedges generate misallocation, we turn to the central question of the paper: how these same distortions affect inequality in equilibrium and hence the tradeoff between allocative efficiency and distributional outcomes.

The owner-side tradeoff is more direct, and is already present in Hsieh and Klenow (2009). Since each firm has one absentee owner and owners are hand-to-mouth, inequality across owners is simply inequality in profits across

[14] In that case, the Lagrange multipliers on the feasibility constraints, λ_H and λ_L, coincide with the wages w_H and w_L in the decentralized equilibrium.

Figure 4: Misallocation vs inequality tradeoff for workers



Notes: The figure shows results from a one-sector version of the model without capital, under a joint lognormal distribution for A_{si}, A_{si}^H, and $1-\tau_{Y_{si}}$. For each parameterization, we solve for both the decentralized equilibrium and the efficient allocation with no dispersion in $1-\tau_{Y_{si}}$. The left panel plots log allocative efficiency, Y^{DE}/Y^{SP}, against $\text{var}(1-\tau_{Y_{si}})$. The right panel plots the skill premium, w_H/w_L, against $\text{corr}(A_{si}^H, 1-\tau_{Y_{si}})$.

firms. In our setting, profits are a constant share of net-of-distortion revenue, $\pi_{si} = \frac{1}{\sigma}(1-\tau_{Y_{si}})P_{si}Y_{si}$. Revenue wedges therefore affect owner incomes by changing both the net profit rate and the distribution of gross profits across firms. When the wedges favor smaller and less profitable firms while implicitly taxing larger and more productive firms, they compress the distribution of owner income but reduce aggregate output. Removing such distortions raises efficiency, but it also reallocates profits toward the owners of larger and more productive firms, thereby increasing inequality across owners. The owner-side tradeoff is therefore strongest when the firms facing the largest distortions are precisely those that would otherwise earn the highest profits.

The worker-side tradeoff depends on how misallocation interacts with the skill premium, w_H/w_L. Using the optimal skill ratio in (11) and the labor market

clearing conditions, the equilibrium skill premium is given by

$$\left(\frac{w_H}{w_L}\right)^{\eta} = \left(\sum_{i,s} \frac{L_{si}}{\overline{L}} \left(A_{si}^H\right)^{\eta-1}\right) \frac{\overline{L}}{\overline{H}}. \tag{12}$$

Equation (12) implies that the skill premium is governed by two components. The first is the relative supply of low-skilled to high-skilled labor, with a higher relative supply of low-skilled labor raising the equilibrium skill premium. The second is the allocation of labor across firms with heterogeneous skill-biased productivity, A_{si}^H. A reallocation of labor toward firms with higher A_{si}^H increases the relative demand for skilled labor, because such firms employ skilled workers more intensively, and therefore raises the skill premium. This expression shows that misallocation affects wage inequality only insofar as it changes how labor is allocated across firms with different levels of skill-biased productivity. In the simulations of the one sector model with no capital (Figure 4), we find this mechanism to be shaped by the covariance between A_{si}^H and $(1-\tau_{Y_{si}})$. When firms with high skill-biased productivity face relatively larger distortions, misallocation compresses the skill premium.

In Figure 4b, we plot the skill premium as a function of the correlation between A_{si}^H and $(1-\tau_{Y_{si}})$. For each parameterization, we solve for the skill premium in the decentralized equilibrium at a given level of dispersion in $(1-\tau_{Y_{si}})$, and compare it to the skill premium under the efficient allocation where $\text{var}(1-\tau_{Y_{si}}) = 0$. The pattern is clear: when $\text{cov}(A_{si}^H, 1-\tau_{Y_{si}}) < 0$, eliminating misallocation raises the skill premium. Intuitively, in this case the distorted equilibrium allocates more labor toward firms with relatively low skill-biased productivity relative to the efficient allocation. Restoring efficiency therefore reallocates production toward skill-intensive firms, increases the relative demand for skilled labor, and raises wage inequality. This condition therefore governs the worker-side efficiency-inequality tradeoff in the model: wedges that misallocate labor away from skill-intensive firms reduce aggregate output, but they also compress wage inequality.

4 Measurement

In this section, we describe how we map the model laid out in Section 3 to the Indian manufacturing data. Using firm-level information on value-added, capital, employment, and the wage bill, we recover the key firm-level objects in the model: skill inputs, factor-neutral productivity, skill-biased productivity, capital shares, and revenue wedges. We then summarize the implied distribution of workers and owners and document the empirical relationships between the recovered productivity measures and wedges that are central to the paper's quantitative results presented in Section 5.

4.1 Inferring wages, skill, productivities and wedges

The model can be inverted firm by firm to recover the latent objects required for the quantitative analysis. Given data on firm value-added, capital, employment, and wage bill, and conditional on values for η, σ, and R, we recover $\{H_{si}, L_{si}, A^H_{si}, \alpha_{si}, \tau_{Y_{si}}, A_{si}\}$ in sequence. Specifically, we first infer economy-wide skilled and unskilled wages. Then we use these wages together with firm wage bills and headcounts to recover skill inputs. We next back out skill-biased productivity from skill ratios. We recover capital shares and revenue wedges from the first-order conditions. Finally, we infer factor-neutral productivity from the consumer demand function and firm value-added.

We identify the economy-wide skilled and unskilled wages, w_H and w_L, using the cross-firm distribution of average wages per worker in each year. In particular, we treat the upper and lower tails of this distribution as informative about the wages paid to predominantly skilled and predominantly unskilled workers. In practice, we anchor w_H and w_L to the highest and lowest observed firm-level average wages per worker, after trimming outliers, which corresponds to interpreting the extreme firms as the limiting cases of fully skilled and fully unskilled employment:

$$w_H \equiv \max_{si} \left\{ \frac{\text{wage bill}_{si}}{\text{workers}_{si}} \right\}, \qquad w_L \equiv \min_{si} \left\{ \frac{\text{wage bill}_{si}}{\text{workers}_{si}} \right\}. \tag{13}$$

Because this strategy relies on the tails of the wage distribution, it is sensitive to outliers. We therefore trim the distribution of average wages per worker at the 5th and 95th percentiles within each survey type and year, as described in Appendix B.1. We report robustness to alternative trimming thresholds in Section 5.4.

Given w_H and w_L, we solve for each firm's skilled and unskilled employment using two equations: the wage bill identity and the headcount identity:

$$w_H H_{si} + w_L L_{si} = \text{wage bill}_{si} \tag{14}$$

$$H_{si} + L_{si} = \text{workers}_{si}. \tag{15}$$

This yields $H_{si} = (\text{wage bill}_{si} - w_L \times \text{workers}_{si})/(w_H - w_L)$ and $L_{si} = \text{workers}_{si} - H_{si}$. Given H_{si}, L_{si}, w_H, w_L, and a value for the elasticity of substitution between skill types η, we invert the optimal skill ratio in (11) to recover A^H_{si} for each firm. Our baseline sets $\eta = 1.5$, following Katz and Murphy (1992). We report robustness to $\eta \in \{2, 3, 4\}$ in Section 5.

The firm's first-order conditions for capital and the labor composite jointly pin down the capital share α_{si} and the revenue wedge $\tau_{Y_{si}}$. Equating the two expressions for the keep rate $(1 - \tau_{Y_{si}})$ from the capital and labor optimality conditions yields

$$\alpha_{si} = \frac{RK_{si}}{w_{si}Z_{si} + RK_{si}}, \tag{16}$$

where $w_{si}Z_{si}$ is total worker compensation. We set $R = 0.10$, following Hsieh and Klenow (2009). The revenue wedge then follows from either first-order

condition:

$$1 - \tau_{Y_{si}} = \frac{1}{\alpha_{si}} \cdot \frac{\sigma}{\sigma - 1} \cdot \frac{RK_{si}}{P_{si}Y_{si}}. \tag{17}$$

We set the elasticity of substitution across varieties within a sector to $\sigma = 3$ and report robustness to $\sigma \in \{4, 5\}$.

Finally, we recover A_{si} by inverting the demand function:

$$A_{si} = \kappa_s \cdot \frac{(P_{si}Y_{si})^{\frac{\sigma}{\sigma-1}}}{K_{si}^{\alpha_{si}} \left[(A_{si}^H H_{si})^{\frac{\eta-1}{\eta}} + L_{si}^{\frac{\eta-1}{\eta}} \right]^{\frac{(1-\alpha_{si})\eta}{\eta-1}}}, \tag{18}$$

where κ_s is a sector-specific constant that we normalize to one and then pin down A_{si} levels by matching aggregate output in the data. Sectoral expenditure shares θ_s are inferred directly from sectoral revenue shares in the data.

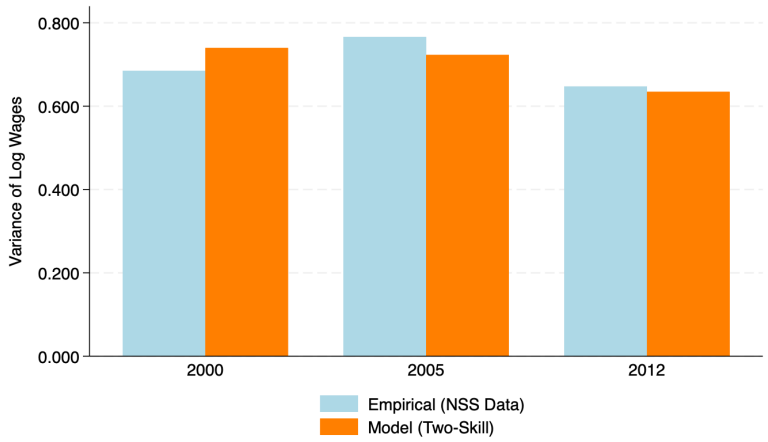
The inversion procedure recovers all firm-level objects required for the quantitative analysis, conditional on the maintained assumption that labor is summarized by two skill types paid common economy-wide wages. A natural concern is whether this parsimonious structure is sufficiently rich to capture the observed dispersion in labor income in the data. As a first pass to assess this, we compare the variance of log wages implied by the model to the empirical variance of log wages in the NSS Employment and Unemployment Survey, restricting to manufacturing workers.[15] As shown in Figure 5, the model generates a variance of log wages that closely matches the empirical variance in each year, suggesting that the binary skill structure captures the first-order variation in labor income across Indian manufacturing workers.

As a further check, in Section 5.3 we verify that our main results are robust to allowing for a continuous distribution of skill types, using the observed wage distribution from the NSS to calculate the share of workers who would lose from reform under a continuum of skill endowments rather than two discrete

[15] We compute $\text{Var}(\log w_i)$ across all workers, where worker i is assigned wage w_H with probability $H_{si}/(H_{si} + L_{si})$ and w_L otherwise, weighted by the firm's sampling weight.

types. In this exercise, each worker is allowed to have a continuous skill type, so earnings are a weighted average of the high-skill and low-skill wages.

Figure 5: Variance of log wages in India NSS versus model with two skill types



Notes: The figure compares the variance of log wages from the baseline model with two skill types to the empirical variance of log wages from the NSS Employment and Unemployment Survey. The empirical measure uses individual-level weekly earnings among manufacturing workers, annualized by multiplying by 52, with NSS sampling weights. The sample is restricted to the same years as the firm data.

4.2 Descriptive statistics

With the recovered latent firm-level objects, we next summarize the key features of the measured economy that matter for our counterfactual analysis. We first describe the distribution of agents across skill types and ownership status, since these determine who gains and who loses from reform. We then examine the relationship between inferred productivities and revenue wedges, which is central to both the magnitude of misallocation and its distributional consequences.

Table 1 summarizes the distribution of agents in the economy implied by our measurement procedure. Panel A reports the inferred stocks of skilled and unskilled labor. Low-skilled workers comprise approximately 90 percent of the manufacturing labor force in each year, with both skill types growing over time.

Panel B reports the split between workers and firm owners, where each firm is assumed to have one absentee owner. Workers account for roughly two-thirds of all agents. The predominance of unskilled workers and small-firm owners will be important for understanding why the counterfactual reform is welfare-reducing for the majority of agents.

Table 1: Distribution of agents

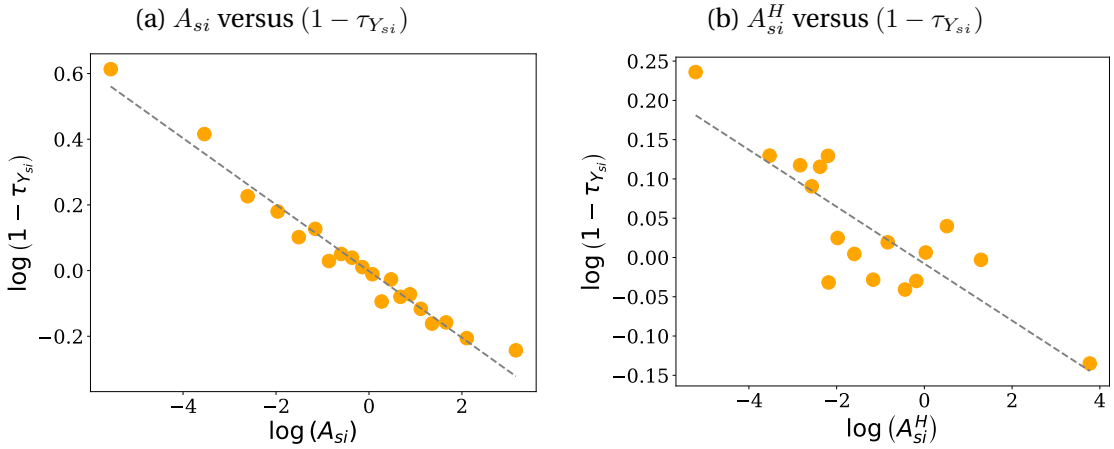
	$\overline{H}$	$\overline{L}$	$\frac{\overline{L}}{\overline{L}+\overline{H}}$	Workers	Owners
	millions			%	%
2000	2.1	18.2	0.90	67	33
2005	2.1	19.4	0.90	66	34
2012	2.5	22.5	0.90	66	34

Notes: Skill types are inferred from the firm data as described in Section 4.1. Workers are counted as the number of workers reported by firms in the ASI and SUM (adjusted for unpaid family labor). Owners are counted as the sample-weighted number of firms, since each firm is assumed to have one absentee owner. India ASI + SUM.

The next set of moments concerns the distribution of distortions across firms and how those distortions line up with firm productivity. Figure 6 plots the relationship between the inferred revenue wedges and the two productivity measures. The left panel shows that factor-neutral productivity A_{si} is strongly negatively correlated with the keep rate $(1 - \tau_{Y_{si}})$: more productive firms face higher effective tax rates. This is the standard misallocation pattern documented by Hsieh and Klenow (2009): inputs are misallocated away from high-productivity firms toward low-productivity firms. The right panel shows that the same pattern holds for skill-biased productivity A_{si}^H, albeit much less tightly: firms with higher skill-biased productivity also face higher wedges. As outlined in Section 3.4, this correlation is what generates the key efficiency-equity tradeoff generated by reform in our model. When size-dependent wedges are removed, production shifts toward firms that are not only more productive but also more skill-intensive, raising the relative

demand for skilled labor and increasing the skill premium.[16]

Figure 6: Relationship between productivities and wedges



Notes: The graphs show binscatters with 20 bins for the relationship between the inferred wedges and productivity measures (see Section 4.1). Each point represents a revenue-weighted bin, residualized on the sector mean. India ASI + SUM 2012.

5 Counterfactual Exercise and Results

This section quantifies the consequences of removing size-dependent distortions. We first construct a counterfactual economy that eliminates the systematic relationship between revenue wedges and firm size while preserving the remaining within-size dispersion in wedges. We then solve for the new general equilibrium and assess its implications for aggregate output, the wage structure, firm profits, and welfare.

5.1 Neutralizing wedges along the firm size dimension

To quantify the tradeoff between size-dependent misallocation and inequality, we construct a counterfactual that removes the component of revenue wedges

[16]The relationship between A^H_{si} and $(1-\tau_{Y_{si}})$ is robust to alternative values of η. See Appendix B.3.

systematically related to firm size while preserving the remaining dispersion in wedges. Specifically, for each year we regress the log keep rate on firm-size-bin indicators and sector fixed effects, weighting observations by firm revenue:

$$\log(1-\tau_{Y_{si}}) = \sum_{b\in\mathcal{B}} \beta_b \mathbb{1}\{\text{size_bin}_{si} = b\} + \gamma_s + \varepsilon_{si}.$$

We then construct counterfactual wedges from the residual component of this regression, recentered so that the mean keep rate is unchanged from baseline. The counterfactual therefore eliminates the systematic size gradient in wedges within sectors while preserving the remaining within-size-bin heterogeneity.

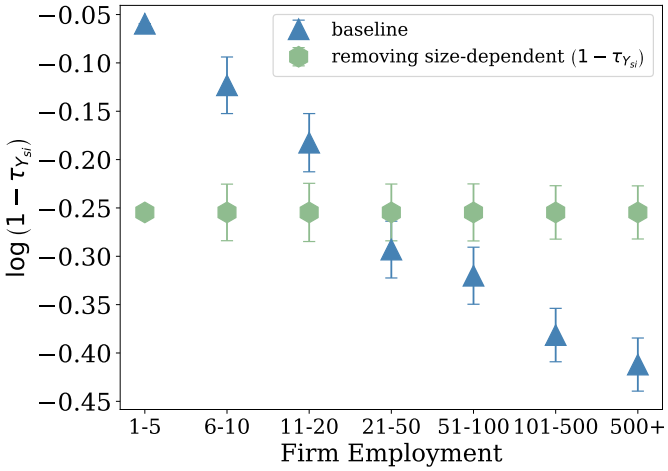
Figure 7 illustrates the exercise. In the baseline, the keep rate declines sharply with firm size, so that smaller firms face lower effective tax rates than larger firms. In the counterfactual, this size gradient is removed. The resulting reform reallocates employment and output from smaller to larger firms, as the implicit relative subsidy to small firms is eliminated and larger, more productive firms expand. Figure 8 shows the share of total employment and output by firm size in the baseline and counterfactual economies. The reform shifts both employment and output from the smallest firms toward larger firms. We then solve for the new general equilibrium, holding fixed all firm-level fundamentals $(A_{si}, A^H_{si}, \alpha_{si})$ and finding the wages w_H and w_L that clear the skilled and unskilled labor markets under the counterfactual wedges.[17]

5.2 Aggregate output, the skill premium, and profits

Table 2 reports the effects of removing size-dependent wedges. Aggregate output rises by 1–4 percent across the three years, confirming that distortions along the firm-size dimension carry meaningful productivity costs. The gains,

[17] The endogenous reallocation of inputs has little effect on the flat gradient, as one can see from Figure 7.

Figure 7: Log revenue wedges by firm size: baseline versus counterfactual

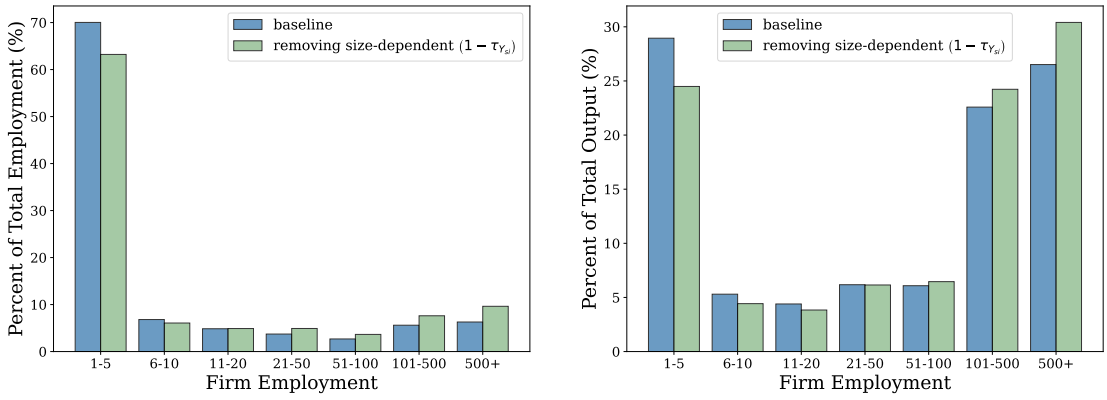


Notes: The figure plots log keep rates by firm size in the baseline and in the counterfactual that removes the size component of wedge dispersion. Baseline wedges are inferred from the data as described in Section 4.1. The regressions are revenue-weighted and include sector fixed effects. Bars denote 95% confidence intervals. India ASI + SUM, 2012.

Figure 8: Employment and output by firm size: baseline versus counterfactual

(a) Employment share by firm size

(b) Output share by firm size



Notes: The figure shows the share of total employment (Panel 8a) and total output (Panel 8b) by the firm's initial size in the baseline economy. The blue bars show this for the baseline economy, and the green bars show it for the counterfactual economy where we neutralize the dispersion in the revenue wedges along the size dimension. India ASI + SUM 2012.

however, are not evenly distributed across skill groups. High-skilled wages rise by 4–7 percent, reflecting the expansion of larger, more skill-intensive firms. By contrast, low-skilled wages fall by 6–7 percent, as the contraction of smaller, skill-light firms reduces the demand for the low-skilled workers concentrated there. As a result, the skill premium increases by 10–15 percent.

Table 2: Change in aggregate output, wages, and skill premium

	Percent Change (%)			
	ΔY	Δw_H	Δw_L	$\Delta \frac{w_H}{w_L}$
2000	1.9	3.9	-5.7	10.2
2005	4.3	6.8	-6.6	14.3
2012	1.4	6.8	-7.1	15.0

Notes: The table reports percent changes from the baseline to the counterfactual economy that removes the size component of wedge dispersion (Section 5.1). India ASI + SUM.

Table 3 shows that the qualitative results are robust to alternative values of the elasticity of substitution between skill types, and highlights the role that imperfect skill substitutability plays in the tradeoff. Higher values of η attenuate the increase in the skill premium and amplify aggregate output gains, as greater substitutability makes it easier for firms to adjust their skill mix. Even at $\eta = 4$, low-skilled wages fall, indicating that the reform amplifies inequality even without low substitutability between skill types.

These results contrast sharply with the standard skill-biased technical change (SBTC) mechanism. Although SBTC also raises the skill premium, it is typically Pareto-improving in the sense that both skilled and unskilled wages rise, with skilled wages rising by more. This is true in our model as well; in Appendix A.2.1 we implement a SBTC counterfactual in which A_{si}^H is increased uniformly across firms, calibrated to match the same percentage increase in aggregate output as the size-based reform. Both low and high skill wages increase in this scenario. By contrast, the size-based reform studied here benefits large, skill-intensive firms at the expense of smaller, skill-light firms.

The welfare losses we find therefore do not reflect a generic consequence of skill-biased change, but arise specifically from the cross-firm reallocation induced by removing size-dependent distortions.

The reform also has heterogeneous effects on owners through firm profits. Table 4 shows that aggregate profits rise by 1–4 percent, but this masks substantial heterogeneity across firms. Conditional on gaining, the average firm's profits increase by 2–5 percent; conditional on losing, they fall by 14–18 percent. Figure 9 makes the source of this heterogeneity clear. The histogram in panel (a) shows that the distribution of profit changes is right-skewed: most firms experience modest losses, while a smaller number of large firms gain substantially. Panel (b) shows that these losses are concentrated among smaller firms, while the gains accrue disproportionately to larger firms. Thus, the reform redistributes profit income away from the owners of small firms and toward the owners of large firms.

Table 3: Change in aggregate output and low-skill wages by skill substitutability

	$\eta = 1.5$	$\eta = 2$	$\eta = 3$	$\eta = 4$
		2000		
ΔY	2.0	2.0	2.2	2.2
Δw_L	-5.7	-4.7	-3.4	-2.5
		2005		
ΔY	3.3	3.7	4.1	4.4
Δw_L	-6.6	-4.8	-2.7	-1.4
		2012		
ΔY	1.4	1.7	2.0	2.2
Δw_L	-7.1	-5.5	-3.5	-2.2

Notes: The table reports percent changes in aggregate output and low-skill wages under the counterfactual that removes the size component of wedge dispersion, for alternative values of the elasticity of substitution between skill types, η. For each η, firm-level objects are re-inferred from the data as described in Section 4.1. India ASI + SUM.

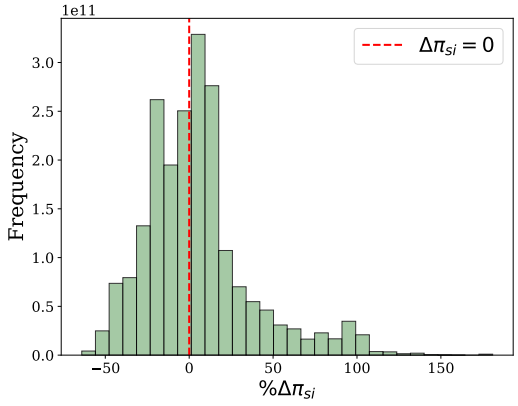
Table 4: Change in aggregate profits and firm profits

	Percent Change (%)		
	$\Delta\Pi$	$\Delta\ \pi_{si}\|\pi_{si}\uparrow$	$\Delta\ \pi_{si}\|\pi_{si}\downarrow$
2000	1.6	2.4	-13.7
2005	3.9	4.6	-15.9
2012	1.1	3.2	-18.1

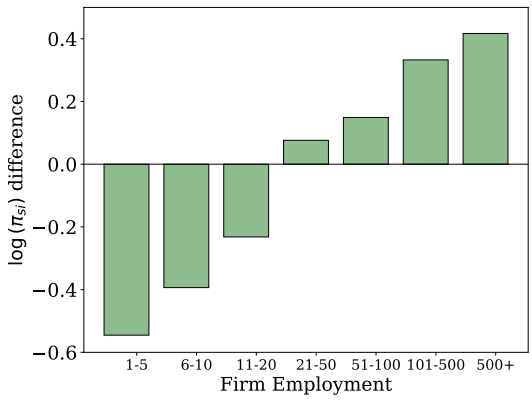
Notes: The table reports the percent change in aggregate profits and the average percent change in firm profits conditional on gaining or losing, from the baseline to the counterfactual economy that removes the size component of wedge dispersion. India ASI + SUM.

Figure 9: Heterogeneity in change in firm profits

(a) Distribution of firm profit changes



(b) Log change in firm profits by firm size



Notes: The figure shows heterogeneity in firm profit changes from the baseline to the counterfactual economy that removes the size component of wedge dispersion. Panel (a) plots the revenue-weighted distribution of percentage profit changes across firms. Panel (b) plots the change in log profits by firm size, measured by baseline employment; the regression is revenue-weighted and residualized on sector means. India ASI + SUM, 2012.

5.3 Political economy and welfare implications

The combination of falling low-skilled wages and falling profits for small-firm owners implies that the vast majority of agents lose from reform. Table 5 shows that only 9–11 percent of all agents experience an income gain in the counterfactual economy. Among workers, roughly 90 percent lose, reflecting the predominance of low-skilled labor in the manufacturing workforce.

Table 5: Percent of agents that gain in counterfactual economy

	Percent of Agents (%)			
	Workers + Owners	Workers	Owners	Owners (revenue-weighted)
2000	11	10	13	59
2005	11	10	12	57
2012	9	10	7	54

Notes: The table reports the share of agents whose income rises in the counterfactual economy that removes the size component of wedge dispersion. For workers, gainers are those whose wage rises; for owners, gainers are those whose profits rise. The final column weights owners by baseline firm revenue. Cells below 50 percent are shaded red and cells above 50 percent green. India ASI + SUM.

This conclusion does not hinge on the assumption of binary skill types. Table 6 relaxes the binary skill assumption by allowing workers to have a continuous skill endowment ϕ, so that a worker's wage is $\phi \cdot w_H + (1-\phi) \cdot w_L$. The indifference threshold ϕ^* solves $\phi^* \cdot \Delta w_H + (1-\phi^*) \cdot \Delta w_L = 0$. We calculate the share of workers worse off using the demeaned wage distribution from the NSS. Even under continuous skills, approximately 72 percent of workers lose from the reform. The share who gain is therefore somewhat larger than under the binary-skill benchmark, but a substantial majority of workers remain worse off.

Among owners, the picture is more nuanced. By headcount, roughly 90 percent of owners lose, because most owners operate small firms. Yet when

Table 6: Percent of workers that gain by number of skill types

	Percent of Workers (%)	
	Binary Skill Types	Continuous Skill Types
2000	10	27
2005	10	28
2012	10	28

Notes: Binary skills assumes all workers earn either w_L or w_H. Continuous skills allows workers to have mixed skill endowments $\phi \cdot w_H + (1-\phi) \cdot w_L$. The indifference threshold ϕ^* solves $\phi^* \cdot \Delta w_H + (1-\phi^*) \cdot \Delta w_L = 0$. The share better off under continuous skills is calculated using the demeaned NSS wage distribution. Cells below 50 percent are shaded red and cells above 50 percent green. India ASI + SUM.

owners are weighted by firm revenue, a majority gain, reflecting the fact that the winners are concentrated among large, high-revenue firms. This distinction highlights an important political economy feature of the reform: it loses under majority rule, but would be supported under a plutocratic voting scheme that weights owners by economic size.

Table 7 translates these distributional changes into *ex ante* consumption-equivalent welfare under CRRA preferences with $\gamma = 2$.[18] In behind-the-veil terms, welfare falls for both workers and owners alike, and for the economy as a whole declines by 8–11 percent. Decomposing this change into level and inequality terms shows why: the positive level effect from higher aggregate output, on the order of 1–4 percent, is more than offset by a negative inequality effect of 9–12 percent. The inequality channel dominates because the gains are concentrated among a relatively small group of high-skilled workers and large-firm owners, while the losses are spread across the much larger population of low-skilled workers and small-firm owners.

[18] See C.1 for details on the consumption-equivalent welfare calculations.

Table 7: Change in consumption equivalent welfare in counterfactual economy

	Percent Change (%)								
	Workers			Owners			All: Workers + Owners		
	Overall	*Level*	*Ineq.*	*Overall*	*Level*	*Ineq.*	*Overall*	*Level*	*Ineq.*
2000	-5.6	0.7	-6.2	-14.5	1.6	-15.9	-8.2	1.1	-9.1
2005	-6.5	2.2	-8.5	-17.9	3.9	-21.0	-9.7	2.9	-12.2
2012	-7.0	2.2	-8.3	-19.8	1.1	-20.7	-10.8	1.3	-11.9

Notes: The table reports the change in behind-the-veil consumption-equivalent welfare from the baseline to the counterfactual economy that removes the size component of wedge dispersion. The overall effect is decomposed into level and inequality components following Flodén (2001). We set the CRRA coefficient to $\gamma = 2$. India ASI + SUM.

5.4 Robustness

Our headline results are qualitatively robust across a range of alternative assumptions about capital supply, preferences, and data construction. We summarize the key findings here; full results are reported in Appendix A.2.2.

Capital supply. Our baseline assumes a small open economy with perfectly elastic capital supply at $R = 0.10$. An alternative is that capital is domestically owned and inelastically supplied, so that the rental rate adjusts in equilibrium. Under this assumption, we allow owners of shrinking firms to lease out their excess capital at the endogenous rate. The results are very similar: the rental rate rises only modestly (from 10% to 10.4%), and the effects on output, wages, and the skill premium are nearly unchanged (Appendix Table 13). This is because capital intensity varies little across the firm size distribution (Appendix Figure 14), so the reform primarily reallocates labor rather than capital. The welfare implications are also essentially unchanged (Table 14).

Substitutability between firms. Higher values of σ amplify both sides of the tradeoff. When firm outputs are closer substitutes, efficient production

requires a bigger reallocation of inputs toward productive firms, generating larger output gains but also a larger increase in the skill premium and larger profit losses for small firms. At $\sigma = 5$, output gains roughly double relative to baseline, but consumption-equivalent welfare losses also nearly double, reaching 20 percent (Appendix Tables 15 and 16).

Risk aversion. The welfare losses vary predictably with the CRRA parameter, but perhaps less than expected. Moving from log preferences ($\gamma = 1$) to $\gamma = 4$ changes the welfare loss by at most 1.5 percentage points in any year (Appendix Table 11). We also verify that the result holds under the assumption that consumption is lognormally distributed when we compute the welfare cost of inequality following Jones and Klenow (2016) (see Appendix Table 12).

Wage trimming. Our baseline trims the tails of average firm wages at the 5th and 95th percentiles within each survey type and year. The welfare results are similar under tighter (6th/94th) or looser (4th/96th) trimming, with consumption-equivalent welfare losses ranging from 8 to 12 percent across specifications and years (Appendix Table 17).

5.5 Implications of removing all misallocation

The results in the preceding sections show that removing size-dependent misallocation would lower welfare for the majority of agents. A natural question is whether this pessimistic conclusion extends to misallocation-reducing reforms more broadly, or whether it is specific to reforms that operate along the size dimension.

Table 8 addresses this question by reporting the effects of three progressively broader counterfactuals: removing all cross-sector dispersion in wedges while preserving within-sector dispersion; removing all within-sector dispersion while preserving cross-sector differences; and removing all wedge dispersion. The results paint a notably different picture from the size-based

reform. When all dispersion in $(1 - \tau_{Y_{si}})$ is removed, output roughly doubles and both skilled and unskilled wages rise substantially. Unlike the size-based reform, low-skilled wages increase by 45–58 percent across the three years.

The comparison across counterfactuals helps clarify the mechanism. Eliminating within-sector dispersion generates large output gains with a modest increase in the skill premium, as it reallocates activity across firms that are relatively similar in skill intensity. By contrast, eliminating cross-sector dispersion produces smaller output gains but a much larger increase in the skill premium because it reallocates activity across sectors with very different skill intensities. Full equalization combines both margins: it raises efficiency enough that both skilled and unskilled wages rise in absolute terms, even though the skill premium increases.

These results suggest that the welfare losses documented above are not an inherent consequence of reducing misallocation, but *are* a feature of reforms that primarily operate along the firm-size dimension. In principle, a sufficiently broad reform that eliminated all wedge dispersion could raise living standards for both skill groups and command broader political support among workers.

Even under full equalization of marginal products across firms, there are important distributional consequences: though low-skilled wages rise, they do so by much less than high-skilled wages and the skill premium increases substantially. Moroever, owners of small firms still lose in absolute terms, and would likely oppose such reforms regardless of the aggregate gains. Second, the full removal of wedge dispersion may not be a feasible policy target. Some of the measured dispersion likely reflects heterogeneous markups or measurement error rather than policy-driven distortions, and the portion that does reflect policy may be difficult to identify and eliminate in practice.

Taken together, these results show that the distributional costs of reform depend critically on which margin of misallocation is being removed. Size-based reforms can lower welfare for the majority even when they raise output, whereas broader equalization of wedges can generate gains large

enough to benefit both skill groups in absolute terms. This comparison helps clarify the central message of the paper. Reducing misallocation need not be welfare-reducing for the majority, but reforms that specifically remove distortions along the firm-size dimension are likely to generate an efficiency-inequality tradeoff because they shift activity toward firms that are more skill-intensive.

Table 8: Change in output, wages, profits from removing all misallocation

	Percent Change (%)						
	ΔY	Δw_H	Δw_L	$\Delta \frac{w_H}{w_L}$	$\Delta \Pi$	$\Delta \pi_{si} \vert \pi_{si} \uparrow$	$\Delta \pi_{si} \vert \pi_{si} \downarrow$
				2000			
no $\tau_{Y_{si}}$ *disp.*	75.9	67.9	43.7	16.8	75.9	242	-73.7
no within sector disp.	60.6	55.8	41.0	10.5	60.6	234	-72.8
no cross sector disp.	8.9	7.8	0.8	6.9	9.0	17.8	-21.4
				2005			
no $\tau_{Y_{si}}$ *disp.*	149	143	58.3	53.2	149	240	-75.3
no within sector disp.	74.4	69.4	51.2	12.0	74.4	256	-73.7
no cross sector disp.	40.3	40.5	4.8	34.1	41.3	34.0	-27.9
				2012			
no $\tau_{Y_{si}}$ *disp.*	96.2	104	45.7	40.0	96.2	218	-71.8
no within sector disp.	53.0	58.6	37.9	15.0	53.0	215	-70.4
no cross sector disp.	25.7	28.2	4.4	22.7	27.1	16.7	-34.4

Notes: The table reports percent changes from the baseline economy under three counterfactuals. "*no* $\tau_{Y_{si}}$ *disp.*" sets all firm-level wedges to the economy-wide revenue-weighted mean. "*no within-sector disp.*" sets each firm's wedge to the revenue-weighted mean wedge in its sector. "*no cross-sector disp.*" preserves within-sector dispersion but recenters each sector's mean wedge to the economy-wide revenue-weighted mean. The table reports changes in aggregate output, skilled and unskilled wages, the skill premium, aggregate profits, and average firm profits conditional on gaining or losing. India ASI + SUM.

6 Conclusion

A large literature has shown that size-dependent policies impose substantial productivity costs on developing economies. In this paper, we asked why such policies have persisted despite their apparent efficiency costs. We argued that the answer could lie in the distributional consequences of reform.

Three empirical facts motivated our analysis. Within Indian manufacturing sectors, revenue productivity is strongly increasing in firm size. Worker wages and household consumption are also steeply increasing in firm size. And larger firms employ substantially more educated workers.

To quantify the tradeoff, we built a general equilibrium model with heterogeneous firms and heterogeneous workers, in which size-dependent revenue wedges simultaneously distorted the allocation of inputs and compressed the skill premium. In this model, the covariance between skill-biased productivity and revenue wedges was crucial for this tradeoff.

We calibrated the model to Indian manufacturing data for 2000, 2005, and 2012, and conducted a counterfactual that removes the component of revenue wedges attributable to firm size. The results revealed an unfavorable tradeoff: removing size-dependent wedges raised aggregate manufacturing output modestly, while increasing the skill premium more steeply. The reform also redistributed profits away from small-firm owners and toward large-firm owners. The upshot was lower average consumption-equivalent welfare for workers and owners alike.

Several important extensions remain for future work. Our model took firm entry, productivity, and worker skill as fixed. Endogenizing these could increase the output gains from size-dependent reforms, especially over time. The impact on wage and consumption inequality is less easy to anticipate. Finally, a fuller analysis could incorporate redistribution policies that could mitigate the inequality effects of size-dependent reforms.

References

Alfaro, Laura and Anusha Chari, "Deregulation, Misallocation, and Size: Evidence from India," *Journal of Law and Economics*, 2014, *57* (4), 897–936.

Atkin, David, Baptiste Bernadac, Dave Donaldson, Tishara Garg, and Federico Huneeus, "The Incidence of Distortions," 2025. Working Paper, Central Bank of Chile No. 1026.

_ , Benjamin Faber, and Marco Gonzalez-Navarro, "Retail Globalization and Household Welfare: Evidence from Mexico," *Journal of Political Economy*, 2018, *126* (1), 1–73.

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

Bau, Natalie and Adrien Matray, "Misallocation and Capital Market Integration: Evidence from India," *Econometrica*, 2023, *91* (1), 67–106.

Bento, Pedro and Diego Restuccia, "Misallocation, Establishment Size, and Productivity," *American Economic Journal: Macroeconomics*, 2017, *9* (3), 267–303.

Bertrand, Marianne, Chang-Tai Hsieh, and Nick Tsivanidis, "Contract Labor and Firm Growth in India," *Review of Economic Studies*, 2021, *88* (3), 1150–1171.

Boar, Corina and Virgiliu Midrigan, "Markups and Inequality," 2024. Working Paper, New York University.

Cai, Wenbiao and Manish Pandey, "Size-Dependent Labor Regulations and Structural Transformation in India," *Economics Letters*, 2013, *119* (1), 67–70.

Dabla-Norris, Era, Laura Jaramillo, Luis Lima, and Alexandre Sollaci, "Size-Dependent Policies, Informality and Misallocation," 2018. IMF Working Paper No. 18/179.

Employees' Provident Fund Organisation, "Compliance Manual," 2025. Official EPFO compliance manual.

_ , "Employees' Provident Funds and Miscellaneous Provisions Act, 1952: Sections 1–16," n.d. Official EPFO document.

_ , "Present Rates of Contribution," n.d. Official EPFO contribution schedule.

Employees' State Insurance Corporation, "Contribution," n.d. Official ESIC contribution page.

Flodén, Martin, "The effectiveness of government debt and transfers as insurance," *Journal of Monetary Economics*, August 2001, *48* (1), 81–108.

Garcia-Santana, Manuel and Josep Pijoan-Mas, "The Reservation Laws in India and the Misallocation of Production Factors," *Journal of Monetary Economics*, 2014, *66*, 193–209.

Garicano, Luis, Claire Lelarge, and John Van Reenen, "Firm Size Distortions and the Productivity Distribution: Evidence from France," *American Economic Review*, 2016, *106* (11), 3439–3479.

Government of India, "The Industrial Disputes Act, 1947 (Act No. 14 of 1947)," March 1947.

_ , "Economic Survey 2005–2006," Technical Report, Ministry of Finance 2006.

Guner, Nezih and Alessandro Ruggieri, "Misallocation and Inequality," 2022. IZA Discussion Paper No. 15174.

_ , Gustavo Ventura, and Yi Xu, "Macroeconomic Implications of Size-Dependent Policies," *Review of Economic Dynamics*, 2008, *11* (4), 721–744.

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

Jones, Charles I. and Peter J. Klenow, "Beyond GDP? Welfare across Countries and Time," *American Economic Review*, 2016, *106* (9), 2426–2457.

Jovanovic, Boyan, "Misallocation and Growth," *American Economic Review*, 2014, *104* (4), 1149–1171.

Katz, Lawrence F. and Kevin M. Murphy, "Changes in Relative Wages, 1963-1987: Supply and Demand Factors," *The Quarterly Journal of Economics*, 1992, *107* (1), 35–78.

Martin, Leslie A., Shanthi Nataraj, and Ann E. Harrison, "In with the Big, Out with the Small: Removing Small-Firm Exemptions from the Value Added Tax," *American Economic Review*, 2017, *107* (4), 1111–1146.

Mehrotra, Santosh and Jajati K. Parida, "India's Employment Crisis: Rising Education Levels and Falling Non-agricultural Job Growth," CSE Working Paper 231, Centre for Sustainable Employment, Azim Premji University October 2019.

National Small Industries Corporation, "Directors' Report, 2023–24," Technical Report, National Small Industries Corporation, Government of India 2024. Official NSIC annual reporting document.

_ , "About Us," n.d. Official NSIC corporate page.

Prasad, Gireesh Chandra, "Future reforms to focus on harnessing potential of the informal sector," *Mint*, March 2024.

PRS Legislative Research, "The Taxation Laws (Amendment) Ordinance, 2019," September 2019. Bill Track page.

PTI, "94.11% registered informal workers' income at Rs 10,000 or below; 74% belong to SC, ST, OBC," *The Economic Times*, May 2022.

Ramaswamy, K. V., "Size Dependent Tax Incentives, Threshold Effects and Horizontal Misallocation: The Case of Indian Manufacturing," Working Paper WP-2016-007, Indira Gandhi Institute of Development Research 2016.

_ , "Size-Dependent Regulations and Firm Size Distribution in Indian Manufacturing," 2021. Working Paper, Indira Gandhi Institute of Development Research.

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

_ and _ , "The Causes and Costs of Misallocation," *Journal of Economic Perspectives*, 2017, *31* (3), 151–174.

Velayudhan, Teevrat, "Misallocation or Misreporting? Evidence from a Value Added Tax," 2018. Working paper.

Appendix

A Additional Tables and Figures

A.1 Additional results for motivating facts

A.1.1 From income inequality to consumption inequality

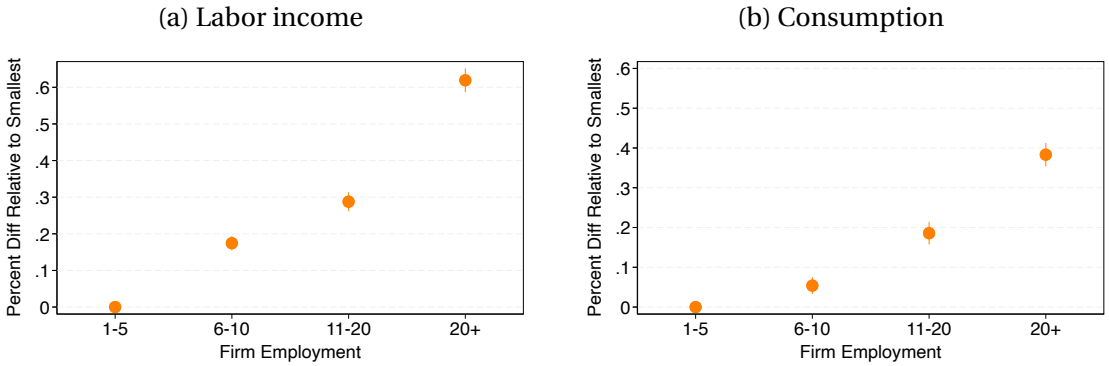
To assess whether the income gradient along firm size documented in Section 2.3 translates into consumption inequality, we estimate regressions of the form

$$\log(y_{si,t}) = \sum_{b \in \mathcal{B}} \beta_b \mathbb{1}\{\text{firm_size_bin}_{si,t} = b\} + \boldsymbol{\theta}\mathbf{X}_{i,t} + \gamma_{s,t} + \varepsilon_{si,t} \tag{19}$$

where i is an individual, t is a year, s is the sector of employment, and $y_{si,t}$ is either labor income or monthly household consumption per capita. Controls $\mathbf{X}_{i,t}$ include indicators for gender, age, and state. Standard errors are clustered at the household-year level.

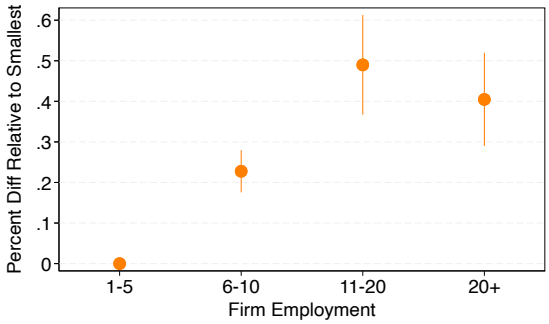
Figure 10 shows the estimated coefficients for workers, with the smallest firm-size bin as the omitted category. Workers at the largest firms have 62 percent higher labor income and 38 percent higher consumption relative to workers at the smallest firms. Figure 11 shows the analogous result for firm owners: consumption is approximately 40 percent higher at the largest firms.

Figure 10: Labor income and consumption by firm size, workers



Notes: Coefficients from equation (19) with the smallest firm size bin as the omitted category. Controls include indicators for gender, age, and state. Includes sector by year fixed effects. Standard errors clustered at the household-year level. Bars show 95% confidence intervals. India NSS.

Figure 11: Consumption by firm size, owners



Notes: Coefficients from equation (19) with consumption as the outcome and the smallest firm size bin as the omitted category. Controls include indicators for gender, age, and state. Includes sector by year fixed effects. Standard errors clustered at the household-year level. Bars show 95% confidence intervals. India NSS.

A.1.2 Elasticity of consumption to labor income

We estimate the elasticity of consumption to labor income using

$$\log(c_{si,t}) = \beta \log(w_{si,t}) + \boldsymbol{\theta}\mathbf{X}_{i,t} + \gamma_{s,t} + \varepsilon_{si,t} \tag{20}$$

where $c_{si,t}$ is monthly household consumption per capita and $w_{si,t}$ is labor income. Controls $\mathbf{X}_{i,t}$ include indicators for gender, age, and state. Table 9 reports the results. The estimated elasticity is approximately 0.27 regardless of whether labor income is measured in cash or in total (including in-kind compensation), confirming that income differences across firm sizes are not fully smoothed through savings or transfers.

Table 9: Elasticity of consumption to labor income

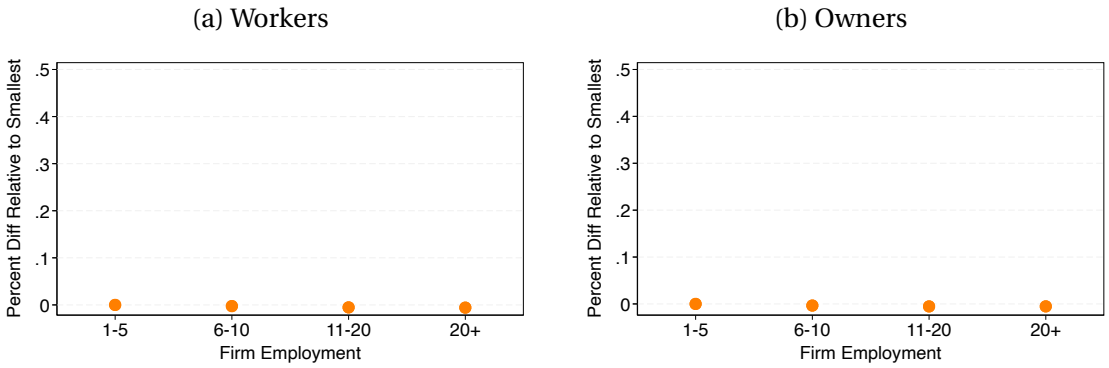
	(1)	(2)
log(Labor Earnings in Cash)	0.262***	
	(0.004)	
log(Labor Earnings in Total)		0.270***
		(0.004)
Adjusted R^2	0.624	0.630
Controls	Yes	Yes
Sector by Year FE	Yes	Yes
Observations	223,175	224,878

Notes: The dependent variable is log monthly household consumption per capita. Controls include indicators for gender, age, and state. Standard errors clustered at the household-year level in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. India NSS.

A.1.3 Leisure by firm size

We estimate equation (19) with leisure hours as the outcome. Leisure is constructed from the NSS time use data as described in Appendix B.1. Figure 12 shows that leisure hours vary by less than 1 percent across firm size bins for both workers and owners, confirming that the consumption differences documented above are not offset by compensating differentials in hours worked.

Figure 12: Leisure by firm size



Notes: Coefficients from equation (19) with log leisure hours as the outcome and the smallest firm size bin as the omitted category. Leisure is constructed as described in Appendix B.1. Controls include indicators for gender, age, and state. Includes sector by year fixed effects. Standard errors clustered at the household-year level. Bars show 95% confidence intervals. India NSS.

A.1.4 Skill sorting versus rent sharing

To assess whether the wage-size gradient documented in Section 2.3 reflects sorting or rent sharing, we augment equation (19) with controls for individual education level. Specifically, we estimate

$$\log(y_{si,t}) = \sum_{b \in \mathcal{B}} \beta_b^c \mathbb{1}\{\text{firm_size_bin}_{si,t} = b\} + \sum_{j \in \mathcal{J}} \beta_j \mathbb{1}\{\text{educ}_{si,t} = j\} + \boldsymbol{\theta} \boldsymbol{X}_{i,t} + \gamma_{s,t} + \varepsilon_{si,t}$$

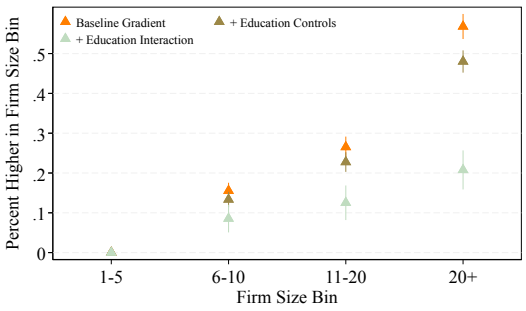
and a version that additionally interacts education with firm-size bins:

$$\log(y_{si,t}) = \sum_{b \in \mathcal{B}} \beta_b^{wi} \mathbb{1}\{\text{firm_size_bin}_{si,t} = b\} + \sum_{j \in \mathcal{J}} \beta_j \mathbb{1}\{\text{educ}_{si,t} = j\} + \sum_{b,j} \beta_{bj} \left(\mathbb{1}\{\text{firm_size_bin}_{si,t} = b\} \times \mathbb{1}\{\text{educ}_{si,t} = j\}\right) + \boldsymbol{\theta} \boldsymbol{X}_{i,t} + \gamma_{s,t} + \varepsilon_{si,t}$$

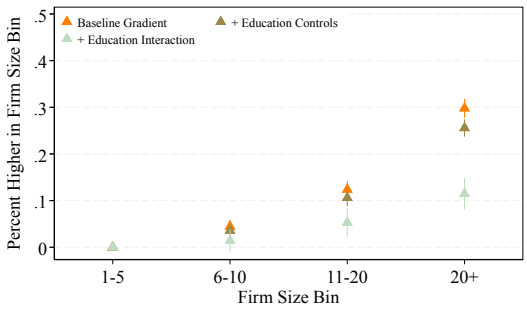
Figure 13 compares the estimated firm-size coefficients across the three specifications (without education controls (β_b), with education controls (β_b^e), and with education interactions (β_b^{wi})) for both labor income and consumption. Including education controls substantially reduces the size gradient, and adding interactions attenuates it further, confirming that the wage-size relationship is primarily driven by the composition of skills rather than rent sharing.

Figure 13: Wage-size gradient with education controls, workers

(a) Labor income



(b) Consumption



Notes: The figure plots the estimated firm-size coefficients from three specifications: without education controls (orange), with education level controls (brown), and with education by firm-size interactions (green). The omitted category is the smallest firm-size bin. Controls include indicators for gender, age, and state. Includes sector by year fixed effects. Standard errors clustered at the household-year level. Bars show 95% confidence intervals. India NSS.

A.2 Counterfactual exercise: additional results

A.2.1 Comparison to skill-biased technical change

Table 10 reports the results of increasing A_{si}^H uniformly across all firms, calibrated to produce the same aggregate output gain as the size-based reform. Unlike the size-based reform, SBTC raises both skilled and unskilled wages. The results are similar whether we use the distorted economy (with the baseline wedges) or an undistorted economy (with wedges set to their revenue-weighted average).

Table 10: Implications of increasing all A_{si}^H

	Percent Change (%)			
	ΔY	Δw_H	Δw_L	$\Delta \frac{w_H}{w_L}$
Distorted economy	1.43	1.67	0.90	0.77
Undistorted economy	1.40	1.64	0.87	0.76

Notes: The table shows the change in aggregate output, wages and the skill premium for an increase in $A_{si}^H \ \forall i$. In "Distorted economy" we run the counterfactual where we use $\tau_{Y_{si}}$ that we backed out from the data. In "Undistorted economy", we run the counterfactual where we impose that all firms have $\tau_{Y_{si}} = \overline{\tau}_Y$ where $\overline{\tau}_Y$ is the revenue-weighted average $\tau_{Y_{si}}$ in the data. $\overline{\tau}_Y = 0.02$ in 2012. India ASI + SUM 2012.

A.2.2 Counterfactual exercise: robustness

Table 11: Change in consumption equivalent welfare varying risk aversion

	Percent Change (%)		
	$\gamma = 1$	$\gamma = 2$	$\gamma = 4$
2000	-8.0	-8.2	-7.2
2005	-9.9	-9.7	-8.2
2012	-10.7	-10.8	-9.4

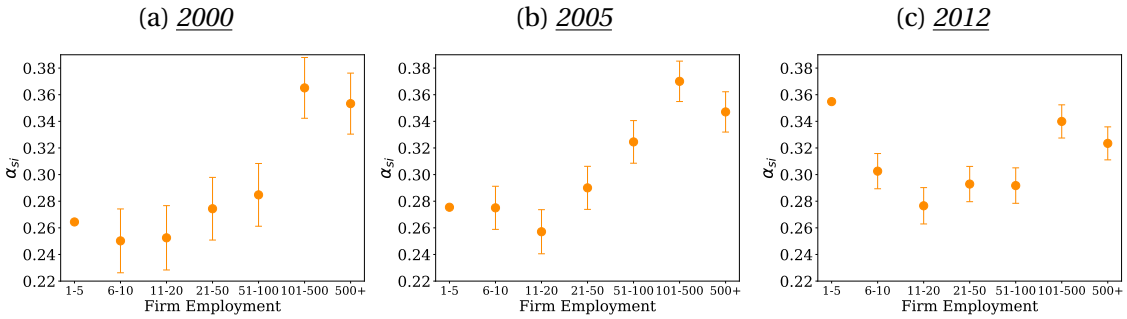
Notes: The table shows the overall percent change in behind-the-veil consumption equivalent welfare for all agents in the economy from the baseline economy to the counterfactual where we remove dispersion in the revenue wedges along the size dimension (see Section 5.1). We calculate the change in welfare for log preferences with $\gamma = 1$, $\gamma = 2$ (our headline results), and $\gamma = 4$. India ASI + SUM.

Table 12: Change in consumption equivalent welfare assuming lognormally distributed consumption

	Standard deviation of log consumption		Loss from consumption inequality (%)	
	Baseline	Counterfactual	$\gamma = 2$	$\gamma = 4$
2000	0.820	0.841	-3.36	-6.60
2005	0.812	0.829	-2.76	-5.44
2012	0.776	0.807	-4.70	-9.17

Notes: The first two columns of the table show the standard deviation of log consumption in the baseline economy and the counterfactual economy where we remove dispersion in the revenue wedges along the size dimension (see Section 5.1). The last two columns show the loss from consumption inequality calculated as in Jones and Klenow (2016) assuming that consumption is lognormally distributed. India ASI + SUM.

Figure 14: Capital intensity by firm size



Notes: The figure shows the capital intensity, α_{si}, by the firm size, measured by the firm employment. The regressions are revenue-weighted and include sector-fixed effects. The bars are the 95% confidence intervals. We use data from the India ASI + SUM for 2000 (14a), 2005 (14b), and 2012 (14c).

Table 13: Change in aggregate output, capital rental rate, wages, and skill premium with elastic vs inelastic capital supply

Percent Change (%)									
ΔY		R		Δw_H		Δw_L		$\Delta \frac{w_H}{w_L}$	
Inelastic	Elastic	Inelastic	Elastic	Inelastic	Elastic	Inelastic	Elastic	Inelastic	Elastic
				2000					
0.7	1.9	0.102	0.100	3.1	3.9	-6.4	-5.7	10.2	10.2
				2005					
1.6	4.3	0.104	0.100	4.9	6.8	-8.3	-6.6	14.4	14.3
				2012					
1.2	1.4	0.100	0.100	6.7	6.8	-7.3	-7.1	15.0	15.0

Notes: The table shows the percent change in aggregate output, rental rate of capital, wages, and skill premium from the baseline economy to the counterfactual economy where we remove dispersion in the revenue wedges along the size dimension (see Section 5.1). We show the change for the economy with elastic capital supply with $R = 0.10$ (headline results), and the economy where we assume inelastic capital supply. India ASI + SUM.

Table 14: Change in consumption equivalent welfare with elastic vs inelastic capital supply

	Percent Change (%)								
	Workers			Owners			All: Workers + Owners		
	Overall	*Level*	*Ineq.*	*Overall*	*Level*	*Ineq.*	*Overall*	*Level*	*Ineq.*
Inelastic	-7.2	1.3	-8.3	-19.9	1.0	-20.7	-10.9	1.1	-11.9
Elastic	-7.0	1.4	-8.3	-19.8	1.1	-20.7	-10.8	1.3	-11.9

Notes: CRRA coefficient $\gamma = 2$. India ASI + SUM 2012.

Table 15: Change in aggregate output, skill premium, and profits by substitutability between firms

	$\sigma = 3$	$\sigma = 4$	$\sigma = 5$
		2000	
ΔY	1.9	2.5	3.3
$\Delta \frac{w_H}{w_L}$	10.2	12.4	14.1
$\Delta \pi_{si} \vert \pi_{si} \uparrow$	2.4	3.4	3.4
$\Delta \pi_{si} \vert \pi_{si} \downarrow$	-13.7	-16.9	-16.9
		2005	
ΔY	4.3	5.7	7.2
$\Delta \frac{w_H}{w_L}$	14.3	18.7	22.9
$\Delta \pi_{si} \vert \pi_{si} \uparrow$	4.6	7.0	7.0
$\Delta \pi_{si} \vert \pi_{si} \downarrow$	-15.9	-21.2	-21.2
		2012	
ΔY	1.4	1.9	2.4
$\Delta \frac{w_H}{w_L}$	15.0	19.1	22.7
$\Delta \pi_{si} \vert \pi_{si} \uparrow$	3.2	4.6	4.6
$\Delta \pi_{si} \vert \pi_{si} \downarrow$	-18.1	-23.1	-23.1

Notes: The table shows the percent change in aggregate output, skill premium, and the average change in firm profits conditional on a gain or loss, from the baseline economy to the counterfactual economy where we remove dispersion in the revenue wedges along the size dimension (see Section 5.1). We show for varying values of the substitutability between firms, where $\sigma = 3$ is our headline results. India ASI + SUM.

Table 16: Change in consumption equivalent welfare varying substitutability between firms

	Percent Change (%)		
	Overall	*Level*	*Inequality*
		2000	
$\sigma = 3$	-8.2	1.1	-9.1
$\sigma = 4$	-11.2	1.2	-12.2
$\sigma = 5$	-14.5	1.3	-15.6
		2005	
$\sigma = 3$	-8.2	1.1	-9.1
$\sigma = 4$	-11.2	1.2	-12.2
$\sigma = 5$	-14.5	1.3	-15.6
		2012	
$\sigma = 3$	-10.8	1.3	-11.9
$\sigma = 4$	-15.4	1.7	-16.8
$\sigma = 5$	-20.1	2.1	-21.7

Notes: The table shows the percent change in behind-the-veil consumption equivalent welfare for workers and owners from the baseline economy to the counterfactual economy where we remove dispersion in the revenue wedges along the size dimension (see Section 5.1). We decompose the overall change into the change coming from the level effect and the inequality effect following Flodén (2001) (details in Appendix C.1)). We show for varying values of the substitutability between firms, where $\sigma = 3$ is our headline results. We set the CRRA coefficient $\gamma = 2$. India ASI + SUM.

Table 17: Change in consumption equivalent welfare varying wage trimming

	Percent Change (%)		
	4th/96th percentile	5th/95th percentile	6th/94th percentile
2000	-8.0	-8.2	-9.0
2005	-11.1	-9.7	-9.5
2012	-11.9	-10.8	-10.0

Notes: The table shows the overall percent change in behind-the-veil consumption equivalent welfare for all agents in the economy from the baseline economy to the counterfactual where we remove dispersion in the revenue wedges along the size dimension (see Section 5.1). We calculate the change in welfare varying the trimming performed on average wages per worker within each survey-year (see Appendix B.1). India ASI + SUM.

B Measurement Appendix

B.1 Data Cleaning

This appendix provides details on the construction and cleaning of the firm and household datasets described in Section 2.1.

B.1.1 Firm data: ASI and SUM

Sample construction. We use the Annual Survey of Industries (ASI) and the Survey of Unorganized Manufacturing (SUM) , restricting both surveys to manufacturing sectors using the NIC code ranges native to each round and pool them using the official sampling weights provided with each survey. State identifiers are harmonized across rounds to account for the creation of new Indian states during our sample period.

Sector classification. Sectors are defined at the four-digit level of the National Industrial Classification, yielding between 100 and 150 manufacturing sectors depending on the year. The NIC classification changed in 1998, 2004, and 2008; we use the classification native to each survey round since our analysis is conducted within year. We drop sectors with fewer than 10 firm-level observations in any year.

Key variables. From both surveys we observe firm revenue (gross value added), the book value of fixed capital at the end of the fiscal year, the number of workers, and total labor compensation. Revenue in the SUM is measured as monthly gross value added, which we annualize. We set the return to capital $R = 0.10$, following Hsieh and Klenow (2009).

Capital. We drop observations with missing or zero reported capital. We do not impute capital for these observations, as the missing values are concentrated among very small informal firms for which book value is unlikely to be meaningful.

Labor and wages. The SUM reports paid and unpaid (family) workers separately. For firms with unpaid family workers, we assume they work half-time and earn half the wage of paid workers in the same firm. Specifically, we adjust the wage bill upward by $0.5 \times (\text{unpaid workers}) \times (\text{wage bill}/\text{paid workers})$ and adjust the labor input to paid workers $+ \; 0.5 \times$ unpaid workers. Where the firm-level wage bill is missing or zero after this adjustment, we impute using the weighted average wage per worker among firms in the same year, sector, survey type (ASI or SUM), and size category. Firms in the SUM that report zero total workers are assigned one unit of labor, representing the owner-operator contributing labor to production.

Trimming. We apply two rounds of trimming. First, we trim the distribution of average wages per worker at the 5th and 95th percentiles within each survey type (ASI or SUM) and year. Second, after computing the revenue wedges $\tau_{Y_{si}}$ (described in Section 4), we trim at the 5th and 95th percentiles of $\tau_{Y_{si}}$ within each sector and year. Both trimming steps use the official sampling weights. We also drop one observation with revenue exceeding 100 billion rupees. We report robustness to alternative trimming thresholds (4th/96th and 6th/94th percentiles) in Section 5.4.

Firm size bins. The firm size in the NSS is reported for the four bins: 1-5, 6-10, 11-20, and 20+ employees. For the counterfactual exercise using the firm data, which requires finer variation in the upper tail, we use seven bins: 1-5, 6-10, 11-20, 21-50, 51-100, 101-500, and 500+ employees.

B.1.2 Household data: NSS Employment and Unemployment

Sample construction. We use the Employment and Unemployment rounds of the National Sample Survey: round 55 (1999-2000), round 61 (2004-2005), and round 68 (2011-2012). We merge individual-level demographic, activity, time use, and wage data with household-level consumption expenditure within each round. We restrict the sample to individuals employed in manufacturing, defined using the same NIC code ranges as the firm data, and to those for whom employer size is reported (dropping observations coded as "unknown").

Worker classification. We classify individuals as owners if their usual principal activity status is self-employed (status 11) or an employer (status 12), and as workers otherwise. Workers include regular salaried employees (status 31), casual wage laborers (status 41 and 51), and unpaid family workers (status 21).

Wages and imputation. Wage and salary earnings are recorded on a weekly basis in the NSS; we annualize by multiplying by 52. For paid workers with missing or zero reported wages, we impute using the weighted average wage among paid workers in the same firm-size bin, sector, and year. Unpaid family workers are assigned half this average wage. To account for the imputation, we construct an effective weight equal to the sampling weight for paid workers and half the sampling weight for unpaid workers, which we use in all NSS regressions. We trim wages at the 1st and 99th percentiles within each year.

Education. Education is recorded in the NSS using detailed categories for general education attainment. We classify individuals as high-skilled if they have completed secondary school or above (general education codes 08, 10, 11, 12, and 13), and low-skilled otherwise.

Consumption. Household consumption is measured as monthly household consumer expenditure (MHCE), which we divide by household size to obtain per capita consumption and annualize by multiplying by twelve.

Time use and leisure. The NSS records the number of half-days and full days worked in the reference week for each activity. We calculate weekly hours worked as $(8 \times \text{full days}) + (2.5 \times \text{half days})$ and construct monthly leisure hours as $\frac{5{,}840}{12} - \text{weekly hours} \times \frac{52}{12}$.

Industry concordance. The NSS records industry of employment using the NIC classification corresponding to each round. For rounds using NIC2008, we concord to NIC2004 and then to NIC98 using official concordance tables to obtain a consistent sector variable for matching with the firm data. The manufacturing restriction is applied using the NIC2008 code range (10101-32909) for round 68, and the equivalent ranges for earlier rounds.

B.2 Measurement details

This appendix provides supplementary details on the measurement procedure described in Section 4.1.

The firm's profit maximization problem yields first-order conditions for capital and the labor composite:

$$(1-\tau_{Y_{si}})\cdot\frac{\sigma-1}{\sigma}\cdot\alpha_{si}\cdot\frac{P_{si}Y_{si}}{K_{si}}=R \tag{21}$$

$$(1-\tau_{Y_{si}})\cdot\frac{\sigma-1}{\sigma}\cdot(1-\alpha_{si})\cdot\frac{P_{si}Y_{si}}{w_{si}Z_{si}}=1. \tag{22}$$

Dividing (21) by (22) eliminates $(1-\tau_{Y_{si}})$ and yields

$$\frac{\alpha_{si}}{1-\alpha_{si}}=\frac{RK_{si}}{w_{si}Z_{si}},$$

which gives equation (16) in the main text. The firm-specific price index for the labor composite is derived from the cost-minimization dual over H_{si} and L_{si}:

$$w_{si}=\left[\left(\frac{w_H}{A^H_{si}}\right)^{1-\eta}+w_L^{1-\eta}\right]^{\frac{1}{1-\eta}},$$

so that $w_{si}Z_{si}$ equals total worker compensation (the wage bill). Given α_{si}, the revenue wedge follows directly from either first-order condition, yielding equation (17).

B.3 Additional descriptive statistics

Table 18 reports the implied skill premium w_H/w_L by year under our baseline and alternative trimming thresholds. The premium ranges from roughly 13 to

19 depending on the year and trimming, reflecting the large gap between the highest- and lowest-paying firms in Indian manufacturing. The premium is sensitive to the trimming threshold because w_H and w_L are identified from the extremes of the wage distribution: tighter trimming brings the tails closer together. Our counterfactual results, which depend on *changes* in the skill premium rather than its level, are robust to these alternatives (Section 5.4).

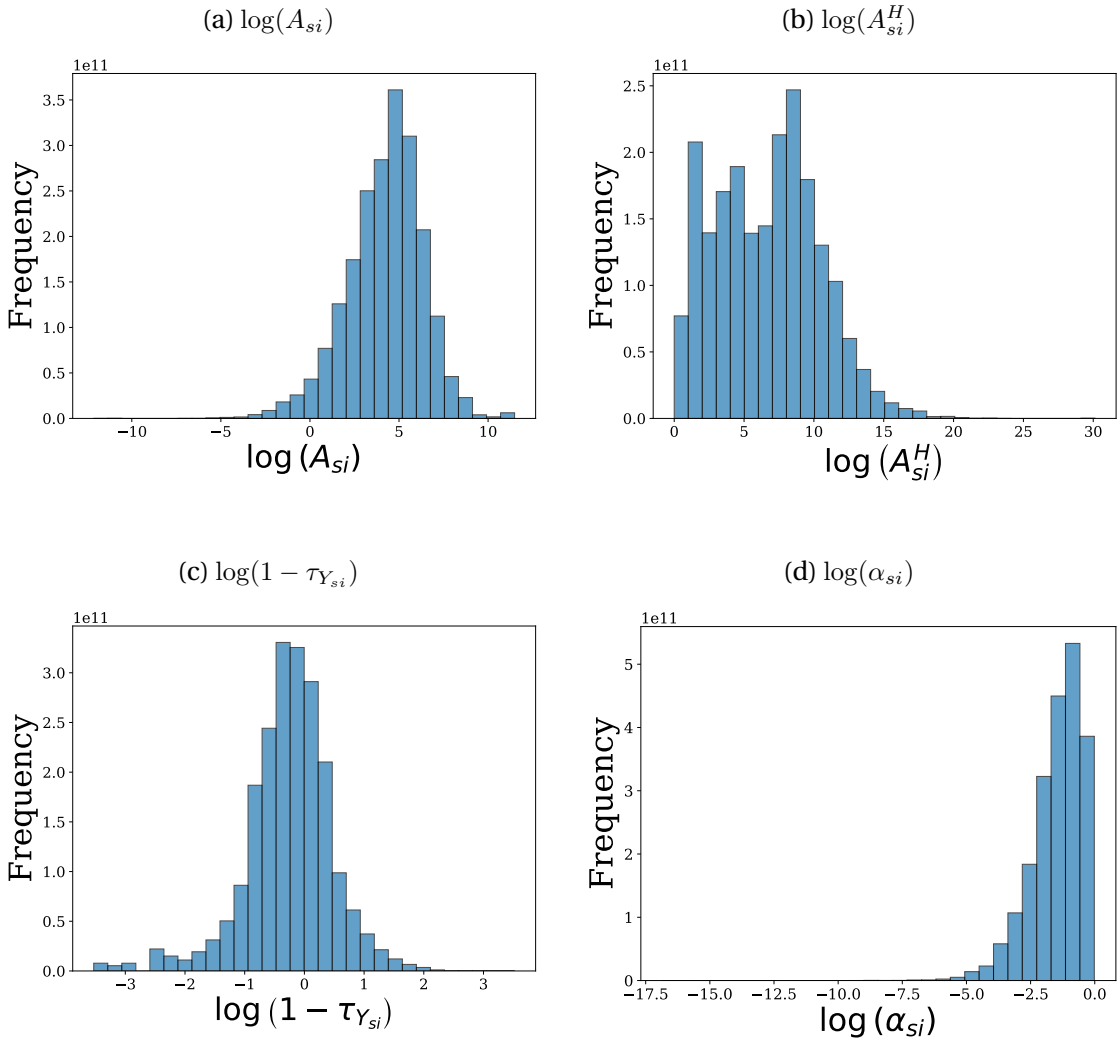
Table 18: Skill premium by year and trimming threshold

	4th/96th	5th/95th	6th/94th
		(headline)	
2000	18.97	17.25	15.84
2005	18.77	17.23	14.51
2012	15.74	13.95	12.75

Notes: The skill premium is w_H/w_L, where w_H and w_L are the maximum and minimum average wage per worker in the firm data after trimming at the indicated percentiles within survey type and year. India ASI + SUM.

Figure 15 shows the revenue-weighted distributions of the four firm-level objects recovered by our measurement procedure: factor-neutral productivity A_{si}, skill-biased productivity A^H_{si}, the keep rate $(1-\tau_{Y_{si}})$, and the capital share α_{si}. All four exhibit substantial dispersion across firms within the manufacturing sector.

Figure 15: Distributions of productivities and wedges

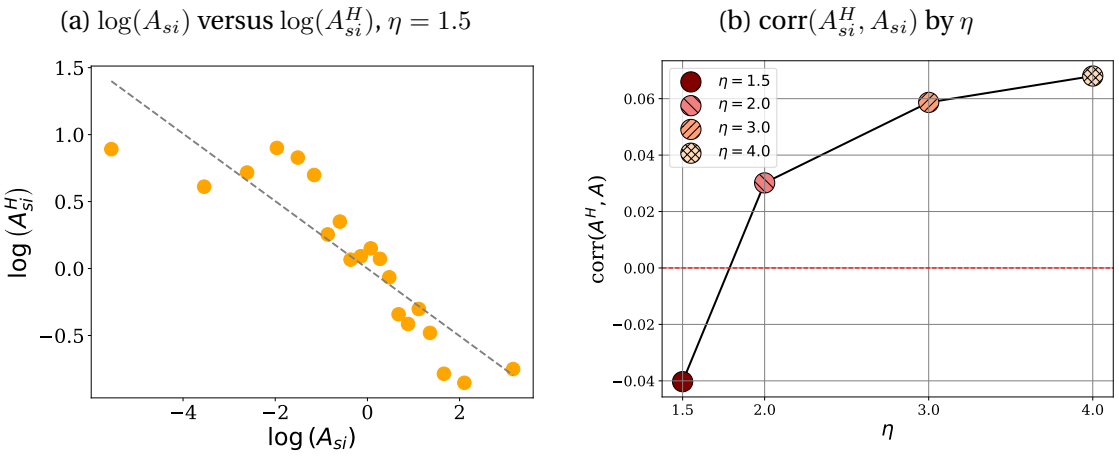


Notes: Revenue-weighted distributions of firm-level productivities and wedges inferred from the measurement procedure in Section 4.1. India ASI + SUM 2012.

Figure 16 shows the relationship between factor-neutral and skill-biased productivity. The two are negatively correlated: firms with high factor-neutral productivity tend to have lower skill-biased productivity, and vice versa. The right panel shows that this negative correlation attenuates and eventually turns slightly positive as η increases, since higher substitutability between skill

types compresses the inferred distribution of A_{si}^H.

Figure 16: Relationship between A_{si} and A_{si}^H

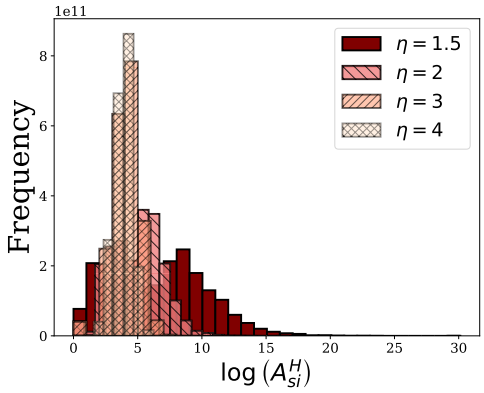


Notes: Left panel: each point represents a revenue-weighted bin. Both axes are in logs and measured relative to the sector mean. Right panel: revenue-weighted correlation between $\log(A_{si}^H)$ and $\log(A_{si})$ for $\eta \in \{1.5, 2, 3, 4\}$. India ASI + SUM 2012.

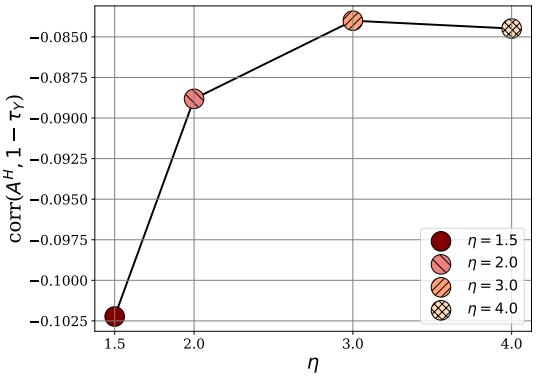
Figure 17 shows that the negative correlation between skill-biased productivity and revenue wedges documented in Figure 6 is robust to alternative values of η. The left panel shows that higher η compresses the distribution of A_{si}^H, since greater substitutability between skill types means less variation in skill-biased productivity is needed to rationalize the observed variation in skill ratios across firms. The right panel confirms that the correlation between A_{si}^H and $(1 - \tau_{Y_{si}})$ remains negative across all values of η, though it attenuates somewhat as η increases.

Figure 17: A_{si}^H versus $(1-\tau_{Y_{si}})$ by η

(a) Distribution of $\log(A_{si}^H)$ by η



(b) corr$(A_{si}^H, 1-\tau_{Y_{si}})$ by η



Notes: Left panel: revenue-weighted distribution of $\log(A_{si}^H)$ for $\eta \in \{1.5, 2, 3, 4\}$. Right panel: revenue-weighted correlation between $\log(A_{si}^H)$ and $\log(1-\tau_{Y_{si}})$ for each value of η. India ASI + SUM 2012.

C Derivations

C.1 Deriving and decomposing behind-the-veil consumption equivalent welfare

Let $x_h = \frac{\overline{H}}{\overline{H}+\overline{L}}$. We specify our social welfare function as

$$SWF = x_h u(c_H) + (1-x_h)u(c_L)$$

where $u(c)$ is CRRA preferences with a coefficient of risk aversion γ. We want to ask "how much change in consumption for every individual in the baseline ("BL") makes them indifferent to living in the counterfactual ("CF") world?". This is given by

$$\text{Overall } \Delta \text{ cons. equiv. welfare} = \left(\frac{x_h u(c_H^{CF}) + (1-x_h)u(c_L^{CF})}{x_h u(c_H^{BL}) + (1-x_h)u(c_L^{BL})}\right)^{\frac{1}{1-\gamma}} - 1 \equiv \Delta \text{ Overall}$$

In our setting, the overall consumption equivalent welfare can change due to (1) level effect — i.e change in average consumption levels, and (2) egalitarian effect — i.e. change in inequality across agents.

We can decompose the overall change in consumption equivalent welfare additively into the level effect and the equality effect.[19]

First define the following:

$$\begin{aligned} \text{average consumption: } \bar{c} \quad &= x_h c_H + (1 - x_h) c_L \\ \text{certainty equivalent consumption: } c_{ce} \quad &= ((1-\gamma)(x_h u(c_H) + (1-x_h)u(c_L)))^{\frac{1}{1-\gamma}} \end{aligned}$$

The change in consumption equivalent welfare that comes from the level change in average consumption is given by

$$\Delta \text{ cons. equiv. welfare from } \Delta \text{ avg. consumption} = \frac{\bar{c}^{CF}}{\bar{c}^{BL}} - 1 \equiv \Delta \text{ Level}$$

We define the cost of inequality as the fraction of average consumption an individual would give up to avoid inequality. In the counterfactual world, this is given by

$$\kappa^{CF} = 1 - \left(\frac{u(c_{ce}^{CF})}{u(\bar{c}^{CF})} \right)^{\frac{1}{1-\gamma}}$$

We define the analogous object in the baseline as

$$\kappa^{BL} = 1 - \left(\frac{u(c_{ce}^{BL})}{u(\bar{c}^{BL})} \right)^{\frac{1}{1-\gamma}}$$

Therefore, the welfare change from the change in equality is given by

$$\Delta \text{ cons. equiv. welfare from } \Delta \text{ inequality} = \frac{1-\kappa^{CF}}{1-\kappa^{BL}} - 1 \equiv \Delta \text{ Inequality}$$

[19] We draw on Flodén (2001).

With $u(.)$ homogeneous, we have that

$$1 + \Delta\text{Overall} = (1 + \Delta\text{Level})(1 + \Delta\text{Inequality}) \approx 1 + \Delta\text{Level} + \Delta\text{Inequality}$$

For the consumption equivalent welfare for owners, we follow a similar procedure. The different owners are put into different groups based on their initial firm size bin and we consider inequality across the outcomes for the representative owner in each size bin.