NBER WORKING PAPER SERIES

THE MICRO AND MACRO PRODUCTIVITY OF NATIONS

Stephen Ayerst
Duc M. Nguyen
Diego Restuccia

Working Paper 32750
http://www.nber.org/papers/w32750

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2024, Revised October 2024

We are grateful for useful comments to Paco Buera, Marcela Eslava, Hugo Hopenhayn, Venky Venkateswaran, and audiences at Toronto, LACEA-LAMES 2023 in Bogota, Hong Kong Macroeconomics Workshop, UNC Chapel Hill, Econometric Society NASM 2024 in Nashville, and ThReD 2024 in Namur. All remaining errors are our own. Restuccia gratefully acknowledges the support from the Canada Research Chairs program and the Bank of Canada Fellowship program. The views expressed herein are those of the authors and should not be attributed to the Bank of Canada or its Governing Council, nor the IMF, its Executive Board, its management, or the National Bureau of Economic Research.

The Micro and Macro Productivity of Nations
Stephen Ayerst, Duc M. Nguyen, and Diego Restuccia
NBER Working Paper No. 32750
July 2024, Revised October 2024
JEL No. O11, O14, O4

ABSTRACT

We examine aggregate productivity differences across nations using cross-country firm-level data and a quantitative model of production heterogeneity with distortions featuring operation decisions (selection) and productivity-enhancing investments (technology). Empirically, less developed countries feature higher distortions and larger dispersion in firm-level productivity, mostly resulting from the higher prevalence of unproductive firms. Quantitatively, measured cross-country differences in the elasticity of distortions with respect to firm productivity generate the bulk of empirical patterns and over two-thirds of cross-country labor productivity differences. Both selection and technology channels are important. Variation in static misallocation also plays an important role, albeit smaller.

Stephen Ayerst
International Monetary Fund
stephen.b.ayerst@gmail.com

Duc M. Nguyen
University of Toronto
Department of Economics
ducm.nguyen@mail.utoronto.ca

Diego Restuccia
University of Toronto
Department of Economics
and NBER
diego.restuccia@utoronto.ca

1 Introduction

Large disparities in aggregate productivity across countries are at the core of international differences in GDP per capita (Klenow and Rodriguez-Clare, 1997; Prescott, 1998; Hall and Jones, 1999). Cross-country productivity differences are linked to distortions in resource allocations across firms within sectors (Restuccia and Rogerson, 2008; Hsieh and Klenow, 2009; Bartelsman et al., 2013). Whereas the misallocation literature emphasizes the aggregate productivity gains from factor reallocation across a given set of producers, producer-level data reveal substantial differences in the firm-level productivity distribution across countries (Hsieh and Klenow, 2009; Gal, 2013; Andrews et al., 2015). In this paper, we take a systematic approach to linking observed firm-level productivity differences with institutions and policies that misallocate resources across firms. Our approach is motivated by empirical evidence from policy reforms that find substantial improvements in firm selection and technology upgrading from reductions in misallocation.[1] To this end, we examine cross-country firm-level data through the lens of a quantitative model of production heterogeneity with distortions featuring operation decisions (selection) and productivity-enhancing investments (technology) that determine the firm-level productivity distribution. We find that differences in selection and technology driven by observed cross-country distortions explain the majority of the micro and macro productivity differences across nations.

We construct a firm-level financial dataset covering many countries using Orbis data, collected and standardized by Bureau Van Dijk, for manufacturing firms over the period 2000-2019. Our final dataset contains 37 countries with an average of around 300,000 firm-year observations. The countries in our sample span a wide range of the world income distribution, with India and Vietnam among the lower-income countries and France and Germany among the higher-income countries. We use these data to examine cross-country differences in firm-level total factor productivity (TFP) and firm-level wedges, a model-based

[1] Some examples of policy reforms with effects on misallocation, technology, and selection include Pavcnik (2002) on trade reform in Chile, Bustos (2011) on technology upgrading by Argentinean firms due to a regional free-trade agreement, and Khandelwal et al. (2013) on export quotas reform in China.

measure of idiosyncratic distortions faced by the firm. We document the following facts: (1) firm-level productivity is more dispersed in less developed countries, (2) productivity tends to fan out—captured by larger gaps between high productivity firms and firms at lower points of the productivity distribution in low income countries, (3) dispersion of firm-level wedges is higher in less developed countries, and (4) firm-level wedges are more highly correlated with firm-level productivity in less developed countries. An important insight of our paper is recognizing that the empirical firm-level productivity and wedge distributions may reflect firm decisions on operation (selection) and technology, an analysis of which requires more structure.

To address the quantitative role of distortions on factor misallocation and differences in the firm-level TFP distribution across countries, we develop a model of production heterogeneity with distortions featuring operation decisions, building on Hopenhayn (1992) and Restuccia and Rogerson (2008). We extend the framework to allow for productivity-enhancing investments, broadly capturing costly activities that firms undertake to improve productivity. Production of a homogeneous good takes place in firms using a decreasing returns technology with labor input. Firms are subject to a fixed operating cost, idiosyncratic distortions, and a transitory productivity shock that becomes known after production decisions are made, which could capture, for example, measurement error in the data as it creates a gap between firm output and input decisions. New firms enter by paying a fixed entry cost, after which they draw an idiosyncratic investment ability and distortion. The productivity of operating firms is determined through costly investment in which higher investment-ability firms face lower investment costs and distortions affect incentives to invest. Importantly, new firms may choose not to invest or operate after drawing their idiosyncratic investment ability and distortion if their expected value is less than the fixed operating cost. This leads to selection in which less productive and more distorted firms exit the economy and are not observed in the firm distribution.

We characterize the four cross-country facts in the model. Firm selection allows the model to replicate the observed higher dispersion and fanning out of firm productivity observed in less developed countries since correlated distortions compress the productivity distribution through the technology channel. Nevertheless, whether the model can replicate the data patterns is a quantitative question that requires us to calibrate the model. The model also highlights how measured distortions may be impacted by firm technology choices, selection, and measurement error, potentially leading to larger or smaller misallocation differences across countries than implied by the data moments alone.

Distortions are parameterized by a systematic component related firm-level productivity, governed by an elasticity of distortions parameter ρ, and a random component drawn from a log normal distribution, which provides an excellent fit of measured distortions and factor allocations within and across countries in the data. We calibrate a distorted benchmark economy to observations for France. The standard deviation of the idiosyncratic distortions, the investment ability, and the transitory component of productivity distributions, the elasticity of distortions ρ, and the fixed operating costs are jointly chosen to match data moments on measured distortions, dispersion in firm-level TFP and employment, and average firm size. To determine a plausible range of variation in distortions and other parameters, we repeat the calibration of the same five parameters for each other country in the dataset. Despite its simplicity, the model replicates for example the firm-level productivity distribution.

We examine whether the model can replicate the cross-country evidence by varying only distortions from the benchmark France values to match distortions of countries across the income distribution. The bulk of the four empirical facts are accounted for by differences in the elasticity ρ of distortions, with variations in the standard deviation of distortions helping improve the model fit. We also show that the transitory productivity shocks, which could capture mismeasurement, can help to explain the full range of outcomes in the data. Notably, high dispersion in transitory productivity can help explain why misallocation appears severe in some high-income countries. We also use the calibrated model to decompose the relationship between the fundamental parameters governing distortions and the observed data moments. After correcting for the impact of technology decisions, selection of operating firms, and noise generated by transitory productivity shocks, distortion parameters tend to be smaller than what would otherwise be implied from the observed data moments. For example, the variance of model distortions tends to be smaller than the measured variance of wedges. Moreover, the bias tends to be stronger in higher-income countries implying that the income gradient of distortions is steeper than implied by cross-country data moments. The upshot being that misallocation tends to be lower than implied by data moments for all countries with a larger gap in misallocation between low- and high-income countries than implied by the data.

The model explains more than half of the cross-country labor productivity gap. Increasing distortions in the benchmark France economy to the fitted values for the lowest income economy would reduce aggregate labor productivity by around 71%, equivalent to around 56% of the observed productivity gap between France and Vietnam. In other words, a policy reform that reduces fitted distortions in the lowest income economy to the level of France would increase aggregate productivity by around 240%. Decomposing the aggregate productivity loss indicates that more than half of the loss arises from the change in the firm-level productivity distribution while the remainder is from static misallocation.

Increasing the elasticity of distortions causes allocative efficiency to decline due to increasing static misallocation and productivity dispersion due to less selection, which amplifies potential gains from reallocation. About three fourths of the decline in allocative efficiency is explained by worsening misallocation for a fixed set of operating firms and about one-fourth by a dynamic channel in which the set of operating firms and chosen technologies adjusts. These results highlight the important interaction between the firm-level productivity distribution and allocative efficiency, an underappreciated cost of misallocation in individual-country survey data. In a separate decomposition, we find that the contribution of the productivity loss from the shift in the productivity distribution is roughly equally divided between selection (the change in operating producers) and technology (firm-level investments).

Our paper makes three main contributions. First, we provide a systematic assessment of the joint firm-level productivity and distortion distribution using cross-country producer-level data, relating our paper to the empirical literature on production heterogeneity and misallocation (Guner et al., 2008; Hsieh and Klenow, 2009; Bartelsman et al., 2013). A novel finding is that the lower productivity dispersion in higher income countries is driven, in part, by a compression of the bottom-end of the productivity distribution. While there are still large gaps, low-productivity firms are much closer to high-productivity firms within high-income countries than within low-income countries. This also connects us with recent efforts to exploit Orbis and other datasets to examine firm-level evidence in many countries (such as by Andrews et al., 2015; Poschke, 2018; Kalemli-Özcan et al., 2024; Alviarez et al., 2023).

Second, we show that the cross-country facts can be reconciled by incorporating producer selection and technology investment into the standard model of misallocation. In this regard, our analysis integrates the quantitative (Restuccia and Rogerson, 2008) and empirical (Hsieh and Klenow, 2009) literatures on misallocation and the associated literature on producer dynamics, technology adoption, and aggregate productivity (Parente and Prescott, 1994; Bhattacharya et al., 2013; Hsieh and Klenow, 2014; Bento and Restuccia, 2017; Comin and Mestieri, 2018; Ayerst, 2025; Buera et al., 2023). The main insights are that selection and technology can theoretically and quantitatively replicate the observed cross-country facts and lead to more substantial aggregate productivity costs of distortions, an effect missing in these literatures. Both selection and technology channels have been emphasized for individual countries, but the cross-country importance of these channels are typically not assessed due to lack of comparable cross-country data.

Third, we derive and measure bias in estimates of the distortions that underlie misallocation. The bias results from selection, technology choices, and transitory productivity shocks (such as mismeasurement). We find that selection is the most important source of bias in higher income countries rendering a flatter relationship between the measured elasticity and income per capita across countries than that of the policy parameter ρ while the standard deviation of wedges tends also to be larger than the underlying policy parameter. The transitory productivity shock bias, while quantitatively significant for some countries, represents a relatively small component of the measured elasticity in less developed countries, indicating that measurement error is not a major factor driving cross-country outcomes.

In a closely related paper, Fattal-Jaef (2022) uses Orbis, and other data sources, to examine the aggregate productivity costs of entry barriers using a model that also features idiosyncratic distortions, endogenous firm exit decisions, and technology investment. While both models feature selection, our focus is on the selection of (ex-ante) more productive firms into markets, in contrast to the exit of ex-post less productive firms examined by Fattal-Jaef (2022). We show ex-ante selection is theoretically and quantitatively important in accounting for cross-country differences in firm-level productivity distributions and impacts the measurement of distortions. The selection channel also relates us to recent papers examining the impact of misallocation on selection (Yang, 2021; Majerovitz, 2023).

The paper proceeds as follows. In the next section, we describe the data and present the main empirical observations from the cross-country data. Section 3 describes the model and characterizes the qualitative role of distortions on firm-level productivity. In Section 4, we calibrate a distorted benchmark economy to observations for France and quantify the effect of distortions on the distribution of firm-level productivity and other outcomes across countries. We conclude in Section 5.

2 Empirical Facts

We describe the cross-country data and details of constructed variables. We then present our main facts on firm-level productivity and distortions across countries.

2.1 Data

We use firm-level financial data from Orbis collected and standardized by Bureau Van Dijk as the main dataset for our analysis. Given that our goal is to characterize cross-country facts on productivity and misallocation, we assemble available data for as many countries as possible in our final dataset. However, we restrict to countries where samples are at least 5,000 observations after cleaning (described below) to allow for enough observations to be representative. We also restrict to the period from 2000 to 2019 since earlier periods tend to have fewer observations in many countries and later periods coincide with the COVID-19 pandemic that may affect cross-firm and cross-country statistics.

Within countries, for comparability with other studies, we focus on firms in the manufacturing sector, nevertheless our results are robust to including all firms. We also drop firm-year observations that are missing sufficient information to construct productivity, are inactive, or are duplicate observations. We trim the remaining variables for extreme values based on output (revenue) at the top and bottom 0.1% and employment greater than 100,000 workers. We drop the bottom 1% of firms based on labor cost as a share of revenue or firms where the wage bill is greater than revenue or value added. We correct employment for firm-year observations that are likely incorrectly reported by replacing employment at the top and bottom 1% of firms based on the wage bill-per-employee with the wage bill-implied employment. We also trim observations at the top and bottom 2% of the productivity and wedge (described below) distribution in each year to limit the influence of outliers. Appendix A provides a detailed description of the data and cleaning procedure.

Our final dataset contains data on 37 countries (from the over 200 countries and regions available in Orbis) with an average of around 300,000 firm-year observations. The number of observations ranges widely across countries, with just over 5,000 observations in Montenegro and Thailand and almost 3 million observations in China. The dataset covers a wide range of the world income distribution with India and Vietnam among the lower-income countries and France and Germany among the higher-income countries.

Discussion of data quality. The Orbis dataset has well-known limitations. Kalemli-Özcan et al. (2024) show that the Orbis firm-level dataset matches well with official country statistics for most European economies. For non-European countries, we restrict to countries with a large number of observations after trimming extreme or unreasonable observations (as discussed above) where coverage is likely to be the best. We find it reassuring that data moments generated in our sample for specific countries match well with other country-specific sources used in the literature. For example, we obtain similar moments for China and India as Hsieh and Klenow (2009) and find that for Vietnam, the lowest income country in our dataset, the Orbis data matches well with the universe of firms in the Vietnam Enterprise Survey (Nguyen, 2025).

2.2 Variable Construction

We describe the firm-level productivity and misallocation distribution across countries using two variables that measure firm-level productivity and distortions. We consider as a starting point a standard decreasing returns to scale production technology for firm i at date t (as in, Hopenhayn, 1992). Output of a homogeneous good is produced using labor according to $y_{i,t} = \mathrm{TFP}_{i,t} n_{i,t}^{\gamma}$, where $\gamma \in (0,1)$. We include labor as the only factor input for parsimony and consistency with the model developed below, although the results extend to capital and intermediates. As in Restuccia and Rogerson (2008), we model misallocation as arising due to idiosyncratic taxes on firm's operating revenue that we label $\tau_{i,t}$. The firm's problem is given by:

$$\max_{n \geq 0} \left\{ (1-\tau_{i,t}) \mathrm{TFP}_{i,t} n^{\gamma} - w_t n \right\}.$$

We refer to firm-level distortions as the firm's wedge since it is a model-based measure of the difference between the firm's realized market allocation and the hypothetical first-best allocation, in which wedges are equalized across firms. In the above specification, firm-level wedges are proportional to $1/(1-\tau_{i,t})$. Following Hsieh and Klenow (2009), we derive empirical measures of firm-level TFP and wedges using the firm's production, inputs, and first-order conditions as:

$$\text{TFP}_{i,t} = \frac{y_{i,t}}{n_{i,t}^{\gamma}}, \qquad\qquad wedge_{i,t} = \frac{y_{i,t}}{n_{i,t}}. \tag{1}$$

While both TFP and the wedge depend on output and employment, the two capture fundamentally different concepts. A simple example is that TFP in an undistorted economy would vary across firms while the wedge would be constant as firms equate the marginal (and average) products of labor. In this section, we simply measure TFP and wedge as defined in Equation (1) and report facts on their distributions. In the next section, we show how these facts relate to model primitives when other factors, such as measurement error, might impact their relationship.

To compute TFP and wedge across firms using data on output and employment, we set the decreasing returns to scale to $\gamma = 0.8$ as is commonly used in the misallocation literature (Guner et al., 2008; Restuccia and Rogerson, 2008). We measure output $y_{i,t}$ as the firm operating revenue, and sales when operating revenue is unreported. We do not use value added because material costs are not widely reported outside of Europe and this limits the final distribution of countries. Employment $n_{i,t}$ is measured as the number of employees hired by the firm. In cases where the number of employees is unavailable, we back out a measure using the wage bill of the firm and a constructed average wage rate for that firm's country sector (two-digit SIC) year. For cross-country comparisons, we remove country-sector-year fixed effects to avoid sectoral and temporal variation from driving our empirical facts.

Robustness of productivity and wedge measures. The wedge in Equation (1) holds for commonly-used production technologies and can be more generally interpreted as the average product of labor. In Appendix B, we show that the main cross-country observations hold if we construct total factor productivity that adjusts for capital inputs, value added as the measure of firm output, if we use a constant elasticity of substitution model as in Hsieh and Klenow (2009), or if we weigh observations on the relative share of firms using national statistics data.

A lingering concern is that using gross output, in place of value added, may lead to differences in production technologies being reflected in the productivity and wedge statistics reported in our stylized facts below. For example, some firms may produce intermediate inputs in-house (high inputs to output) while others may process near-complete intermediates (low inputs to output) leading productivity and the wedge in (1) to appear small for in-house intermediate producers and larger for near-complete intermediate using producer. In Appendix B.4, we explore whether systematic differences in production technologies may drive the data moments in France and Vietnam, the countries spanning the range of aggregate labor productivity in our analysis. To this end, we construct firm groups by the relative importance of intermediate inputs using the value added share of revenues in France and cost of goods sold share of revenues in Vietnam. We do not find systematic differences in moments for the constructed groups and the pooled data used in the baseline analysis. We take this as suggestive evidence that differences in production technologies are not systematically related to underlying firm fundamentals in a way that would substantially alter our conclusions.

2.3 Stylized Facts

We assess how the firm productivity and wedge distributions vary across countries at different stages of development. We document four facts describing systematic differences in these distributions across countries.

Illustration. We start with an example comparing the productivity distributions of three European countries in 2013. Figure 1 reports the firm productivity distribution without residualizing productivity, such that the levels are comparable across countries.

Figure 1 highlights differences between the countries. First, the relatively high income

Figure 1: Productivity Distribution of Operating Firms

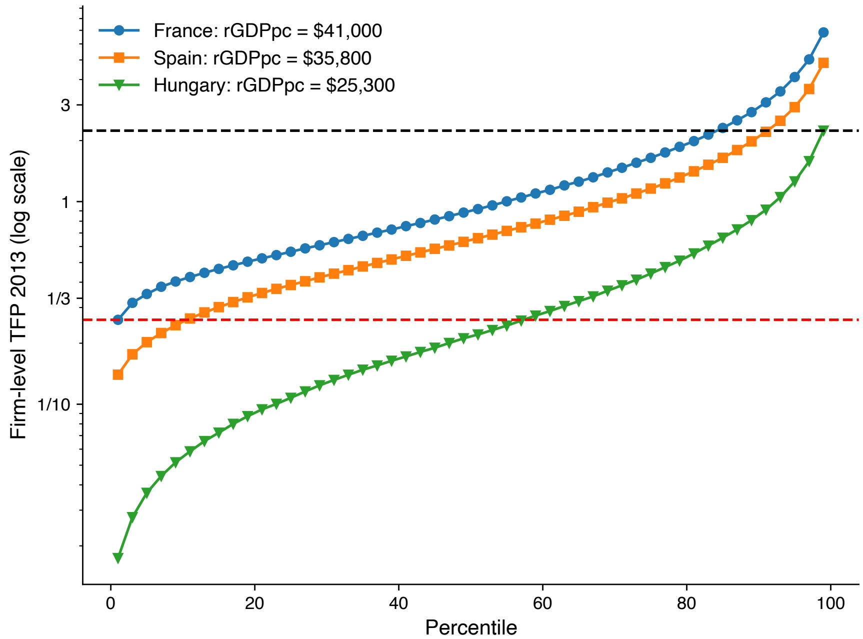

Notes: Values for 50 percentile points in each country of the firm-level TFP distribution in 2013 (in log scale), from percentile one (p1) to percentile 99 (p99). Each point reports the productivity of the indicated percentile firms within the productivity distribution.

France has less dispersed productivity. Second, the productivity distributions in other countries (Spain and Hungary) appear to "fan out" from the top end of the distribution. The most productive firms in France, Spain, and even Hungary have relatively similar productivity while the productivity gap at the bottom percentile of the distribution is much larger. For instance, around 58 percent of firms in Hungary have lower TFP than the bottom one percentile firm in France, whereas only 15 percent of firms in France has higher productivity than the top one percentile of firms in Hungary.

Cross-country productivity distribution. We use the cross-country data to document broader facts. In the following comparisons, we demean productivity and wedges by regressing each variable (in logs) on country-year-sector fixed effects. In this regard, we are comparing the distribution of relative firm productivity excluding levels, unlike Figure 1. We report statistics against aggregate labor productivity in 2015 (real GDP per worker) from the Penn World Table (Feenstra et al., 2015).

Figure 2: Cross-Country Productivity Distribution

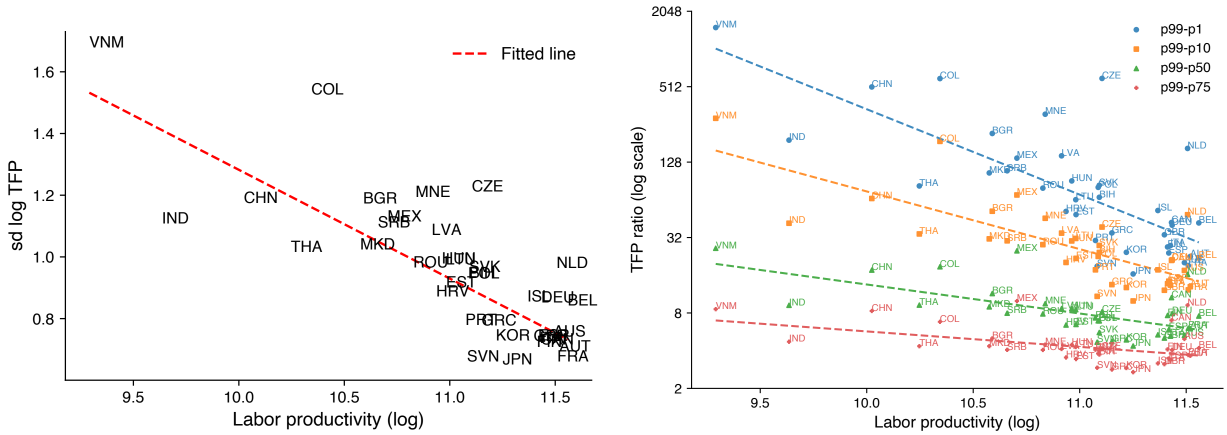

(a) Standard deviation

(b) Fanning out of productivity

Notes: Panel (a) reports firm-level productivity dispersion measured by the standard deviation of log TFP across firms in each country. Panel (b) reports the firm-level TFP ratio at different points of the productivity distribution in log scale. Dashed lines represent the best fit. Each observation is the estimated value for the indicated country. Aggregate labor productivity in logs is based on 2015 data from the Penn World Table (Feenstra et al., 2015).

Fact 1. *Productivity dispersion tends to be higher in lower income countries.*

Figure 2a reports the standard deviation of firm-level TFP. Productivity tends to be more dispersed in lower income countries. This is also confirmed by similar patterns in other measures of dispersion, such as the inter-quartile or inter-decile range, that place less weight on outliers (Appendix B.3).

Fact 2. *Productivity tends to fan out in lower income countries, captured by increasing gaps between high productivity firms and lower productivity firms at different points of the productivity distribution.*

Figure 2b reports the comparison of high-productivity firms with firms at different percentiles of the productivity distribution to provide information on the overall shape of the distribution across countries, and whether, in general, the productivity distribution fans out or not in countries at lower levels of aggregate labor productivity. We compare firms at the top of the distribution (p99) to firms at different points of the distribution (p75, p50, p10, and p1) in all countries. Noting the log scale, uniform shifts in firm-level productivity dis-

Figure 3: Cross-Country Wedge Distribution

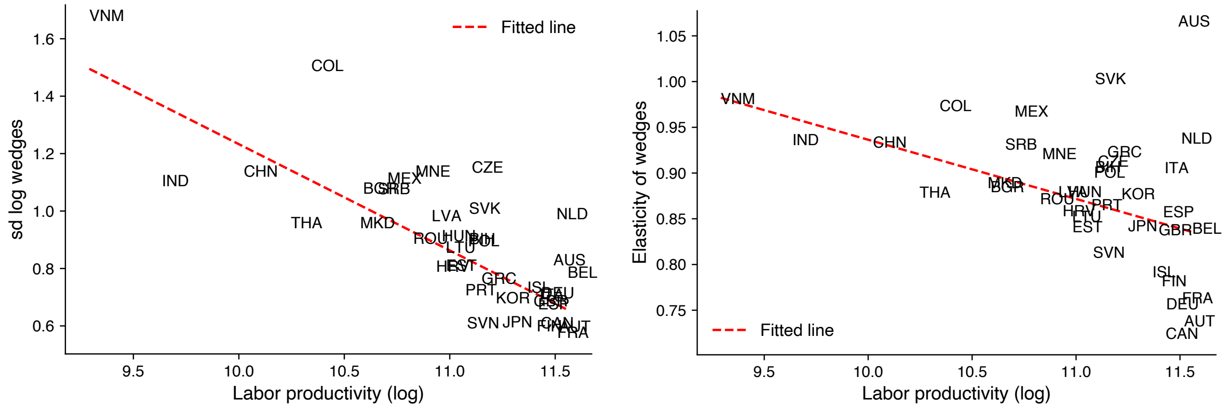

(a) Standard deviation of distortions

(b) Elasticity of distortions

Notes: Panel (a) reports the standard deviation of log wedge across firms in each country. Panel (b) reports the elasticity of distortions measured by the slope coefficient of a regression between log wedge and log TFP in each country. Aggregate labor productivity in logs is based on 2015 data from the Penn World Table (Feenstra et al., 2015).

tribution would result in flat percentile ratios across countries while shifts in productivity dispersion alone would result in the same slopes of the four percentile ratios across countries. As with the Hungary and France comparison, the productivity distribution tends to fan out in lower income countries. The productivity gap of firms at higher productivity percentiles is flatter across countries. This implies that lower productivity dispersion in higher income countries is driven by a relative improvement of low-productivity firms. Our interpretation of this fact is that stronger firm selection in higher income countries discourages entry and operation by low-productivity firms, compressing the productivity distribution.

Cross-country wedge distribution. We next examine how the wedge distribution varies across countries. We focus on two commonly reported moments: the standard deviation of log wedges and the measured elasticity of wedges to productivity.

Fact 3. *Wedge dispersion tends to be higher in lower income countries.*

Figure 3a reports the standard deviation in the log wedge across firms for all the countries in the data. Consistent with previous evidence for individual countries, lower income countries tend to have more dispersed wedges. This indicates that firms in lower income countries tend to have more dispersion in their marginal products of factors and output could be increased by reallocating factors across firms. The extent of the gains from reallocation depend on the joint distribution of wedges and productivity, which we document in the next fact.

Fact 4. *The measured elasticity of wedges to productivity is higher in lower income countries.*

Figure 3b reports the relationship between firm-level wedges and productivity. Lower income countries tend to have distortions that are more correlated with productivity. This captures that more productive firms in lower income countries face more constraints to increasing size, such as through financial frictions that prevent productive firms from increasing scale (Buera et al., 2013) or size-dependent regulations that encourage firms to remain small (Guner et al., 2008). While the measured elasticities reported in Figure 3b are sensitive to modeling and data choices, the cross-country relationship is robust to these choices (reported in Appendix B.2). Reassuringly, we find similar elasticities for Vietnam as in Nguyen (2025) and for China and India as reported by Hsieh and Klenow (2009) using the same model and parameterization.

We emphasize that the pattern of distortions across countries ultimately depends on the relationship between firm-level productivity and firm size since the wedge simply captures a change in return required for a firm to choose the observed employment as an equilibrium of the model. Hence, in light of the systematic relationship between distortions and firm productivity across countries, it is illustrative to document more directly the relationship between firm productivity and employment. Figure 4 shows that lower income countries feature a much lower elasticity of employment to productivity across firms compared to higher income countries.

Other supporting evidence. We document other stylized facts in Appendix B.3 that support our main conclusions. We find that employment dispersion is relatively similar between high- and low-income countries. Since productivity and wedges impact firm employ-

Figure 4: Cross-Country Elasticity of Employment

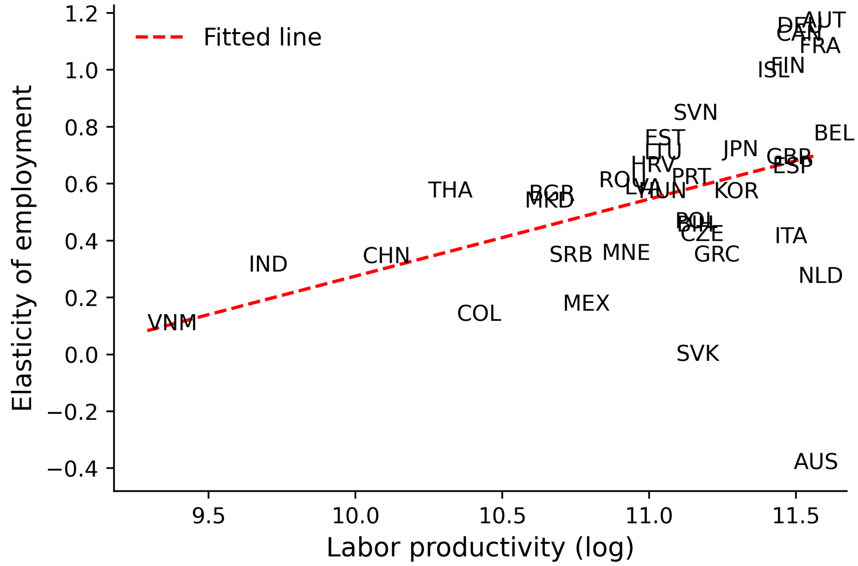

Notes: Elasticity of employment measured by the slope coefficient of a regression between log employment and log TFP across firms in each country. Aggregate labor productivity in logs is based on 2015 data from the Penn World Table (Feenstra et al., 2015).

ment decisions and low-income countries feature greater dispersion in both firm productivity and wedges, the relatively constant employment dispersion is indicative of offsetting factors, such as the negative correlation between firm-level productivity and wedges in Fact 4. Consistently, we find employment is less elastic to TFP in lower income countries as documented in Figure 4. We also find lower correlations between output and employment across firms in lower income countries, which may indicate that higher output firms face difficulties scaling up operations.

The four facts document that the productivity and wedge distributions change systematically across countries. An important insight of our paper is recognizing that the empirical firm-level productivity and wedge distributions may reflect firm technology decisions and selection through operation choices. Analyzing technology and selection requires more structure. In the next sections, we develop and employ a model to examine the roles of technology and selection in driving the above facts and their impact on allocative efficiency and aggregate productivity. We also use the model to examine measurement issues that may create gaps between measured and fundamental notions of productivity and distortions.

3 Model

We develop a model of production heterogeneity with distortions featuring operation decisions, building on Hopenhayn (1992) and Restuccia and Rogerson (2008). We extend the framework to allow for productivity-enhancing investment and highlight selection from operation decisions of firms. We present a parsimonious model that highlights mechanisms that can explain the documented stylized facts and quantify their aggregate implications, abstracting from firm life-cycles (as in, for example, Akcigit et al., 2022; Ayerst et al., 2023). A broad set of distortions potentially underlie the empirical wedges as countries differ substantially in terms of institutional and policy distortions, leading us to take an indirect approach to modeling wedges, rather than a direct approach focusing on specific market distortions.

3.1 Economic Environment

Technologies. At each date, a homogeneous good is produced by firms indexed by i. Firms have access to the decreasing-return-to-scale technology:

$$y_i = v_i {z_i}^{1-\gamma} {n_i}^{\gamma}, \quad \gamma \in (0,1),$$

where y_i is output, n_i is the labor input, and $v_i {z_i}^{1-\gamma}$ is the firm total factor productivity. The term $z_i^{1-\gamma}$ is a permanent component of total factor productivity that is the result of the firm's investment decision and v_i is a transitory component of total factor productivity with $\mathbb{E} v_i = 1$ that is drawn each period from an iid cumulative distribution function $H(v)$ after production decisions are made (similar to the timing in Boar et al., 2022). The transitory component v_i could also capture measurement error, which would appear as a disconnect between the reported output and labor inputs.[2] Firms are subject to an operating fixed cost c_f per period in units of labor and may choose to exit, before v_i is realized, to avoid incurring this cost.

Entry and exit. Entering firms incur an entry cost c_e in units of labor. After entry, firms draw an innovation ability χ_i from an iid cumulative distribution $G(\chi)$ and choose to invest in productivity z_i at cost $\psi z_i^{\phi}/\chi_i$, where $\psi > 0$ and $\phi > 1$. We find similar results if we instead model investment costs in labor (Appendix D.2). Collapsing firm dynamics to a one-time investment decision upon entry allows us to parsimoniously capture the stationary equilibrium relationship between the firm productivity distribution and distortions.[3]

Firms face a random exit shock with probability λ each period. Firms may also choose not to invest and exit the market after learning their type. We denote the mass of entrants by E and the total mass of operating firms by N.

Households. There is a representative household of measure one with log preferences for consumption, $u(C) = \log(C)$, that discounts future periods at rate r. The household is endowed with one unit of productive time each period that is supplied inelastically to the market.

3.2 Market Structure

Firms face idiosyncratic distortions which we model as proportional revenue taxes τ_i as in Restuccia and Rogerson (2008). Following Bento and Restuccia (2017) and Restuccia (2019), we assume that idiosyncratic distortions feature a systematic component related with firm's productivity $z_i^{-\rho}$ and a firm-specific random component ϵ_i. Specifically, we assume that firm-

[2] We could also model a similar transitory component on measured labor inputs but the implications for measurement would be similar and hence for simplicity we only include the output component v_i.

[3] In a similar model, Ayerst (2025) shows how gradual productivity improvements by firms can imply a similar stationary distribution and aggregate properties as when firms make one-time productivity investments.

level distortions $\tau_i(z_i, \epsilon_i)$ are equal to:

$$\log(1-\tau(z_i,\epsilon_i)) = (1-\gamma)\left[-\rho \log z_i - \log \epsilon_i\right],$$

where ρ is the elasticity of distortions with respect to the firm's permanent TFP, determining the systematic component of distortions, and ϵ_i is the random component of distortions drawn from an iid cumulative distribution function $F(\epsilon)$. Intuitively, ρ distorts the productivity gradient of firm size and ϵ captures the effect of distortions on firm size that are independent of productivity. Taxes are collected by a government that redistributes revenues as a lump-sum transfer T to households.

We model τ_i as a catch-all of the myriad of policies and institutions that affect business operation, abstracting from specific drivers of distortions. Our goal is to examine differences across a wide range of countries where sources of distortions may vary widely. We emphasize, however, that the literature has identified numerous policies and institutions creating wedges in marginal products across firms in many different contexts (Hopenhayn, 2014; Restuccia and Rogerson, 2017). Prominent examples of specific policies and institutions creating systematic wedges across firms include firing taxes (Hopenhayn and Rogerson, 1993; Hopenhayn, 2014), financial frictions (Buera et al., 2013), and size-dependent regulations (Guner et al., 2008). Further, the link between distortions and productivity is supported by numerous trade reforms that have reduced misallocation, improved selection, and encouraged technology upgrading (Pavcnik, 2002; Bustos, 2011; Khandelwal et al., 2013).

3.3 Equilibrium

We focus on a stationary competitive equilibrium to examine factors driving long-term cross-country differences. Households and firms take prices as given, prices are constant, and the distribution of resource allocations and firm types are stationary. The price of the output good is normalized to one and the price of labor is denoted by w.

Incumbent firms. An incumbent firm is characterized by productivity z and distortion τ. The firm chooses the optimal labor n to maximize expected per-period profit $\pi(z, \tau)$:

$$\begin{aligned}\pi(z_i, \tau_i) &= \max_{n \geq 0} \quad \mathbb{E}_v \left[v_i (1 - \tau_i) z_i^{1-\gamma} n^\gamma - wn - c_f w \right], \\ &= \max_{n \geq 0} \quad (1 - \tau_i) z_i^{1-\gamma} n^\gamma - wn - c_f w.\end{aligned}$$

In the above expression, the transitory productivity shock v drops out of the firm's problem since $\mathbb{E}v = 1$. The firm's problem implies that the labor demand and optimal output are:

$$\begin{aligned}n(z_i, \tau_i) &= (1 - \tau_i)^{\frac{1}{1-\gamma}} z_i \left(\frac{\gamma}{w} \right)^{\frac{1}{1-\gamma}} = \frac{z_i^{1-\rho}}{\epsilon_i} \left(\frac{\gamma}{w} \right)^{\frac{1}{1-\gamma}}, \\ y(z_i, v_i, \tau_i) &= (1 - \tau_i)^{\frac{\gamma}{1-\gamma}} v_i z_i \left(\frac{\gamma}{w} \right)^{\frac{\gamma}{1-\gamma}} = \frac{z_i^{1-\rho\gamma}}{\epsilon_i^\gamma} v_i \left(\frac{\gamma}{w} \right)^{\frac{\gamma}{1-\gamma}}.\end{aligned}$$

We expand the distortion τ_i into its components in the right-most terms to demonstrate the impact of the two dimensions of distortions and the transitory productivity shock on allocations and output. The systematic component of distortions $z_i^{-\rho}$ determines the gradient of output and inputs to firm-level productivity with higher elasticities ρ indicating a flatter gradient. The random component of distortions ϵ_i leads to firms scaling up or down labor inputs results in more or less output. The transitory productivity shock v_i only impacts output since firms learn of this shock after hiring labor and cannot adjust labor to reflect higher or lower productivity. In this respect, v_i functions similar to mismeasurement that would reflect gaps between measured output and inputs.

Expected operating profits are equal to:

$$\pi(z_i, \tau_i) = \Omega (1 - \tau_i)^{\frac{1}{1-\gamma}} z_i - c_f w, \quad \text{where} \quad \Omega \equiv \left(\frac{\gamma}{w} \right)^{\frac{\gamma}{1-\gamma}} (1 - \gamma).$$

The expected value of a firm can be written as the expected discounted per-period profit stream. The expected value of an incumbent firm $W(z_i, \tau_i)$ is:

$$W(z_i, \tau_i) = \max\left\{\pi(z_i, \tau_i) + (1-\lambda)\frac{W(z_i, \tau_i)}{1+r}, 0\right\},$$
$$= \max\left\{\frac{\Omega(1-\tau_i)^{\frac{1}{1-\gamma}} z_i - c_f w}{1-R}, 0\right\},$$

where $R = (1-\lambda)/(1+r)$, noting that firms with negative (expected) profit would not operate and hence have zero value. We characterize operation decisions below at the time of entry.

Entering firms. Entrants draw an idiosyncratic innovation ability χ_i from distribution $G(\chi)$ and a random distortion component ϵ_i from distribution $F(\epsilon)$. The firm then chooses productivity to maximize its incumbent firm value net of productivity investment cost:

$$V(\chi_i, \epsilon_i) = \max_{z \geq 0}\left[W(z, \tau(z, \epsilon_i)) - \psi\frac{z^{\phi}}{\chi_i}\right],$$

where $W(z, \tau)$ is the value of an incumbent firm with productivity z and $\tau(z, \epsilon_i)$ is the distortion faced by the firm given the choice of z and the random component ϵ_i. We denote the optimal productivity level from this problem by the function $z(\chi, \epsilon)$. Even though there is an optimal productivity level associated with every type (χ, ϵ), only a fraction of firm types operate in the market.

Optimal productivity z for an entrant drawing (χ_i, ϵ_i) is given by:

$$z(\chi_i, \epsilon_i) = \left(\frac{(1-\rho)\tilde{\Omega}\chi_i}{\psi\phi\epsilon_i}\right)^{\frac{1}{\phi+\rho-1}}, \quad \text{where} \quad \tilde{\Omega} \equiv \frac{\Omega}{1-R}. \tag{2}$$

Both χ and ϵ affect productivity proportionally. The elasticity of distortions ρ flattens the relationship between firm productivity and firm type (χ, ϵ). Intuitively, firms with higher innovative ability should operate at a larger employment scale, but distortions associated with ρ discourage it by reducing profits and hence, distortions of this type represent an effective tax to productivity-enhancing investment.

Using this optimal productivity and substituting for the value of an incumbent firm, the value of an entrant firm drawing (χ_i, ϵ_i) is given by:

$$\begin{aligned} V(\chi_i, \epsilon_i) &= \max\left\{\tilde{\Omega} z(\chi_i, \epsilon_i)^{1-\rho}\epsilon_i - \psi\frac{z(\chi_i, \epsilon_i)^{\phi}}{\chi_i} - \frac{c_f w}{1-R}, 0\right\}, \\ &= \max\left\{\Gamma(w, \rho)\chi_i^{\frac{1-\rho}{\phi+\rho-1}}\epsilon_i^{-\frac{\phi}{\phi+\rho-1}} - \frac{c_f w}{1-R}, 0\right\}, \end{aligned}$$

where

$$\Gamma(w, \rho) \equiv \frac{\phi+\rho-1}{\phi}\tilde{\Omega}\left(\frac{(1-\rho)\tilde{\Omega}}{\psi\phi}\right)^{\frac{1-\rho}{\phi+\rho-1}}.$$

As firms only operate when their value is non-negative, the decision to operate for a firm drawing (χ_i, ϵ_i) can be characterized as:

$$o(\chi_i, \epsilon_i) = \begin{cases} 1 \quad \text{if} \quad \Gamma(w, \rho)\chi_i^{\frac{1-\rho}{\phi+\rho-1}}\epsilon_i^{-\frac{\phi}{\phi+\rho-1}} \geq \frac{c_f w}{1-R}, \\ 0 \quad \text{otherwise.} \end{cases} \tag{3}$$

We emphasize from the technology and operation decisions of firms that even for economies with the same distribution of innovation ability $G(\chi)$, the distribution of productivity z of operating firms may differ if the economic environment features different distortions across firms (ρ, $F(\epsilon)$). To put it differently, in economies with no distortions, there is a simple mapping from the exogenous distribution of innovation ability $G(\chi)$ to the distribution of firm productivity, but distortions affect this mapping though effects on the operation and technology decisions of firms.

At the beginning of each period, the entry value V_e is given by:

$$V_e = \mathbb{E}_{\chi,\epsilon}V(\chi, \epsilon) - c_e w \leq 0.$$

The entry condition requires that firms enter to the point where further entry is no longer valuable, and is equal to zero in an equilibrium with positive entry.

Firm distribution. The firm distribution is straightforward to characterize since we abstract from firm dynamics other than entry, exit, and the random productivity shock v. We denote by $\mu(\chi, \epsilon)$ the mass of producers with investment ability χ and random distortion component ϵ. The law of motion for $\mu(\chi, \epsilon)$ is given by:

$$\mu'(\chi, \epsilon) = (1 - \lambda)\mu(\chi, \epsilon) + Eo(\chi, \epsilon)dF(\epsilon)dG(\chi).$$

The implied stationary distribution is given by:

$$\mu(\chi, \epsilon) = \frac{E}{\lambda} o(\chi, \epsilon)dF(\epsilon)dG(\chi). \tag{4}$$

The mass (number) of firms in a stationary equilibrium is

$$N = \int_\chi \int_\epsilon d\mu(\chi, \epsilon) = \frac{E}{\lambda} \int_\chi \int_\epsilon o(\chi, \epsilon)dF(\epsilon)dG(\chi). \tag{5}$$

Definition of equilibrium. A stationary competitive equilibrium comprises a wage w; decision functions for firms: labor demand $n(z, \tau)$, profits $\pi(z, \tau)$, value of incumbent firm $W(z, \tau)$, productivity $z(\chi, \epsilon)$, operating decision $o(\chi, \epsilon)$, net value of firm $V(\chi, \epsilon)$, value of entry V_e, a distribution of firms $\mu(\chi, \epsilon)$, mass of firms N and entrants E; lump-sum transfer T; and allocation C for households such that:

(i) Given w and T, the allocation C solves the household's problem.

(ii) Given w, decision function $n(z, \tau)$ solves the incumbent's firm problem, determining per-period profits $\pi(z, \tau)$ and the value of incumbent firms $W(z, \tau)$.

(iii) Given w, entrants choose productivity $z(\chi, \epsilon)$ and operating decision $o(\chi, \epsilon)$ to maximize the net value of the firm $V(\chi, \epsilon)$.

(*iv*) Zero profit entry condition $V_e = 0$.

(*v*) Invariant distribution of firms μ given by equation (4), which implies the mass of firms is constant and given by equation (5).

(*vi*) The government's budget is balanced:

$$0 = T + \int_\chi \int_\epsilon \tau(\chi, \epsilon)(1 - \tau(\chi, \epsilon))^{\frac{\gamma}{1-\gamma}} z(\chi, \epsilon) \left(\frac{\gamma}{w}\right)^{\frac{\gamma}{1-\gamma}} \mu(d\chi, d\epsilon).$$

(*vii*) The goods and labor markets clear:

$$\int_\chi \int_\epsilon z(\chi, \epsilon)^{1-\gamma} n(\chi, \epsilon)^\gamma \mu(d\chi, d\epsilon) = C + E \int_\chi \int_\epsilon \psi \frac{z(\chi, \epsilon)^\phi}{\chi} o(\chi, \epsilon) dG(\chi) dF(\epsilon),$$

and

$$1 = \int_\chi \int_\epsilon n(\chi, \epsilon) o(\chi, \epsilon) \mu(d\chi, d\epsilon) + E c_e + N c_f.$$

Equilibrium solution. Given a wage rate w, all firm decision functions can be solved and since V_e is a strictly decreasing function of w, the zero profit entry condition solves for w (Proposition 1). The labor market clearing condition solves for the mass of entry E which in turn determines all other variables such as the invariant distribution and number of firms.

Proposition 1. *The equilibrium wage rate w is determined by the zero profit entry condition:*

$$\int_\chi \int_\epsilon \max\left\{\Gamma(w, \rho) \chi_i^{\frac{1-\rho}{\phi+\rho-1}} \epsilon_i^{-\frac{\phi}{\phi+\rho-1}} - \frac{c_f w}{1 - R}, 0\right\} G(d\chi) F(d\epsilon) - c_e w = 0. \tag{6}$$

The left-hand side (LHS) of Equation (6) is the expected value of potential entrants net of entry costs which must be zero in an equilibrium with positive entry as in our framework.

3.4 Model Implications

In the empirical section, we documented distributional patterns of the firm-level productivity and wedges across countries. We now discuss how the model relates with these observations. The measured firm-level productivity and wedge of firm i are given by:

$$\mathrm{TFP}_i = \frac{y_i}{n_i^{\gamma}} = z(\chi_i, \epsilon_i)^{1-\gamma} v_i, \qquad \qquad wedge_i = \frac{y_i}{n_i} = \left(\frac{w}{\gamma}\right) \frac{v_i}{1-\tau_i}. \tag{7}$$

We note that both technology and selection channels affect the dispersion of firm-level productivity via firm choices of technology z and operation. Similarly, the transitory productivity component v_i (mismeasurement) also drives dispersion in firm-level productivity. Moreover, the transitory component implies a mechanical relationship between measured productivity and wedges. Therefore, it is essential to quantify the importance of these different channels.

Cross-country productivity distribution. The impact of technology on the productivity distribution can be assessed through technology choice z in Equation (7), rewritten as:

$$\log\left(z(\chi_i, \epsilon_i)\right) = \frac{1}{\phi+\rho-1} \log\left(\frac{(1-\rho)\tilde{\Omega}}{w\psi\phi}\right) + \frac{1}{\phi+\rho-1}\left[\log(\chi_i) - \log(\epsilon_i)\right]. \tag{8}$$

In an undistorted economy, technology dispersion is driven by firm investment ability χ. Productivity choice is affected by the elasticity of distortions ρ through its impact on firm incentives to invest in productivity, where a higher elasticity of distortions ρ lowers the χ-gradient of technology. The random distortion component ϵ also affects technology choice in the same proportion as investment ability χ, generating dispersion in z for a given χ.

The impact of selection can be assessed in the operation cutoff condition in Equation (3), rewritten as:

$$\frac{1-\rho}{\phi+\rho-1}\log(\chi) - \frac{\phi}{\phi+\rho-1}\log(\epsilon) \geq \log\left[\frac{c_f w}{(1+R)\Gamma(w,\rho)}\right]. \tag{9}$$

It follows from the above expression that a higher right-hand side (i.e., more selection) implies that, for a given χ, the operation cutoff for the random distortion component ϵ (equivalently, a lower τ) is smaller. Consequently, more selection implies a more positive measured covariance between χ and ϵ, denoted by $cov(\chi, \epsilon|o)$, as low χ, high ϵ firms do not operate. Additionally, more selection implies lower variances of χ and ϵ conditional on operating, denoted by $\sigma^2_{\chi|o}$ and $\sigma^2_{\epsilon|o}$. While we cannot directly map the extent of distortions into selection, looking ahead, we find in our quantitative analysis that more distorted economies tend to have less selection. Intuitively, more elastic distortions (higher ρ) decrease the average size of more productive firms, lowering wages and freeing resources for lower productivity firms to remain active in the economy.

Proposition 2 derives the measured variance of productivity and percentile gaps reported in Facts 1 and 2.

Proposition 2. *The measured variance of log TFP is given by:*

$$\begin{aligned} var(TFP) &= (1-\gamma)^2 \sigma^2_{z|o} + \sigma^2_v, \\ &= \left(\frac{1-\gamma}{\phi + \rho - 1} \right)^2 (\sigma^2_{\chi|o} + \sigma^2_{\epsilon|o} - 2\ cov(\chi, \epsilon|o)) + \sigma^2_v. \end{aligned}$$

Productivity dispersion is decreasing in the elasticity ρ and the extent of selection.

The measured TFP gap between the percentile i and j firm is given by:

$$gap_{i,j} = \log TFP_i - \log TFP_j = \frac{1-\gamma}{\phi + \rho - 1} \log \left(\frac{\chi_i / \epsilon_i}{\chi_j / \epsilon_j} \right) + \log \left(\frac{v_i}{v_j} \right).$$

The measured TFP variance is decreasing in the elasticity of distortions ρ, other factors equal. In particular, in the absence of selection (i.e., when $cov(\chi, \epsilon|o) = 0$, $\sigma^2_{\chi|o} = \sigma^2_\chi$, and $\sigma^2_{\epsilon|o} = \sigma^2_\epsilon$), more distorted economies have lower productivity dispersion via the technology channel. Consequently, for the model to generate the observed higher dispersion in productivity documented in Fact 1, the selection channel is needed and quantitatively must counteract the technology channel that works in the opposite direction. As a result, on the basis of technology and selection alone it is possible for the model to fail in generating an increase in productivity dispersion in lower income countries. Larger transitory shocks (e.g., mismeasurement) σ_v in lower income countries could also account for the higher dispersion in firm-level productivity, but this channel alone would not account for the fanning out of productivity at the bottom of the distribution in Fact 2. Disentangling the quantitative significance of each of these channels is the purpose of our quantitative analysis.

Cross-country variation in the elasticity of distortions ρ or in the variances of σ_χ^2 and σ_ϵ^2 would result in proportional changes in the productivity gap between the percentiles i and j firms across countries, in the case without selection and if transitory shocks are small. This follows because the relative ordering of firms within the country remains constant while adjustments to ρ, σ_χ^2, σ_ϵ^2 proportionally stretch or compress productivity gaps between firms. Selection also provides a channel through which the ordering of firms within the country changes (as low percentile firms exit), potentially accounting for the observed fanning out pattern in Fact 2.

Cross-country wedge distribution. Proposition 3 derives the measured wedge variance and elasticity of wedges reported in Facts 3 and 4.

Proposition 3. *The measured variance of the log wedge is equal to:*

$$var(wedge) = (1-\gamma)^2\rho^2\sigma_{z|o}^2 + (1-\gamma)^2\sigma_{\epsilon|o}^2 + 2\rho(1-\gamma)^2 cov(z,\epsilon|o) + \sigma_v^2. \tag{10}$$

The measured elasticity of wedges is given by:

$$elas(TFP, wedge) = \frac{\rho(1-\gamma)^2\sigma_{z|o}^2 + \sigma_v^2 + (1-\gamma)^2 cov(z,\epsilon|o)}{(1-\gamma)^2\sigma_{z|o}^2 + \sigma_v^2}. \tag{11}$$

The standard deviation of wedges reflects dispersion in the random distortion component σ_ϵ^2 and dispersion in the firm productivity through the correlated component of distortions

ρ and the covariance between productivity and the random distortion component. The standard deviation of wedges may be overestimated when factors that firms cannot adjust inputs to, such as measurement error, affect firms output (or inputs), captured by the standard deviation in the transitory productivity component σ_v.

The proposition also shows that there are three potential sources of bias between the measured elasticity of wedges and the fundamental parameter ρ. First, the transitory productivity component v creates a mechanical relationship between the measured wedge and TFP that increases the elasticity. Since firms cannot adjust inputs in response to this shock, the shock has the same impact on both the measured wedge and productivity creating a positive bias. Second, selection creates a positive bias in the estimated elasticity. This follows because unproductive (low χ), high distortion (high τ) firms do not operate, making it more likely that observed relatively unproductive firms have relatively low distortions (low τ), biasing upwards the measured elasticity. Third, firms choose technology z based on their draw of random distortions ϵ creating a mechanical negative relationship between measured wedges and productivity, biasing downwards the estimated elasticity (see also, Ayerst, 2025).

The actual magnitude of these biases and their net impact on outcomes across countries is a quantitative question that we examine in the next section. Depending on the magnitude, the biases could imply either smaller or larger cross-country differences in misallocation.

4 Quantitative Analysis

We proceed in three steps. First, we calibrate the model to observations for France. Second, we re-calibrate key parameters to each country in our dataset to provide a reasonable range of parameter values. We show that differences in distortions in the range of cross-country estimates can account for the facts documented in Section 2, with the elasticity of distortions ρ accounting for the bulk of differences. Third, we decompose the productivity losses from varying distortions into static misallocation, selection, and technology choice.

4.1 Calibration

We calibrate a distorted benchmark economy to micro and aggregate data for France. We parameterize the distributions of $\log \chi$, $\log v$, and $\log \epsilon$ to be normal with mean zero. There are 11 parameters to calibrate in the model: the decreasing returns to scale γ, the exogenous firm exit rate λ, the real interest rate r, the dispersion in innovation ability σ_χ, the dispersion transitory ex-post productivity shock σ_v, the level and curvature parameters of innovation cost function ϕ and ψ, the fixed costs of entry c_e and operation c_f, the productivity elasticity of distortions ρ, and the dispersion of the random wedge component σ_ϵ.

Six parameters are either normalized or assigned values from outside evidence. We set the decreasing returns to scale to $\gamma = 0.8$ as in the empirical section, the exit rate to $\lambda = 0.10$ (Davis et al., 1998), the real interest rate to $r = 0.04$, and the curvature of investment cost function to $\phi = 2$ (Acemoglu et al., 2018). We normalize the investment cost $\psi = 1$ and the entry cost $c_e = 1$.

The remaining five parameters ρ, σ_ϵ, σ_χ, σ_v, and c_f are jointly calibrated to match the following moments: (1) the distortion-productivity elasticity, (2) the standard deviation of log distortions, (3) the standard deviation of log employment, (4) the standard deviation of log TFP, and (5) average firm size. Table 1 reports the calibrated parameter values and the model and data moments. The benchmark economy calibrated to the France data replicates the targeted moments.

Table 1: Calibration of Distorted Benchmark Economy to France

Parameter	Value	Targeted moments	Model	Data
ρ	0.478	Measured elasticity of distortions	0.75	0.76
σ_ϵ	1.57	sd log distortions	0.55	0.56
σ_χ	9.94	sd log employment	1.40	1.40
σ_v	0.25	sd log TFP	0.67	0.66
c_f	0.10	Average firm size	14.8	14.9

We also highlight how the calibration moments provide identification for the model parameters. Table 2 reports the sensitivity of the France calibration moments to a 10 percent increase in each calibrated parameter, highlighting the extent of interconnectedness between the model moments and parameters. The measured elasticity of distortions is most sensitive to the model parameter ρ, but also depends on the other parameters due to direct or indirect impacts of these parameters on the measured moment. The standard deviation of distortions mostly reflects the model distortions themselves through ρ and σ_ϵ, but is also impacted by the transitory productivity dispersion σ_v because this is measured as part of the wedge, and by the other parameters through their impact on productivity dispersion and selection. The dispersion parameters σ_v and σ_χ and the fixed cost c_f all directly impact either the dispersion in potential productivity or firm productivity types that select into operating, which impacts the dispersion in firm productivity and employment as well as the average firm size. Similarly, the distortion parameters ρ and σ_ϵ directly impact the dispersion of firm productivity and employment and the mass of firms through firm technology and operating choices.

Table 2: Effects of 10% Changes in Calibrated Parameter Values

	ρ	σ_ϵ	σ_v	σ_χ	c_f
Measured elasticity of distortions	3.8	1.7	0.9	−0.9	0.3
sd log distortions	6.7	3.9	2.1	2.1	−0.3
sd log TFP	4.6	2.2	1.5	3.0	−0.5
sd log employment	−4.7	1.4	0.0	3.4	−0.8
Average firm size	−24.4	11.2	0.0	18.5	5.1

Notes: The values indicate the percentage changes in each moment when the indicated parameter is increased by 10 percent relative to the benchmark value for France and all other parameters are fixed to benchmark values.

Figure 5 reports the percentile distribution of firm-level TFP in the model compared to the Orbis data for France. To highlight that the model is able to replicate the distribution of other countries in the data, we also report the fit of the economy calibrated to Vietnam, the lowest income per capita country in our sample. Despite the calibration assuming a log normal distribution of innovation abilities and only targeting the standard deviation of log TFP, the resulting distribution of firm-level TFP matches closely for both countries. We also note that the calibrated France economy features strong selection of firms in operation with around 95 percent of firms not operating in equilibrium.

Figure 5: Firm-level TFP Distribution in France and Vietnam

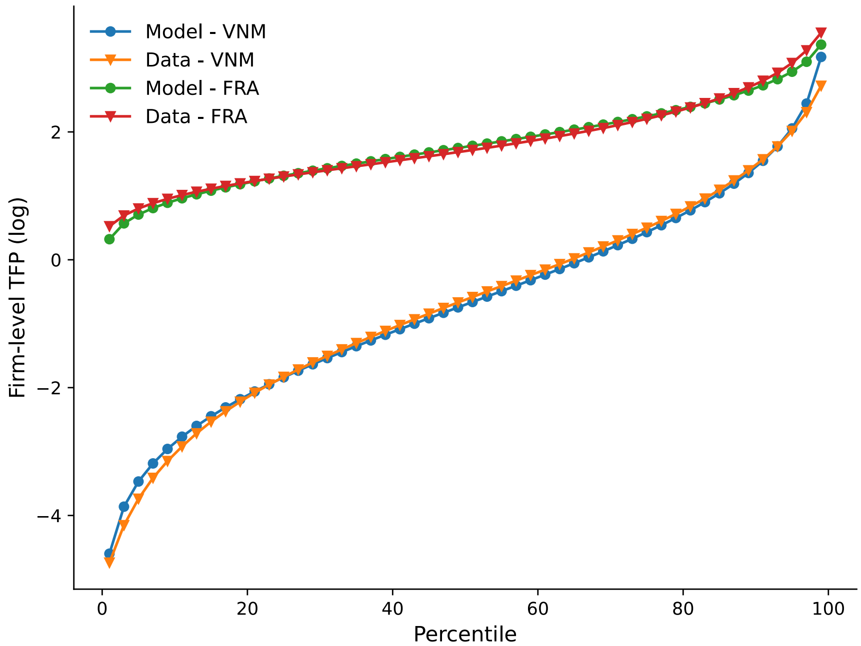

Notes: Data refers to the distribution of firm-level TFP in France (FRA) and Vietnam (VNM) from Orbis, whereas Model refers to the calibrated distribution of productivity in the distorted economies. For ease of illustration, the figure only plots 50 percentile points of the distributions, from percentile one (p1) to percentile 99 (p99). Means of the data distributions are normalized to be equal to the model counterparts.

4.2 Cross-Country Experiments

To provide a plausible range of parameter values across countries, we repeat the calibration procedure of the benchmark economy for every country in our dataset. Figure 6 reports the estimated values of the distortion parameters (ρ, σ_ϵ) and the remaining parameters $(\sigma_v, \sigma_\chi, c_f)$ are reported in Appendix D.1. We find a negative relationship between calibrated parameters ρ, σ_ϵ, and σ_v and aggregate labor productivity across countries, whereas the calibrated parameters σ_χ and c_f show much less systematic variation across countries.

We examine whether measured distortions can generate the cross-country patterns on productivity and wedges documented in Section 2. In these experiments, we start with the

Figure 6: Cross-Country Calibration Results

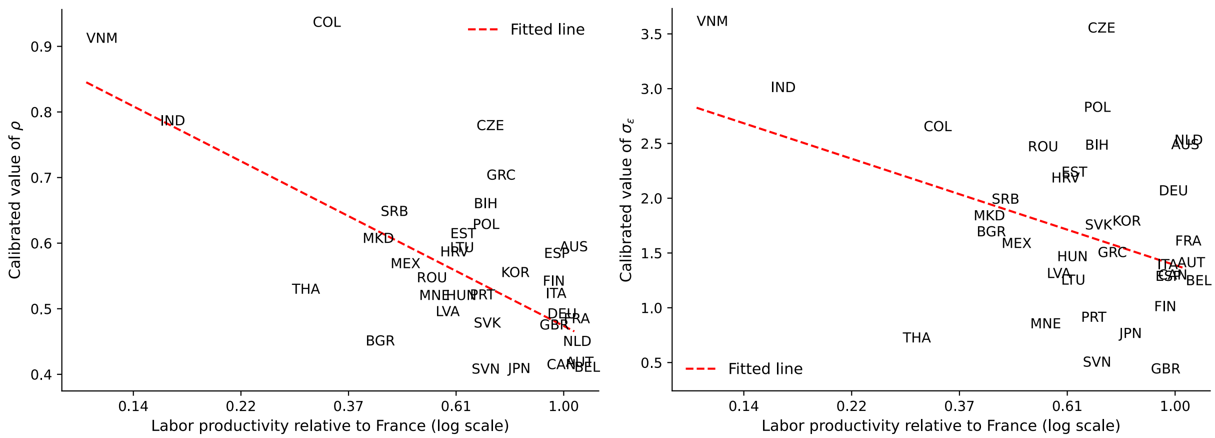

(a) Elasticity of distortions ρ (b) Dispersion of distortions σ_ϵ

Notes: Country codes indicate the parameter values for the calibration of the economy setting the five parameters $\{\rho, \sigma_\epsilon, \sigma_v, \sigma_\chi, c_f\}$ to match the country moments reported in Table A.1. The red dashed line is the linear fit of parameters and log aggregate labor productivity.

benchmark economy calibrated to France and then adjust distortions captured by (ρ, σ_ϵ) to match other countries along the aggregate productivity distribution. In this regard, the counterfactual economies represent the effect of greater distortions in France matching more distorted economies. Rather than focusing on individual countries, we adjust distortions along the fitted line in Figure 6 such that the counterfactual economies capture average features of each country along the income distribution. For example for ρ_i, we set $\rho_i = \rho_{FRA} + \beta \log LP_i$ for a counterfactual economy i with labor productivity LP_i relative to France and β being the estimated cross-country slope coefficient in Figure 6.

We assess parameter adjustments in stages, starting with just the elasticity of distortions (Model w/ $\Delta\rho$) and then adjusting both the elasticity of distortions and the standard deviation of distortions (Model w/ $\Delta\rho, \sigma_\epsilon$). Since we do not target the level of income or aggregate productivity in the calibration, the model also provides a quantification of the relative strength of the different channels on aggregate productivity. The transitory productivity shock v, which captures noise or mismeasurement, can also be important for explaining the range of outcomes observed across countries, although we emphasize that variation in σ_v

does not generate effects on aggregate productivity. To give a sense of the importance of the dispersion in the noise term, σ_v, we also report moments with a higher value of σ_v than in the benchmark economy (Model w/ ρ, σ_ϵ + higher σ_v) that we set equal to $\sigma_v = 0.76$, near the calibrated value for the Netherlands.

Model fit. We compare the cross-country model results against the facts from Section 2 and other relevant moments. We report each moment against aggregate labor productivity and normalize aggregate labor productivity in France in the data and the model.[4] In the model, aggregate labor productivity is aggregate output since aggregate labor is normalized and constant across economies.

As is expected, the model under-accounts for the cross-country labor productivity gap, which is not targeted in the calibration. To provide a sense of the range, we compare labor productivity in France with Vietnam, where labor productivity in Vietnam is around 11% of France in the data. The model generated gap is around 33% when only ρ varies and around 29% when ρ and σ_ϵ vary, indicating that the model accounts for between 50% and 56% ($\approx \log(0.29)/\log(0.11)$) of the gap. Note that the fitted elasticity $\rho = 0.85$ in these experiments is lower than the $\rho = 0.90$ found for Vietnam in Figure 6, which would instead account for around 63% of the aggregate productivity gap. Taken together, the model accounts for around half to two-thirds of the productivity gap between France and the lowest productivity country, consistent with other metrics (e.g., the gradient of aggregate productivity of the calibrated economies is around half of the same gradient in the data). We note that the remaining aggregate productivity gaps could be accounted for by other channels not emphasized in the model, such as aggregate barriers to technology adoption (e.g., Parente and Prescott, 1994; Ayerst, 2025), potentially captured by cross-country differences in the parameter ψ. We leave the exploration of this possibility for future work.

Figure 7 reports the changes in both the standard deviation of log TFP and the produc-

[4] We avoid using cross-country measures of aggregate TFP because of the important measurement issues associated with this statistic in the data, especially related to measures of capital, but we emphasize that our main conclusions would hold.

Figure 7: Dispersion Measures of Firm-Level TFP (Facts 1 and 2)

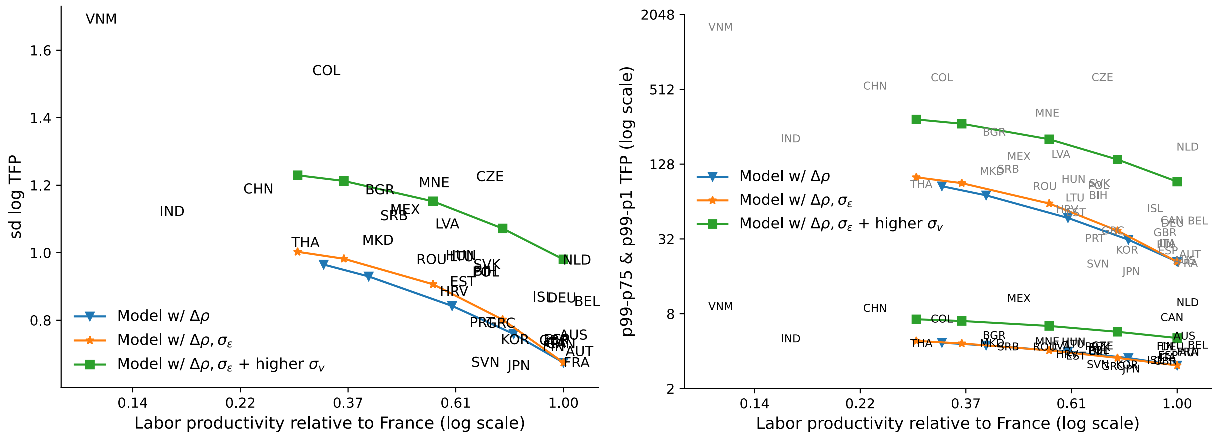

(a) Standard deviation of log TFP (b) p99-p75 & p99-p1 of log TFP

Notes: The blue-triangular, orange-starred, and green-squared lines correspond to experiments: (1) varying ρ, (2) varying ρ and σ_ϵ, and (3) varying ρ and σ_ϵ with a higher value of σ_v, respectively. Country codes are data values for the indicated country. Panel (a) reports the standard deviation of measured TFP. Panel (b) reports the ratio p99/p75 and p99/p1 against aggregate labor productivity. Aggregate labor productivity in logs is based on 2015 data from the Penn World Table (Feenstra et al., 2015).

tivity gaps between the p99 and p75 and p99 and p1 firms (Facts 1 and 2). We exclude other gaps reported in Figure 2b for clarity. Varying distortions fit the cross-country data relatively well. The model elasticity of distortions ρ generates the pattern of the cross-country differences while the standard deviations of distortions σ_ϵ is relatively less important. For example, the model changing only ρ explains most of the empirical gaps between France and India, China, and Thailand. The transitory productivity shock σ_v is also important for explaining outliers in the data. For example, the model with a higher transitory shock σ_v over explains the gaps between France and most countries but helps the model's fit of some countries (e.g., Colombia, Czechia, Netherlands) that are also visual outliers in Figures 2. The p99 and p75 and p99 and p1 gaps are not targeted in the calibration and act as a goodness-of-fit check on the model.

Figure 8 documents the relationship between the standard deviation and measured elasticity of distortions across countries (Facts 3 and 4). As with the dispersion in firm-level productivity, the models fit the cross-country data well. Similarly, the model with only dif-

Figure 8: Standard Deviation and Elasticity of Wedges (Facts 3 and 4)

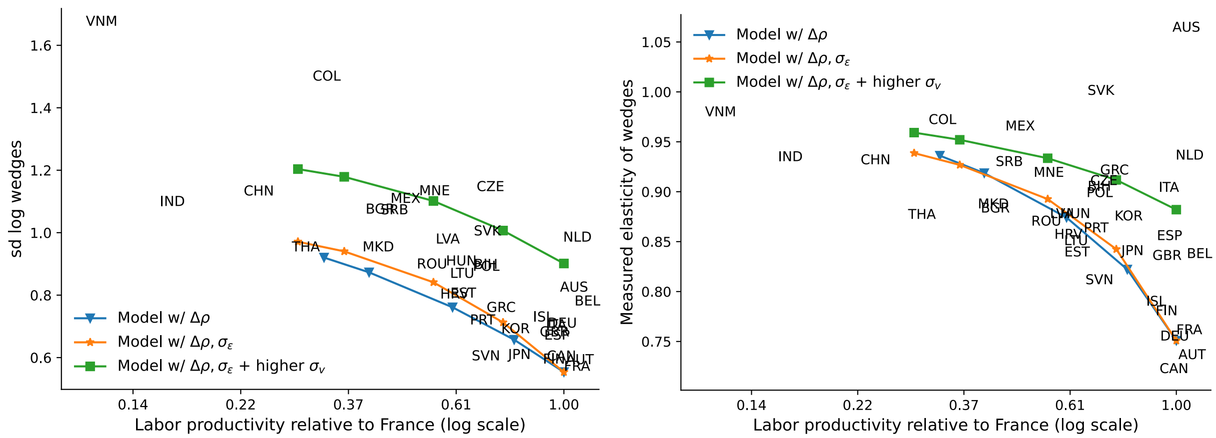

(a) Standard deviation of log wedges

(b) Measured elasticity of wedges

Notes: The blue-triangular, orange-starred, and green-squared lines correspond to experiments: (1) varying ρ, (2) varying ρ and σ_ϵ, and (3) varying ρ and σ_ϵ with a higher value of σ_v, respectively. Country codes are data values for the indicated country. Panel (a) reports the standard deviation of the measured log distortion. Panel (b) reports the measured elasticity of distortions. Aggregate labor productivity in logs is based on 2015 data from the Penn World Table (Feenstra et al., 2015).

ferences in the elasticity of distortions ρ explains most of the gap in the standard deviation of wedges between France and individual countries, excluding outliers where the transitory productivity shock can help generate enough dispersion. The bulk of the measured elasticity of distortions are generated by differences in the model elasticity of distortions ρ across economies, even after accounting for potential biases in the measured elasticity. It is also worth highlighting that mismeasurement (higher σ_v) represents a much smaller component of the measured elasticity of distortions in lower income countries.

Figure 9 reports the dispersion in firm-level employment, measured by the standard deviation of log employment across countries. The data does not indicate a systematic relationship between the dispersion in firm size and aggregate productivity across countries. Variation in the elasticity of distortions ρ alone results in lower dispersion in firm size in higher ρ economies since higher values of ρ directly compress the employment distribution. Including variation in σ_ϵ realigns the model with the data. This suggests that the dispersion in firm size is driven more by productivity dispersion in low elasticity ρ economies, where employment

Figure 9: Employment Dispersion across Firms

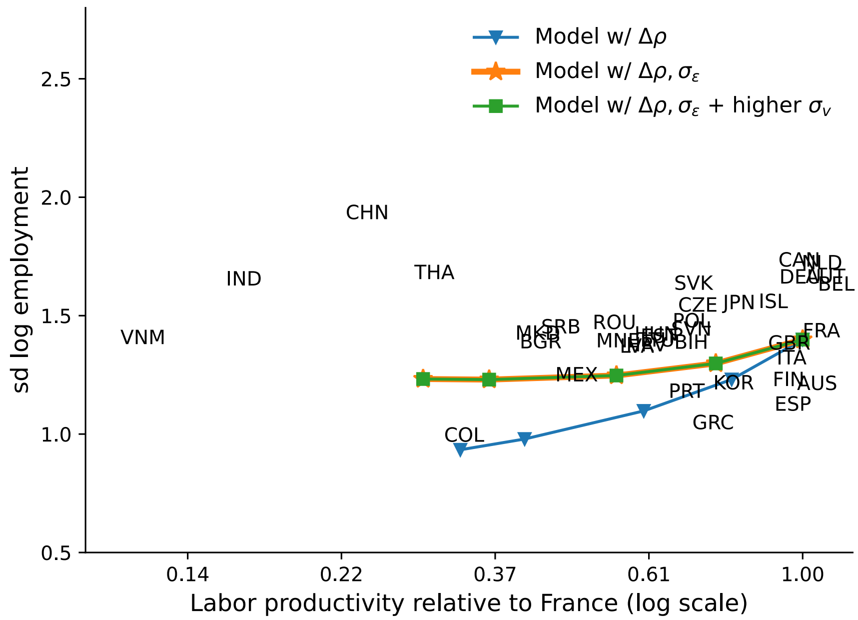

Notes: Employment dispersion is measured by the standard deviation of log employment across firms. The blue-triangular, orange-starred, and green-squared lines correspond to experiments: (1) varying ρ, (2) varying ρ and σ_ϵ, and (3) varying ρ and σ_ϵ with a higher value of σ_v, respectively. Country codes are data values for the indicated country. Experiments (2) and (3) are perfectly aligned because v does not affect employment decisions. Aggregate labor productivity in logs is based on 2015 data from the Penn World Table (Feenstra et al., 2015).

is closer to efficient levels, and more by dispersion in distortions in lower income countries.

In Appendix D.3, we show that models with selection or technology alone would not generate the relatively flat profile of cross-country employment dispersion and would generate too much or too little dispersion in productivity. In this regard, the quantitative fit of the moments also serves as a check on the need for both technology choice and selection mechanisms in the model.

Figure 10 reports allocative efficiency in the model and the data, which is also not targeted in the calibration. Allocative efficiency is the ratio of aggregate output in an economy to the aggregate efficient output that could be achieved if labor were efficiently allocated, for the given set of operating firms and technologies in each economy. As a result, allocative efficiency depends on the joint productivity and wedge distribution. The model implies an allocative efficiency of 0.76 in the France calibration economy, close to to 0.65 in the France

Figure 10: Allocative Efficiency in Model and Data

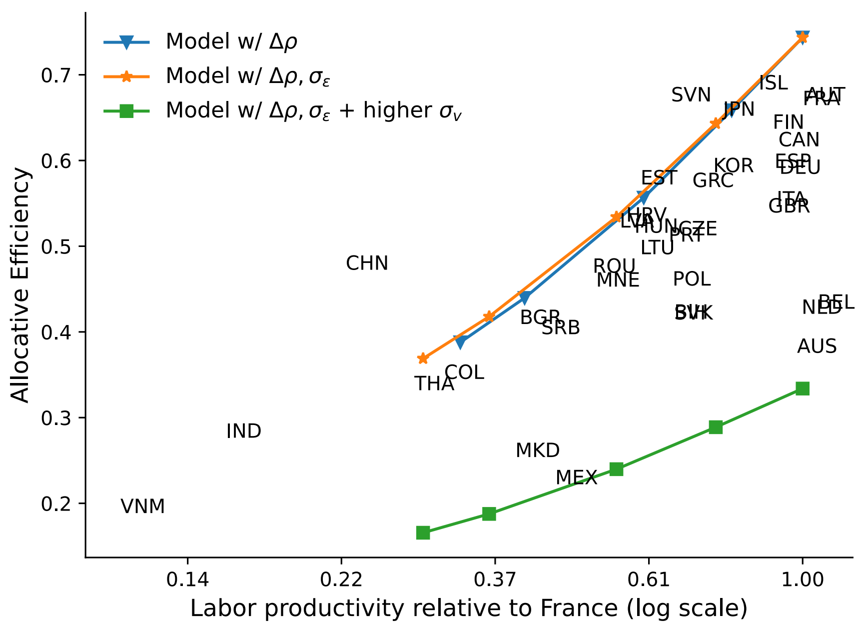

Notes: The blue-triangular, orange-starred, and green-squared lines correspond to experiments: (1) varying ρ, (2) varying ρ and σ_ϵ, and (3) varying ρ and σ_ϵ with a higher value of σ_v, respectively. Country codes are data values for the indicated country. Allocative efficiency measures output divided by output in the efficient allocation economy. Aggregate labor productivity in logs is based on 2015 data from the Penn World Table (Feenstra et al., 2015).

data. Varying the elasticity ρ and dispersion σ_ϵ of distortions to match the most distorted economy implies allocative efficiency of 0.37, compared to the 0.18 in the Vietnam data, and a gap between France and the most distorted economy of around 0.50, compared to the 0.28 observed for the France-Vietnam gap in the data. More generally, the model with only differences in ρ and σ_ϵ fits well the measured allocative efficiency across countries in the data, albeit higher in some cases, which is to be expected given the stylized nature of the calibrated model. The higher random productivity components σ_v leads the model to over-explain the empirical allocative efficiency for most countries, similar to the other empirical moments. However, as noted before, the inclusion of σ_v can be important for explaining the relatively low measured allocative efficiency in some countries. An interpretation is that high measurement error in some countries can lead to low measures of allocative efficiency, potentially explaining why some high-income countries (e.g., Australia, Belgium, Netherlands) appear to have low allocative efficiencies.

Measured distortions and model mechanisms. We close the cross-country discussion by examining the relationship between the model fundamentals and the measured moments related to distortions. The model implies three mechanisms that impact the measurement of distortions: technology, selection, and noise. The impact of the three channels can be decomposed as:

$$
\begin{aligned}
var(wedge) = \rho^2 var(TFP) + (1-\gamma)^2\sigma_\epsilon^2 & \underbrace{- \frac{2\rho(1-\gamma)^2}{\phi+\rho-1}\sigma_\epsilon^2}_{technology} \\
& \underbrace{+ \frac{(1-\gamma)^2}{\phi+\rho-1}\left[2\rho cov(\chi,\epsilon|o) + (\phi-\rho-1)[\sigma_{\epsilon|o}^2 - \sigma_\epsilon^2]\right]}_{selection} + \underbrace{(1-\rho^2)\sigma_v^2}_{noise},
\end{aligned}
$$

and

$$
\begin{aligned}
elas(TFP, wedge) = \rho & \underbrace{- \frac{(1-\gamma)^2}{\phi+\rho-1}\frac{\sigma_\epsilon^2}{var(TFP)}}_{technology} \\
& \underbrace{+ \frac{(1-\gamma)^2}{\phi+\rho-1}\left[\frac{cov(\chi,\epsilon|o)}{var(TFP)} + \frac{(\sigma_\epsilon^2 - \sigma_{\epsilon|o}^2)}{var(TFP)}\right]}_{selection} + \underbrace{(1-\rho)\frac{\sigma_v^2}{var(TFP)}}_{noise}.
\end{aligned}
$$

where the expressions assume that $var(TFP)$ is fixed as this is observable in the data. The expressions highlight how each model mechanism impacts the relationship between the model parameters σ_ϵ and ρ governing the extent of distortions and the measured moments. In both expressions, the first terms are what would be inferred from the data moments in the absence of the the technology, selection, and noise channels in the model. In this case, the data moments directly correspond to the fundamental dispersion—either through the random $(1-\gamma)^2\sigma_\epsilon^2$ or systematic $\rho^2 var(TFP)$ component of distortions—and elasticity ρ of distortions. Technology creates a negative relationship between distortions and productivity implying that distortions need to be more dispersed and elastic to match the measured dispersion and elasticity of wedges. Selection has an ambiguous impact on the wedge dispersion and increases the measured elasticity of wedges as high distortion, low productivity firms select out of the economy, creating a positive relationship between productivity and wedges. Noise increases both the measured dispersion and elasticity of wedges as it creates correlated variation in measured TFP and wedges not explained by distortions. However, the impact of the noise channel is reduced by the elasticity ρ of distortions, as noise only matters to the extent that it is not captured by TFP. Figure 11 reports the decomposition.

Figure 11: Measured Distortions and Model Mechanisms

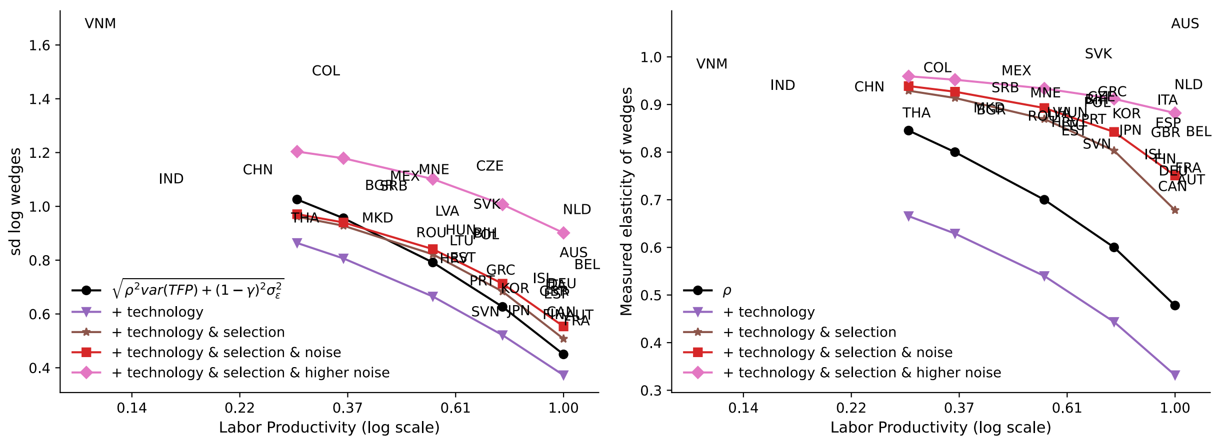

(a) Standard deviation of log wedges (b) Measured elasticity of wedges

The black line with circles reports the dispersion in distortions, $\rho^2\text{var}(\text{TFP}) + (1-\gamma)^2\sigma_\epsilon^2$, and the elasticity of wedges, $elas(\text{TFP}, \text{wedge})$, that would be observed in the absence of other model channels. The purple line with triangles, brown line with stars, red line with squares, and pink line with diamonds show how the measured moments are affected by, respectively: (1) technology choice, (2) selection into operation, (3) noise in the measure of output (set to the benchmark value), and (4) higher noise (set to the level observed in the Netherlands, a high-noise country). Country codes are data values for the indicated country. Aggregate labor productivity in logs is based on 2015 data from the Penn World Table (Feenstra et al., 2015).

The gap between the measured standard deviation and elasticity of distortions tends to be larger than the first term. Put differently, a model without technology, selection, or noise channels would tend to overestimate the severity of distortions and misallocation. The gap for the elasticity tends to be larger in higher income countries. This implies larger cross-country differences in the severity of distortions and misallocation as measured through their systematic component than the data would imply without accounting for these channels. Mechanically, selection and noise drives measured distortions to be larger. Selection is important in explaining why the measured elasticity of distortions is larger than the model elasticity ρ in higher-income countries. Selection, which is stronger in higher-income countries, creates a positive relationship between the productivities and the random distortion ϵ of observed firms (i.e., those that operate), creating a positive bias in the measured elasticity of distortions.

4.3 Decomposing Productivity Losses

The previous experiments highlight that the model can account for key cross-country moments, that there are substantial losses stemming from distortions, and that the calibrated elasticity of distortions ρ accounts for the bulk of cross-country differences, even after correcting for measurement issues. We now use the France calibration (benchmark) economy to provide insights into the sources of productivity losses.

We focus on increasing the elasticity ρ from the benchmark value of 0.48 to 0.60, 0.70 and 0.85, consistent with the implied cross-country range of measured elasticities. We also report an economy with $\rho = 0$ for reference, although we note that this economy is not an undistorted economy since the random component of distortions σ_ϵ is fixed at its calibrated value in the benchmark economy in all the experiments.

Table 3 reports the results. The top row documents aggregate labor productivity in each economy relative to the benchmark $\rho = 0.48$ economy. Increasing ρ from 0.48 to 0.85 reduces aggregate productivity substantially by 67 percent, from a normalized value of 1.00 in the benchmark economy to 0.33 in the $\rho = 0.85$ economy. Put differently, if the $\rho = 0.85$ economy were to implement policy reforms to reduce ρ to 0.48 as in the benchmark economy, aggregate productivity would increase in this economy by 3-fold (an increase of 200 percent).

We decompose the change in aggregate labor productivity into its sources and to relate with existing approaches in the literature. We start with the most narrow form of misallocation: static misallocation (second row, Table 3) that measures the effect of higher ρ on

Table 3: Experiments with Alternative ρ Values

	Value of ρ				
	0.00	0.48	0.60	0.70	0.85
Aggregate labor productivity	1.36	1.00	0.79	0.60	0.33
A. Static versus dynamic misallocation					
Static misallocation	1.05	1.00	0.91	0.80	0.61
Contribution (%)	*16*	–	*42*	*44*	*44*
Dynamic misallocation					
Firm-level productivity	1.26	1.00	0.90	0.80	0.63
Contribution (%)	*75*	–	*47*	*44*	*42*
Firm productivity with distortions	1.03	1.00	0.98	0.94	0.86
Contribution (%)	*9*	–	*10*	*12*	*14*
Allocative efficiency	1.08	1.00	0.89	0.75	0.52
Contribution (%)	*25*	–	*53*	*56*	*58*
B. Technology versus selection					
Technical efficiency	2.18	1.00	0.78	0.63	0.44
Technology	1.30	1.00	0.90	0.80	0.64
Contribution (%)	*34*	–	*44*	*47*	*54*
Selection	1.68	1.00	0.87	0.78	0.69
Contribution (%)	*66*	–	*56*	*53*	*46*

Notes: Static misallocation is aggregate output with distortions when operating firms i and technologies z_i are held fixed at the benchmark economy. Firm-level productivity is aggregate output in the efficient allocation in each economy. Allocative efficiency is aggregate output relative to efficient aggregate output. Technical efficiency is the average of firm-level productivity in the efficient allocation. Technology is technical efficiency keeping the operation decisions of firms constant at the benchmark economy, whereas Selection is calculated as a residual from technical efficiency and technology.

aggregate output when the set of producers and technologies are fixed to match the benchmark economy. This type of misallocation is the focus in Restuccia and Rogerson (2008) and a large quantitative literature. Static misallocation generates a reduction in aggregate productivity of 39 percent in the $\rho = 0.85$ economy, or around 44 percent ($= \log(0.61)/\log(0.33)$) of the total loss in aggregate productivity. The remaining 56 percent results from changes in firm-level productivities, implying that the dynamic channels of selection and technology increase the productivity loss from static misallocation by a factor of 2.2 times

($\approx \log(0.33)/\log(0.61)$) (consistent with Ayerst, 2025; Ayerst et al., 2023). The change in the firm-level productivity distribution (third row, Table 3) reduces aggregate productivity by 37 percent, accounting for 42 percent ($\log 0.63/\log 0.33$) of the loss in aggregate productivity in the $\rho = 0.85$ economy. The remaining 14 percent is an interaction of changes in firm-level productivity with increased distortions (fourth row, Table 3). Intuitively, distortions that compress employment among firms have larger effects on aggregate output when there are more high-productivity firms to distort.

Table 3 also reports allocative efficiency (fifth row) that measures aggregate output divided by the efficient allocation aggregate output holding the set of operating firms and technologies in each economy fixed. The inverse of allocative efficiency is the efficiency gain used by Hsieh and Klenow (2009). In contrast to static misallocation, allocative efficiency calculates the output losses given the counterfactual firm distribution and not the benchmark firm distribution. Consistent with the substantial changes in the firm-level productivity distribution, allocative efficiency drops much more than static misallocation, to around 0.52 in the $\rho = 0.85$ economy, representing about 58 percent ($= \log(0.52)/\log(0.33)$) of the aggregate output loss. This finding helps rationalize why quantitative analyses of misallocation when the producer productivity distribution is constant lead to smaller aggregate losses than in empirical analyses of allocative efficiency (Restuccia and Rogerson, 2017). Since static misallocation accounts for about 44 percent of the aggregate productivity loss and allocative efficiency for 58 percent, the static component is three-fourths of the loss attributed to allocative efficiency, with the remaining one-fourth due to the change in the productivity distribution.

We also decompose the contributions of the selection channel (operation decision of firms) versus the technology channel (productivity-enhancing investment decision of firms) to aggregate productivity. To do so, we calculate technical efficiency in each economy, defined as the average of firm-level productivity. We measure the contribution of the technology channel as the change in technical efficiency when the set of operating firms is the same as in the benchmark economy. This is equivalent to setting the operating decisions $o(\chi,)$ to the benchmark economy. The remaining change in technical efficiency is due to the selection channel. We report this counterfactual in Panel B of Table 3. We find that the effect on technical efficiency is roughly equally shared between selection and technology channels. However, note that for smaller values of ρ the role of technology differences is smaller and hence the contribution of selection is larger.

To summarize, plausible variations in the elasticity of distortions ρ from the benchmark economy generate implications on misallocation and firm-level TFP distributions consistent with the cross-country data. Additionally, the aggregate output loss is 2.2-fold larger than from static misallocation alone.

5 Conclusions

We examine micro and macro productivity differences across nations using cross-country firm-level data and a quantitative model of misallocation featuring operation decisions (selection) and productivity-enhancing investments (technology). Empirically, we find that less developed countries feature higher distortions and larger dispersion in firm-level productivity, mostly resulting from the prevalence of unproductive firms compared to developed countries. Quantitatively, we find that the aggregate productivity cost of distortions extends beyond static misallocation. Measured distortions in the form of higher productivity elasticity of distortions generate large aggregate productivity losses, over half of which are accounted for by changes in the productivity distribution. About one-fourth of the change in allocative efficiency is attributed to the change in misallocation due to changes in the productivity distribution.

Our analysis connects policies and institutions that underlie static misallocation with cross-country differences in the firm-level productivity distribution. The analysis highlights the importance of the technology and selection channel, which may both amplify the costs of misallocation and bias summary measures of misallocation. Further work would benefit from examining within-country institutional differences (e.g., Ayerst et al., 2023) or reform periods (e.g., Bustos, 2011) that can tightly link the technology and selection channels with distortions. While the technology channel has received more attention, the selection channel is relatively underexplored. Lessons from linking distortions with the productivity distribution may also inform further work on specific drivers of misallocation or firm dynamics. In particular, understanding how distortions impact specific firm investments—e.g., management practices, technology adoption, multi-national FDI—are important paths forward.

References

Acemoglu, D., Akcigit, U., Alp, H., Bloom, N., and Kerr, W. (2018). Innovation, Reallocation and Growth. *American Economic Review*, 126(4):1374–1443.

Akcigit, U., Akgunduz, Y. E., Alp, H., Cilasun, S. M., and Quintero, J. M. (2022). Cost of Size-dependent Regulations: The Role of Informality and Firm Heterogeneity. Technical report, University of Chicago.

Alviarez, V., Cravino, J., and Ramondo, N. (2023). Firm-embedded productivity and cross-country income differences. *Journal of Political Economy*, 131(9):2289–2327.

Andrews, D., Criscuolo, C., and Gal, P. N. (2015). Frontier Firms, Rechnology Diffusion and Public Policy: Micro Evidence from OECD Countries.

Ayerst, S. (2025). Distorted technology adoption. *The Economic Journal*, 135(668):1167–1190.

Ayerst, S., Brandt, L., and Restuccia, D. (2023). Distortions, Producer Dynamics, and Aggregate Productivity: A General Equilibrium Analysis. Technical report, National Bureau of Economic Research.

Bartelsman, E., Haltiwanger, J., and Scarpetta, S. (2013). Cross-Country Differences in Productivity: The Role of Allocation and Selection. *American Economic Review*, 103(1):305–334.

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

Bhattacharya, D., Guner, N., and Ventura, G. (2013). Distortions, Endogenous Managerial Skills and Productivity Differences. *Review of Economic Dynamics*.

Boar, C., Gorea, D., and Midrigan, V. (2022). Why are returns to private business wealth so dispersed? Technical report, National Bureau of Economic Research.

Buera, F., Hopenhayn, H., Shin, Y., and Trachter, N. (2023). Big Push in Distorted Economies. Technical report.

Buera, F. J., Moll, B., and Shin, Y. (2013). Well-intended policies. *Review of Economic Dynamics*, 16(1):216–230.

Bustos, P. (2011). Trade liberalization, exports, and technology upgrading: Evidence on the impact of mercosur on argentinian firms. *American Economic Review*, 101(1):304–340.

Comin, D. and Mestieri, M. (2018). If Technology has Arrived Everywhere, Why Has Income Diverged? *American Economic Journal: Macroeconomics*, 10(3).

Davis, S. J., Haltiwanger, J. C., Schuh, S., et al. (1998). Job creation and destruction. *MIT Press Books*, 1.

Fattal-Jaef, R. (2022). Entry Barriers, Idiosyncratic Distortions, and the Firm Size Distribution. *American Economic Journal: Macroeconomics*, 14(2):416–68.

Feenstra, R. C., Inklaar, R., and Timmer, M. P. (2015). The next generation of the penn world table. *American Economic Review*, 105(10):3150–3182.

Gal, P. N. (2013). Measuring Total Factor Productivity at the Firm Level using OECD-ORBIS.

Guner, N., Ventura, G., and Xu, Y. (2008). Macroeconomic Implications of Size-Dependent Policies. *Review of economic Dynamics*, 11(4):721–744.

Hall, R. E. and Jones, C. I. (1999). Why Do Some Countries Produce So Much More Output Per Worker Than Others? *The Quarterly Journal of Economics*, 114(1):83–116.

Hopenhayn, H. and Rogerson, R. (1993). Job turnover and policy evaluation: A general equilibrium analysis. *Journal of political Economy*, 101(5):915–938.

Hopenhayn, H. A. (1992). Entry, Exit, and Firm Dynamics in Long Run Equilibrium. *Econometrica: Journal of the Econometric Society*, pages 1127–1150.

Hopenhayn, H. A. (2014). Firms, misallocation, and aggregate productivity: A review. *Annu. Rev. Econ.*, 6(1):735–770.

Hsieh, C.-T. and Klenow, P. J. (2009). Misallocation and Manufacturing TFP in China and India. *The Quarterly Journal of Economics*, 124(4):1403–1448.

Hsieh, C.-T. and Klenow, P. J. (2014). The Life Cycle of Plants in India and Mexico. *The Quarterly Journal of Economics*, 129(3):1035–1084.

Kalemli-Özcan, Ş., Sørensen, B. E., Villegas-Sanchez, C., Volosovych, V., and Yeşiltaş, S. (2024). How to construct nationally representative firm-level data from the orbis global database: New facts on smes and aggregate implications for industry concentration. *American Economic Journal: Macroeconomics*, 16(2):353–374.

Khandelwal, A. K., Schott, P. K., and Wei, S.-J. (2013). Trade Liberalization and Embedded Institutional Reform: Evidence from Chinese Exporters. *American Economic Review*, 103(6).

Klenow, P. J. and Rodriguez-Clare, A. (1997). The Neoclassical Revival in Growth Economics: Has It Gone Too Far? *NBER Macroeconomics Annual*, 12:73–103.

Majerovitz, J. (2023). Misallocation and the Selection Channel. Technical report.

Nguyen, D. (2025). Openness to Foreign Firms, Industrialization and Aggregate Growth. Technical report, Amherst College.

Parente, S. L. and Prescott, E. C. (1994). Barriers to Technology Adoption. *Journal of Political Economy*, 102(2):298–321.

Pavcnik, N. (2002). Trade Liberalization, Exit, and Productivity Improvements: Evidence from Chilean Plants. *The Review of Economic Studies*, 69(1):245–276.

Poschke, M. (2018). The firm size distribution across countries and skill-biased change in entrepreneurial technology. *American Economic Journal: Macroeconomics*, 10(3):1–41.

Prescott, E. C. (1998). Lawrence R. Klein Lecture 1997: Needed: A Theory of Total Factor Productivity. *International Economic Review*, pages 525–551.

Restuccia, D. (2019). Misallocation and aggregate prodcutivity across time and space. *Canadian Journal of Economics*, 52(1):5–32.

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

Restuccia, D. and Rogerson, R. (2017). The Causes and Costs of Misallocation. *Journal of Economic Perspectives*, 31(3):151–74.

Yang, M.-J. (2021). Micro-level Misallocation and Selection. *American Economic Journal: Macroeconomics*, 13(4):341–368.

Online Appendix

A Data Details

We describe the details of the construction of the final dataset.

A.1 Variables, Sample Selection, and Data Cleaning

We use the Historical Product of the Orbis data that provides detailed firm-level financial data for a number of countries. Kalemli-Özcan et al. (2024) discuss the advantages and disadvantages of the Orbis dataset and provide a comparison of aggregate outcomes with national statistics.

Our final dataset is constructed over the period 2000 to 2019 and observations are at the firm-year (i, t) level. In most countries, the number of observations increases substantially around 2000 and starts to decline in more recent periods. We additionally drop 2020 and later periods to avoid including the COVID-19 pandemic. We define a sector (denoted by s) as the firm's primary two-digit SIC code.

Variables and cleaning. The main variables used in our analysis are constructed as follows:

- **Output.** Our baseline measure of output $y_{i,t}$ is set equal to the reported operating revenues, or sales when operating revenues is unavailable. We use this measure instead value added since material costs are not systematically available in most non-European countries in our dataset. Value added is constructed as operating revenues or sales minus the material costs. Value added is used in the data cleaning and in the robustness exercises as an alternative measure of output.

- **Employment** We measure employment $n_{i,t}$ as the reported count of employees at the firm. We also use capital $k_{i,t}$ to construct an alternative measure of firm-level TFP in a robustness analysis. The wage rate $w_{i,t}$ is constructed as the wage bill divided by the employment count ($w_{i,t} = W_{i,t}/n_{i,it}$). We construct the average wage bill of firms within a sector year as $\bar{w}_{s,t} = \sum_{i \in \mathcal{I}_{s,t}} W_{i,t}/n_{i,t}$ for firms that report both the wage bill $W_{i,t}$ and employment, where $\mathcal{I}_{s,t}$ is the set of firms in sector s and year t. We construct an alternative wage-bill-implied employment using the average sectoral wage rate as $n_{i,t}^{alt} = W_{i,t}/\bar{w}_{s,t}$. We replace employment for the top and bottom one percent of the employee count distribution, as these are most likely driven by misreporting.

- **Capital.** Firm capital $k_{i,t}$ is set equal to the total book value of firm fixed assets. Capital is used in the robustness exercises to construct an alternative measure of firm-level productivity.

Dropped observations. We drop observations based on the following criteria:

- **Missing data.** We drop firm-year observations without output (operating revenue or sales) or employment count data, needed to construct the baseline measures of productivity and wedges.

- **Inactive firms.** We exclude firms that are listed with unknown status, in bankruptcy, dissolved, in liquidation, or inactive in the current period and each following period.

- **Sectors.** We focus on the manufacturing sector and only include firms with four-digit NACE code between 1000 and 3300. We also exclude firms identified as a non-corporate identity (e.g., bank).

- **Data trimming.** We trim the remaining observations based on the following criteria. (i) The top and bottom 0.1% of firm observations within a year based on output. (ii) Firms with more than 100,000 employees. (iii) For countries with sufficient data, firms with a reported labor share (labor cost divided by value added) of more than one or in the bottom 1% of the sample. For other countries, firms with a reported cost of good sold to output or more than one or less than 0.1. (iv) Firms in the top or bottom 1%

of cost of goods sold to employment. (v) Firms in the top or bottom 2% based of the productivity and wedge distribution after removing year and sector differences. (vi) Firms with a capital-output or capital-labor ratio that are in the top or bottom 1% of the sample (only for robustness results that include capital).

Multiple observations. Many firms report multiple filings within a year for various reasons. We remove multiple observations based on:

- **Consolidated financial records**. Firms may report financial records for either unconsolidated, consolidated, or both. In the case where both are reported, we default to using the unconsolidated records.
- **Filing type**. Firms may report financial records as "annual reports" or "local registry filings". In the case where both are reported, we use the annual reports.
- **Other duplicates**. Other instances of firms reporting multiple filings are relatively rare and for the most part represent duplicated data. For these duplicates, we choose between the maximum and minimum observed values based on which values minimize the absolute error with output and employment in the previous period.

Time and sector trends. We regress measures of productivity and wedges on country-by-year-by-sector fixed effects and report summary statistics for the residualized variable. That is, for variable $\tilde{X}_{f,t}$ we estimate $\log \tilde{X}_{f,t} = \Gamma_{s,t} + \log X_{f,t}$ and then construct the residualized variable as $\log X_{i,t} = \log \tilde{X}_{i,t} - \Gamma_{s,t}$.

A.2 Firm Weights

An issue with the Orbis data is that it tends to over-sample large firms and under-sample small firms in some countries. Kalemli-Özcan et al. (2024) report that European countries tend to reflect the size distribution of firms reported in Eurostat (based on national statistics).

However, less is known about the coverage outside of European countries. We construct firm weights using national statistics to allow us to re-weight the data to match the actual distribution. We show that our results are robust to this re-weighting in Appendix B.

We denote our final firm weights as $\omega_{n,t}$ and denote $h^{D}_{[\underline{n},\overline{n}],c,t}$ as the share of firms with between $\underline{n}$ and $\overline{n}$ employees from dataset D. For example, the share of firms in France in 2013 with between 10 and 19 employees in the Orbis dataset is denoted by $h^{Orbis}_{[10,19],FR,2013}$. We construct the final firm weights as

$$\omega_{[\underline{n},\overline{n}],c,t} = \frac{h^{D}_{[\underline{n},\overline{n}],c,t}}{h^{Orbis}_{[\underline{n},\overline{n}],c,t}}.$$

We discuss the construction of $h^{D}_{[\underline{n},\overline{n}],c,t}$ using nationally representative data below.

Eurostat data. We use the distribution of firms by employment size as reported by Eurostat in the business demography (BD) and structural business statistics (SBS) datasets to construct observation weights. The BD dataset reports more granular data for smaller business sizes and separates non-employer businesses. The BD dataset reports firms in employment bins {0,1-4,5-9,10+}. The SBS dataset is more granular at higher employment levels but lumps non-employers into the smallest size bin. The SBS dataset reports firms in employment bins {0-9,10-19,20-49,50-249,250+}.

We exclude non-employer businesses from the final dataset. We construct the final Eurostat bins as:

$$
\begin{aligned}
h^{ES}_{[1,4],c,t} &= \frac{h^{SBS}_{[0,9],c,t} - h^{BD}_{0,c,t}}{1 - h^{BD}_{0,c,t}} \times \frac{h^{BD}_{[1,4],c,t}}{h^{BD}_{[1,4],c,t} + h^{BD}_{[5,9],c,t}}, \\
h^{ES}_{[5,9],c,t} &= \frac{h^{SBS}_{[0,9],c,t} - h^{BD}_{0,c,t}}{1 - h^{BD}_{0,c,t}} \times \frac{h^{BD}_{[5,9],c,t}}{h^{BD}_{[1,4],c,t} + h^{BD}_{[5,9],c,t}}, \\
h^{ES}_{[10,19],c,t} &= \frac{h^{SBS}_{[10,19],c,t}}{1 - h^{BD}_{0,c,t}}, \\
h^{ES}_{[20,49],c,t} &= \frac{h^{SBS}_{[20,49],c,t}}{1 - h^{BD}_{0,c,t}}, \\
h^{ES}_{[50,249],c,t} &= \frac{h^{SBS}_{[20,49],c,t}}{1 - h^{BD}_{0,c,t}}, \\
h^{ES}_{[250,\infty],c,t} &= \frac{h^{SBS}_{[250,\infty)}}{1 - h^{BD}_{0,c,t}}.
\end{aligned}
$$

We extend the firm shares $h^{ES}_{[\underline{n},\overline{n}],c,t}$ to earlier and late periods by assuming that firm shares are the same as in the closest period. For example, if the earliest period with sufficient data for Austria is 2005 then we assume the weights $h^{ES}_{[\underline{n},\overline{n}],AT,2005}$ also apply to the period 2000 to 2004. Firm shares tend to be relatively stable over time and this allows us to maximize the usable data. We also interpolate data missing in intermediate periods as a linear combination of the two surrounding periods.

OECD data. The OECD database reports the firm size distribution divided into either three or five size bins. The three size bin categories are {1-19,20-249,250+} and the five size bin categories are {1-9,10-19,20-49,50-249,250+}. We follow same procedure as with the Eurostat data to fill in missing data.

We also use the OECD data to construct weights for countries without alternative sources. For these countries, we first construct the expected share of firms in the 20-249 size bin by regressing $h_{[20,249],c,t} = \alpha \ln LP_{c,t} + F_t + \epsilon_{c,t}$, where $LP_{c,t}$ is aggregate labor productivity and F_t is a year fixed effect. The coefficient α captures the relationship between aggregate productivity and the size distribution across countries, where $\alpha > 0$ ($\alpha < 0$) implies that wealthier countries have more (fewer) firms in this size bin.

Individual country data. We supplement the above information with data on Vietnam, Mexico, and Korea. The Vietnam data, from the Vietnamese Statistical Yearbook, group firms into the size categories {1-4,5-9,10-49,50-199,200-299,300+} and are available from 2004 to 2019. The Mexico data group firms into the size categories {1-10,11-50,51-250,250+} and are available every five years between 2004 to 2019. The Korea data group firms into the size categories {1-4,5-9,10-49,50-99,100-199,200-299,300+} and are available from 2011 to 2019.

A.3 Overview of Final Dataset

Table A.1 reports the number of observations in the final dataset along with the source of the firm distribution used to construct firm weights. We drop countries without at least 5,000 observations after the previously described cleaning and trimming is done in order to reduce sample size issues. That said, countries have a wide range of observations in the final dataset from just over 5,000 in Montenegro and Thailand to almost 3 million observations in China.

B Empirical Analysis

We examine the robustness of the main empirical results to alternative constructions of productivity and wedges. In addition, for each version of the model, we compare outcomes with labor input (tfp^{lo}) and a Cobb-Douglas aggregate of capital and labor (tfp^{cd}).

- **Value added**: The baseline results use gross output to construct statistics since this improves the representation across countries. In the alternative model, we construct value added by subtracting material costs $m_{i,t}$ from revenues $s_{i,t}$. The measures of

Table A.1: Final Dataset

	Obs	sd log TFP	sd log Wedge	sd log Emp	Meas Elasticity
AUS	131,259	0.74	0.81	1.19	1.06
AUT	28,657	0.69	0.58	1.64	0.73
BEL	109,157	0.84	0.77	1.61	0.83
BGR	342,105	1.17	1.06	1.36	0.88
BIH	51,644	0.93	0.88	1.36	0.90
CAN	5,364	0.72	0.59	1.71	0.72
CHN	2,883,637	1.18	1.12	1.91	0.93
COL	21,592	1.53	1.49	0.97	0.97
CZE	237,561	1.21	1.13	1.52	0.91
DEU	268,178	0.85	0.70	1.64	0.75
ESP	1,219,398	0.73	0.66	1.10	0.85
EST	52,307	0.90	0.79	1.39	0.84
FIN	174,124	0.71	0.58	1.20	0.78
FRA	586,150	0.66	0.56	1.41	0.76
GBR	191,000	0.73	0.67	1.36	0.83
GRC	74,933	0.78	0.75	1.02	0.92
HRV	140,171	0.87	0.79	1.35	0.85
HUN	242,811	0.98	0.90	1.39	0.87
IND	12,351	1.11	1.09	1.63	0.93
ISL	5,681	0.85	0.72	1.53	0.79
ITA	1,701,564	0.72	0.69	1.29	0.90
JPN	542,260	0.65	0.60	1.53	0.84
KOR	531,262	0.73	0.68	1.19	0.87
LTU	55,830	0.97	0.86	1.37	0.85
LVA	60,948	1.07	0.97	1.34	0.87
MEX	29,953	1.11	1.10	1.22	0.96
MKD	46,456	1.02	0.94	1.40	0.88
MNE	5,235	1.19	1.12	1.37	0.91
NLD	18,294	0.96	0.97	1.70	0.93
POL	132,100	0.93	0.88	1.45	0.89
PRT	367,970	0.78	0.71	1.15	0.86
ROU	450,985	0.97	0.89	1.44	0.87
SRB	197,803	1.10	1.06	1.42	0.93
SVK	148,323	0.95	0.99	1.61	1.00
SVN	90,069	0.66	0.59	1.42	0.81
THA	5,771	1.02	0.94	1.65	0.87
VNM	174,893	1.68	1.66	1.38	0.98

productivity are then:

$$tfp_{i,t}^{lo} = \frac{s_{i,t} - m_{i,t}}{n_{i,t}^{\gamma}}, \qquad tfp_{i,t}^{cd} = \frac{s_{i,t} - m_{i,t}}{(k_{i,t}^{\alpha} n_{i,t}^{1-\alpha})^{\gamma}}.$$

- **Constant elasticity of substitution**: Hsieh and Klenow (2009) construct a model in which firms have constant returns to scale and face constant elasticity of substitution with other firms' goods. We examine our main results using this framework and also taking their calibrated elasticity of substitution $\sigma = 3$. In this version of the model, productivity can be constructed as

$$tfp_{i,t}^{lo} = \frac{(p_{i,t} y_{i,t})^{\frac{\sigma}{\sigma-1}}}{n_{i,t}}, \qquad tfp_{i,t}^{cd} = \frac{(p_{i,t} y_{i,t})^{\frac{\sigma}{\sigma-1}}}{k_{i,t}^{\alpha} n_{i,t}^{1-\alpha}}.$$

- **Population weighting**: The final version of the model that we report is identical to the baseline model but we weight the results by the population weights constructed in Appendix A.

Wedges are the same in each version of the model since wedges do not rely on the structure of the production function. We construct wedges based on the Cobb-Douglas inputs, the labor input, and the capital input. In the case where the distortion is on firm revenues, as opposed to factor inputs, these three wedges are theoretically equivalent. This is not necessarily the case when distortions tend to impact one factor more than the other, such as, if credit constraints limit capital inputs more than employment.

$$wedge_{i,t}^{y} = \frac{y_{i,t}}{k_{i,t}^{\alpha} \ell_{i,t}^{1-\alpha}}, \qquad wedge_{i,t}^{\ell} = \frac{y_{i,t}}{\ell_{i,t}}, \qquad wedge_{i,t}^{k} = \frac{y_{i,t}}{k_{i,t}},$$

where in the value added approach the numerator is revenues minus material costs.

B.1 Firm Productivity Distribution

Figure B.1 reports the standard deviation of firm productivity across the six different models described above. We find a similar decreasing pattern with aggregate labor productivity across all six models as in the main text. In the value added measures there are fewer lower income countries included in the sample because material costs are not reported in these countries. In general, the Cobb-Douglas production function tends to reduce some of the variation in productivity as it controls for cross-firm differences in capital intensity. The CES assumption on the production and demand functions implies higher productivity dispersion. While Hsieh and Klenow (2009) show that this model is isomorphic to our baseline model, the two models make different assumptions on the mapping of data to productivity, potentially explaining the differences. We find that weighting observations increases the measured dispersion of productivity, which could reflect larger dispersion in productivity of small, under-sampled firms.

Figure B.2 reports the p99 to p75 and p99 to p1 ratios for each of the six models. As in the baseline results, we find that the gap between the p99 and p1 firms increase more than the gaps between the p99 and p75 firms moving from higher to lower labor productivity countries. Similar to the previous set of figures, the magnitude of the results depends on the model used, with the CES model having a notably much larger gap in the percentile ratios.

B.2 Elasticity of Wedges

Figure B.3 reports the elasticity of distortions with respect to the constructed measure of firm-level productivity in each of the six models. As with the baseline results, we find that the higher income countries tend to have lower measured elasticity of distortions. We find that this moment varies by model specification and parameter choice with values ranging between 0.3 and 0.6 in the CES models and 0.8 and 1.1 in the weighted Cobb-Douglas model. This could also be due to the sources of bias having different impacts depending on the model (e.g., if capital is less accurately measured or due to the model-implied measures

Figure B.1: Standard Deviation of Firm Productivity

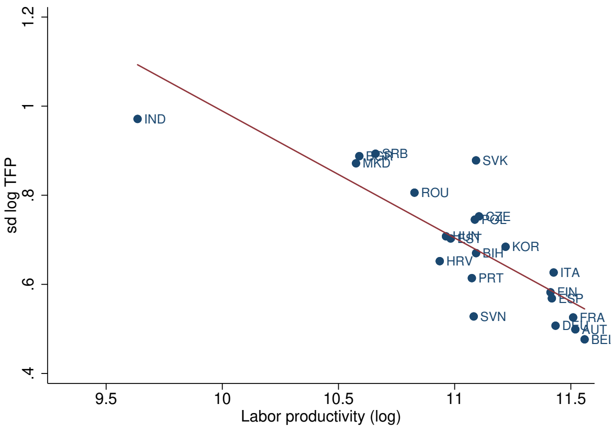

(a) Value added labor only tfp^{lo}

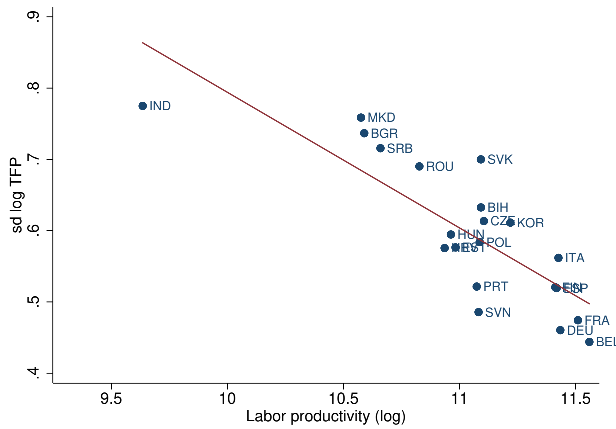

(b) Value added Cobb-Douglas tfp^{cd}

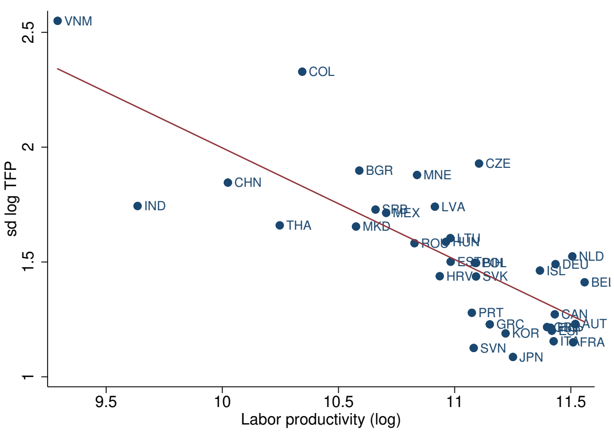

(c) CES labor only tfp^{lo}

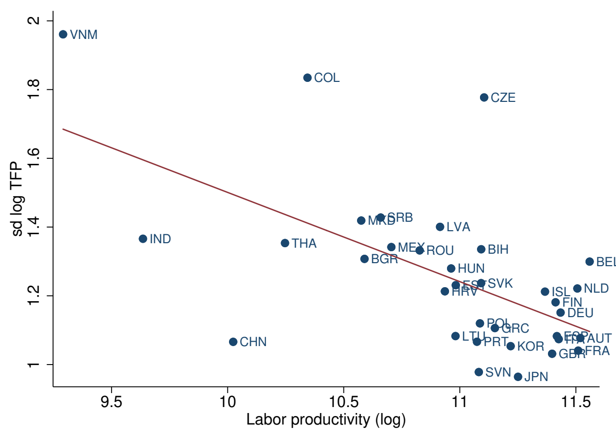

(d) CES Cobb-Douglas tfp^{cd}

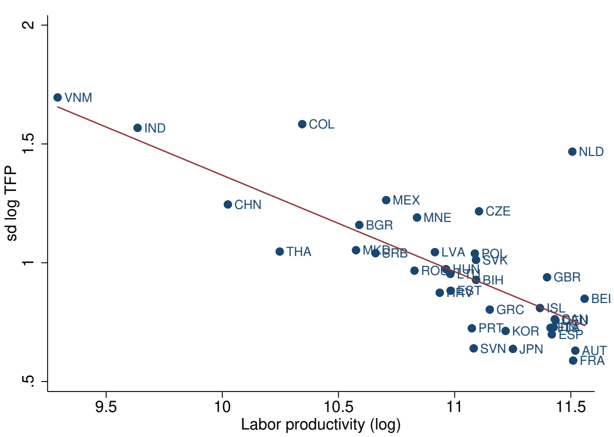

(e) Weighted labor only tfp^{lo}

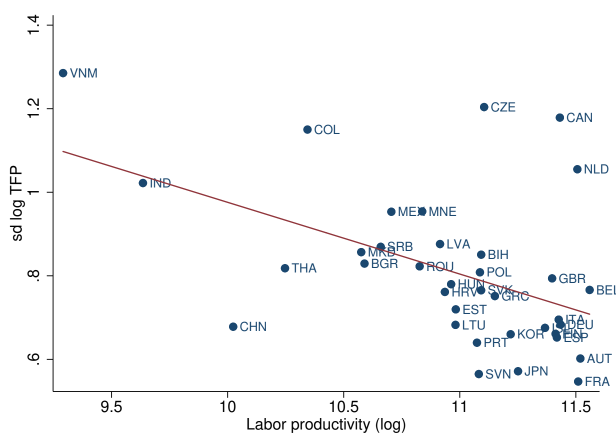

(f) Weighted Cobb-Douglas tfp^{cd}

Figure B.2: Firm Productivity Distribution

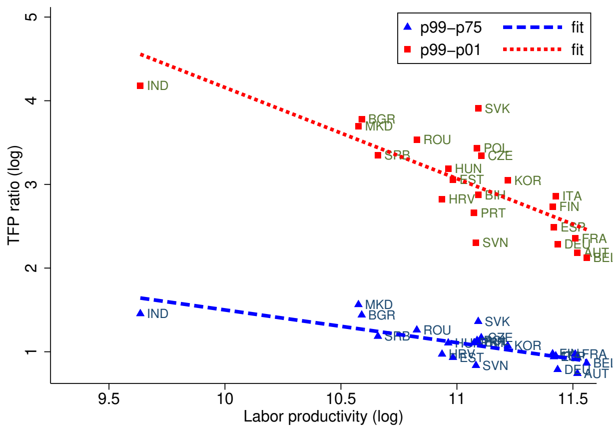

(a) Value added labor only tfp^{lo}

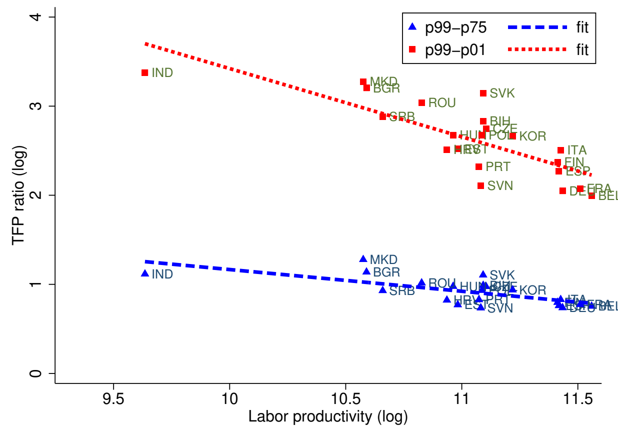

(b) Value added Cobb-Douglas tfp^{cd}

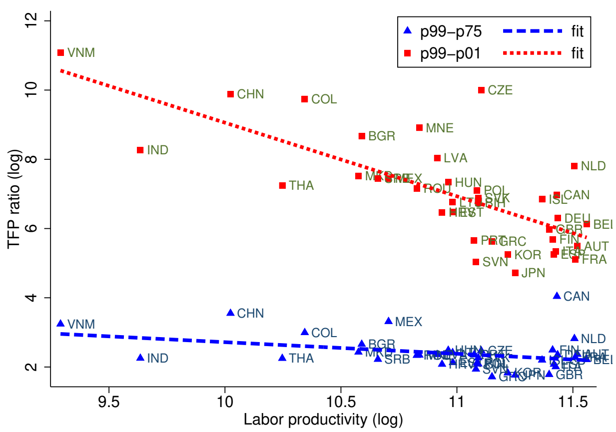

(c) CES labor only tfp^{lo}

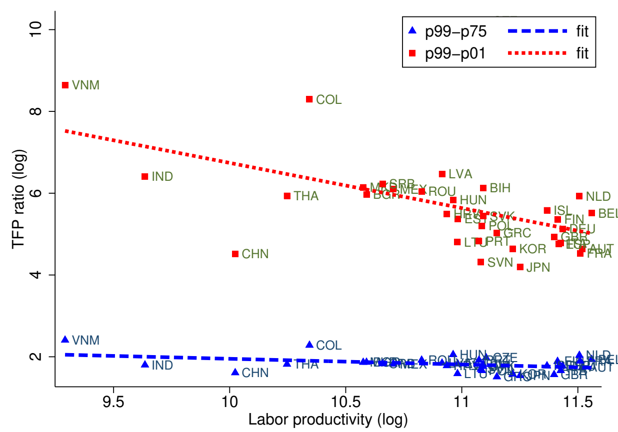

(d) CES Cobb-Douglas tfp^{cd}

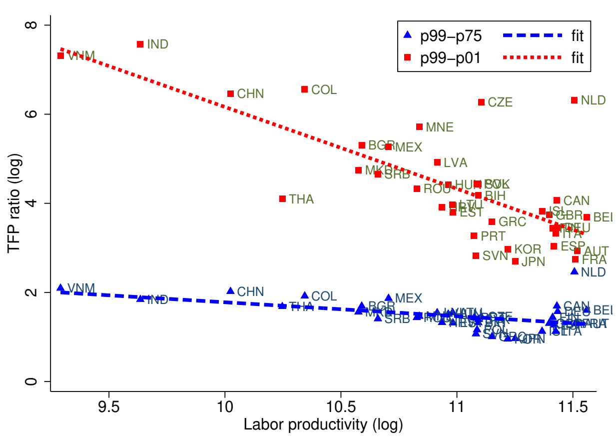

(e) Weighted labor only tfp^{lo}

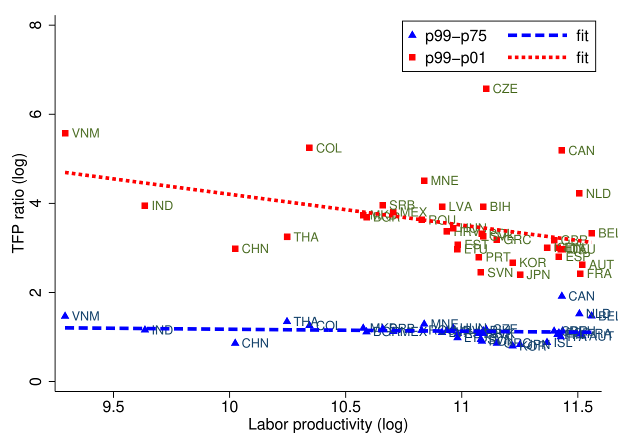

(f) Weighted Cobb-Douglas tfp^{cd}

Figure B.3: Elasticity of Wedges

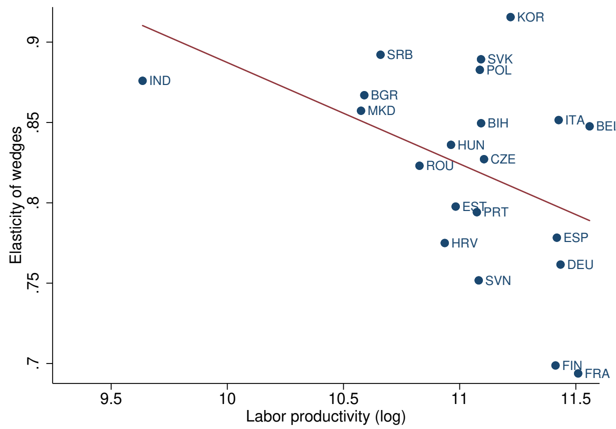

(a) Value added labor only tfp^{lo}

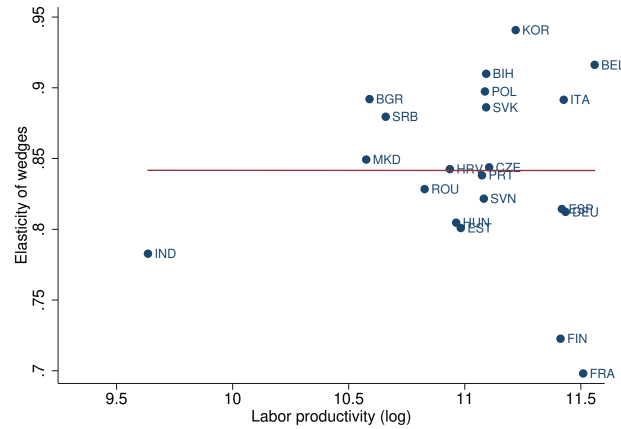

(b) Value added Cobb-Douglas tfp^{cd}

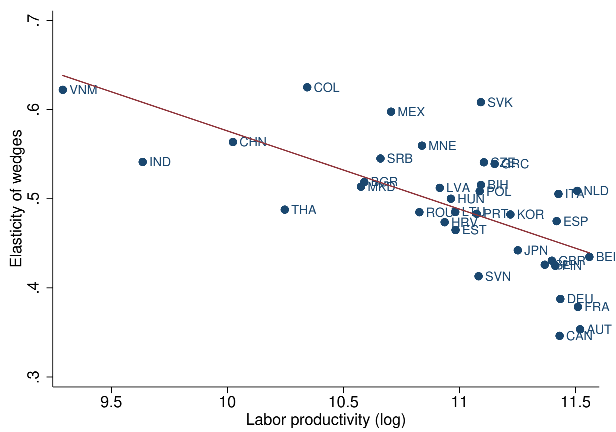

(c) CES labor only tfp^{lo}

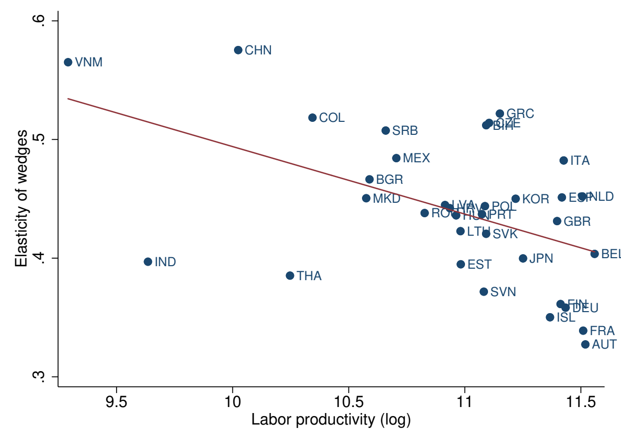

(d) CES Cobb-Douglas tfp^{cd}

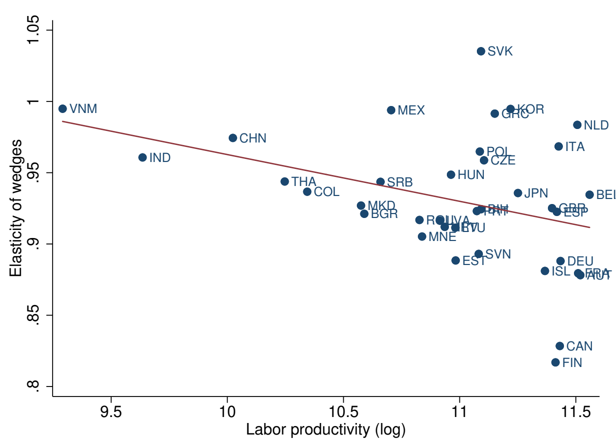

(e) Weighted labor only tfp^{lo}

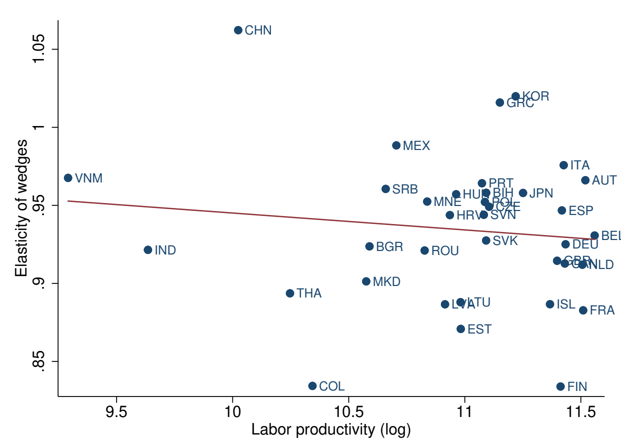

(f) Weighted Cobb-Douglas tfp^{cd}

of productivity). We also find similar elasticity as in Hsieh and Klenow (2009) for India and China using the same model and parameterization (Panel d).

B.3 Additional Empirical Results

Figure B.4 reports the inter-quartile and inter-decile ratios of productivity in each country and the standard deviation of employment. Panel (a) shows the the increased dispersion of productivity reported in Fact 1 is not being driven by outliers. Panel (b) highlights that despite more dispersed productivity and wedges–both driven dispersion of employment–in lower income countries, employment dispersion is only slightly higher, or even flat. the flatness of employment dispersion reflects that dispersion in productivity and wedges must have some negatively correlation in their impact on employment–e.g., productive firms receive lower wedges–consistent with the systematic component of distortions in the model.

Figure B.4: Dispersion in Firm-Level TFP and Employment

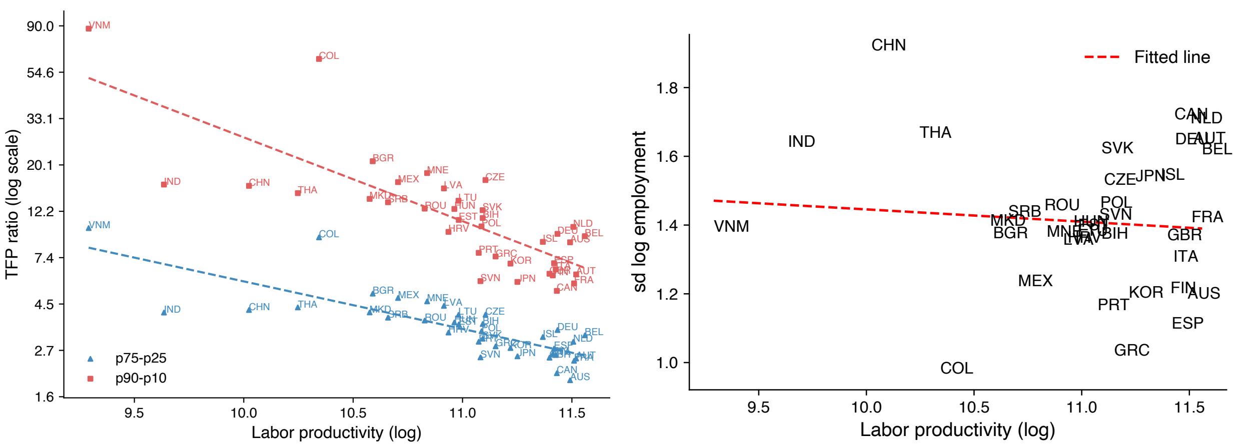

(a) TFP inter-quartile, decile ranges

(b) Standard deviation of employment

Figure B.5 reports two additional moments related to the elasticity of wedges to productivity. Panel (a) reports the elasticity of firm-level employment to productivity. In high income countries, employment is more positively related to productivity indicating high allocative efficiency of markets. Panel (b) reports the correlation of output and employment. The lower correlation implies that higher output firms tend to have lower employment in lower income countries. This is consistent with higher productive firms facing more obstacles in accumulating resources.

Figure B.5: Employment Elasticity and Correlation with Output

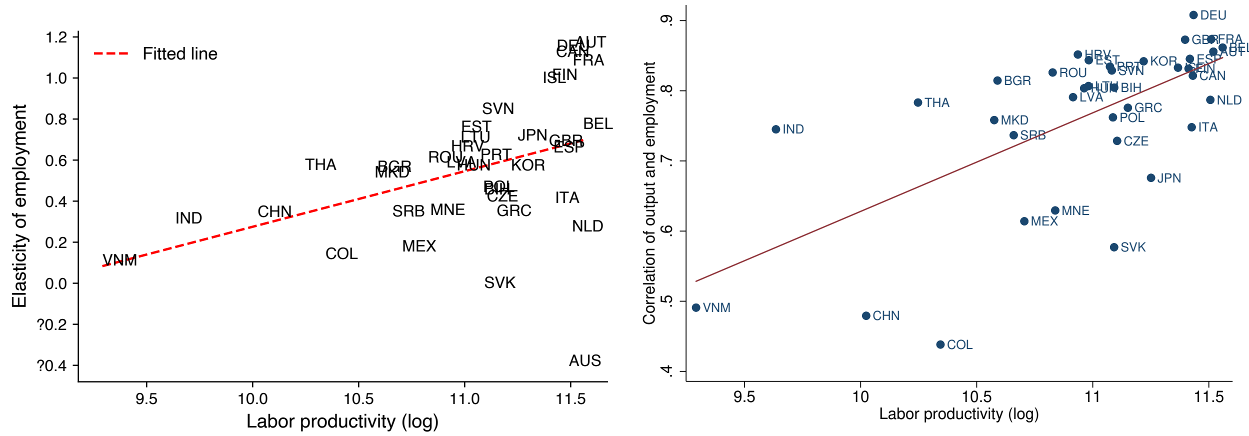

(a) Elasticity of employment

(b) Correlation of Output and Employment

B.4 Production Technologies in France and Vietnam

Our baseline results measure productivity and wedges using gross output rather than value added to accommodate that many countries do not report the data necessary to construct value added. We show in Appendix B that our main results hold using value added for the subset of countries with sufficient data. In this Appendix, we take a different approach by grouping firms by their relative value added and cost of goods sold shares of revenue. This allows us to compare the moments in France and Vietnam, the countries that represent the range in aggregate labor productivity in our quantitative analysis.

The concern with using gross output is that differences in production technologies may imply differences in the intermediate share of production. Higher intermediate share firms would appear relatively productive and distorted based our measures of productivity and wedges. This could bias our results if production technologies are systematically related to other firm fundamentals (e.g., productivity). To explore this possibility, we use a k-means clustering algorithm to construct five groups of firms based on the value added share of revenue in France and the cost of goods sold share of revenue in Vietnam. We use the cost of goods sold share in Vietnam because value added is not available for a large number of firms and the cost of goods sold should capture similar mechanisms. Similarly, cost of goods sold is not available for a large number of firms in France.

Table B.2: Data Moments Grouping Firms by Shares of Value Added and Cost of Goods Sold

A: France

	Meas Elasticity	sd log TFP	sd log Wedge	sd log Emp
baseline	0.76	0.66	0.56	1.41
g1	0.75	0.52	0.46	1.26
g2	0.71	0.60	0.49	1.44
g3	0.69	0.63	0.50	1.52
g4	0.68	0.61	0.49	1.50
g5	0.76	0.52	0.47	1.28

B: Vietnam

	Meas Elasticity	sd log TFP	sd log Wedge	sd log Emp
baseline	0.98	1.68	1.66	1.38
g1	0.99	1.58	1.59	1.36
g2	0.97	1.60	1.57	1.46
g3	0.96	1.65	1.61	1.38
g4	0.97	1.69	1.65	1.26
g5	0.99	1.61	1.61	1.13

Notes: Baseline moments refer to the pooled data including all firms in both countries. Moments for the groups g1-g5 are calculated using only the indicated group of firms. France groups are constructed using a k-means clustering algorithm with five groups on firm value added shares (revenues minus material costs divided by revenues). Vietnam groups are constructed using a k-means clustering algorithm with five groups on firm cost of goods sold share (cost of good sold divided by revenues).

Table B.2 reports moments for the five constructed groups along with the baseline results that uses the full final dataset. For both France and Vietnam, the moments appears relatively similar between the baseline and the five groups. While this does not rule out the possibility that underlying production technology differences drive some of the differences in the data, it is reassuring that grouping firms based on these closely related variables does not substantially alter the data moments that are the focus of our analysis.

C Model Details

We provide the proofs of propositions in the paper.

Proof of Proposition 1. Simply follows from the zero-profit entry condition.

Proof of Proposition 2. Productivity can be written as

$$\begin{aligned}\ln TFP_i &= (1-\gamma)\ln z_i + \ln v_i, \\ &= \frac{1-\gamma}{\phi+\rho-1}\left[\ln \chi_i + \ln \epsilon_i\right] + \ln v_i.\end{aligned} \tag{C.1}$$

Then, the standard deviation of ln TFP across firms is given by

$$\sigma_{TFP}^2 = (1-\gamma)^2 \sigma_{z|o}^2 + \sigma_v^2,$$

from equation (C.1) and the fact that $cov(z, v) = 0$. Define a variable $\ln \tilde{x}$ as $\ln x - \ln \bar{x}$ where $\ln \bar{x} = \mathbb{E} \ln x$. Then, going further

$$\begin{aligned}\sigma_{z|o}^2 &= \mathbb{E}\left[\frac{1}{\phi+\rho-1}(\ln \chi_i + \ln \epsilon_i) - \ln \bar{z} \mid o\right]^2, \\ &= \left(\frac{1}{\phi+\rho-1}\right)^2 \mathbb{E}\left[(\ln \tilde{\chi}_i + \ln \tilde{\epsilon}_i) \mid o\right]^2, \\ &= \left(\frac{1}{\phi+\rho-1}\right)^2 \mathbb{E}\left[(\ln \tilde{\chi}_i)^2 + (\ln \tilde{\epsilon}_i)^2 + \ln \tilde{\chi}_i \ln \tilde{\epsilon}_i \mid o\right], \\ &= \left(\frac{1}{\phi+\rho-1}\right)^2 \mathbb{E}\left[\sigma_{\chi|o}^2 + \sigma_{\epsilon|o}^2 + cov(\chi, \epsilon|o)\right].\end{aligned}$$

Substituting into the expression for σ_{TFP}^2 confirms the result in the main text:

$$\begin{aligned}var(\text{TFP}) &= (1-\gamma)^2 \sigma_{z|o}^2 + \sigma_v^2, \\ &= \left(\frac{1-\gamma}{\phi+\rho-1}\right)^2 (\sigma_{\chi|o}^2 + \sigma_{\epsilon|o}^2 + 2\ cov(\chi, \epsilon|o)) + \sigma_v^2.\end{aligned}$$

Proof of Proposition 3. The variance of the log wedge is given by

$$
\begin{aligned}
var(\log(wedge)) &= \mathbb{E}\left[-\rho(1-\gamma)\log\tilde{z} - (1-\gamma)\log\tilde{\epsilon} + \log v\right]^2 \\
&= \mathbb{E}[\rho^2(1-\gamma)^2(\log\tilde{z})^2 + (1-\gamma)^2(\log(\tilde{\epsilon})^2) + 2\rho(1-\gamma)^2(\log\tilde{z})(\log\tilde{\epsilon}) + (\log v)^2] \\
&= \rho^2(1-\gamma)^2\sigma^2_{z|o} + (1-\gamma)^2\sigma^2_{\epsilon|o} + 2\rho(1-\gamma)^2 cov(z,\epsilon|o) + \sigma^2_v
\end{aligned}
$$

The elasticity of distortions (wedge) with respect to firm-level productivity is given by

$$
elas(wedge_i, TFP_i) = \frac{\mathbb{E}[\ln \tilde{wedge}_i \ln \tilde{TFP}_i]}{\sigma^2_{TFP}}.
$$

The numerator is equal to

$$
\begin{aligned}
\mathbb{E}[\ln \tilde{wedge}_i \ln \tilde{TFP}_i] &= \mathbb{E}\left[(\ln\tilde{v}_i + \rho(1-\gamma)\ln\tilde{z}_i + (1-\gamma)\ln\tilde{\epsilon}_i)((1-\gamma)\ln\tilde{z}_i + \ln\tilde{v}_i)\right], \\
&= \mathbb{E}\begin{bmatrix} (1-\gamma)\ln\tilde{z}_i\ln\tilde{v}_i + \rho(1-\gamma)^2(\ln\tilde{z}_i)^2 - (1-\gamma)^2\ln\tilde{z}_i\ln\tilde{\epsilon}_i \\ +(\ln v_i)^2 + \rho(1-\gamma)\ln\tilde{z}_i\ln\tilde{v}_i - (1-\gamma)\ln\tilde{v}_i\ln\tilde{\epsilon}_i \end{bmatrix}, \\
&= \rho(1-\gamma)^2\sigma^2_{z|o} + \sigma^2_v - (1-\gamma)^2 cov(z,\epsilon|o).
\end{aligned}
$$

The last line follows from $\mathbb{E}\ln\tilde{z}_i\ln v_i = 0$ and $\mathbb{E}\ln v\ln\tilde{\epsilon}_i = 0$. Along with $\sigma^2_{TFP} = (1-\gamma)^2\sigma^2_{z|o} + \sigma^2_v$, the above expression confirms the result in the main text:

$$
elas(\text{TFP}, wedge) = \frac{\rho(1-\gamma)^2\sigma^2_{z|o} + \sigma^2_v - (1-\gamma)^2 cov(z,\epsilon|o)}{(1-\gamma)^2\sigma^2_{z|o} + \sigma^2_v}.
$$

D Quantitative Details

D.1 Other Calibration Parameters

Figure D.6 reports the calibrated values for each country of the dispersion in the transitory productivity σ_v, the dispersion in innovation ability σ_χ, and the fixed cost of operation c_f.

These variables are fixed at the values of the benchmark economy calibrated to France in the main experiments.

Figure D.6: Cross-Country Calibration Results

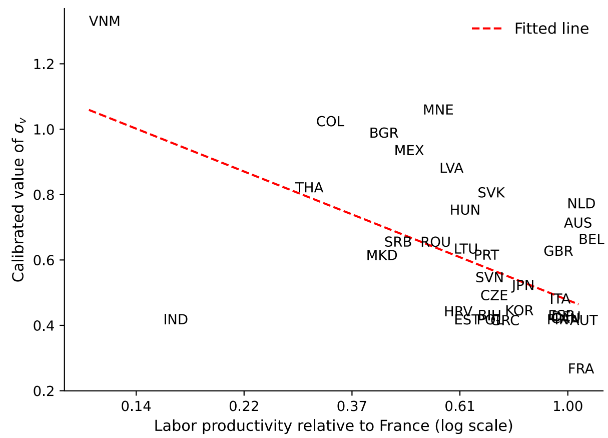

(a) Dispersion of transitory productivity σ_v

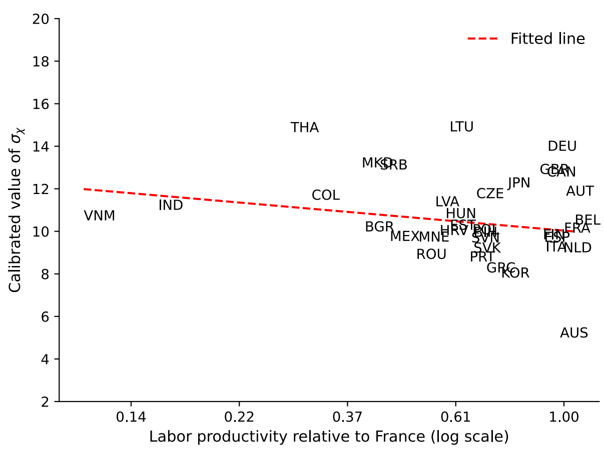

(b) Dispersion of innovation ability σ_χ

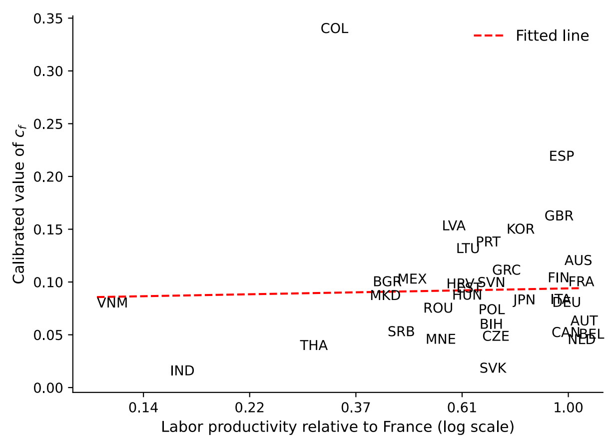

(c) Fixed cost of operation c_f

D.2 Innovation Costs in Units of Labor

We consider an alternative version of the model in which innovation costs are measured in units of labor, rather than units of output as in the main text. The problem for entering firms are as follows.

Entering firms. Entrants draw an idiosyncratic innovation ability χ_i from distribution $G(\chi)$ and a random distortion component ϵ_i from distribution $F(\epsilon)$. The firm then chooses productivity to maximize its incumbent firm value net of productivity investment cost:

$$V(\chi_i, \epsilon_i) = \max_{z \geq 0} \left[W(z, \tau(z, \epsilon_i)) - \psi \frac{z^\phi}{\chi_i} w \right],$$

where $W(z, \tau)$ is the value of an incumbent firm with productivity z and $\tau(z, \epsilon_i)$ is the distortion faced by the firm given the choice of z and the random component ϵ_i. We denote the optimal productivity level from this problem by the function $z(\chi, \epsilon)$. Even though there is an optimal productivity level associated with every type (χ, ϵ), only a fraction of firm types operate in the market.

Optimal productivity z for an entrant drawing (χ_i, ϵ_i) is given by:

$$z(\chi_i, \epsilon_i) = \left(\frac{(1-\rho)\tilde{\Omega}\chi_i}{\psi\phi\epsilon_i w} \right)^{\frac{1}{\phi+\rho-1}}, \quad \text{where} \quad \tilde{\Omega} \equiv \frac{\Omega}{1-R}.$$

Using this optimal productivity and substituting for the value of an incumbent firm, the value of an entrant firm drawing (χ_i, ϵ_i) is given by:

$$\begin{aligned} V(\chi_i, \epsilon_i) &= \max \left\{ \tilde{\Omega} z(\chi_i, \epsilon_i)^{1-\rho} \epsilon_i - \psi \frac{z(\chi_i, \epsilon_i)^\phi}{\chi_i} w - \frac{c_f w}{1-R}, 0 \right\}, \\ &= \max \left\{ \Gamma(w, \rho) \chi_i^{\frac{1-\rho}{\phi+\rho-1}} \epsilon_i^{-\frac{\phi}{\phi+\rho-1}} - \frac{c_f w}{1-R}, 0 \right\}, \end{aligned}$$

where

$$\Gamma(w, \rho) \equiv \frac{\phi+\rho-1}{\phi} \tilde{\Omega} \left(\frac{(1-\rho)\tilde{\Omega}}{\psi\phi w} \right)^{\frac{1-\rho}{\phi+\rho-1}}.$$

As firms only operate when their value is non-negative, the decision to operate for a firm drawing (χ_i, ϵ_i) can be characterized as:

$$o(\chi_i, \epsilon_i) = \begin{cases} 1 & \text{if} \quad \Gamma(w,\rho)\chi_i^{\frac{1-\rho}{\phi+\rho-1}}\epsilon_i^{-\frac{\phi}{\phi+\rho-1}} \geq \frac{c_f w}{1-R}, \\ 0 & \text{otherwise.} \end{cases}$$

Benchmark calibration. We find that this version does not affect the targeted moments in the model. The benchmark economy calibrated to the France data yields similar parameter values to the main text and is able to replicate the targeted moments.

Table D.3: Calibration of Distorted Benchmark Economy to France

Parameter	Value	Targeted moments	Model	Data
ρ	0.478	Measured elasticity of distortions	0.75	0.76
σ_ϵ	1.57	sd log distortions	0.55	0.56
σ_χ	9.94	sd log employment	1.40	1.40
σ_v	0.25	sd log TFP	0.67	0.66
c_f	0.10	Average firm size	14.8	14.9

Quantitative results. The micro-level moments are identical across the two versions. The alternative version also matches the key cross-country moments highlighted in the main analysis. We next compare the aggregate outcomes when varying the elasticity parameter ρ from the benchmark value of 0.48 to $0.60, 0.70$, and 0.85, consistent with the empirically observed range of cross-country elasticities.

Table D.4 reports results for both the baseline specification (innovation cost in units of output) and the alternative specification (innovation cost in units of labor). The top panel shows aggregate labor productivity relative to the benchmark economy with $\rho = 0.48$. Increasing ρ from 0.48 to 0.85 lowers aggregate productivity by about 67% in the baseline version and 65% in the alternative version. The two specifications also yield identical results for allocative efficiency across experiments. Overall, modeling innovation costs in labor rather than in output has limited impacts on the main quantitative results of the paper. Buera et al. (2023) find similar results when comparing investment costs in units of labor and

Table D.4: Experiments with Alternative ρ Values

	Value of ρ				
	0.00	0.48	0.60	0.70	0.85
Aggregate labor productivity					
Baseline	1.36	1.00	0.79	0.60	0.33
Alternative	1.23	1.00	0.82	0.64	0.35
Allocative efficiency					
Baseline	1.08	1.00	0.89	0.75	0.52
Alternative	1.08	1.00	0.89	0.75	0.52

Notes: Baseline is the model where cost of productivity investment is in units of output, whereas alternative refers to the model where the cost is in units of labor. Allocative efficiency is aggregate output relative to efficient aggregate output.

output between economies with estimated distortions in India and an undistorted economy calibrated to the US.

D.3 Technology and Selection

Figure D.7 presents the results from experiments decomposing the effects on the standard deviation of log TFP and log employment of the transitory productivity shocks v, technology choice $z(\chi, \epsilon)$, and selection $o(\chi, \epsilon)$. These experiments vary the value of ρ while keeping other parameters constant, including: (1) a baseline experiment with the transitory productivity shock v, changes in operating decision $o(\chi, \epsilon)$, and technology choice $z(\chi, \epsilon)$ across economies (Model w/ technology & selection & noise); (2) a baseline experiment but without transitory productivity shock v (Model w/ technology & selection); (3) an experiment without transitory productivity shock v, fixing operating decisions $o(\chi, \epsilon)$ to the benchmark economy values and changing only technology choices $z(\chi, \epsilon)$ across economies (Model w/ technology); and (4) an experiment without the transitory productivity shock v, fixing technology choice $z(\chi, \epsilon)$ to the benchmark economy values and changing only operating decisions $o(\chi, \epsilon)$ across economies (Model w/ selection).

The transitory productivity shock v does not generate important quantitative differences

Figure D.7: Effect of Productivity Shocks, Technology and Selection in Sample Moments

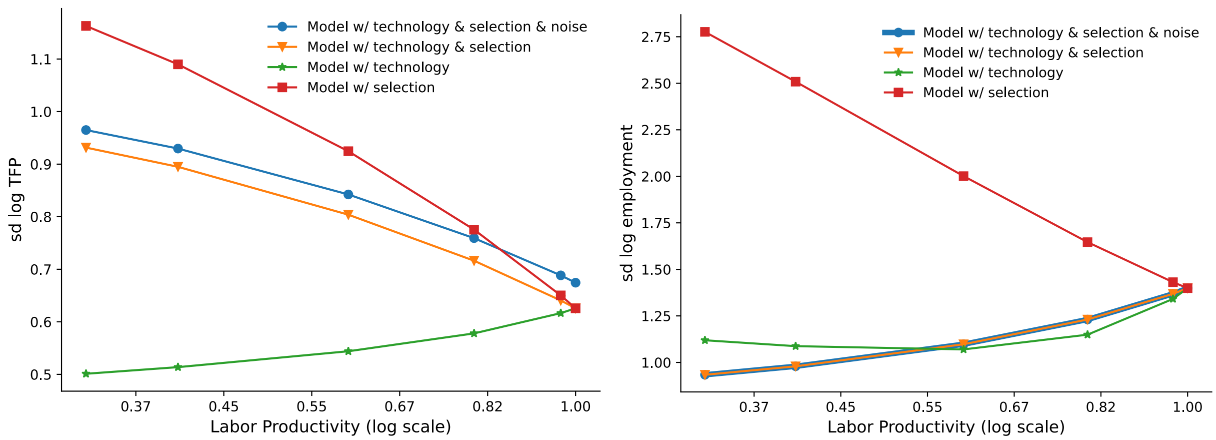

(a) Standard deviation of log TFP (b) Standard deviation of log employment

Notes: The circle-blue line represents the sample moments in the model with productivity shock v, changes in technology choice $z(\chi, \epsilon)$, and changes in operating decision $o(\chi, \epsilon)$. The triangle-orange line reports the same but assumes no productivity shock v. The star-green and square-red lines report the same (no v) but further assume no changes in technology choice $z(\chi, \epsilon)$ and in operating decision $o(\chi, \epsilon)$ relative to the benchmark economy, respectively.

in the pattern of the standard deviation of log TFP and have no impact on the standard deviation of log employment. However, the technology and selection channels produce opposite patterns in both the standard deviation of log TFP and log employment. While the experiment with changes in only the technology choice channel generates a positive relationship between the standard deviation of log TFP and aggregate productivity, contradicting the pattern in the data, the experiment with changes in only the selection channel produces patterns consistent with the data.

Nevertheless, changes in the selection channel create a strong negative relationship between the standard deviation of log employment and aggregate productivity, in contrast with the data pattern where there is no systematic relationship between the dispersion of log employment and aggregate productivity. The experiment with changes in technology choice effectively accounts for this dimension. These experiments highlight the important role of both technology and selection channels in generating the patterns in key micro moments consistent with the cross-country data.