Productivity Volatility and the Misallocation of Resources in Developing Economies*

John Asker, Allan Collard-Wexler and Jan De Loecker
NYU Stern & NBER, NYU Stern & NBER, and Princeton, NBER & CEPR

July 16, 2012

Abstract

We investigate the role of dynamic production inputs and their associated adjustment costs in shaping the dispersion of total factor productivity and static measures of capital misallocation within a country. Using firm-level production data sets from Chile, Colombia, India, Mexico, Slovenia, Ghana, Kenya, Tanzania and the World Bank's Enterprise Research Data (covering a further 33 countries), we find that countries exhibiting greater time-series volatility of productivity have greater cross-sectional dispersion of both productivity and marginal revenue product of capital. We use a standard model of investment with adjustment costs to show that variation in the volatility of productivity across these developing economies is sufficient to quantitatively replicate the cross-country variation in the dispersion of productivity and marginal revenue product of capital.

*JEL Code: D24, D92, O12. We would like to thank Dave Backus, for introducing us to the World Bank Data, and Nick Bloom, Gian Luca Clementi, John Fernald, Alessandro Gavazza, Chang-Tai Hsieh, Panle Jia Barwick, Pete Klenow, Marc Melitz, Richard Rogerson, Daniel Xu, and participants at the NBER SI Macro Productivity Workshop 2011, NYU Stern Micro Lunch, University of Minnesota, University of Wisconsin, Harvard University, University of Chicago, MIT and the CFSP (Financial Systems, Industrial Organization, and Economic Development) Workshop for their comments. Financial Assistance from the NYU-Stern Japan Center is greatly appreciated. Contact details: Allan Collard-Wexler (corresponding author), wexler@stern.nyu.edu; John Asker, jasker@stern.nyu.edu; and Jan De Loecker, jdeloeck@princeton.edu. The usual caveat applies.

1 Introduction

This paper considers the interpretation of well-documented cross-country differences in the dispersion of firm-level productivity, marginal products of inputs and other measures of performance.[1] Viewed through a static model, variation in marginal products across firms suggests some market distortion that impedes the efficient allocation of resources. The natural implication is that policies directed at reducing distortions or reallocating resources can realize significant welfare gains.

In this paper, we examine these cross-country differences in the dispersion of marginal products through the lens of a dynamic model. Specifically, we consider the dynamic optimization problem faced by firms that must choose capital stocks, which last for multiple periods, subject to adjustment costs. In our model, firms also face a productivity shock in each period that is determined by some known stochastic process. Importantly, in this model, there are no distortions in the output and input markets. We show that this model can explain, qualitatively and quantitatively, much of the cross-country variation in the dispersion of marginal products of capital, and of productivity. A literal implication is that resource allocation, while appearing inefficient in a static setting, may well be efficient in a dynamic sense. Clearly, we are not the first to make this point, but our paper goes beyond this by empirically quantifying the extent to which differences in dispersion can be generated from a dynamic model of investment.

This finding contributes to the discussion of the welfare implications of, and appropriate policy response to, dispersion in productivity and marginal products in developing countries. If the (admittedly extreme) view is taken that the stochastic process governing productivity shocks is exogenous (invariant to policy), then dispersion can be welfare-irrelevant, in the sense that firms appear to allocate resources efficiently given the shock process. A more balanced perspective would be that policies that endeavor to reduce the volatility of the process generating productivity shocks can be useful in increasing economic performance and welfare. That is, policies that seek to adjust the shock process for productivity may have some value. This contrasts with a static view, in which the notion of volatility cannot be discussed. Further, redistributive policies aimed at reducing dispersion in marginal products may not be welfare-improving in a dynamic setting, since dispersion by itself need not indicate any need for reallocation.

Our central contribution, then, is to highlight the usefulness of considering policies that influence this shock process of productivity, especially in developing coun-

[1] Marginal products, throughout the paper, should be understood as referring to the marginal revenue product of an input in a static model of production. To emphasize this, we will, at times, refer to them as "static marginal products."

tries. These policies complement policies aimed at easing any input market frictions that may exist.[2]

The starting point for our paper, and for much of the accompanying literature, is the fact that firms differ in performance — more specifically, total factor productivity (TFP), or simply productivity. Cross-sectional dispersion in firm-level productivity is even observed within narrowly defined industries.[3] Across countries, the extent of this dispersion varies considerably, particularly when comparing countries at different stages of economic development. A recent literature has considered the welfare effect of this TFP dispersion, and has tried to identify the degree of misallocation of resources from the variation in marginal products of inputs across producers. For example, Hsieh and Klenow (2009) find that if producers in the manufacturing sectors of India and China had the same degree of misallocation as the manufacturing sector in the United States, output would increase by thirty and sixty percent, respectively. Recently, largely spurred by this set of facts, a number of papers have tried to identify specific mechanisms to explain why TFP differences do not get eliminated by market-based resource reallocation.[4] Underscoring the potential macroeconomic gains from increased allocative efficiency, studies done at the industry level have shown that undoing misallocation can have first-order welfare effects. A well-known example is Olley and Pakes' (1996) study of productivity growth in the telecommunications equipment industry. They find that the reallocation of output to more-productive firms accounts for a large fraction of aggregate productivity growth.[5]

We begin by writing down a variant of a standard dynamic investment model in which firms: a) face costs when adjusting one factor of production (capital); b) can acquire all inputs in a frictionless spot market and; c) get a firm-specific productivity shock in each period generated by an AR(1) process. We show that, when firms are making decisions in this setting, dispersion in productivity and (more interestingly) in the marginal product of capital arises naturally.[6] In particular, we show that

[2]Given that a dynamic model with no input market distortions can not capture all the variation in the data, there is ample room for some of the dispersion in marginal products to arise due to input market distortions—that is, taking a dynamic view does not imply being blind to the potential for static input market frictions to distort allocations.

[3]See Bartelsman and Doms (2000), Bartelsman, Haltiwanger, and Scarpetta (2009) and references therein.

[4]See, for instance, Restuccia and Rogerson (2008), Collard-Wexler (2009), Midrigan and Xu (2009), and Moll (2010) for some recent work.

[5]Bartelsman, Haltiwanger, and Scarpetta (2009) rely on the reallocation measure introduced by Olley and Pakes (1996)—the covariance term between output and productivity and find it to play a key role in accounting for aggregate productivity growth across a wide range of countries.

[6]Midrigan and Xu (2009) use a similar dynamic model to investigate the role of capital frictions in explaining misallocation. They show that credit constraints alone, as opposed to other capital-adjustment

(in the range consistent with our data) as the volatility of productivity increases so does the the cross-sectional dispersion in productivity and the marginal product of capital.

We then confront this model with data, drawing from two types of data. The first are country-specific data on establishment/firm production in each of Chile, Colombia, India, Mexico, Slovenia, Ghana, Kenya and Tanzania (all of which have been widely used in the development and productivity literatures). The second are the World Bank's Enterprise Research Data, which allows us to exploit production data on firms in 33 countries. Each type of data has different strengths: The country-specific data sets have many more observations and somewhat tighter data collection protocols, while the World Bank data allows us access to a broader set of countries.

The basic reduced-form pattern implied by the model—that as volatility increases, so does dispersion—is strongly supported in the data across all data sets. After documenting this, we then take a more structural approach to see how well the model does at capturing cross-country variation in dispersion, and other moments. For this exercise, we first estimate capital adjustment costs. These adjustment-cost estimates, along with each country's AR(1) shock process, are used to generate model predictions (that is, we hold all other parameters constant). The model does surprisingly well: when confronted with cross-country data on dispersion in the marginal product of capital it generates a measure of fit equivalent to an uncentered R^2 of 0.7 (note that none of the model's parameters are estimated by matching this moment, making this a demanding empirical test). This suggests that the dynamic process of productivity is important, both empirically and theoretically, in determining the patterns observed in the cross-section.

These macro-level findings sit well with micro-level studies of the myriad challenges facing firms in developing countries. What we model as the productivity shock process is a reduced-form for a range of time-varying shocks to production, including (but not limited to): demand shocks (Collard-Wexler, 2008); natural disasters, (such as floods or landslides De Mel, McKenzie, and Woodruff (2012)); infrastructure shocks, such as power-failures or transportation links being established (Fisher-Vanden, Mansur, and Wang, 2012); variation in the incidence of corruption or nepotism (Fisman and Svensson, 2007); changes in mark-ups due to demand shocks or market-structure changes (De Loecker, Goldberg, Khandelwal, and Pavcnik, 2012); and changes in local or regional politics that affect productive outcomes

costs, cannot rationalize the extent of misallocation observed in Korean and Colombian plant-level data. If productivity is persistent enough, then productive firms will quickly save enough to escape their credit constraints, yielding a first-best outcome.

(Fisman, 2001). This paper can be viewed as suggesting a channel through which these micro effects can have macro implications.

The remainder of the paper is organized as follows. In Section 2, we present our dynamic model of investment. Section 3 discusses the measurement of productivity across several countries and consider reduced-form empirical evidence. Section 4 confronts the predictions of the dynamic investment model with the data using a structural approach. We discuss a few outstanding issues in Section 5 and conclude in Section 6.

2 Theoretical framework

In this section, we posit a simple model that allows us to consider how the time-series process of productivity should affect the cross-sectional dispersion of productivity, (static) marginal revenue products of capital and other variables. Central to the model is the role of capital adjustment costs, and a one-period time-to-build, in making optimal capital-investment decisions the solution to a dynamic problem: These adjustment frictions create links between the time-series process generating firm-level productivity shocks and firm-level heterogeneity in the adjustment of capital stocks.

2.1 Modeling preliminaries

We begin by providing an explicit model of productivity, in the context of a profit-maximizing firm (since we assume that establishments operate as autonomous units, firms and establishments, for our purposes, are synonymous). A firm i, in time t, produces output Q_{it} using the following (industry-specific) technology:[7,8]

$$Q_{it} = A_{it} K_{it}^{\alpha_K} L_{it}^{\alpha_L} M_{it}^{\alpha_M} \tag{1}$$

where K_{it} is the capital input, L_{it} is the labor input, and M_{it} is materials. This production function is industry-specific and throughout the paper, the coefficients β and α are kept industry-specific unless noted otherwise. The demand curve for

[7] We adopt a gross output approach to productivity in the model's exposition. In those instances when we use a value-added approach, the model presented here should be adapted such that $\alpha_M = 0$. We discuss associated measurement issues in the data section and appendix.

[8] To avoid the use of myriad subscripts, we omit subscripts that would indicate the country, and we omit industry subscripts on the α's and β's despite these coefficients being allowed to vary across country-industries in the empirical work.

the firm's product is given by a constant elasticity of demand curve:

$$Q_{it} = B_{it} P_{it}^{-\epsilon} \tag{2}$$

Combining these two equations, we obtain an expression for the sales-generating production function:

$$S_{it} = \Omega_{it} K_{it}^{\beta_K} L_{it}^{\beta_L} M_{it}^{\beta_M} \tag{3}$$

where $\Omega_{it} = A_{it}^{1-\frac{1}{\epsilon}} B_{it}^{\frac{1}{\epsilon}}$, and $\beta_X = \alpha_X(1 - \frac{1}{\epsilon})$ such that $X \in \{K, L, M\}$. For the purposes of this paper, productivity is defined as ω_{it} such that $\omega_{it} = \ln(\Omega_{it})$.

A fact that we will use repeatedly is that, in a static model with no frictions, profit maximization implies that the marginal revenue product (MRP) of an input should be equal to its unit input cost. For capital, the static marginal revenue product is given by

$$\frac{\partial S_{it}}{\partial K_{it}} = \beta_K \frac{\Omega_{it} K_{it}^{\beta_K} L_{it}^{\beta_L} M_{it}^{\beta_M}}{K_{it}} \tag{4}$$

Our notion of productivity is a revenue-based productivity measure, or TFPR as introduced by Foster, Haltiwanger, and Syverson (2008). As is common in this literature, we do not separately observe prices and quantities at the producer level, and, therefore, we can only directly recover a measure of profitability or sales per input precisely.[9] This implies that all our statements about productivity refer to TFPR, and, therefore, deviations across producers in our measure of productivity, or its covariance with firm size, could reflect many types of distortion, such as adjustment costs, markups or policy distortions, as Hsieh and Klenow (2009) discuss in detail.

2.2 A dynamic investment model

We now articulate a dynamic investment model that allows us to examine the link between productivity volatility and dispersion in both the static marginal revenue product of capital and productivity. Our model follows, and builds on, a standard model of investment used in the work of Bloom (2009), Cooper and Haltiwanger (2006), Dixit and Pindyck (1994), and Caballero and Pindyck (1996).

Taking the structure in section 2.1 as given, we begin by assuming that firms can hire labor in each period for a wage p_L and acquire materials in each period at a price p_M. Both of these inputs have no additional adjustment costs. Thus, we can optimize out labor and materials, conditional on Ω_{it} and K_{it}. This leads to a 'period-profit' (ignoring capital costs for the moment) of:

[9] See De Loecker (2010) for a detailed discussion and implications for actual productivity analysis.

$$\pi(\Omega, K) = \lambda \Omega^{\frac{1}{\beta_K+\epsilon^{-1}}} K^{\frac{\beta_K}{\beta_K+\epsilon^{-1}}} \tag{5}$$

where $\lambda = \left(\beta_K + \epsilon^{-1}\right) \left(\frac{\beta_L}{p_L}\right)^{\frac{\beta_L}{\beta_K+\epsilon^{-1}}} \left(\frac{\beta_M}{p_M}\right)^{\frac{\beta_M}{\beta_K+\epsilon^{-1}}}$.[10]

Capital depreciates at rate δ so $K_{it+1} = (1-\delta)K_{it} + I_{it}$ where I_{it} denotes investment. These investment decisions are affected by a one-period time to build and a cost of investment $C(I_{it}, K_{it}, \Omega_{it})$.[11] We employ an adjustment cost function composed of: 1) a fixed disruption cost of investing and 2) a convex adjustment cost expressed as a function of the percent investment rate and, therefore, $C(I_{it}, K_{it}, \Omega_{it})$:

$$C_K^F 1(I_{it} \neq 0)\pi(\Omega_{it}, K_{it}) + C_K^Q K_{it} \left(\frac{I_{it}}{K_{it}}\right)^2 \tag{6}$$

Next, let $\omega_{it} \equiv \ln(\Omega_{it})$ follow an AR(1) process given by:[12]

$$\omega_{it} = \mu_c + \rho_c \omega_{it-1} + \sigma_c \nu_{it} \tag{7}$$

where $\nu_{it} \sim \mathcal{N}(0,1)$ is an i.i.d. standard normal random variable. Note that we allow the mean of productivity, as measured by μ_c, the volatility σ_c, and the persistence coefficient, ρ_c to vary by country c. When we present results from computing our model, we will vary the volatility and persistence parameters $(\mu_c, \sigma_c, \rho_c)$.[13]

A firm's value function V is given by the Bellman equation:

$$\begin{aligned} V(\Omega_{it}, K_{it}) = &\max_{I_{it}} \pi(\Omega_{it}, K_{it}) - I_{it} - C(I_{it}, K_{it}, \Omega_{it}) \\ &+ \beta \int_{\Omega_{it+1}} V(\Omega_{it+1}, \delta K_{it} + I_{it}) \phi\left(\Omega_{it+1} \,|\Omega_{it}, \rho_c, \mu_c, \sigma_c\right) d\Omega_{it+1} \end{aligned} \tag{8}$$

and, thus, a firm's policy function $I^*(\Omega_{it}, K_{it})$ is just the investment level that

[10]It is worth noting that the λ term operates as a scaling term on the profit function. That is, with the flexibility to set the input prices, λ can be calibrated to any value that the researcher desires. An implication of this is that the qualitative predictions of the model do not depend on the number of variable inputs (labor, materials, energy...) in the production function. For instance, a value-added formulation (that is, with $Q_{it} = A_{it} K_{it}^{\alpha_K} L_{it}^{\alpha_L}$) can generate exactly the same patterns after adjusting the λ parameter appropriately.

[11]This time to build assumption is, in itself, a friction that we can easily shut down by allowing investment to become productive within a period (equivalent to one month in our implementation). As an indication of the economic effect of adjustment costs, if we set these to zero, then dispersion in the MRPK is reduced by 50 percent.

[12]Throughout the paper, lower case denotes logs, such that $x = \ln(X)$.

[13]Note that the specification in equation (7) rules out aggregate-level shocks to productivity growth. However, a regression of changes in productivity on country-year dummies in the World Bank data yields an R^2 of only six percent. Similarly, for the eight individual country data sets we find R^2's between 0.001, for Mexico, and 0.023, for Chile, when running productivity growth against year dummies. Thus, there appears to be only a small aggregate component to productivity change.

maximizes the firm's continuation value.

Note that since there is neither entry nor exit in this model, there is no truncation of the productivity distribution.[14] Thus, given the AR(1) structure above, the cross-sectional standard deviation of productivity is mechanically given by the ergodic distribution of Ω_{it}. Hence,

$$\text{Std.}(\omega_{it}) = \frac{\sigma_c}{\sqrt{1-\rho_c^2}} \tag{9}$$

where, as earlier, $\omega_{it} = \ln(\Omega_{it})$.

We analyze the model using computation. The parameters we use are found in Table 1. Parameters for the elasticity of demand, depreciation rate, and discount rate follow those adopted by Bloom (2009). Bloom uses a model in which investment decisions are made each month, with the model's predictions aggregated to the year level to fit the data. Modeling decisions on a monthly level is an attractive approach, as the model incorporates the likely time aggregation embedded in annual data. We follow this approach in computing the model and interpret a period in the model as equivalent to a month in data.[15] The results we report here are in terms of what one would see in annual data — that is, we aggregate up from monthly decision making to annualized data.

The coefficients of the sales-generating function we use correspond to the average over the World Bank Sample. We implicitly normalize the prices of non-capital inputs by setting $\lambda = 1$.[16]

The last set of parameters we need to fix are the σ_c, ρ_c and μ_c terms in the AR(1) process, which governs the evolution of productivity over time. In Section 4.1, we estimate this process using the firm data from the World Bank Enterprise Survey. For the moment, however, we merely note that the range of σ_c observed in the data lies in the interval $[0.11, 1.04]$. As a result, we compute the model for values of σ_c between 0.1 and 1.4. For ρ_c we pick three values that span the bulk of the estimated values, 0.78, 0.86 and 0.97. Lastly, we set $\mu_c = 0$. For more details on these estimated values, see the subsequent discussion in Section 4.1 and Table 8.

[14]The absence of entry and exit is a consequence of the decreasing returns to scale in the revenue equation (yielded by constant returns to scale in the production function and an elastic demand curve) and the absence of fixed costs, which make it profitable for any firm to operate at a small enough scale. See Midrigan and Xu (2009) for a discussion of the role of entry and exit in a similar model environment. However, our principal data source, does not cover a long enough time period to credibly get at the net-entry mechanism.

[15]This interpretation requires transforming the AR(1) process–which is quoted to reflect, and empirically estimated off, annual data–into its monthly equivalent. After noting that the sum of normal random variables with the same mean is distributed normally, this reduces to a straightforward algebraic exercise.

[16]More precisely, what were are normalizing is λ, a function of these non-capital input prices. The functional form of λ puts structure on the relative prices of non-capital inputs. Subject to this structure, normalizing λ is equivalent to a normalization of one of the non-capital input prices.

We compute the optimal investment policies for the value function in equation (8). We solve this model using a discretized version of the state space (Ω_{it}, K_{it}). Specifically, we use a grid of capital states ranging from log capital 3 to log capital equal to 20, in increments of 0.03. Moreover, we use a grid of productivity with 30 grid points, whose transition matrix and grid points are computed using Tauchen (1986)'s method. The model is solved using policy iteration with a sparse transition matrix (since there are 17,000 states). Using the computed optimal policies, we simulate the evolution of a country, or industry, for 10,000 plants over 1,000 periods. We use the output from the 1,000th and 988th periods to compute the reported results (corresponding to years t and $t-1$; recall that we interpret a period as a month).

2.3 Computational comparative static results

Figure 1 shows the output of the model. Panel (a) puts values of σ_c on the horizontal axis, and computed values of Std.$[\log(\beta_K) + s_{it} - k_{it}]$ are on the vertical axis. That is, it examines the way dispersion in the static marginal revenue product of capital changes as σ_c, the volatility of productivity, changes. In the figure there are three bold lines and three grey dashed lines. The bold lines correspond to the model with both a one period time to build and the adjustment costs. The dashed grey lines show the model without adjustment costs. Each set of bold and dashed lines has three lines stacked one above the other. In all panels, from top to bottom these correspond to ρ equal to 0.97, 0.86 and 0.78 respectively. In panel (a) for instance, this means that, for any specification and any level of σ, as ρ increases so does dispersion in the static marginal revenue product of capital.

Panels (b) through (e) have the same format, showing the computed dispersion in productivity (Std.$[\omega_{it}]$), the computed Olley-Pakes covariance (Cov.(ω_{it}, s_{it}), a measure of misallocation first suggested by Olley and Pakes (1996)), the computed volatility in the static marginal revenue product of capital over time (Std. $[(\log(\beta_K) + s_{it} - k_{it}) - (\log(\beta_K) + s_{it-1} - k_{it-1})]$), and the volatility in firms' capital over time (Std.$[k_{it} - k_{it-1}]$). Panel (f) focuses on the volatility in firms' capital over time (Std.$[k_{it} - k_{it-1}]$) in the full model.

Panel (b) is the most mechanical of the relationships reported in Figure 1. As volatility (σ_c) increases, so does the cross-section dispersion in productivity. As noted above, the dispersion in productivity is given by $\frac{\sigma_c}{\sqrt{1-\rho_c^2}}$. That is, it is given by the ergodic distribution of Ω_{it}. A further implication of this is that, if ρ_c and σ_c were constant over countries, there would be no cross-country differences in productivity dispersion.

Panel (a) contains the dispersion of the static marginal revenue product of capital (Std.[$\log(\beta_K) + s_{it} - k_{it}$]). Again, as productivity volatility (σ_c) increases, so does dispersion in the static marginal revenue product of capital. To further understand the pattern in Panel (a), note that this dispersion reflects the optimal investment choices of firms facing different productivity shocks over time and, hence, different state variables. To make the effect of this clear, note that if all plants had the same capital stock, this graph would replicate the relationship found in Panel (b). Yet the relationship between Std.[$\log(\beta_K) + s_{it} - k_{it}$] and σ_c is not linear and has a slope change at $\sigma_c = 0.5$ for $\rho_c = 0.97$ and at $\sigma_c = 0.7$ for $\rho_c = 0.86$. There is no readily discernible slope change in this range of σ_c for $\rho_c = 0.78$.

To see why this is happening, examine Panels (e) and (f). These Panels show the relationship between Std.[$k_{it} - k_{it-1}$] and σ_c.[17] As volatility increases, plants will engage in more investment and disinvestment. Since greater volatility leads to larger changes in productivity, it is natural that plants respond by altering their capital stock more frequently. However, for at least some values of the state space, plants begin to reduce their response to productivity shocks after σ_c reaches 0.5 for $\rho_c =$ 0.97 and 0.7 for $\rho_c = 0.86$, while for $\rho_c = 0.78$, the same pattern exists but is more gradual. At these high levels of volatility, current productivity is a weaker signal of the future marginal revenue product of capital. In the limit, where the productivity process is an i.i.d. draw, current productivity provides no information about future profitability. Firms would choose an optimal level of capital and stick to it forever, resulting in no variance in investment across firms. Thus, the flattening out of capital-adjustments to volatility is due to the changing trade-off in determining the value of investment today, between the size of shocks experienced today and the likelihood that they will be swamped by future shocks.

The results in panels (e) and (f) help explain the relationship between misallocation and volatility in Panel (a). As volatility increases above 0.5, the capital adjustment mechanism starts to shut down, and this speeds up the dispersion of the static marginal revenue product of capital.

Finally, Panels (c) and (d) show the Olley-Pakes covariance and the relationship between the standard deviation of the change in [$\log(\beta_K) + s_{it} - k_{it}$] and σ_c. Both relationships are essentially linear. The former is driven primarily driven by the dispersion in productivity, while the latter is driven primarily by year-to-year changes in productivity, rather than by large year-to-year changes in capital stock.

[17]Panel (f) focuses in on the adjustment cost part of Panel (e)

3 Data and preliminary analysis

In the rest of the paper, we work with a variety of data sets to understand the extent to which the framework developed in the preceding section is helpful in organizing the observed patterns in cross-country firm-level productivity differences.

3.1 Production data

The firm-level production data that we use come from two sources. The first is individual-country-level production data from eight countries. The second is the World Bank's Enterprise Research Data, which gives us access to data collected in a coordinated way across 33 countries. Each data source has tradeoffs: each individual-country-level production data set provides more-exhaustive coverage of the establishments/firms in just one country, together with tighter data collection protocols; while the World Bank data provide a sample of firms across many countries. Our introduction to each data set is brief, and we refer the reader to Appendix A for more details.

3.1.1 Individual-country-level Production Data

The first set of data is "high-quality" producer-level data from eight countries: Chile, Colombia, Ghana, India, Kenya, Mexico, Slovenia, and Tanzania. Each of these data sets has been used extensively in the literature; with a strong focus on the analysis of productivity.[18] The data sets differ in the time period covered, and in how producers are sampled. Table 1 summarizes the main features of the eight datasets. In Appendix A, we discuss each country data set in more detail and refer to a selective list of published work relying on these data.

3.1.2 The World Bank Enterprise Research Data

The second data source is the World Bank Enterprise Research Data, which gives us access to data collected in a coordinated way across 33 countries. Table 2 lists the countries we are able to use, together with the number of observations on each country.

These data were collected by the World Bank across 41 countries and many different industries between 2002 and 2006. Standard output and input measures are reported in a harmonized fashion. In particular, we observe sales, intermediate

[18]See, for instance, Tybout and Westbrook (1995), Roberts (1996), Pavcnik (2002), Rankin, Söderbom, and Teal (2006), Van Biesebroeck (2005), De Loecker and Konings (2006); De Loecker (2007), and Goldberg, Khandelwal, Pavcnik, and Topalova (2009).

inputs, various measures of capital, and employment, during (and covering up to) a three-year period, which allows us to compute changes in productivity and capital. Out of the 41 countries in the data, 33 have usable firm-level observations. This is primarily because, for many years and countries, the World Bank did not collect multi-year data on capital stock.

To construct data on productivity and the change in productivity we need two years of information on sales, assets, intermediate inputs and employment. 5,558 firms across our 33 countries meet this criterion.[19] In the data appendix we provide further details on sample construction and compare the firms in our sample, with the universe of sampled firms.

3.2 Measuring productivity

To guide the measurement of productivity, we build on the explicit model of productivity in Section 2.1 and, in particular, rely on the sales-generating production function in equation (3). In order to recover a measure of log productivity, ω_{it}, we need to impute the value of β_L, β_M and β_K by industry-country. Profit maximization implies that for each input facing no adjustment costs, the revenue production function coefficient equals the share of the input's expenditure in sales, or formally:

$$\beta_X = \frac{P_{it}^X X_{it}}{S_{it}} \qquad \text{for } X \in \{L, M\} \tag{10}$$

As mentioned before, we allow β_X to vary at the industry level within a country, thereby allowing the production function to vary across industries and countries. Thus, our approach to measuring productivity is to compute, for each individual firm:

$$\omega_{it} = s_{it} - \beta_K k_{it} - \beta_L l_{it} - \beta_M M_{it} \tag{11}$$

We recover the capital coefficient, for each industry-country observation, assuming constant returns to scale in physical production function (equation (1)); that is, $\beta_K = (1 - \frac{1}{\epsilon}) - \beta_L - \beta_M$. In order to compute β_K for each firm, we need to assign a value to the elasticity parameter, ϵ. We follow Bloom (2009) and set it equal to four.[20,21] Importantly, this approach in inferring β_K allows capital to have adjust-

[19] We also drop countries with fewer than 25 observations. This has little effect on our results.

[20] Alternatively, we could estimate the output elasticity directly from production data. We follow the standard in this literature and rely on cost shares to compute TFP and thereby avoid the issues surrounding identification of output elasticities (in our case, across many industries and countries).

[21] Table 10 reports results with either a lower or higher elasticity of demand ($\epsilon = 2$, $\epsilon = 6$), using plant-level input-shares instead of industry-level input shares to compute productivity, and using an OLS regression to estimate production function coefficients instead of using the results of a first-order condition. The relationship described in the paper is essentially unchanged.

ment costs, since it does not rely on a static first-order condition for the capital.[22] Appendix A2 provides further implementation details.

3.2.1 Summary statistics

Table 3 presents summary statistics for each of the data sets we use at the firm and country levels (the African data sets, having a common collection protocol, are consolidated). The data sets differ somewhat in the size of the firms that are included. The largest firms are in the Mexican data (and likely in the Indian data, although in that data set, the number of workers is not separately reported from the wage bill). The firms in the World Bank data also appear to have a relatively large number of employees, lying between the Mexican firms and the firms in the remaining data sets.

We next report the logs of value added, materials, capital and labor relative to log sales. This allows a unit free metric of the size of firm characteristics. To aid interpretation, consider the World Bank data: ln(Value Added) - s_{it} is equal to -0.9, which is equivalent to a value added to sales ratio of 40 percent; $m_{it} - s_{it}$ is -0.6, which is equivalent to a materials to sales ratio of 55 percent; $k_{it} - s_{it}$ is -0.1, which is equivalent to a capital to sales ratio of 90 percent; and $l_{it} - s_{it}$ is -1.8, which is equivalent to a labor to sales ratio of 17 percent.

We then summarize the year-to-year changes in log sales, capital, labor and productivity. Lastly, we report the capital share coefficients for both gross output and value added measures of productivity. Overall, all data sets have similar characteristics on these dimensions, a fact that is interesting in itself.

We then turn to productivity dispersion and volatility measures computed at the country level. These show the the extent to which the data sets differ according to how productivity is measured. Value-added measures tend to magnify differences across countries in their productivity dispersion and, to a lesser extent, volatility. In what follows value-added measures are used only to check if results obtained using gross output measures are robust.

3.3 Dispersion and volatility

After measuring TFP for each firm using data on sales and input usage, we construct the standard deviation of ω_{it} as a measure of productivity dispersion in each country.

[22]See De Loecker and Warzynski (2013) for more discussion. In addition, our alternative measure of productivity,using value-added (ω^{VA}), is obtained similarly using $\omega_{it}^{VA} = va_{it} - \beta_L^{VA} l_{it} - \beta_K^{VA} k_{it}$, where va_{it} is log of value-added for a firm-(country)-year, and the coefficients are now the share of input expenditures in value added. Again, we obtain similar results using value-added production functions.

We rely on the standard deviation of $(\omega_{it} - \omega_{it-1})$ as a measure of productivity volatility for each country. We then examine the correlation between these measures of a country's productivity dispersion and productivity volatility. The result of this process is shown in Figure 2(a) (depicting specification I in Panel A of Table 5). Figure 2(b) replaces the standard deviation of ω_{it} with the standard deviation of $(\log(\beta_K) + s_{it} - k_{it})$, the log of the static marginal revenue product of capital.

Figure 2(a) illustrates the positive correlation between productivity dispersion and productivity volatility. Indeed, in the World Bank data (depicted by the empty circles), cross-country variation in productivity volatility explains 64 percent of the cross-country variation in productivity dispersion in an OLS regression with a constant as the only other regressor. To examine the extent to which this pattern is an artifact of the World Bank data set, we superimpose the dispersion and volatility for each of the eight countries for which we have extensive high-quality production data. These are indicated by the solid circles. As the figure shows, these countries coincide with the World Bank data.

Table 5(A), presents regressions of productivity dispersion on productivity volatility, using the World Bank data. Specification I (depicted in Figure 1(a)) shows the OLS regression, using observations at the country level, weighted by the number of productivity observations per country. This weighting is used to give more importance to countries whose measurements of productivity dispersion and productivity volatility are relatively precise. In this specification, productivity volatility accounts for 64 percent of the (appropriately weighted) variation in within-country productivity dispersion. Specification II shows the results from an unweighted regression. Across both specifications, we find coefficients of 0.86 and 0.75, with standard errors of 0.21 and 0.23. Thus, the data appear consistent with the hypothesis that dispersion and volatility are related.

In specifications III and IV, the unit of observation changes from the country to the firm. The standard deviation of ω_{it} is common for all firms in a country, but we now control for firm size using total assets and the industrial activity of the firm–i.e., we include industry fixed effects. The coefficients are similar to those found without these controls. The standard errors, which are clustered by country, are also comparable. The results from these regressions eliminate the concern that dispersion and volatility are co-generated by a third variable, such as a country's industrial composition or the size of plants within a country. In specification V, we replicate specification I, but use a value-added measure of productivity. The results are robust to this specification change.

3.4 Capital misallocation and volatility

Productivity dispersion is economically relevant, to the extent that it reflects movements away from an optimal feasible resource allocation. This is most often examined in the context of productive inefficiency within an economy by inspecting differences in the static marginal revenue product of capital across firms. The static marginal revenue product of capital should, in the absence of adjustment costs (or other frictions), be equal across firms. As noted in the model in Section 2.1, the static marginal revenue product of capital (MRPK) is given by:

$$MRPK = \frac{\partial S_{it}}{\partial K_{it}} = \beta_K \frac{S_{it}}{K_{it}} \tag{12}$$

Thus, the dispersion (measured in standard deviations) of $\log(MRPK)$ is:[23]

$$\text{Std.}\left(\log(MRPK)\right) = \text{Std.}\left(\log(\beta_K) + \log(S_{it}) - \log(K_{it})\right) = \text{Std.}\left(\log(\beta_K) + s_{it} - k_{it}\right) \tag{13}$$

We use this as our measure of dispersion of the marginal revenue product of capital.

Table 5(B) presents regressions of static misallocation, $\text{Std.}\left(\log(\beta_K) + s_{it} - k_{it}\right)$, on productivity volatility, $\text{Std.}(\omega_{it} - \omega_{it-1})$. We use the same controls and estimation procedures as before, and, as such, the only difference between Panels A and B of Table 5 is the dependent variable. Figure 2(b) illustrates the positive correlation between dispersion in the static marginal revenue product of capital and productivity volatility (corresponding to specification I in Table 5(B)).

The coefficients in each specification of Table 5(B) are 0.67, 0.75, 0.64, and 0.63, respectively. All coefficients are statistically significant. Moreover, the R^2 is 0.31 in specification I, where no other controls are included. This increases to 0.36 when industry fixed-effects and log assets are included. Thus, a substantial fraction of cross-country differences in misallocation can be attributed to differences in country-specific productivity volatility. This suggests the existence of a link between the volatility of productivity in a country and the extent of (static) capital misallocation in that economy, and, is consistent with the model predictions presented in Section 2.

We also consider an alternative measure of misallocation first suggested by Olley and Pakes (1996): the covariance between a firm's market share and its TFP level.[24] Table 6 shows regressions of the Olley-Pakes covariance on productivity

[23] We allow β_K to vary at the industry-country level.

[24] The covariance measure can also be computed as the difference between from the market share weighted TFP average from the unweighted TFP average. See Olley and Pakes (1996) for more details, and also see Bartelsman, Haltiwanger, and Scarpetta (2009) for a discussion and application of this measure in the context of explaining productivity differences across countries.

volatility. Specification I presents a country-level regression analogous to specification I in Table 5. Specifications II and III present firm regressions analogous to specifications IV and III in Table 5. As in the earlier regressions, a statistically significant relationship emerges that is consistent with the model prediction.

3.5 Robustness: Industry-level analysis

The model presented in Section 2 applies equally to cross-industry dispersion differences and cross-country differences. Building on this, we check the robustness of our findings by taking each of the data sets that we have access to, and replicating the analysis done above at the country level, using industry-country-level observations. Panels A and B in Table 7 show the results.

Panel A takes each data set (we pool the three African countries' data) and, in specification I, projects the industry-country dispersion in productivity onto the volatility in productivity (again at the industry-country level) and a constant.

All coefficients are positive and significant, consistent with the model presented in Section 2. It is notable that, aside from the Colombian data, the coefficients have comparable magnitudes. As might be expected, when country and industry effects are added to the World Bank data, the coefficient moves toward zero – reflecting the likely introduction of some attenuation bias. Specification II controls for firm productivity and capital, but this makes no qualitative difference.

Panel B does the same exercise, using the standard deviation of the log sales to capital ratio (the static marginal revenue product of capital) at the industry level as the dependent variable. All coefficients are positive and all but one are significant.

The results presented in Tables 7 are consistent with the model predictions in Section 2. Since these results are generated using multiple independent data sets, they suggest that the phenomena illustrated earlier using just the World Bank data are found in most, if not all, productivity data sets drawn from developing economies (and even from those, like Slovenia, that are more properly termed developed).

4 Structural analysis: model and empirics

The previous section established a relationship between productivity volatility and various measures of static misallocation in the data. Further, the relationships we observe in the data match those predicted by the dynamic model presented in Section 2. In this section, we investigate quantitative aspects of the link between dispersion and volatility in a more structural setting, employing a calibrated model.

We adopt an approach in which we rely on the calibrated model to predict the cross-country pattern of various moments, which we then compare to the same moments in the data. The exercise is demanding in that, to predict a country's moments, we use only that country's data to estimate the AR(1) process determining productivity shocks and the production function coefficients β.[25] All other parameters are held constant across countries. This means that the moments we seek to match are not used to estimate the parameters that we use to generate predictions. This is the sense in which we are engaging in demanding quantitative evaluation of the extent to which productivity volatility, together with adjustment costs, can account for cross-country variation in the distribution of firm-level productivity and static misallocation measures.

The exercise focuses on the World Bank data in order to compare the various moments across a large set of countries for which the underlying data was collected in a uniform fashion.[26]

We proceed in four steps. The first step is to estimate the AR(1) process for each country. The second step is to obtain estimates of adjustment costs. We choose to use Chilean data from the World Bank, rather than from any other country or data source, as Chile has the largest number of observations in the World Bank data and is a fairly typical country in terms of its AR(1) process.[27] We use the same adjustment parameters to estimate all the countries's moments. We do this for two reasons: First, it focuses attention on the productivity volatility process (mirroring the reduced-form analysis in Section 3); and, second, we find (both empirically and in the computational results in Section 2) that several of the moments are relatively unaffected by the level of adjustment costs, provided that some adjustment friction exists (e.g., the one-period time to build). By fixing adjustment costs across countries, we make it harder for the model to fit the data. The third step is to arrive at a way to evaluate fit: We use an adaptation of the familiar uncentered R^2 statistic. The fourth step is to evaluate the results.

[25] Except for Chile, as the data used to estimate adjustment cost parameters is Chilean.

[26] Guyana, Kyrgyzstan, the Philippines, Poland, Tajikistan and Tanzania are excluded, since their estimated ρ_c's exceed 1. This means that producing a stationary distribution in the simulation is not feasible. Ecuador is also excluded as its estimated ρ_c is so close to one that computing a stationary equilibrium is not feasible within machine precision. See Table 8 and Section 4.1 for more information on the estimated AR(1) coefficients.

[27] In a previous version of this paper, we used the capital adjustment costs estimated by Bloom (2009) off of COMPUSTAT firms in the United States, and we find results similar to those in this paper.

4.1 Estimating the productivity AR(1) process

The first step is to estimate the AR(1) productivity process for each country. The specification used is:[28]

$$\omega_{it} = \mu_c + \rho_c \omega_{it-1} + \sigma_c \nu_{it} \tag{14}$$

The data is a short panel, as we only have two years of data per firm, and thus μ_c also captures any aggregate shocks at the country level. The unit of observation is the firm. Identification of the AR(1) relies on the assumption that ρ and σ are constant over time and across firms.[29] This allows us to identify the model using cross-sectional variation in firm-specific productivity pairs, $< \omega_{it}, \omega_{it-1} >$.

Tables 8 summarizes the results of this exercise, and also provides comparison estimates using the other country-specific data sets we have access to, for which we can rely on a much longer panel, as described in Table 2. Chile and Tanzania feature in both the World Bank and country-specific data. The Tanzanian estimates for the ρ and σ parameters are almost identical in both data sets. The Chilean estimates for ρ are very similar and the estimate of σ is considerably higher using the World Bank data. Note that the samples in each data set are different, since the World Bank data are for 2002-2004, while the Census data for Chile are for 1979 to 1986.

More-detailed information on the specification reported in Table 8 is given in Table 10, specification V (in the appendix) together with comparisons with alternate specifications.

4.2 Estimation of Adjustment Costs

To estimate adjustment costs, we rely on the Chilean data from the World Bank Enterprise Data, employing both the Chilean specific production function coefficients and the AR(1) coefficients. Recall that the adjustment cost specification is given by:

$$C_K^F 1(I_{it} \neq 0)\pi(\Omega_{it}, K_{it}) + C_K^Q K_{it} \left(\frac{I_{it}}{K_{it}}\right)^2 \tag{15}$$

We estimate C_K^F and C_K^Q using a minimum-distance procedure very similar to that in Cooper and Haltiwanger (2006). That is, we seek parameters that minimize the distance between moments predicted by the model and those found in the data. The moments we use are: the proportion of firms with less than a 5 percent year-on-year

[28]The 'c' subscript indicates the country.

[29]This restriction is driven only by the data, and our framework could handle various forms of time-specific persistence and volatility if the data had a longer time dimension. We have estimated this model on a longer panel (of about 7-12 years) for two countries, India and Slovenia, and find that the AR(1) coefficient is stable over time.

change in capital; the proportion of firms with more than a 20 percent year-on-year change in capital; and, the standard deviation of the year-on-year change in log capital.

Denote the predicted moments from the model as $\mathbf{\Psi}(\theta)$, found by solving for the firm's optimal policies and simulating the model forward for 1000 months for 10,000 firms, and computing moments based on the last two years of the simulated data set, as in Section 2.2 . The moments from the data are denoted $\hat{\mathbf{\Psi}}$. We estimate the model's adjustment costs using minimum distance with a criterion function given by the usual quadratic form, with weighting matrix $\mathbf{W}$:

$$\mathcal{Q}(\theta) = \left(\hat{\mathbf{\Psi}} - \mathbf{\Psi}(\theta)\right)' \mathbf{W} \left(\hat{\mathbf{\Psi}} - \mathbf{\Psi}(\theta)\right) \tag{16}$$

As the moments in the data are similarly scaled, we pick the identity matrix as a weighting matrix ($\mathbf{W} = \mathbf{I}$). We find the minimized value of the criterion using a grid search.[30]

We obtain the following estimates: Fixed Adjustment Costs (C_K^F), 0.17 with standard error equal to 0.05; Convex Adjustment Cost (C_K^Q), 0.75 with standard error equal to 0.20. The fixed cost of adjustment is equivalent to two months of output, while the convex adjustment costs are such that when a firm doubles its capital, this component of cost is equal to 0.75 of the value of its investment. These parameters are comparable to those found in Bloom (2009) (Table 3, column 2) who obtains fixed adjustment costs of 0.01 and convex adjustment costs of 1.00.

4.3 Computation and evaluating model fit

To compute country-specific predictions, we use the country's estimated AR(1) process, as well as the country-specific production function parameters β reported in Table B. Adjustment costs are common for all countries, as well as all other parameters (discount and depreciation rates and the like) reported in Table 1. The computation of the model follows that described in Section 2.2. The only departure is that we compute a prediction for each country-industry, as these have different β's, and then aggregate up to the country level.[31]

To assess the fit of the model, we compute the sum of squared errors, scaled by

[30]Standard errors are computed using the usual formula for minimum-distance estimators:

$$Cov(\hat{\theta}) = \left(\frac{\partial \mathbf{\Psi}'}{\partial \theta}\mathbf{W}\frac{\partial \mathbf{\Psi}}{\partial \theta}\right)^{-1} \left(\frac{\partial \mathbf{\Psi}'}{\partial \theta}\mathbf{W}\mathbf{Var}(\hat{\mathbf{\Psi}})\mathbf{W}\frac{\partial \mathbf{\Psi}}{\partial \theta}\right) \left(\frac{\partial \mathbf{\Psi}'}{\partial \theta}\mathbf{W}\frac{\partial \mathbf{\Psi}}{\partial \theta}\right)^{-1} \tag{17}$$

We bootstrap the data to obtain estimates of the covariance of the moments in the data $\mathbf{Var}(\hat{\mathbf{\Psi}})$.

[31]Note that the variation between industries is far smaller than the variation within industries.

the sum of the squared 'dependent' variable (data). That is, if the data are a vector $\mathbf{x}$ that is predicted by a variable $\hat{\mathbf{x}}$, then we compute

$$S^2 = 1 - \frac{(\mathbf{x}-\hat{\mathbf{x}})'(\mathbf{x}-\hat{\mathbf{x}})}{\mathbf{x}'\mathbf{x}} \tag{18}$$

as our measure of fit. This measure of fit is closely related to the uncentered R^2 measure of fit familiar from regression analysis. However, because our model's prediction does not come from a regression, but from a parameterized model, nothing in the structure restricts S^2 to lie between 0 and 1, though, by definition, it must be less than or equal to one. That being said, to map our measure of fit into a context equivalent to the R^2, it is correct to interpret S^2 as the proportion of the observed data captured by the model's prediction, with the caveat that it is possible for this number to be negative.[32]

4.4 Results

Table 9 presents S^2 statistics for five moments of interest. Statistics are present for both the full model (including the adjustment costs) and a model in which the only adjustment friction is the one-period time to build. Figure 3 examines three moments in more detail, showing each model's prediction for each country, plotted against the measure of the variable of interest found in the data. Circles indicate countries, and circle size is proportional to the number of firms per country. Each country is plotted using an (x, y) coordinate, where the x -coordinate indicates the model's prediction and the y - coordinate indicates the value in the data. The closer the country lies to the 45^o line, the more accurate the model's prediction.

The first moment of interest is the dispersion in the static marginal revenue product of capital (Std.$[\log(\beta_K) + s_{it} - k_{it}]$). As shown in Table 9, the S^2 of the full model is 0.704, while the model with only the one-period time to build (the partial model) has an S^2 of 0.863. This indicates that 70 percent of the observed static capital misallocation is captured by the full model's prediction. It is of some interest that the partial model outperforms the full model. This is shown more clearly in Panels (a) and (b) of Figure 3. When adjustment costs are added to the model, the model tends to over-predict for most countries. The two countries in which over-prediction is greatest are (from right to left) Mauritius and Indonesia.[33]

[32] We use an uncentered measure of fit, as our model does not incorporate anything analogous to the estimated constant commonly found in a regression specification.

[33] Note that the AR(1) process substantially overpredicts the steady-state productivity dispersion for Mauritius and Indonesia. Thus, it is the failure of the AR(1) process, rather than the investment model itself, that leads to poor predictions for these two countries.

The second moment of interest is the dispersion in productivity. As shown in Table 9, the S^2 of the both models is 0.863 (recall that the ergodic distribution is independent of adjustment costs), indicating that 87 percent of the the data on cross-country differences in productivity dispersion is captured by the model's prediction.

The third moment of interest is the Olley-Pakes covariance. The S^2 of the full model is 0.722, while the partial model has an S^2 of 0.722. While the models do have slightly different performance, they are very very similar, which mirrors the results obtained from the computational simulations (see Figure 1). This, in itself, is an informative feature consistent with the model. Figure 3 shows the results in more detail, with Panels (c) and (d) being, essentially, interchangeable.

The fourth moment of interest is the dispersion in the change in the static marginal revenue product of capital (Std.$[(\log(\beta_K) + s_{it} - k_{it}) - (\log(\beta_K) + s_{it-1} - k_{it-1})]$). The S^2 of the full model is -0.318, while the partial model has an S^2 of 0.709. Here, as in the cross-sectional dispersion (moment one), the full model is over-predicting, albeit by a greater magnitude. Again, the two countries for which over-prediction is greatest are Mauritius and Indonesia.

The last moment, the dispersion in the change in capital, shows a different feature of the model. Here, the S^2 of the full model is 0.585, while the partial model has an S^2 of -7.10. That is, the full model does a good job of capturing cross-country variation in the data, while the partial model barely relates to the data at all. Figure 3, Panels (e) and (f), compare the partial and full models. As can be seen, without the full array of adjustment costs, the partial model dramatically over-predicts changes in capital. On this dimension, the full model is clearly preferred to the partial model.

The full and partial models, taken together, display the usefulness of the dynamic model presented in Section 2. This model, appropriately parameterized, does a good job at predicting a range of moments in the data.

It is also interesting to compare the full and partial models. The fact that each does poorly on one dimension (spectacularly so, in the case of the partial model), suggests that reconciling the two may require a re-casting of how adjustment costs are parameterized. There are two components that give adjustment frictions: the one-period time to build; and the 'monetary' adjustment costs. The one-period time to build introduces a state-dependency to the adjustment friction, since a high-productivity state, with a high persistence process will be affected differently as compared to a low-productivity state, or a process with little persistence. Thus, the one-period time to build introduces an 'uncertainty-cost' feature to adjustment costs that is parameterized by the length of time taken to build (in our case, one month). To accommodate the apparent trade-offs between fitting moments better captured by the full and partial models, we conjecture that allowing the length of the time to build to vary may be important. This would make uncertainty have a more significant role in creating adjustment frictions, a theme being explored in contemporaneous research. See, for example, Bloom, Floetotto, Jaimovich, Saporta-Eksten, and Terry (2011).

4.5 Discussion

The model does a surprisingly good job of fitting the data, despite the fact that all parameters other than the μ_c, ρ_c and σ_c terms in the AR(1) productivity process and the production function parameters, β, are either taken as standard or estimated using Chilean data. This has a series of implications for our thinking about resource allocation, productivity differences across countries, and the welfare and policy implications therein.

First, the dynamic process underlying productivity can generate cross-sectional dispersion of productivity and capital allocation. Our contribution is to demonstrate the empirical importance of this mechanism. Our findings suggest that the dynamic process governing productivity shocks is a first-order determinant of differences in productivity and, hence, income across countries.

Second, the conclusions one draws regarding welfare and policy depend on the model one has in mind regarding this dynamic process. If one has the view that the productivity process is an exogenous, or primitive, feature of the model, then our findings suggest that, in an aggregate sense, the firms in the countries we studied are acting much as the social planner in our model would have them act (assuming that the social planner takes the capital adjustment costs as given). This suggests that there are few welfare implications for differences in productivity dispersion and static misallocation across countries. On the other hand, if government policy can affect the productivity process, then there may be welfare dividends to policy interventions aimed at moving toward some socially-optimal productivity process. However, characterization of what this optimal process is likely requires a more subtle modeling approach than that offered here.

The position one takes about the extent to which the productivity process is exogenous requires careful consideration of what is captured by the "revenue" measure of productivity we use. In particular, productivity is not just technological in nature. The fact that the sales function is used, means that our measure of productivity volatility captures changes in managerial and physical technology. It also captures year-on-year variation in the intensity of corruption (and the implicit tax therein); regulatory frictions; environmental factors (e.g., floods) and the efficacy of infrastructure used to cope with them; and year-on-year variation in markups and product market competition. Many of these elements of measured productivity volatility may be effectively influenced by appropriate policy aimed at providing a stable business environment. By offering an alternative view that is strongly supported by the data, this paper sharpens the debate on the role of policy interventions that are geared towards eliminating resource misallocation in developing countries.

5 Robustness Checks

The relationship we have documented between volatility and the dispersion of both productivity and the static marginal revenue product of capital is well founded as a point of theory. However, there remains a concern that the empirical relationship we have documented might be contaminated by measurement error that affects both the volatility and dispersion of productivity. Below, we briefly discuss three robustness checks that speak to concerns that measurement error is distorting our findings.[34] Table 10 reports additional robustness and specification checks.

5.1 Additional Data

Our main approach is to replicate our basic findings using many data sets, which are collected separately by different collection agencies in different countries using different collection methods. As discussed earlier in the paper, we observe the same relationships, between productivity volatility, dispersion and dispersion in the (static) marginal revenue product of capital, in eight individual-country data sets and the World Bank Enterprise data (see Figure 2, for example). Moreover, these relationships hold not only across countries, but also within country, using industry-level variation, both for the World Bank Enterprise Data and for each of our eight high-quality datasets, as predicted by the theory (Tables 7 and 8). This suggests that measurement error arising from the idiosyncrasies of any one data set is unlikely

[34]For the entirety of the relationships we document to be an artifact of the data would require a very specific formulation of the measurement error process. We have not been able to derive what the process would have to be. For instance, the most natural model of measurement-error that would replicate our dispersion-volatility relationship would be a model in which serially correlated measurement error in productivity would cause both greater dispersion in productivity and greater volatility of productivity (since the World Bank Enterprise survey is a retrospective survey rather than a true panel, we find it more plausible to believe that worse measurement is associated with less rather than more volatility in capital or output). Yet this model would fail to reproduce the relationship we found between the volatility of productivity and the Olley-Pakes covariance, since greater measurement error should be associated with a lower correlation between output and productivity.

to be at the heart of the relationships we document.

5.2 Outliers

All of our results are robust with respect to eliminating outliers in terms of productivity, recalling that we dropped observations with productivity above 6 in absolute value. In practice, we looked at whether our results are affected by trimming the top and bottom ten percent, and by the interquartile range (which is less sensitive to outliers) rather than the standard deviation as a dependent variable. Both of these robustness checks are presented in Table 10, and our results still hold with these two changes. This suggests that our findings are not driven by common data problems in the tails of the samples or by the way the samples are trimmed.

5.3 Predictive power of productivity measures

In order to test whether our results could still be plagued by remaining measurement error, we follow Hsieh and Klenow (2009) and relate our measure of productivity to decision variables that plausibly have little room for measurement error.

Regardless of the ultimate source of measurement error, if measured productivity were mere measurement error, we would not expect actual behavior to be correlated with measured productivity. With this in mind, we ran a probit with an indicator for positive investment as the dependent variable, and productivity, log capital and country fixed effects as the explanatory variables (using the World Bank data). The average marginal effect on productivity was estimated to be 0.11 with a standard error of 0.01, making it significant at better than one percent. The pseudo-R-squared was 0.16. We also ran an OLS regression with the log investment to capital ratio as the dependent variable, and (again) productivity, log capital and country fixed effects as the explanatory variables (using the World Bank data). The coefficient on productivity was 0.23, again significant at better than one percent. The R-squared was 0.19.

The indicator for positive investment is likely to be well measured and is positively, and significantly, correlated with productivity. The log investment to capital ratio, while arguably more prone to measurement error, displays the same pattern. This constitutes evidence that plausibly well-measured decision variables are correlated with productivity.

6 Conclusion

We have focused on the adjustment costs in capital, coupled with productivity shock processes, to interpret the large dispersion in marginal revenue product of capital. In doing so, we shut down many other economically relevant features of a firm's environment that could lead to differences in the measured marginal revenue product of inputs, including, for instance, differences in the factor prices for inputs, differences in the size of adjustment costs, and heterogeneity in market power. This keeps our model parsimonious and makes the approach in this paper directly comparable with the approach taken in the existent literature on cross-country productivity differences. However, the dynamic model of misallocation proposed in this paper is clearly compatible with additional sources of heterogeneity between producers.

A natural alternative starting point would be to include additional heterogeneity in market power and interpret the differences in marginal revenue product differently–i.e., as a reflection of differences in market power that vary over time.[35] We note this to underscore the fact that observed productivity differences can have many underlying drivers. We focus on just one.

The primary contribution of this paper is to establish the link between the dynamic process governing productivity changes over time and cross-sectional measures of productivity dispersion and (static) capital misallocation. In particular, we show that a parsimonious model of the country-specific productivity process explains much of the variation in the dispersion of productivity and of the static marginal revenue product of capital across countries. Thus, commonly used static measures of misallocation are difficult to compare across environments that have different processes for productivity. We provide evidence to support the claim that the dynamic process of productivity is important, both empirically and theoretically, in determining the patterns observed in the cross-section.

Our findings reinforce the point that the country-specific stochastic process of productivity is sufficient to explain a significant proportion of cross-country variation in productivity and static capital misallocation. They suggest that producers in countries that experience larger uncertainty in the operating environment (i.e., higher volatility in productivity) make different investment decisions than those producers active in less volatile environments. This leads to different levels of capital and output and, moreover, means that the welfare gains from policies inducing reallocation of factors of production are likely to be lower than otherwise implied

[35] De Loecker and Warzynski (2013) provide a way to obtain producer-level markups using standard production data, while allowing explicitly for dynamic inputs of production, such as capital. Pairing their approach with our framework could, in principal, allow for a decomposition of "static" and "dynamic" components of the perceived misallocation from the standard model. This we leave for future work.

by static models.

An alternative suite of policy options, aimed at making the productivity process more benign, may be attractive as a complement to the redistributive measures featured in the counterfactuals considered in other studies. It is likely that at least some component of the stochastic process of productivity is influenced by government policy. To the extent that this is true, our findings imply that, if government policies can provide a more predictable business environment, then this will benefit the economy and help producers allocate resources in more-productive ways. This raises the issue of the sources of adjustment costs and productivity volatility, a topic on which we are silent in this paper. Our aim here is to merely cast light on the importance of dynamics in assessing the welfare relevance of productivity dispersion and in evaluating an appropriate policy response.

A Appendix

In this Appendix, we discuss briefly the data and elaborate on how we measure TFP in our data. It is important to stress that we measure TFP such that it is consistent with the theoretical model introduced in the paper.

A.1 Data

A.1.1 Slovenia

The data are taken from the Slovenian Central Statistical Office and are the full company accounts of firms operating in the manufacturing sector between 1994 and 2000. We have information on 7,915 firms: an unbalanced panel with information on production, employment, investment, intermediate input, and balance-sheet variables. We would like to thank Joze Damijan at Ljublijana University for sharing the data. We refer the reader to De Loecker and Konings (2006) and De Loecker (2007) for more on the data.

A.1.2 Mexico

Annual plant-level data on manufacturing plants are recorded by Mexico's Annual Industrial Survey and are provided by Mexico's Secretary of Commerce and Industrial Development (SEC-OFI). These data, which cover the period 1984-1990, include various production, employment, investment, intermediate input, and balance-sheet variables. The sample of plants represents approximately eight percent of total output, where the excluded plants are the smallest ones. The data were generously provided by Jim Tybout through a license at IES Princeton University. Please see Tybout and Westbrook (1995) for more details.

A.1.3 Colombia

Annual plant-level data on all manufacturing plants were provided by Colombia's Departmento Administrativo Nacional de Estadistica (DANE). The census covers all plants, but after 1982, it only covers those plants with ten or more workers. These data, which cover the period 1978-1991, include various production, employment, investment, intermediate input, and balance-sheet variables. The data were generously provided by Jim Tybout through a license at IES Princeton University. Please see Roberts (1996).

A.1.4 Chile

Annual plant-level data on all manufacturing plants with at least ten workers were provided by Chile's Instituto Nacional de Estadistica (INE). These data, which cover the period 1979-1986, include various production, employment, investment, intermediate input, and balance-sheet variables. The data were generously provided by Jim Tybout through a license at IES Princeton University. See Pavcnik (2002) for a productivity study using these data.

A.1.5 India

Annual firm-level data on manufacturing firms were provided by Prowess, and are collected by the Centre for Monitoring the Indian Economy (CMIE). Prowess is a panel that tracks firms over time for the period 1989-2003. The data contain mainly medium and large Indian firms. These data include various production, employment, investment, intermediate input, and balance-sheet variables. The data are used in De Loecker, Goldberg, Khandelwal, and Pavcnik (2012), and more details on the data are discussed therein.

A.1.6 African Survey Data: Ghana, Kenya, and Tanzania

These data were made available by the Centre for the Study of African Economics at the University of Oxford. Permission for use of the data for academic research is given by the Centre and its collaborating and funding institutions and can be found here: `http://www.csae.ox.ac.uk/datasets/cfld/cfld-main.html`. The data come from various surveys of African manufacturing firms, see Rankin, Söderbom, and Teal (2002) for background on the Ghanaian data set, Söderbom (2001) for information on the Kenyan data set, and Harding, Söderbom, and Teal (2002) for the Tanzanian data. Söderbom (2001), section 3, is a useful reference for the survey details relevant to all these data. All data sets contain multiple observations on the same firms over time–i.e., panel data. For an application of the data on Ghana, Kenya, and Tanzania, see Rankin, Söderbom, and Teal (2006). We would like to thank Jo Van Biesebroeck for pointing us to the data source. Also see Van Biesebroeck (2005) for a productivity study using some of these data.

A.1.7 World Bank Data

The data are available from `http://www.enterprisesurveys.org`, accessed on December 15th, 2010. Extensive documentation is available from the same website.

The survey documentation describes the sampling universe as follows: "6. The population of industries to be included in the Enterprise Surveys and Indicator Surveys, the Universe of the study, includes the following list (according to ISIC, revision 3.1): all manufacturing sectors (group D), construction (group F), services (groups G and H), transport, storage, and communications (group I), and subsector 72 (from Group K). Also, to limit the surveys to the formal economy the sample frame for each country should include only establishments with five (5) or more employees. Fully government owned establishments are excluded as the Universe is defined as the non-agricultural private sector." from page 3 in 'Enterprise Survey and Indicator Surveys Sampling Methodology' August 29th, 2009 at `http://www.enterprisesurveys.org/Documents/Sampling_Note.pdf` downloaded 23 April, 2011.

The survey used a stratified sampling procedure, in which firms were sampled randomly within groups based on the firm's sector of activity, firm size, and geographical location. The structure of the sampling leads to an oversampling of larger firms (relative to random sampling of all firms in the economy). The exact structure of the stratification varies by the size of the economy in question. We have chosen to not do any sampling correction, preferring to maintain as much transparency as possible as to the mapping from data to findings, being mindful of the fact that we can use data from only 7 percent of the sampled firms in any case and, most importantly, considering the absence of a well-defined criterion that could be used to guide any such correction. It is an open question whether, ideally, we would weight by, say, contribution to GDP or would weight each firm equally. More likely, weighting by some measure of activity makes more sense for our purpose, but to the extent that, say, any Eritrean government statistics we would use to do this would have measurement error, this may merely contribute to attenuation bias in the results. This is especially so, given that the sampling structure used in the surveys overweights large firms and, hence, already moves in the direction of weighting by contribution to economic activity. In any case, the results in the paper are robust to controlling for differences in the size and industrial composition of firms across countries.

The firms in the data are drawn from the manufacturing, construction, services, and transport, storage, and communications sectors. As would be expected, the precise industry composition (defined at the two-digit ISIC level) varies by country.

Firms were surveyed between 2002 and 2006. The majority of firms within a country were surveyed in the same year. The survey asked questions about activity in the current year and the previous two years. Thus, the panel-data aspect of these data, relating to activity in year $t-1$, comes from the recollections and records of managers in year t.

While there are over 41,000 observations in the data, only 5,558 have information on capital over several years, which is needed to compute productivity volatility. Table C presents summary statistics of the data, where for each variable, the first line refers to the data that we use, while the second presents the data that we dropped since there was not enough information to compute changes in productivity. The dropped observations are usually smaller plants with lower sales and fewer employees. However, changes in inputs (such as changes in capital or labor) are comparable across the data we did and did not use. Notice that the dispersion of productivity is similar between the two data sets, with a standard deviation of 1.0 (our data) versus 1.2 (dropped data), as well as the dispersion of the sales to capital ratio which is 1.1 (our data) versus 1.3 (dropped data). Thus, the sampling bias will slightly understate the level of productivity and sales to capital dispersion, but this effect is small relative to the large differences in dispersion across countries.

A.2 Measuring Productivity

As discussed in the main text we rely on a standard production function where a firm i, in country c, in time t, produces output Q_{it} using the following (industry specific) technology:

$$Q_{it} = A_{it} K_{it}^{\alpha_K} L_{it}^{\alpha_L} M_{it}^{\alpha_M} \tag{19}$$

where K_{it} is the capital input, L_{it} is the labor input, and M_{it} is materials. The demand curve for the firm's product is given by a constant elasticity of demand curve:

$$Q_{it} = B_{it} P_{it}^{-\epsilon} \tag{20}$$

Combining these two equations, we obtain an expression for the sales-generating production function:

$$S_{it} = \Omega_{it} K_{it}^{\beta_K} L_{it}^{\beta_L} M_{it}^{\beta_M} \tag{21}$$

where $\Omega_{it} = A_{it}^{1-\frac{1}{\epsilon}} B^{\frac{1}{\epsilon}}$, and $\beta_X = \alpha_X(1 - \frac{1}{\epsilon})$ such that $X \in \{K, L, M\}$.

Our base results rely on TFP measured as described in Section 3.2. As mentioned in the main text, we allow for the coefficients (β) to vary at the industry-country level. In practice, in order to obtain a robust measure of these shares, we rely on the median of the expenditure share for labor and intermediate inputs, in a given industry-country (sc), or

$$\beta_X^{sc} = median(\frac{P_{it}^X X_{it}}{S_{it}}) \qquad \text{for } X \in \{L, M\}\,, i \in sc \tag{22}$$

To recover the coefficient on capital, β_K, we use our assumption of constant returns to scale in production–i.e., $\sum_x \alpha_x = 1$, such that:

$$\beta_K^{sc} = \frac{\epsilon - 1}{\epsilon} - \beta_L - \beta_K \tag{23}$$

To compute (gross output) productivity, we simply plug in the coefficients obtained above into:[36]

$$\omega_{it} = s_{it} - \beta_K k_{it} - \beta_L l_{it} - \beta_M M_{it} \tag{24}$$

When we consider value-added-based productivity, we apply the same procedure as above, and replace sales by value added in all the above. That is, we first compute the share of input's expenditures in value added, and obtain a measure of productivity subtracting the weighted inputs from (log) value added.

To measure TFP, we require a measure of plant-level sales (S_{it}), employment (L_{it}), intermediate input use (M_{it}) and the capital stock (K_{it}). We follow the standard practice and refer to Bartelsman, Haltiwanger, and Scarpetta (2009) for an excellent clear overview and discussion on the measurement of TFP using similar data sources. In particular, we handle the data from the eight individual data sets in a standard way.

For some of the countries in the World Bank Enterprise Data, a number of issues emerged in the calculation of productivity. In particular, labor use is typically reported as the number of employees or a wage bill converted to the number of employees with no correction for hours worked. Moreover, sales and gross output data are not corrected for inventories, and the capital stock is based on book values. These are standard data restrictions researchers face using this type of data.

Sales are directly measured in the data, whereas labor is measured by the total number of workers active in a plant, or, alternatively, we convert the total wage bill of a plant into the number of workers using a plant-specific wage. The latter is corrected for aggregate wage trends using the median wage trend in a given industry-country pair. Finally, we rely on the book value of capital as measured by either total assets or net book value. We experimented with both measures and our results are invariant. When we consider a measure of value added, we compute it by netting the sales variable from the use of intermediate inputs.

[36]Here, we revert back to suppressing industry and country subscripts and superscripts.

For many firm-years in the data, we can compute productivity directly. However, for some firm-years, we observe only the firm's wage bill and not the number of workers. To address this issue, we use the median country-industry wage, $\tilde{w}$, (imputed from observations with both the wage bill and the number of workers) as a deflator and apply it to the wage bill to give a measure of labor. That is, to compute L_{it} we use $L_{it} = \frac{wL_{it}}{\tilde{w}}$. In what is presented in this paper, we use this measure for all firm-year observations.

Finally, we convert all relevant variables into real values using detailed producer price and input price deflators where available. In particular, for each of the eight individual countries, we rely on industry-specific price indices to convert the data on sales, input expenditures and input use. For the 33 countries covered in the World Bank data, these price indices are, unfortunately, not available. Therefore, we rely on the World Bank deflators to convert all monetary variables into USD. As mentioned in the text, the data are converted from local currency units into U.S. dollars. We use the World Bank's measure of purchasing power parity (PA.NUS.PPP). Note that we account for differences in the rate of inflation across countries by using a year-specific measure of PPP. Since productivity is a ratio, these PPP conversions get netted out in many specifications, but they are useful when, for instance, we use controls for firm size.

References

Bartelsman, E., and M. Doms (2000): "Understanding productivity: Lessons from longitudinal microdata," *Journal of Economic Literature*, 38(3), 569–594.

Bartelsman, E., J. Haltiwanger, and S. Scarpetta (2009): "Cross-country differences in productivity: the role of allocation and selection," Discussion paper, NBER Working Paper 15490.

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

Bloom, N., M. Floetotto, N. Jaimovich, I. Saporta-Eksten, and S. Terry (2011): "Really uncertain business cycles," Working Paper Stanford University.

Caballero, R., and R. Pindyck (1996): "Uncertainty, investment, and industry evolution," *International Economic Review*, 37(3), 641–662.

Collard-Wexler, A. (2008): "Demand Fluctuations in the Ready-Mix Concrete Industry," *Working Paper, New York University.*

——— (2009): "Productivity Dispersion and Plant Selection in the Ready-Mix Concrete Industry," *Working Paper, New York University.*

Cooper, R., and J. Haltiwanger (2006): "On the nature of capital adjustment costs," *Review of Economic Studies*, 73(3), 611–633.

De Loecker, J. (2007): "Do exports generate higher productivity? Evidence from Slovenia," *Journal of International Economics*, 73(1), 69–98.

——— (2010): "Product differentiation, multi-product firms and estimating the impact of trade liberalization on productivity," *Econometrica.*

De Loecker, J., P. Goldberg, A. Khandelwal, and N. Pavcnik (2012): "Prices, Markups and Trade Reform," NBER Working Paper.

De Loecker, J., and J. Konings (2006): "Job reallocation and productivity growth in an emerging economy. Evidence from Slovenian manufacturing," *European Journal of Political Economy*, 22(2006), 388–408.

De Loecker, J., and F. Warzynski (2013): "Markups and firm-level export status," *American Economic Review.*

De Mel, S., D. McKenzie, and C. Woodruff (2012): "Enterprise Recovery Following Natural Disasters*," *The Economic Journal*, 122(559), 64–91.

Dixit, A., and R. Pindyck (1994): *Investment Under Uncertainty.* Princeton University Press.

Fisher-Vanden, K., E. Mansur, and Q. Wang (2012): "Costly Blackouts? Measuring Productivity and Environmental Effects of Electricity Shortages," Discussion paper, National Bureau of Economic Research.

Fisman, R. (2001): "Estimating the value of political connections," *The American Economic Review*, 91(4), 1095–1102.

Fisman, R., and J. Svensson (2007): "Are corruption and taxation really harmful to growth? Firm level evidence," *Journal of Development Economics*, 83(1), 63–75.

Foster, L., J. Haltiwanger, and C. Syverson (2008): "Reallocation, firm turnover, and efficiency: Selection on productivity or profitability?," *American Economic Review*, 98(1), 394–425.

Goldberg, P., A. Khandelwal, N. Pavcnik, and P. Topalova (2009): "Trade liberalization and new imported inputs," *American Economic Review*, 99(2), 494–500.

Harding, A., M. Söderbom, and F. Teal (2002): *The Tanzanian manufacturing enterprise survey 2002*. Centre for the Study of African Economies.

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

Midrigan, V., and D. Y. Xu (2009): "Finance and Misallocation," Discussion paper, NYU.

Moll, B. (2010): "Productivity Losses from Financial Frictions: Can Self-Financing Undo Capital Misallocation?," Discussion paper, Princeton University.

Olley, G. S., and A. Pakes (1996): "The dynamics of productivity in the telecommunications equipment industry," *Econometrica*, 64(6), 1263–1297.

Pavcnik, N. (2002): "Trade liberalization, exit, and productivity improvements: Evidence from Chilean plants," *Review of Economic Studies*, pp. 245–276.

Rankin, N., M. Söderbom, and F. Teal (2002): "The Ghanaian manufacturing enterprise survey," *Center for the Study of African Economies, Oxford, UK*.

——— (2006): "Exporting from manufacturing firms in sub-Saharan Africa," *Journal of African Economies*, 15(4), 671–687.

Restuccia, D., and R. Rogerson (2008): "Policy distortions and aggregate productivity with heterogeneous establishments," *Review of Economic Dynamics*, 11(4), 707 – 720.

Roberts, M. (1996): "Colombia, 1977-1985: Producer Turnover, Margins, and Trade Exposure," in *Industrial evolution in developing countries: micro patterns of turnover, productivity and market structure*, ed. by M. Roberts, and J. Tybout, chap. 10, pp. 227–259. Oxford University press, New York.

Söderbom, M. (2001): "Constraints and Opportunities in Kenyan Manufacturing: Report on the Kenyan Manufacturing Enterprise Survey 2000," .

Tauchen, G. (1986): "Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions," *Economics Letters*, 20(2), 177–81.

Tybout, J., and M. Westbrook (1995): "Trade liberalization and the dimensions of efficiency change in Mexican manufacturing industries," *Journal of International Economics*, 39(1), 53–78.

Van Biesebroeck, J. (2005): "Firm size matters: Growth and productivity growth in African manufacturing," *Economic Development and Cultural Change*, 53(3), 545–583.

Figure 1: Model simulation results

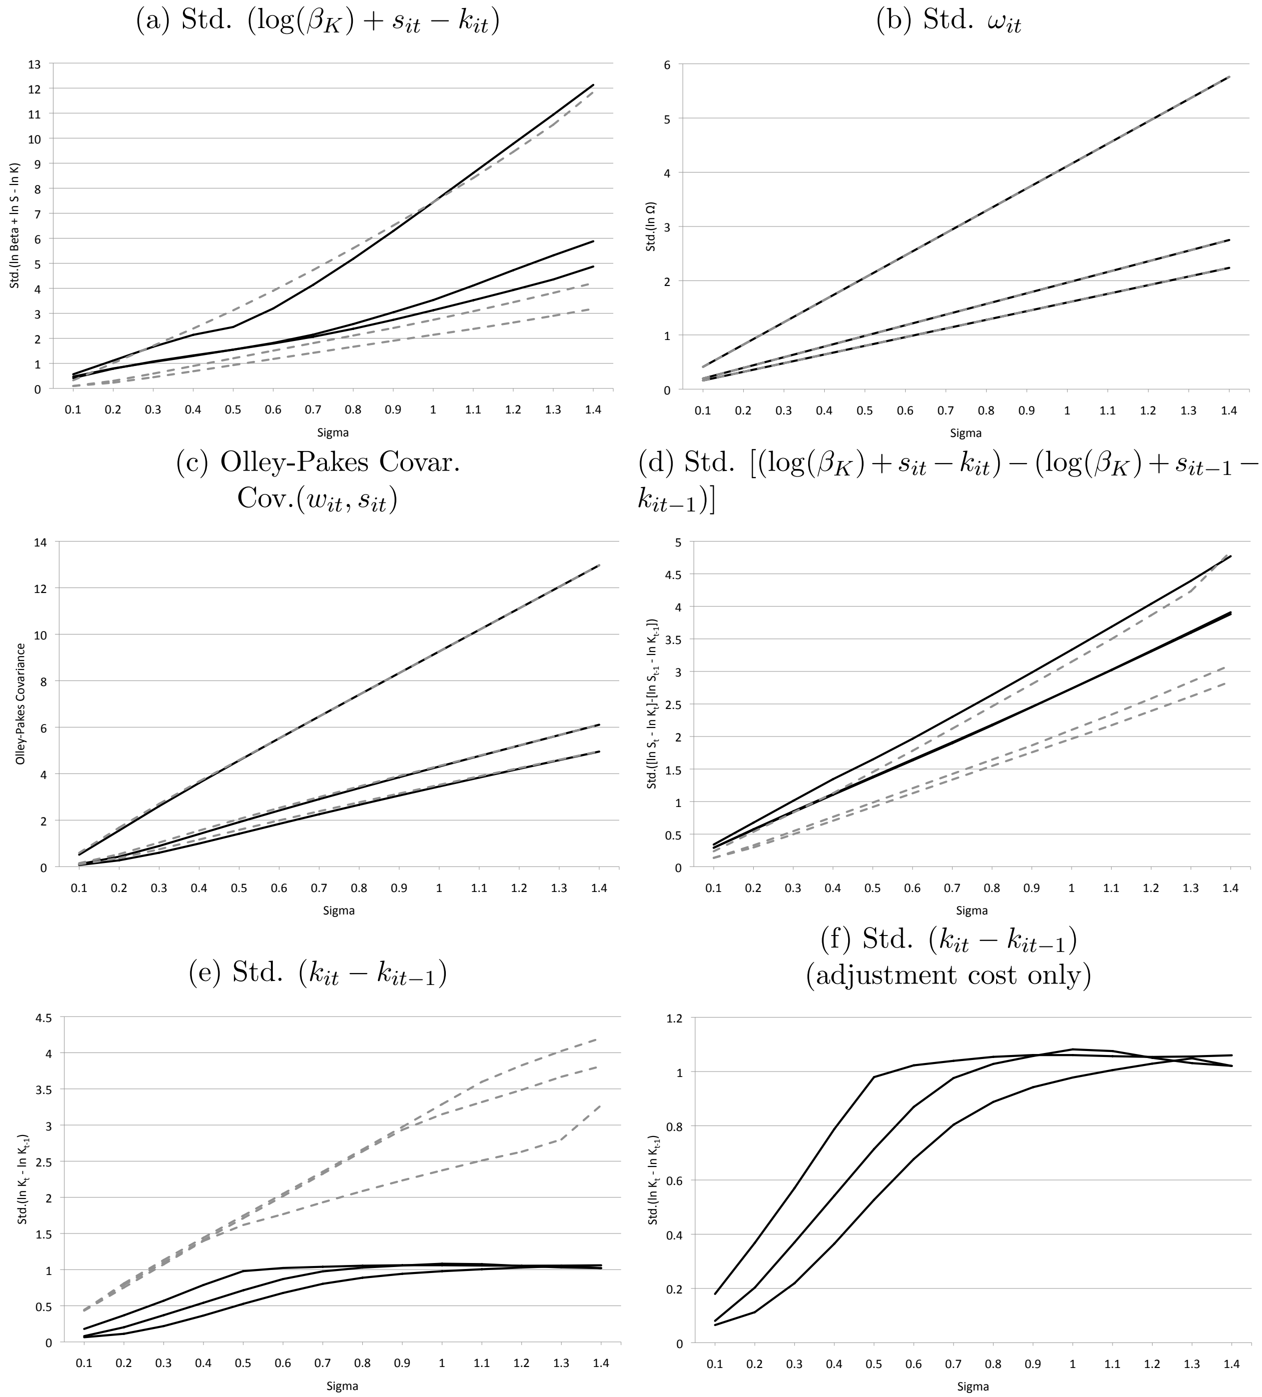

Note: In each of the figures, there are three bold lines and three grey dashed lines. The bold lines correspond to the model with both a one-period time to build and the adjustment costs. The grey dashed line shows the model without adjustment costs. Each set of bold and dashed lines has three lines stacked one above the other. In all panels, from top to bottom, these correspond to ρ equal to 0.97, 0.86 and 0.78, respectively. In panel (a), for instance, this means that, for any specification and any level of σ, as ρ increases, so does dispersion in the static marginal revenue product of capital.

Figure 2: Productivity dispersion, static misallocation and productivity volatility

(a) Productivity dispersion and productivity volatility

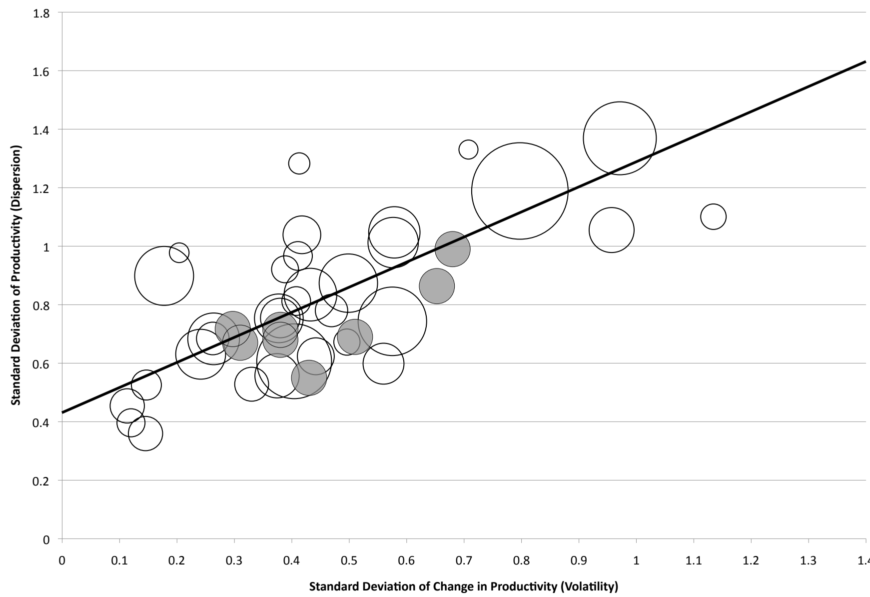

(b) Static misallocation and productivity volatility

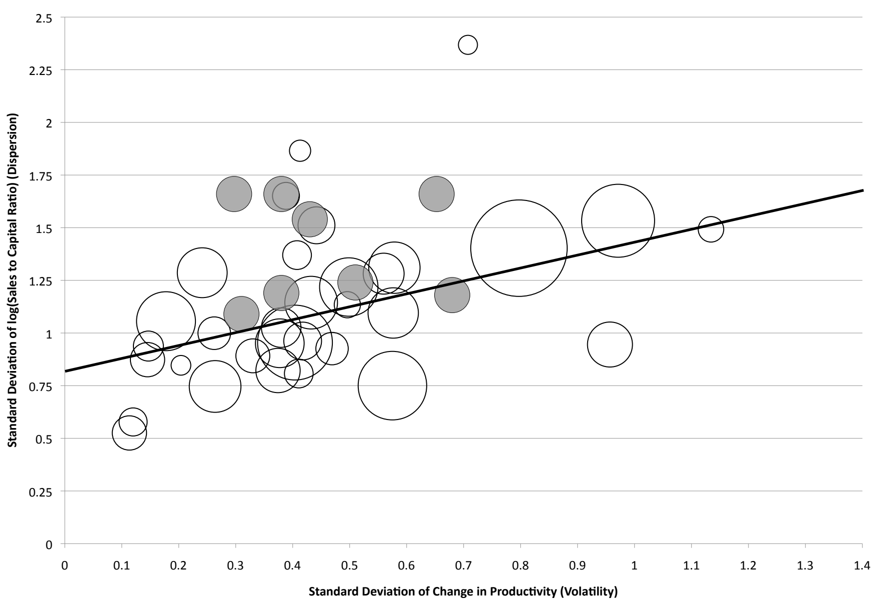

Note: Circles indicate countries. Unfilled circles are from the World Bank data. Unfilled-circle size is proportional to the number of firms per country. Filled circles are generated from country-specific data sets (size does not indicate number of firms). The bold straight line is the line–of–best–fit for the World Bank data (computed using OLS with a constant term, as per specification I in Table 5). The horizontal axis indicates the value of the standard deviation of $[\omega_{it} - \omega_{it-1}]$ in both panels. The vertical axis indicates the standard deviation in ω_{it}, where $\omega_{it} = \ln(\Omega_{it})$, and Ω_{it} is defined as in equation (3) in panel (a) (top) and the standard deviation in $\log(\beta_K) + s_{it} - k_{it}$ (the log of the sales-to-capital ratio) in panel (b) (bottom).

Figure 3: Data vs. Model: Selected moments from structural modeling

(a) Std. $(\log(\beta_K) + s_{it} - k_{it})$
One period time to build

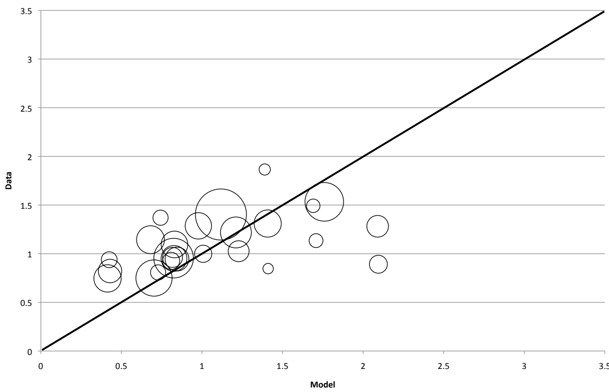

(b) Std. $(\log(\beta_K) + s_{it} - k_{it})$
plus adjustment costs

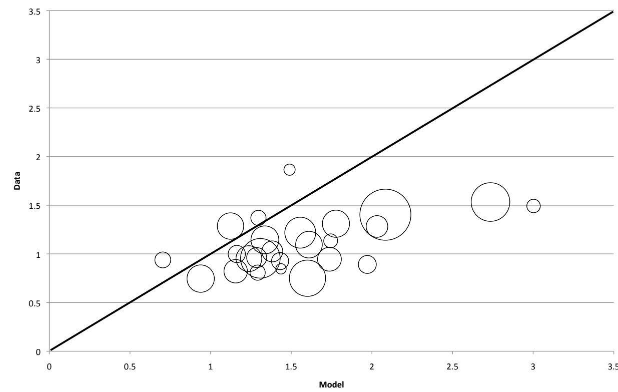

(c) Olley-Pakes Covar.
One period time to build

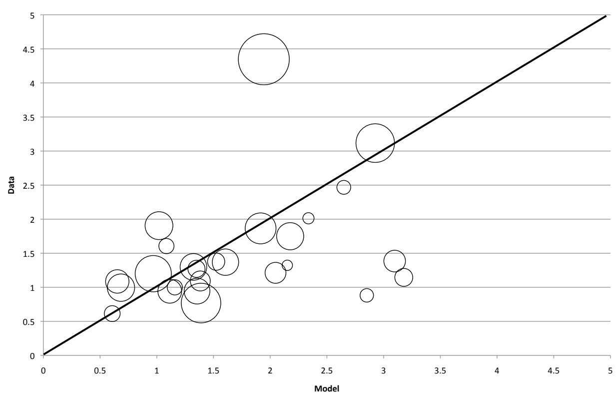

(d) Olley-Pakes Covariance
plus adjustment costs

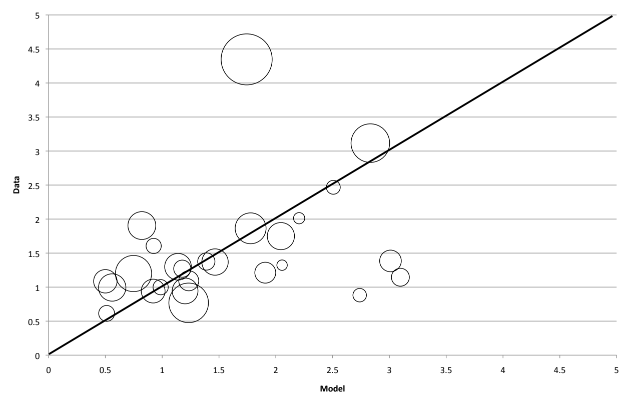

(e) Std. $(k_{it} - k_{it-1})$
One period time to build

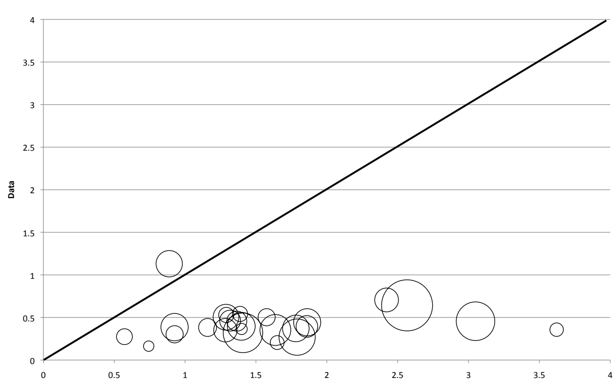

(f) Std. $(k_{it} - k_{it-1})$
plus adjustment costs

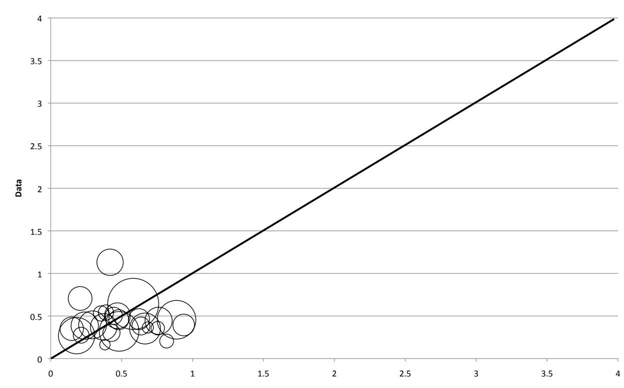

Note: Circles indicate countries. The size of the circle reflects the number of observations for the country.

Table 1: Simulation parameters

Parameter	Comments
$\epsilon = -4$ $\delta = 10\%$ $\beta = \frac{1}{1+6.5\%}$	Values drawn from Bloom (2009).
$\beta_K = 0.12$ $\beta_M = 0.47$ $\beta_L = 0.16$	Mean values in World Bank data.
$C_K^F = 0.17$ (fraction of annual sales) $C_K^Q = 0.75$	Estimated using World Bank data on Chile, see section 4.2 .
$\rho_c \in \{0.78, 0.86, 0.97\}$ $\mu = 0$ $\sigma_c \in [0.1, 1.4]$	Selected to fall within range of estimated values. See section 4.1 .
$\lambda = 1$	Scaling parameter that normalizes the price of non-capital inputs.

Table 2: Individual-country level data sets

Country	Time period	Obs	Survey Type and Criterion
Colombia	1978-1991	55,740	Census, establishments $\geq$10 workers
Chile	1979-1986	37,600	Census, establishments $\geq$10 workers
India	1989-2003	32,588	Prowess, medium & big firms
Mexico	1984-1990	22,526	Representative sample, medium & big establishments
Slovenia	1994-2000	38,856	Census, establishments firms $\geq$10 workers
Ghana	1991-2003	3,390	Stratified random sample of firms
Kenya	1992-1999	3,240	Stratified random sample of firms
Tanzania	1992-1999	2,625	Stratified random sample of firms

Note: The unit of observation in each case is the establishment/firm-year. See Appendix A for more details.

Table 3: Countries in the World Bank data sample

Region	**Country**	**Standard Deviation of TFP**	**Firms**
North Africa			
	Morocco	0.74	376
Sub-Saharan Africa			
	Benin	0.97	66
	Ethiopia	1.05	211
	Madagascar	0.78	84
	Malawi	0.75	125
	Mauritius	1.10	52
	South Africa	0.63	199
	Tanzania	0.92	58
	Zambia	0.56	157
Central Asia			
	Kyrgyzstan	0.45	94
	Tajikistan	0.36	94
	Uzbekistan	0.53	92
Middle East			
	Syria	0.67	55
South Asia			
	Bangladesh	0.60	134
	Sri Lanka	1.04	114
South East Asia			
	Indonesia	1.37	426
	Philippines	0.90	278
	Thailand	0.68	214
	Vietnam	0.68	448
Central America			
	Costa Rica	0.87	273
	Ecuador	0.62	109
	El Salvador	0.75	190
	Guatemala	1.06	162
	Honduras	1.01	203
	Nicaragua	0.83	222
South America			
	Brazil	0.68	85
	Chile	1.19	745
	Guyana	1.33	29
	Peru	0.98	31
Europe			
	Moldova	0.53	72
	Lithuania	0.81	66
	Poland	0.40	63
	Turkey	1.28	36

Note: Productivity is measured using gross output.

Table 4: Firm-level summary statistics by data set

	Chile		**Colombia**		**India**		**Mexico**		**Slovenia**		**African Countries**		**World Bank**	
	Mean	Std.	Mean	Std.	Mean	Std.	Mean	Std.	Mean	Std.	Mean	Std.	Mean	Std.
Workers	44	89	71	185	n.a.	n.a.	339	379	49	202	831	215	285	875
ln(Value Added) - s_{it}	-0.90	0.67	-0.87	0.57	-0.63	0.46	-0.67	0.64	-1.27	0.71	-10.65	2.43	-0.9	0.7
$m_{it} - s_{it}$	-0.74	0.48	-0.73	0.55	-0.98	0.70	-0.70	0.60	-0.44	0.46	-0.77	0.57	-0.6	1.0
$k_{it} - s_{it}$	-1.38	1.22	-2.12	1.27	-0.57	1.14	-3.78	1.40	-1.32	1.42	-0.80	0.57	-0.1	1.1
$w_{it} - s_{it}$	-1.92	0.72	-5.51	0.84	-2.83	1.06	-1.84	0.89	-1.66	0.92	-2.17	1.04	-1.8	1.1
Change in s_{it}	0.02	0.47	0.15	1.39	0.08	0.54	0.00	0.46	0.12	0.67	0.01	0.76	0.1	0.6
Change in k_{it}	-0.01	0.28	0.16	1.71	0.04	0.29	-0.15	0.66	0.11	0.73	0.03	0.26	0.1	0.5
Change in l_{it}	0.01	0.35	0.00	0.91	0.10	0.33	-0.01	0.30	0.05	0.41	-0.01	0.42	0.2	0.7
Change in ω_{it}^{GO}	0.014	0.35	0.11	0.50	0.05	0.30	0.02	0.40	0.02	0.35	0.01	0.44	0.0	0.6
Change in ω_{it}^{VA}	-0.02	0.82	0.02	1.15	0.07	0.51	0.05	0.73	0.09	0.64	-0.01	0.58	0.1	0.7
β_{Kict} (GO)	0.09	0.05	0.06	0.04	0.26	0.07	0.08	0.10	0.04	0.02	0.10	0.05	0.12	0.04
β_{Kict} (VA)	0.44	0.03	0.41	0.05	0.63	0.03	0.58	0.07	0.16	0.06	0.48	0.06	0.37	0.13
Statistics computed at the country level:														
std. $(\log(\beta_K) + s_{it} - k_{it})$	1.19		1.24		1.09		1.18		1.54		1.66		1.1	0.4
k_{it} Dispersion	1.98		2.05		1.61		2.13		2.51		3.15		2.1	0.4
ω_{it}^{GO} Dispersion	0.68		0.74		0.67		0.93		0.54		0.77		0.8	0.3
ω_{it}^{GO} Volatility	0.33		0.50		0.30		0.40		0.42		0.44		0.4	0.2
ω_{it}^{VA} Dispersion	1.09		1.44		1.03		1.54		0.94		1.94		1.1	0.3
ω_{it}^{VA} Volatility	0.86		1.15		0.51		0.72		0.64		0.58		0.6	0.2

Note: The African countries are Ghana, Kenya and Tanzania. To be included in the final data set, a firm needed to have at least two years of information on sales, materials, assets, and salaries. We exclude firms with productivity (ω_{it}) greater than 6 log points away from the mean to remove the effect of outliers. The results reported in the paper are qualitatively unchanged if the threshold on ω_{it} is set to be 2 or 9. "GO" indicates Gross Output, "VA" indicates Value Added.

Table 5: Productivity dispersion, static misallocation, and volatility: using the World Bank data

Panel A: Productivity dispersion and volatility					
Specification	I	II (unweighted)	III	IV	V
Dependent Var:	Standard Deviation of ω_{it}, by country				
Std.$[\omega_{it} - \omega_{it-1}]$	0.86***	0.74***	0.84***	0.84***	1.06***
	(0.21)	(0.23)	(0.22)	(0.22)	(0.21)
Log Assets $(t-1)$				0.01	
				(0.01)	
Industry FE			X	X	
Constant	0.43***	0.49***	0.43***	0.43***	0.47***
	(0.09)	(0.09)	(0.09)	(0.10)	(0.14)
Firms	5563	5563	5563	5563	5274
Countries	33	33	33	33	33
R^2	.64	.45	.66	.66	.58

Panel B: Static misallocation and volatility					
Specification	I	II (unweighted)	III	IV	V
Dependent Var:	Standard Deviation of $\log(\beta_K) + s_{it} - k_{it}$, by country				
Std.$[\omega_{it} - \omega_{it-1}]$	0.67***	0.75**	0.64***	0.63***	0.83***
	(0.21)	(0.28)	(0.22)	(0.21)	(0.20)
Log Assets $(t-1)$				0.00	
				(0.01)	
Industry FE			X	X	
Constant	0.78***	0.79***	0.79***	0.77***	0.57***
	(0.10)	(0.12)	(0.10)	(0.10)	(0.12)
Firms	5563	5563	5563	5563	5274
Countries	33	33	33	33	33
R^2	.31	.22	.36	.36	.40

Note: Productivity is measured using gross output, except in column V, where it is measured using value-added. Column I and II run regressions on country-level aggregates. Column I runs a weighted OLS with weights equal to the number of firms per country, whereas Column II has equal weights for each country. Columns III and IV run regressions at the firm level (where the dependent variable and Std.$[\omega_{it} - \omega_{it-1}]$ only vary at the country level). Error induced from the use of estimated dependent and independent variables is accounted for using a bootstrap procedure where Std.ω_{it} and Std.$[\omega_{it} - \omega_{it-1}]$ are recomputed for each bootstrap replication (200 bootstrap replications are used). These standard errors are clustered by country by having the bootstrap resample countries rather than individual firms.

Table 6: Olley-Pakes Covariance Regressions: using the World Bank data

Specification	I	II	III
Dependent Var:	Olley-Pakes Covariance		
Std.$[\omega_{it} - \omega_{it-1}]$	3.16**	3.12**	2.83**
	(1.33)	(1.24)	(1.13)
Log Capital		0.02	0.01
		(0.04)	(0.03)
Industry FE			X
Constant	0.15	0.02	-0.13
	(0.45)	(0.52)	(0.47)
Firms	5563	5563	5563
Countries	33	33	33
R^2	.44	.44	.53

Note: All columns show regressions on country-level aggregates weighted by the number of firms per country. Error induced from the use of estimated dependent and independent variables is accounted for using a bootstrap procedure where the dependent variable and Std.$[\omega_{it} - \omega_{it-1}]$ are recomputed for each bootstrap replication (200 bootstrap replications are used).

Table 7: Industry-country dispersion and volatility: By data set

Dependent Var:	Panel A: Standard deviation of (ω_{it})				Panel B: Standard deviation of $(\log(\beta_K) + s_{it} - k_{it})$				
Specification:	I		II		I		II		
	Volatility	R^2	Volatility	R^2	Volatility	R^2	Volatility	R^2	Obs.
Colombia	0.04* (0.02)	0.03	0.06** (0.03)	0.17	0.15** (0.07)	0.14	0.17** (0.06)	0.19	[21,345 - 80]
Chile	0.42*** (0.10)	0.15	0.42*** (0.10)	0.16	0.24 (0.17)	0.06	0.23 (0.16)	0.06	[28,845 - 72]
India	0.46*** (0.09)	0.12	0.46*** (0.09)	0.12	0.80*** (0.19)	0.15	0.78*** (0.19)	0.15	[31,574 - 510]
Mexico	0.51*** (0.05)	0.29	0.41*** (0.07)	0.22	0.20*** (0.07)	0.03	0.24*** (0.06)	0.08	[17,679 - 903]
Slovenia	0.70*** (0.17)	0.29	0.69*** (0.16)	0.30	0.77*** (0.24)	0.13	0.76*** (0.24)	0.13	[33,602 - 133]
African countries (FE)	0.38*** (0.09)	0.47	0.38*** (0.06)	0.54	0.09** (0.04)	0.47	0.08 (0.07)	0.50	[3,055 - 78]
World Bank data	0.53** (0.09)	0.30	0.52** (0.09)	0.31	0.43*** (0.08)	0.12	0.42*** (0.08)	0.12	[5537 - 249]
World Bank data (FE)	0.23** (0.08)	0.67	0.23** (0.10)	0.67	0.28** (0.10)	0.53	0.28** (0.10)	0.53	[5537 - 249]

Note: Observations are reported as [# of obs - #industry-time observations]. Standard errors are clustered by industry-country. *, ** and *** denote significant at the 10%, 5%, and 1% or lower, respectively. Error induced from the use of estimated dependent and independent variables not accounted for. All TFP measures use TFP Gross-Output. We supress the constant term in specifications I and II, and the coefficients on productivity and capital in specification II. The results for Chile are based on value-added productivity dispersion and volatility. All results on the individual countries, and the group of African countries, are robust to the inclusion of industry and time fixed effects. (FE) refers to the inclusion of both country and industry fixed effects.

Table 8: Country-specific AR(1) coefficients

Specification: $\omega_{it} = \mu_c + \rho_c \omega_{it-1} + \sigma_c \eta_{it}$						
Country	ρ_c	se(ρ_c)	σ_c	se(σ_c)	μ_c	se(μ_c)
Individual-Country data						
Chile	0.76	0.00	0.32	0.00	0.92	0.01
Colombia	0.79	0.00	0.48	0.00	1.30	0.00
India	0.84	0.00	0.28	0.00	0.30	0.00
Mexico	0.94	0.00	0.40	0.00	0.01	0.03
Slovenia	0.55	0.00	0.41	0.00	-0.59	0.01
Ghana	0.87	0.01	0.37	0.01	0.54	0.06
Kenya	0.74	0.03	0.62	0.02	1.20	0.15
Tanzania	0.90	0.02	0.29	0.01	0.46	0.08
World Bank data						
Bangladesh	0.92	0.08	0.56	0.03	0.19	0.23
Benin	0.80	0.05	0.36	0.03	0.54	0.12
Brazil	0.94	0.04	0.26	0.02	0.26	0.12
Chile	0.68	0.02	0.70	0.02	1.08	0.07
Costa Rica	0.85	0.03	0.48	0.02	-0.09	0.03
Ecuador	0.99	0.07	0.44	0.03	0.02	0.19
El Salvador	0.86	0.03	0.36	0.02	0.14	0.05
Ethiopia	0.84	0.04	0.55	0.03	0.36	0.09
Guatemala	0.30	0.04	0.60	0.03	1.81	0.12
Guyana	1.05	0.10	0.69	0.09	-0.06	0.50
Honduras	0.71	0.03	0.50	0.02	0.66	0.10
Indonesia	0.74	0.03	0.90	0.03	0.81	0.11
Kyrgyzstan	1.00	0.03	0.11	0.01	0.01	0.05
Lithuania	0.81	0.06	0.37	0.03	0.58	0.16
Madagascar	0.79	0.06	0.44	0.03	0.66	0.20
Malawi	0.92	0.04	0.37	0.02	0.29	0.12
Mauritius	0.61	0.13	1.04	0.10	1.08	0.41
Moldova	0.94	0.03	0.14	0.01	0.14	0.08
Morocco	0.56	0.03	0.47	0.02	1.34	0.10
Nicaragua	0.76	0.03	0.38	0.02	0.54	0.08
Peru	0.98	0.04	0.20	0.03	0.11	0.12
Philippines	1.01	0.01	0.18	0.01	-0.01	0.03
Poland	1.03	0.04	0.12	0.01	-0.05	0.10
South Africa	0.95	0.03	0.24	0.01	0.28	0.10
Sri Lanka	0.85	0.03	0.38	0.03	0.41	0.10
Syria	0.92	0.10	0.49	0.05	0.12	0.21
Tajikistan	1.03	0.04	0.14	0.01	-0.13	0.08
Tanzania	1.00	0.06	0.38	0.04	0.06	0.16
Thailand	0.84	0.02	0.24	0.01	0.57	0.08
Turkey	0.93	0.05	0.40	0.05	0.27	0.16
Uzbekistan	0.97	0.07	0.33	0.02	-0.04	0.13
Vietnam	0.84	0.03	0.39	0.01	0.50	0.08
Zambia	0.68	0.05	0.33	0.02	0.89	0.12

Note: Productivity is measured using gross output. Note that the μ coefficients will not be comparable across data sets due to the use of different measurement units.

Table 9: Model fit, reported as S^2, of different specifications

Moment	Definition	Full Model	Partial Model
Dispersion in static MRPK	Std.$(\log(\beta_K) + s_{it} - k_{it})$	0.704	0.863
Productivity dispersion	Std.(ω_{it})	0.863	0.863
Olley-Pakes covariance	Cov.(w_{it}, s_{it})	0.722	0.722
Std. of change in static MRPK	Std.$[(\log(\beta_K) + s_{it} - k_{it})$		
	$-(\log(\beta_K) + s_{it-1} - k_{it-1})]$	-0.318	0.709
Std. of change in capital	Std.$(k_{it} - k_{it-1})$	0.585	-7.10

Notes: The S^2 statistic is as described in equation 18. "Full Model" has parameters as described in Table 1 and $(\mu_c, \rho_c, \sigma_c)$ as estimated in Table 8. "Partial Model" sets the adjustment cost parameters equal to zero, leaving the one-period time to build as the only adjustment friction.

Table 10: Robustness Checks: Productivity Measurement

	Dependent Variable	
	std. ω_{it}	std. $\log(\beta_K) + s_{it} - k_{it}$
Baseline	0.86***	0.67**
	(0.21)	(0.21)
Plant Level Input Shares	0.62***	0.47*
	(0.20)	(0.23)
Less Elastic Demand ($\epsilon = 2$)	0.70*	0.65**
	(0.28)	(0.18)
More Elastic Demand ($\epsilon = 6$)	0.72***	0.69***
	(0.17)	(0.15)
Productivity Estimated via OLS	0.88***	0.77***
(with industry-country fixed effects)	(0.14)	(0.13)
Drop top and bottom decile for each country	0.32*	1.10***
	(0.12)	(0.22)
Interquartile Range	0.31*	0.54**
	(0.14)	(0.16)

Note: All regressions share a common specification: y_{it} = constant + Std.$(\omega_{it} - \omega_{it-1})$ run using a weighted OLS with weights equal to the number of firms per country. 'Baseline' refers to Column I of panel A and B in Table 5. 'Plant Level Input Shares' uses plant-level labor and material shares to compute plant-level production function coefficients β_{it}. 'Less and More Elastic' computes productivity assuming either $\epsilon = 2$ or $\epsilon = 6$ (the results in the Baseline specification assume $\epsilon = 4$). 'Productivity estimated via OLS' computes production function coefficients as the coefficients of an OLS regression of log sales on log labor, materials and capital. These coefficients are allowed to vary by country-industry pair, and include a country-industry specific intercept. 'Interquartile Range' computes the dependent variables as interquartile ranges rather than standard deviations.

A Appendix Tables

Table A: Time series process, AR(1), for productivity: Using the World Bank data

Dependent Var: Productivity ω_{it}	I	II	III	IV	V
ω_{it-1}	0.88***	0.92***	0.79**	0.91***	0.91***
	(0.05)	(0.12)	(0.30)	(0.04)	(0.04)
(ω_{it-1}) •(Country Dummy)				X	X
Var.				0.07	0.16
ω_{it-1}^2		0.13*	0.15		
		(0.06)	(0.14)		
ω_{it-1}^3		-0.04**	-0.03		
		(0.01)	(0.03)		
ω_{it-1}^4		0.00**	0.00		
		(0.00)	(0.00)		
Constant	0.33*	0.01	0.09	0.22	0.49
	(0.15)	(0.11)	(0.24)	(0.11)	
Country Specific Constant					X
Var.					0.06
Variance σ					
Constant	0.56***	0.43***	0.56***	0.56***	0.45
	(0.06)	(0.04)	(0.00)	(0.00)	
Country Specific Variance			X	X	X
Var.			.23	.24	.23
Log Assets		0.02*			
		(0.01)			
Observations	5563	5563	5563	5563	5274
Countries	33	33	33	33	33
Log-Likelihood	-4636	-4366	-3355	-3352	-3352

Note: Productivity is measured using gross output. Standard Errors (in parentheses) clustered by country. 'Var.' indicates the standard deviation of the set of parameters indicated in the row above. For Column V, averages from country-level regressions are presented. The full set of coefficients and standard errors, together with those estimated using the other data sets, are presented in Table 8.

Table B: Production function coefficients: Mean estimates by country

	Labor Coefficient β_l	Material Coefficient β_m	Capital Coefficient β_k
Bangladesh	0.14	0.50	0.11
Benin	0.17	0.48	0.10
Brazil	0.17	0.48	0.11
Chile	0.15	0.44	0.16
Costa Rica	0.17	0.47	0.12
Ecuador	0.15	0.48	0.12
El Salvador	0.15	0.48	0.12
Ethiopia	0.18	0.46	0.11
Guatemala	0.17	0.47	0.11
Guyana	0.12	0.50	0.13
Honduras	0.16	0.47	0.12
Indonesia	0.15	0.48	0.12
Kyrgyzstan	0.16	0.47	0.12
Lithuania	0.17	0.44	0.14
Madagascar	0.17	0.46	0.12
Malawi	0.14	0.48	0.12
Mauritius	0.14	0.48	0.12
Moldova	0.16	0.47	0.12
Morocco	0.16	0.48	0.11
Nicaragua	0.16	0.47	0.11
Peru	0.17	0.47	0.11
Philippines	0.14	0.49	0.12
Poland	0.15	0.48	0.12
South Africa	0.16	0.47	0.12
Sri Lanka	0.15	0.48	0.11
Syria	0.16	0.48	0.11
Tajikistan	0.17	0.47	0.11
Tanzania	0.14	0.49	0.11
Thailand	0.15	0.49	0.11
Turkey	0.13	0.49	0.13
Uzbekistan	0.16	0.48	0.12
Vietnam	0.16	0.47	0.12
Zambia	0.13	0.50	0.12

Table C: Selection Bias due to Missing Data in World Bank Data

Variable	Mean	Std. Dev.	N
Log Sales	7.0	3.1	5579
	6.7	3.3	51043
Log Value Added	6.0	3.1	4719
	5.9	3.3	42230
Log Materials	6.4	3.3	5579
	5.2	3.5	46642
Log Capital	6.9	3.1	5579
	7.5	3.0	12728
Log Labor	5.2	2.9	4715
	4.8	3.1	23696
Workers	284	874	5579
	145	1010	50891
Productivity (GO)	2.3	1.0	5579
	2.4	1.2	4750
Sales to Capital Ratio	0.1	1.1	5579
	0.2	1.3	12528
Sales to Labor Ratio	2.9	2.2	5579
	3.1	3.2	37918
Change in Capital	0.1	0.5	5579
	0.1	0.5	11268
Change in Labor	0.2	0.7	4626
	0.1	0.6	14360
Change in the Sales to Capital Ratio	0.0	0.7	5579
	0.0	0.7	11017

Note: The first row shows the data used in the paper, and the second row indicates data that we dropped due to some missing observation.