Chapter Title: Plant Size Distribution and Cross-Country Income Differences

Chapter Author: Laura Alfaro, Andrew Charlton, Fabio Kanczuk

Chapter pages in book: (p. 243 - 272)

6

Plant Size Distribution and Cross-Country Income Differences

Laura Alfaro, *Harvard Business School and NBER*
Andrew Charlton, *London School of Economics*
Fabio Kanczuk, *Universidade de São Paulo*

I. Introduction

Cross-country differences in income per worker are widely known to be enormous. Per capita income in the richest countries exceeds that in the poorest countries by more than a factor of 50 (see Klenow and Rodriguez-Clare 1997; Prescott 1998; Hall and Jones 1999; Caselli 2005). An important strand of the literature trying to understand cross-country differences in per capita incomes has focused on the role of aggregate factor accumulation by abstracting from heterogeneity in the production units.[1] But there is an emerging and growing body of research that takes a different approach, focusing instead on the misallocation of resources across plants (Restuccia and Rogerson 2007; Bartelsman, Haltiwanger, and Scarpetta 2008; Hsieh and Klenow 2009). Policies' and institutions' differential effects on the business climate broadly defined might significantly influence the allocation of resources across establishments. The working hypothesis in this literature is that not only the level of factor accumulation but also how these factors are allocated across heterogeneous production units matters in trying to understand income differences.

Our paper contributes to this literature by performing a development accounting exercise using a new data set of more than 20 million establishments in 79 developing and industrialized countries. Specifically, we develop a simple model of heterogeneous production units that follows Melitz (2003).[2] Plants' dynamics and policy distortions are as in Restuccia and Rogerson (2007), but we assume that production units have constant returns to scale technologies and some degree of market power, as in Hsieh and Klenow (2009). We calibrate the model to match our data set. Our calibration exercise consists in finding the profile of output taxes and subsidies needed to match each country's plant size distribution. These distortions should be interpreted as the different

types of policies that might generate these effects such as noncompetitive banking systems, product and labor market regulations, corruption, and trade restrictions. For example, specific producers might be offered, by governments, special tax deals and contracts financed by taxes on other production activities, and by noncompetitive banking systems, favorable interest rates on loans based on noneconomic factors, leading to misallocation of credit across establishments. Corruption and trade restrictions might also result in less productive firms obtaining a larger share of the market.

In practice, our calibration exercise amounts to making each artificial economy's plant size distribution match the observed plant size distribution for each country. Taking the United States as a supposedly undistorted benchmark economy, we find the distribution of plant-specific productivities needed to generate its plant size histogram.[3] We then find, for each country, the plant size specific distortions needed to match its plant size histogram, assuming that it faces the same distribution of productivity as the U.S. economy. This enables us to calculate how much aggregate output is being wasted as a result of misallocation attributable to distortions.

To make them directly comparable, our results are reported using the same framework as Caselli (2005). We measure the success of our model by computing cross-country income dispersion under the assumption that all countries have the same productivity. In other words, we calculate the extent to which differences in the misallocation of resources (as well as differences in the amount of physical and human capital resources) explain dispersion in income per worker.

We find misallocation of resources across plants to be a powerful explanatory factor of cross-country differences in income. For our benchmark calibration, the model explains 58% of the log variance of cross-country income dispersion. This figure should be compared to the 42% success rate of the usual model, which considers physical and human capital (average years of schooling).

We redo the basic experiment with subsamples of the data that are more reliable, choose different parameter calibrations, and truncate the data at different thresholds. We find that the results are not particularly driven by the parameter calibration or sample differences or biases. We conclude with a discussion of the limitations of, and possible extensions to, our exercise. The acknowledged shortcomings notwithstanding, our results suggest that misallocation of resources is a crucial determinant of income dispersion.

As noted above, the papers closest to the present study are those of Restuccia and Rogerson (2007), Bartelsman et al. (2008), and Hsieh and Klenow (2009).

Restuccia and Rogerson (2007) consider idiosyncratic policies that have the direct effect of engendering heterogeneity in the prices to individual producers and reallocation of resources across plants but that do not change aggregate capital accumulation and aggregate relative prices. Nonetheless, the authors find substantial effects of these policies on aggregate output and measured efficiency. In their benchmark model, they find that the reallocation of resources implied by such policies can lead to reductions of as much as 30% in output, even though the underlying range of available technology is the same.

As these authors show through simulations in their model, as distortions increase, more resources are shifted toward subsidized plants, implying higher drops in output and measured productivity. The source of the measured productivity differences is that subsidized plants become larger and taxed plants become smaller. That is, whereas in the undistorted economy all plants with the same productivity are of the same size, in the distorted economy there is a nondegenerate distribution of plant size within a plant-level productivity class. This entails an efficiency loss, which shows up in the aggregate measure of productivity.

In one set of simulations, the authors consider the case in which plants with low productivity receive a subsidy and plants with high productivity are taxed. Specifically, assuming that half the plants with low productivity are subsidized and the other half are taxed, they find that a tax rate equal to 10% entails an output and aggregate productivity loss of 13%. They also study the case in which large, presumably productive plants are subsidized. This policy is usually associated with the view that larger, productive plants need to assume a larger role in the development process. In the context of their model, however, policies that subsidize high-productivity plants, because they also distort optimal plant size, also have negative effects on output and productivity. Quantitatively, subsidizing 10% of the highest-productivity plants and taxing the rest at 40% implies a drop in output and productivity of 3%, according to their model.

Hsieh and Klenow (2009) use plant-level information from the Chinese and Indian manufacturing census data to measure dispersion in the marginal products of capital and labor within four-digit manufacturing sectors. When capital and labor are hypothetically reallocated to equalize marginal products to the extent observed in the United States, the authors find efficiency gains of 25%–40% in China and 50%–60% in India.

An analogous exercise performed on those two countries using our data yields similar results.

Our work also relates to that of Bartelsman et al. (2008), who, using an Olley and Pakes (1996) decomposition of industry-level productivity, find evidence of considerable cross-country variation in allocative efficiency in a sample of 24 industrial, transition, and emerging economies. They show that their simulated model cannot fully match some key aspects of the firm dispersion observed in the actual data, which should be seen as a limitation of this type of model.

The rest of the paper is organized as follows. Section II describes the data set and its characteristics. The model is presented in Section III and its calibration detailed in Section IV. The results are discussed in Section V. In Section VI, we carry out a number of robustness tests, and in Section VII we discuss some limitations and unaddressed extensions of our analysis and present tentative conclusions.

II. Data Description

Recent theoretical work in macroeconomics, trade, and development has emphasized the importance of heterogeneity in production units and the level of their dynamism to economic activity. Cross-country empirical investigations at the firm and establishment levels, however, are notoriously challenging because of the lack of data and difficulty of comparing the few data sets that are available.[4] There are few high-quality time-series data sets (mostly in rich countries), but there is a clear need to combine data from multiple countries (in particular, developing ones) in order to understand, for example, the role of institutional policy differences. The paucity of data is particularly acute for developing countries, and selection problems tend to be associated with biases in, and the potential endogeneity of, the cross-country samples. The reason for the data constraint is simple: economic censuses of establishments are infrequently collected because of high cost and institutional restrictions that impose an "upper bound" on research, especially in poor countries. No institution has the capacity or resources to overcome the limitation of "lack of census data" for a wide range of countries and periods. Hence, most methodologies face this restriction. Because the implications of firm heterogeneity warrant going forward despite existing data limitations, researchers have sought other sources of business "compilations" (registries, tax sources) such as data from the United Nations Industrial Development Organization, Amadeus, and the WorldBase data set used in this paper.

Dun & Bradstreet's (D&B) WorldBase is a database of public and private companies in 205 countries and territories.[5] The data, compiled from a number of sources including partner firms in dozens of countries, telephone directory records, Web sites, and self-registration, are meant to provide clients with contact details for and basic operating information about potential customers, competitors, and suppliers. Information from local insolvency authorities and merger and acquisition records are used to track changes in ownership and operations. All information is verified centrally via a variety of manual and automated checks.

It is important to note that the unit of observation in WorldBase is the establishment rather than the firm. Establishments (also referred to as plants), like firms, have their own addresses, business names, and managers but might be partly or wholly owned by other firms. It is therefore possible to observe new enterprises spawned from existing firms or, by aggregating to the firm level, examine only independent new firms. Our unit of observation in this paper is the establishment. WorldBase reports each establishment's age, number of employees, and the four-digit SIC-1987 code of the primary industry in which it operates as well as sales and exports (albeit with much less extensive coverage).[6]

The main advantage of our database is its size. Our original sample included nearly 24 million private establishments in 2003/4. Excluding territories and countries with fewer than 10 observations and those for which the Penn World Table 6.1 provides no data left us with observations in 79 countries that exhibited significant variation in international wealth and resource misallocation, precisely what we wanted for a study of development accounting.[7]

In most of the countries considered, our data set provides highly satisfactory coverage. To give some sense, we compared our data with the Statistics of U.S. Businesses collected by the U.S. Census Bureau. The U.S. 2001–2 business census records 7,200,770 "employer establishments" with total sales of $22 trillion; our data include 4,293,886 establishments with more than one employee with total sales of $17 trillion. The U.S. Census records 3.7 million small-employer establishments (fewer than 10 employees); our data include 3.2 million U.S. establishments with more than one and fewer than 10 employees.

We also compare the U.S.-owned subsidiaries in the WorldBase data with information on U.S.-owned plants from the U.S. Bureau of Economic Analysis (see fig. 1a and b). The BEA's U.S. Direct Investment Abroad: Benchmark Survey, a confidential census conducted every 5 years, covers

virtually the entire population of U.S. multinational companies (MNCs). The firm-level data are not readily available, but the BEA reports aggregate and industry-level information. In 2004, the BEA (http://bea.gov/bea/di/usdop/all_affiliate_cntry.xls) reported sales (employment) by foreign affiliates of U.S. MNCs totaling $3,238 billion (10.02 million employees). According to D&B data for 2005, the sum of all sales (employment) by foreign establishments that reported U.S. parents was $2,795 billion (10.07 million employees). Not only are the totals similar, but the distributions across countries are also consistent. Figure 1a plots the total sales (by country) of the foreign affiliates of U.S. MNCs as reported in the BEA's Benchmark Survey 2004 against the total sales (by country) of all plants in the D&B data that reported a U.S.-based parent. The correlation is striking, suggesting that the cross-country distribution of multinational activity in the D&B data matches that found in the U.S. BEA's Benchmark Survey.[8]

Although we consider the WorldBase data to be highly informative with respect to the question we posed, we are nevertheless aware of their limitations. In our final sample, the number of observations per country ranges from more than 7 million establishments in the United States to fewer than 20 in Malawi. That this variation reflects differences not only in country size but also in the intensity with which D&B samples in different countries raises the concern that our measures of size might be affected by cross-country differences in the sample frame.

For example, in countries in which coverage is lower, more established, often older, and larger enterprises might be overrepresented in the sample, which could bias our results. In particular, we know that poorer countries typically have a sizable informal sector populated
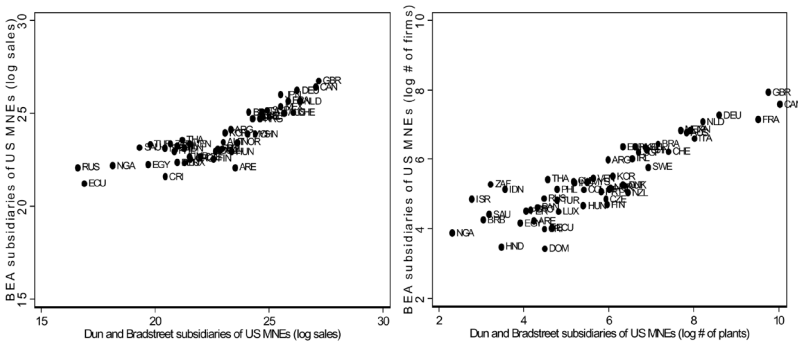

Fig. 1. Comparison of U.S. multinationals—BEA versus D&B. a, Sales of U.S. multinationals. b, Number of U.S. subsidiaries.

by small production units (Schneider and Enste 2000). Because it probably does not capture most of the informal sector, the D&B sample tends to underreport the number of smaller establishments in poor countries. To address this concern, we slice the data in different ways and redo our calculations for many possible cases.

To mitigate the potential for bias resulting from not having small establishments in poor countries represented, we truncate the data for all countries. In our benchmark exercise, we use information only for establishments with at least 20 employees, but we also work with other thresholds to test the robustness of the results. Similarly, although in our benchmark exercise we use all countries for which there are more than 10 observations, we also work with subsamples in which countries have large numbers of observations.[9]

We depict the main features of the data set in figures 2–5, in which we measure the size of an establishment in terms of the logarithm of the number of employees. Income per worker, from the Penn World Table version 6.1, refers to purchasing power parity adjusted dollars. Figures 2, 3, 4, and 5 plot, respectively, the mean, "coworker mean,"[10] variance, and skewness of the establishment size distributions of each country against income per worker (in logarithm). Note that mean size, coworker mean size, and variance size are negatively related to income, with correlations equal to −0.73, −0.73, and −0.62, respectively (significant at the 1% level). Skewness, in contrast, is positively correlated with income (0.52, also significant at the 1% level).
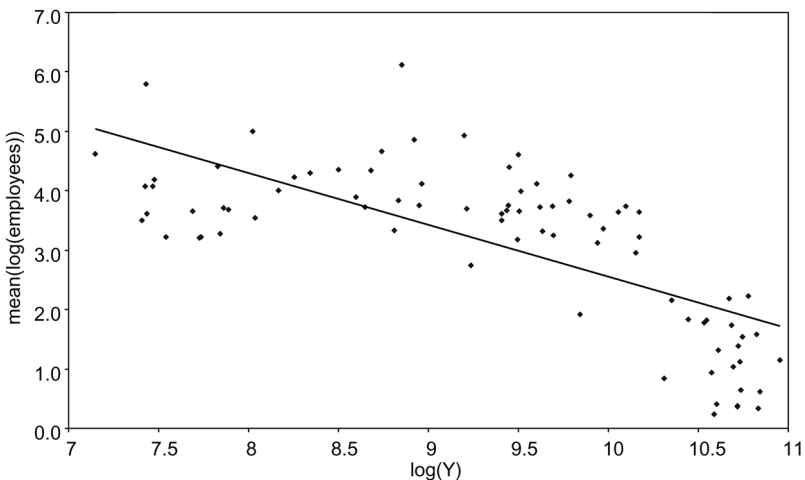

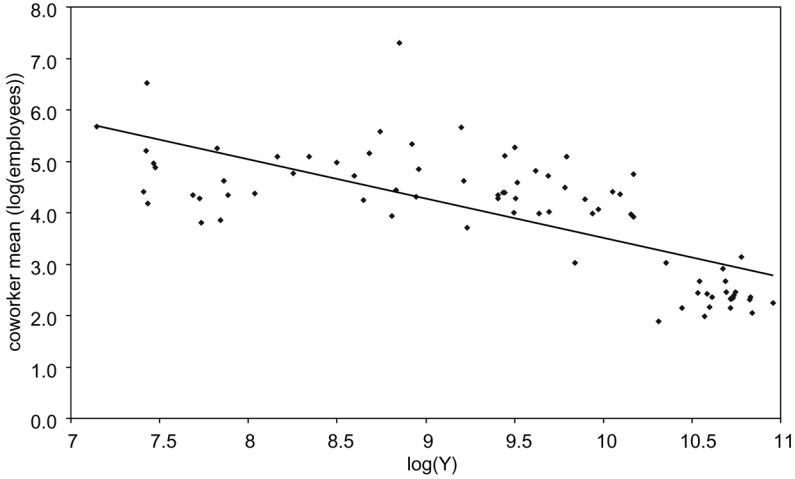**Fig. 2.**   Mean of establishment size against income per worker

Fig. 3.   Coworker mean of establishment size against income per worker

Figure 6 depicts the relation between mean size of the establishment and size of the market, measured in terms of the number of employees (as reported in Penn World Table 6.1). Note that these two variables are not correlated (the correlation is equal to 0.02, which is not significant at the 5% level).

Some additional intuition can be grasped by directly comparing selected countries' histograms. In figure 7, the horizontal axis refers to
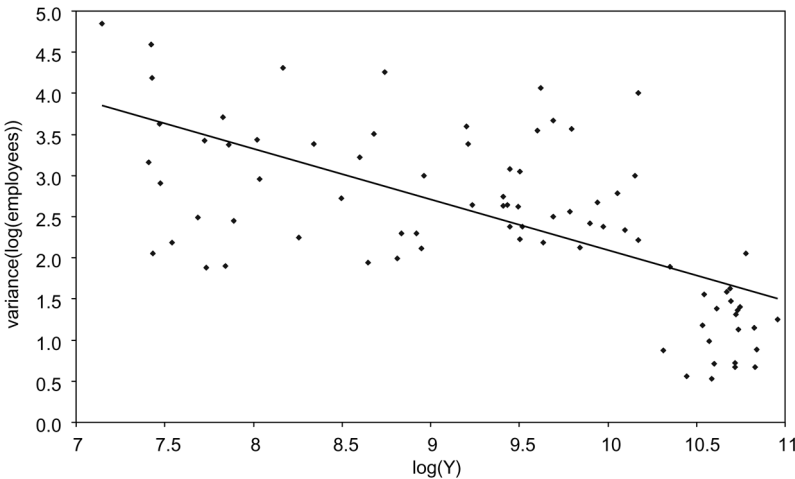

Fig. 4.   Variance in establishment size against income per worker
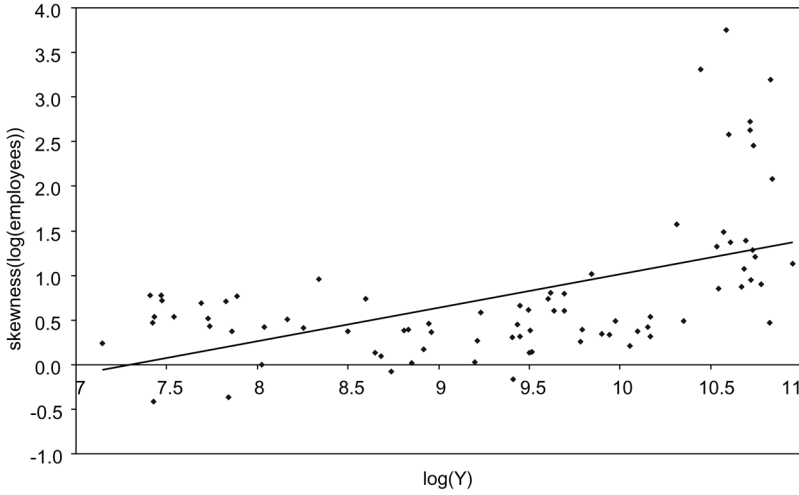
Fig. 5. Skewness of the establishment size distribution against income per worker

intervals for the number of employees, and the vertical axis reports the frequency of firms within each bracket. Notice that, starting with the United States, the curve is always decreasing, indicating that small plants constitute the most common form of production unit. Typically, histograms tend to be slightly flat, more like those of Austria or Brazil than that of the United States. India, in which plants tend to be much larger, is an extreme case; Norway, in which the plant size distribution

Fig. 6. Mean of establishment size against market size
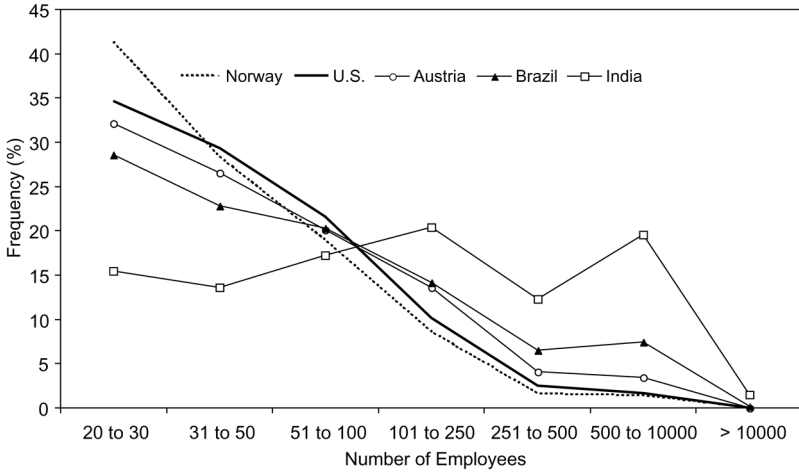
Fig. 7. Histograms of selected countries

is more concentrated in small firms than even in the United States, is a
curious case.
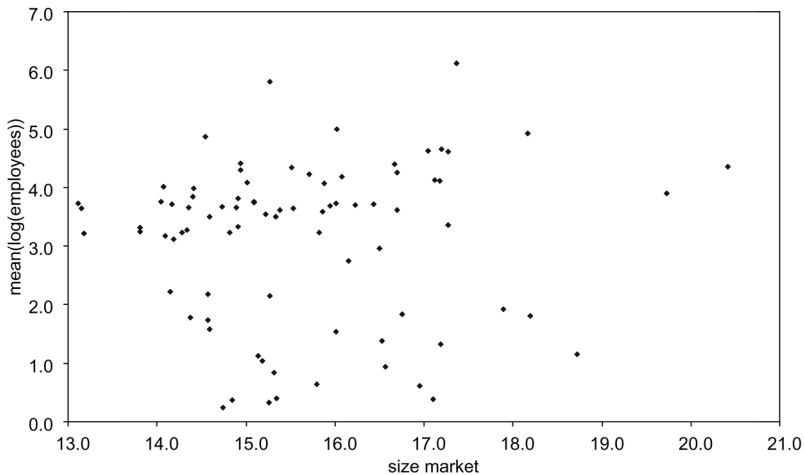
III. Model

Our model draws heavily on the work of Melitz (2003), Restuccia and
Rogerson (2007), and Hsieh and Klenow (2009). Plant dynamics and
policy distortions are as in Restuccia and Rogerson's work, but we assume
that production units have constant returns to scale technologies and some
degree of market power, as in Hsieh and Klenow's work. Because of the
degree to which our model borrows from these previous works, we at-
tempt to be as concise as possible.

Assume that the final output is a constant elasticity of substitution
aggregate of a continuum of differentiated goods, indexed by ω:

$$Y = \left[\int_{\omega \in \Omega} y_i^{(\sigma-1)/\sigma} d\omega \right]^{\sigma/(\sigma-1)}, \tag{1}$$

where the measure of the set Ω represents the mass of available goods.
This implies that demand for good ω is given by

$$y_\omega = \frac{Y}{p_\omega^\sigma}, \tag{2}$$

where p_ω denotes the price of good ω and the price of final output is normalized to one.

The production unit is the plant (establishment). There exists a continuum of plants, each of which chooses to produce a different variety ω. Plants' technologies share the same Cobb-Douglas functional form but might differ in their productivity factors, which are indexed by φ:

$$y_\varphi = AA_\varphi k_\varphi^\alpha l_\varphi^{1-\alpha}, \tag{3}$$

where A is the economy-wide productivity factor, A_φ is the plant-specific productivity factor, k_φ and l_φ are, respectively, the capital rented and labor hired by such a plant, and α is the usual capital share parameter. Conditional on remaining in operation, an incumbent plant maximizes its period profit, which is given by

$$\pi_\varphi = (1 - \tau_\varphi) p_\varphi y_\varphi - rk_\varphi - wl_\varphi, \tag{4}$$

where τ_i denotes a plant-specific output tax (or subsidy), and r and w denote the rental rates of capital and labor, respectively. Note that we assume taxes to be a function of a plant's productivity. Following Restuccia and Rogerson (2007), one should understand τ_i to be not literally a tax, but rather a general distortion. Among the different types of policies that might generate these effects are noncompetitive banking systems, product and labor market regulations, corruption, and trade restrictions.

Profit maximization, subject to the demand curve, implies the following expressions:

$$l_\varphi = \frac{y_\varphi}{AA_\varphi} \left[\frac{(1-\alpha)r}{\alpha w} \right]^\alpha, \tag{5}$$

$$k_\varphi = \frac{y_\varphi}{AA_\varphi} \left[\frac{\alpha w}{(1-\alpha)r} \right]^{1-\alpha}, \tag{6}$$

and

$$\left\{ p_\varphi - \frac{r^\alpha w^{1-\alpha} \left[\left(\frac{1-\alpha}{\alpha} \right)^\alpha + \left(\frac{\alpha}{1-\alpha} \right)^{1-\alpha} \right]}{(1-\tau_\varphi) AA_\varphi} \right\} \Big/ p_\varphi = \frac{1}{\sigma}, \tag{7}$$

which correspond to the labor and capital allocation and pricing equation (Lerner's formula). Plugging the last expression back into demand (2) gives the amount produced by each plant:

$$
y_\varphi = Y \left\{ \frac{(\sigma-1)(1-\tau_\varphi)AA_\varphi}{\sigma r^\alpha w^{1-\alpha} \left[\left(\dfrac{1-\alpha}{\alpha} \right)^\alpha + \left(\dfrac{\alpha}{1-\alpha} \right)^{1-\alpha} \right]} \right\}^\sigma . \tag{8}
$$

The equilibrium will be characterized by a mass M of plants (and thus M goods) and a distribution μ_φ of plant productivity factors over a subset of $(0, \infty)$. In such an equilibrium, the aggregate levels of capital and labor are given by $K = \int_0^\infty M k_\varphi \mu_\varphi d\varphi$ and $L = \int_0^\infty M l_\varphi \mu_\varphi d\varphi$. Plugging (8) into (6) and (7) yields expressions for K and H as functions of Y. Combining these expressions with (1) and (8), we obtain

$$
Y = A \frac{\left[\int_0^\infty (1-\tau_\varphi)^{\sigma-1} A_\varphi^{\sigma-1} M \mu_\varphi d\varphi \right]^{\sigma/(\sigma-1)}}{\int_0^\infty (1-\tau_\varphi)^\sigma A_\varphi^{\sigma-1} M \mu_\varphi d\varphi} K^\alpha L^{1-\alpha}. \tag{9}
$$

This equation will constitute the backbone expression for our calculations. As we will see in the next section, it is not necessary to specify the rest of the economic environment to use this equation. We do so, however, to gain a better understanding of the interplay of the different effects of the hypothesis on the results.

Following Restuccia and Rogerson (2007), we consider the economy to be populated by an infinitely lived representative household with preferences over streams of consumption goods that does not care about leisure. There is also a large (unbounded) pool of plants prospectively entering the industry. To enter, however, incurs a cost; prospective entrants must make their entry decision knowing that they face a distribution of potential draws for A_φ (and thus τ_φ). Although a plant's productivity and tax remain constant over time, in any given period each plant faces a constant probability of death.

The steady-state equilibrium of this model is obtained as follows. As usual, the consumer problem determines the rental rate of capital, which is a function of the time discount factor and capital depreciation rate. Given the rental rate of capital, the zero profit condition for plant entry determines the steady-state wage rate. Labor supply being inelastic, in equilibrium, total labor demand must be equal to one. It turns out that labor market clearing determines the equilibrium mass of plants.

IV. Calibration

As noted above, our data set consists of establishment size histograms for each country. The fundamental step in our calibration is thus to find plant-specific tax distortion profiles that make each country's artificial economy plant size histogram match the data. That is, we must find the distortions profile that would make the histogram of the U.S. economy, which is presumably undistorted, the histogram of another country. To do this, we need to map the plants of each country to the plants of the U.S. economy, which involves dividing each country's histogram into a certain (large) number of cells denoted by N. To achieve this mapping as simply and directly as possible, we make this division such that all countries' histograms have the same number of cells. We further give the cells of each histogram the same mass. As we shall see, however, calibration requires that across countries the cells have different masses.

Note that N denotes the number of cells, not the number of establishments, in the sample. In the case of N much greater than the number of establishments in the sample, there will be many cells for each establishment.[11] The relevant implication here is that we are capturing (at most) N moments of each distribution.

Having completed the division of the histograms, we can begin to find plant-specific productivity factors and taxes. Plugging equation (8) into equation (5) and comparing the resulting labor input for two different plants gives us

$$\frac{l_i}{l_j} = \frac{(1 - \tau_i)^\sigma A_i^{\sigma-1}}{(1 - \tau_j)^\sigma A_j^{\sigma-1}}, \tag{10}$$

where i and j refer to two plants (i.e., two different cells of the histogram). As noted earlier, we assume the U.S. economy to be sufficiently undistorted to provide a good benchmark against which to assess establishment-specific productivities. More precisely, we assume $\tau_j = 0$ for all U.S. establishments and use the U.S. data to determine the A_j factors. We do this by normalizing $A_1 = 1$ and using equation (10) to determine A_i, for $i = 2, 3, \ldots, N$.

The next step is to find the distortions for each country, which we accomplish by mapping the histogram cells of each country to the U.S. histogram cells. This is done the natural way, by sorting the histogram cells by the number of employees (see fig. 8). The mapping between any two countries' histograms is thus such that the nth smaller cell
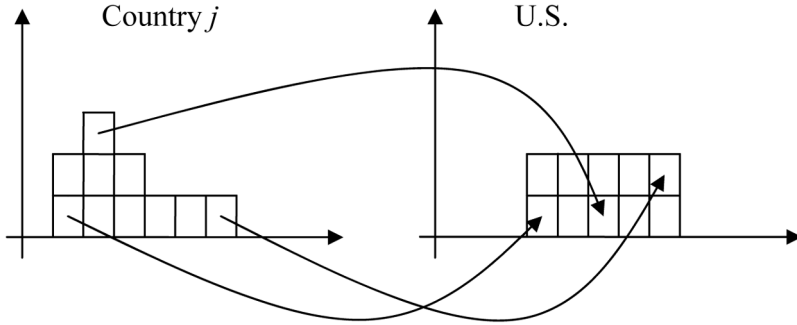
Fig. 8. Mapping histograms

of one corresponds to the *n*th smaller cell of the other.[12] This approach engenders the *minimum distortion possible* to our economies; that is, tax distortions affect establishment size but do not change the size ordering of a country's establishments. In other words, distortions never result in more productive establishments having fewer input factors than less productive establishments.[13]

Returning to equation (10), we can use the previously determined A_i's to obtain τ_i as a function of τ_1, for $i = 2, 3, \ldots, N$, for each country. More precisely, using

$$1 - \tau_i^* = (1 - \tau_i)/(1 - \tau_1), \tag{11}$$

we can obtain τ_i^* for $i = 1, 2, 3, \ldots, N$. Note that we do not need τ_1 to employ equation (9). If we plug equation (11) into equation (9), the terms on $1 - \tau_1$ in the numerator and the denominator cancel out, giving equation (9) with τ_i^* replacing τ_i.

To calibrate the mass of each country's plant distribution, we resort to the labor market–clearing equation, $L = \int_0^\infty M l_\varphi \mu_\varphi d\varphi$. In practice, after the histogram divisions, and remembering that we are normalizing the labor force to unity, this becomes

$$M = N \Big/ \sum_{i=1}^N l_i. \tag{12}$$

We borrow the technology parameters from the literature. As usual, we assume $\alpha = 1/3$. In our benchmark calibration, we set $\sigma = 6$. This parameter value delivers a 20% markup in price over marginal cost, which is in line with Rotemberg and Woodford (1992).

To relate our model to Caselli's (2005) calculations, we substitute labor for "quality-adjusted" workforce. With some abuse of notation, we rewrite equation (9) to include the human capital factor h:

$$Y = A \frac{\left[\sum_{i=1}^{N} M(1 - \tau_i^*)^{\sigma-1} A_i^{\sigma-1} \right]^{\sigma/(\sigma-1)}}{\sum_{i=1}^{N} M(1 - \tau_i^*)^{\sigma} A_i^{\sigma-1}} K^{\alpha} (Lh)^{1-\alpha}. \tag{13}$$

We then follow Caselli (2005) in calibrating the remaining parameters, the values for Y, K, and h, for each country. Briefly, Y and K are from the 6.1 version of the Penn World Tables. Capital is calculated by the perpetual inventory method, with depreciation rate equal to 6% and the initial capital determined by the initial investment rate and its geometric growth over the period. Following Hall and Jones (1999), h is measured by the formula $h = \exp[\psi(s)]$, where, following Barro and Lee (2001), ψ is piecewise linear and s denotes average number of years of schooling. We continue to normalize the size of the labor force to $L = 1$, as we evaluate output in terms of number of workers.

An interesting aspect of the calibration is that it did not require that we specify many economy parameters such as the household's preference discount factor, plants' entry costs, and the probability that a plant exits the market. Such a specification would have been necessary to obtain a complete characterization of the equilibrium including determination of factor prices and tax distortions (i.e., τ_1). We would also have had to use the "free-entry" condition, which was not required for our purposes.

For the purposes of this paper, the distribution of plant size is a summary statistic of the resource misallocation for each country. That is, in order to employ the chosen model to determine cross-country income differences, we do not need to know entry costs or the probability that a plant dies, which presumably depends on the bankruptcy laws. But an analysis of these characteristics is probably essential for drawing explicit policy implications related to interplant misallocation of resources.

V. Results

To perform the calibration, we make the number of cells, N, equal to 100,000. With that, the artificial histogram becomes a good approximation of the real data histogram even when establishment size distribution is extremely skewed. We then obtain establishment productivity for the United States, A_i (fig. 9), and, for each country, the distortions τ_i (fig. 10).[14]

Figure 9 presents the U.S. plant size distribution and figure 10 the type of distortion needed to transform that distribution into another country's plant size distribution. Remember that τ_1, the distortion of the smallest plants, was normalized to zero for all countries. Thus, one should not understand τ_j to be indicative of the aggregate distortion. Rather, figure 10 indicates, for each country, the magnitude of distortions over plants *relative to* the distortions over small plants.

Note the considerable variety in τ profiles. For some countries, τ is not monotonic in the size of the plant; for some countries it is positive, for others negative. The cloud of τ profiles also indicates that the "median" distortion corresponds to negative values for τ, which become more negative with plant size. That is, the most typical distortion corresponds to subsidies to big plants (or taxes to small plants) that increase (decrease) with the size of the plant.

After obtaining the distortions, we calculate the impact of resource misallocation on each country's productivity. Analogously to Caselli (2005), we make

$$y = ADk^{\alpha}h^{1-\alpha}, \tag{14}$$

where $y = Y/L$ and $k = K/L$ are output per worker and capital per worker, respectively, and D is the misallocation factor, defined as

$$D \equiv \frac{\left[\sum_{i=1}^{N} M(1 - \tau_i^*)^{\sigma-1} A_i^{\sigma-1}\right]^{\sigma/(\sigma-1)}}{\sum_{i=1}^{N} M(1 - \tau_i)^{\sigma} A_i^{\sigma-1}}. \tag{15}$$
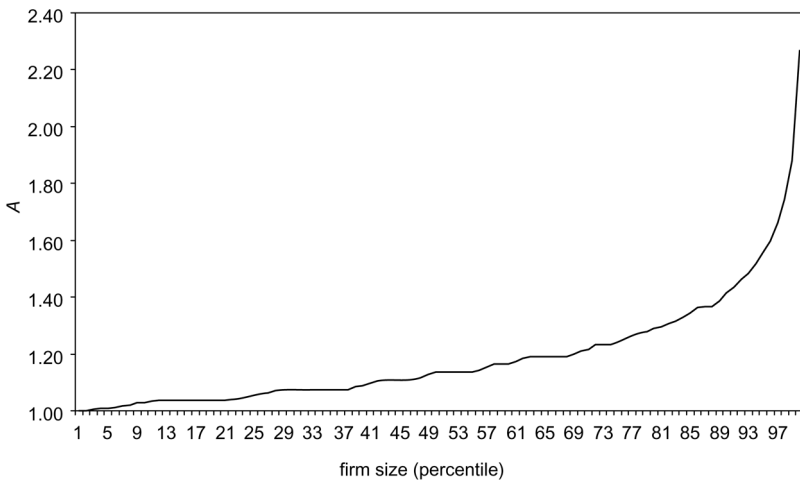

Fig. 9. Establishments' productivities
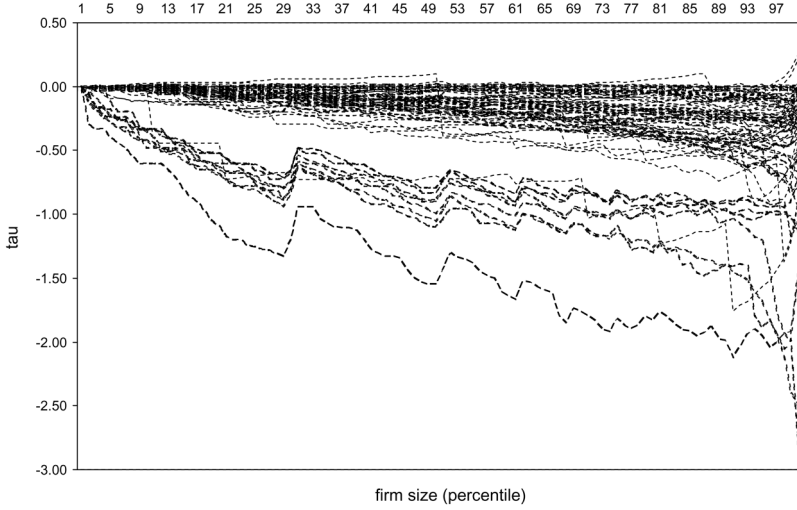
Fig. 10. Tax distortions for sample countries

To calculate the measure of the success of our exercise and compare to previous work, we define the *factor-only model*, y_{KH}, and the *misallocation model*, y_{DKH}, as, respectively,

$$y_{KH} \equiv k^\alpha h^{1-\alpha} \tag{16}$$

and

$$y_{DKH} = Dk^\alpha h^{1-\alpha}. \tag{17}$$

Our measure of success is based on the following question: What would the dispersion of incomes be if all countries had the same A? That is, we define the measure of success of the factor-only model and misallocation model, respectively, as

$$\text{Success}_{KH} \equiv \frac{\text{Var}[\log(y_{KH})]}{\text{Var}[\log(y)]} \tag{18}$$

and

$$\text{Success}_{DKH} \equiv \frac{\text{Var}[\log(y_{DKH})]}{\text{Var}[\log(y)]}. \tag{19}$$

Before we analyze the results of our benchmark experiment, summarized in table 1, it is useful to check their consistency with Caselli's

Table 1
Success in Benchmark Experiment

Subsample	Observations	Var[log(y)]	Var[log(y_{KH})]	Var[log(y_{DKH})]	Success$_{KH}$	Success$_{DKH}$
All	79	1.26	.52	.73	.42	.58
OECD	22	.047	.030	.049	.65	1.0
Non-OECD	57	.91	.36	.47	.39	.52
Africa	22	.75	.25	.27	.33	.36
Americas	23	.39	.20	.32	.50	.82
Asia and Oceania	18	.64	.32	.59	.49	.91
Europe	16	.16	.036	.061	.23	.40

(2005). Because our sample contains 79 countries and his 94, the success of the factor-only model in our case is 0.417, slightly greater than the 0.385 obtained in his calculation. Analogous observations apply to each subsample of countries.

The first and foremost observation about the misallocation model is that it displays a success measure of 0.58. The 16-percentage-point difference, about 39% higher than the corresponding figure for the factor-only model, indicates that dispersion of misallocation across countries is quantitatively important in addition to physical and human capital.

We further observe that the misallocation model displays high correlation with countries' income, equal to 0.95, as shown in figure 11. In fact, the correlation between the misallocation factor D and countries' income (in logarithms) is equal to 0.54. That is, misallocation is not only
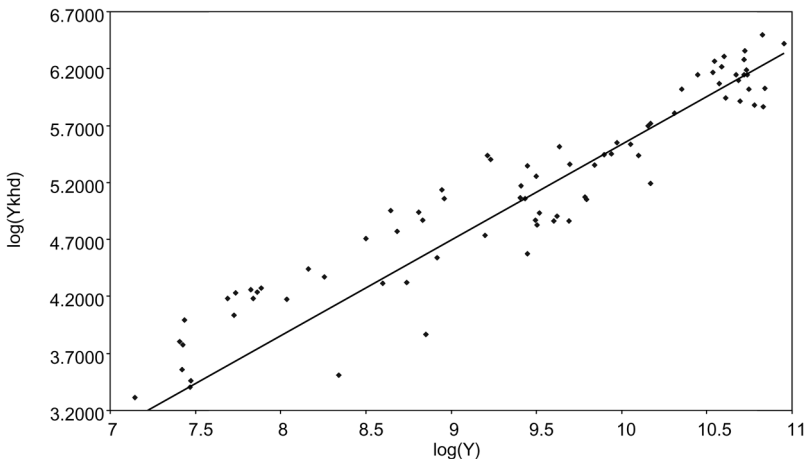

Fig. 11.   Misallocation model against income per worker

adding noise to the model but contributing to our understanding of income differences.

The subsamples of countries mostly yield conclusions that parallel those reached in Caselli's (2005) discussion of the factor-only model. Variation in log income per worker is higher in subsamples that are, on average, poorer (non-OECD, Africa). Moreover, it is more difficult to explain precisely the income differences in subsamples in which poor countries are involved, which is where a model is most needed. Unfortunately, as table 1 indicates, the misallocation model does not help much in this dimension. Although it can fully account for the income variation of OECD countries, the misallocation model can explain only about half the dispersion in the non-OECD group.

For the Asia and Oceania subgroup, however, the misallocation model performs much better than the factor-only model. It is important to note that this was not driven by a single country. Rather, the misallocation factor increased the dispersion of incomes in a homogeneous way, displaying a correlation with actual incomes of 0.93.

VI. Robustness

We test the sensitivity of our results by conducting a series of robustness checks. Specifically, we change some of the model's hypotheses and slice the data in different ways. The main results are unchanged by these manipulations.

A. Sampling Intensity

Our benchmark experiment includes all countries with a sample size greater than 10 observations (i.e., 10 establishments). This enabled us to study a large group of countries but might raise concerns about the reliability of the data and of results for countries with fewer observations. We report here the results obtained when we select only countries with sample sizes greater than 100 establishments, 1,000 establishments, and 10,000 establishments. Reducing the data set in this way has two effects: (1) it restricts the sample to countries with higher sample intensity, and (2) it excludes countries in which D&B collected little information. The latter countries tend to be poor countries in which smaller establishments are likely to be underrepresented. Coincidently, these are the countries in which the dual market (black market) operates, making the collection of data more difficult. Remember, however, that,

knowing that this could bias our results, we truncated the data in the benchmark exercise.

Table 2 displays the results. Note that as we reduce the sample of countries, the misallocation model exhibits similar performance. Although it improves for the subsample of more than 1,000 establishments, its explanatory power returns to 0.58 for more restrictive (and reliable) data sets. We take from this that our original experiment is probably a good quantitative point of reference.

B. Elasticity of Substitution

In our benchmark experiment, we calibrate the elasticity of substitution as $\sigma = 6$. Although this is our preferred calibration, there is considerable uncertainty about this parameter. In this subsection, we redo the entire experiment using $\sigma = 3.8$ and $\sigma = 10$. The former figure, used by Bernard et al. (2003), implies a 36% markup in price over marginal cost; the latter figure delivers an 11% markup in line with Basu and Fernald's (1997) findings.

Table 3 presents the results for $\sigma = 3.8$. The first line of the table, which reports results for the entire sample, gives a favorable first impression. For this calibration, success increases to 0.755. Such a reading is, however, misleading. The subsamples' analysis indicates that the misallocation model tends to overexplain the data in many cases. This might be a consequence of the small number of observations in each subsample or a problem with the data. We believe it to be more reasonable to conclude that the problem is with the calibration; in our view, with $\sigma = 3.8$, the model overestimates the effect of misallocation.

Table 4 reports the results for $\sigma = 10$. For this calibration, the misallocation model's success is 0.500, and the subsamples do not indicate any

Table 2
Success in More Reliable Subsamples

Subsample	Observations	Var[log(y)]	Var[log(y_{KH})]	Var[log(y_{DKH})]	Success$_{KH}$	Success$_{DKH}$
More than 10 establishments	79	1.26	.53	.72	.42	.58
More than 100 establishments	65	.97	.41	.58	.42	.60
More than 1,000 establishments	41	.44	.18	.32	.41	.73
More than 10,000 establishments	27	.23	.087	.13	.38	.58

Table 3
Success in Experiment with Elasticity of Substitution $\sigma = 3.8$

Subsample	Observations	Var[log(y)]	Var[log(y_{KH})]	Var[log(y_{DKH})]	Success$_{KH}$	Success$_{DKH}$
All	79	1.26	.52	.95	.42	.75
OECD	22	.047	.030	.086	.65	1.8
Non-OECD	57	.91	.36	.61	.39	.68
Africa	22	.75	.25	.32	.33	.42
Americas	23	.39	.20	.46	.50	1.16
Asia and						
Oceania	18	.64	.32	.87	.49	1.35
Europe	16	.16	.036	.11	.23	.70

inconsistencies. The question here is whether this calibration is more adequate than our benchmark. Although it is hard to know for sure, casual observation suggests that industries in developing countries tend to be less competitive than those in developed countries and are thus characterized by higher markups. With this in mind, and given that $\sigma = 10$ was obtained for the U.S. economy, we think that it is reasonable to consider the table 4 results to be a lower bound for the actual success of the misallocation model.

C. Truncation

In the benchmark experiment, we truncate our data set. The rationale for considering only establishments larger than 20 employees was that differences in the intensity with which D&B samples establishments might be affected by cross-country differences and, thus, bias our results. Poorer countries typically have large, informal sectors characterized by small establishments that might not be captured in the data set. Consequently, the data set might erroneously indicate, for example, a low mean for the plant size distribution in poor countries.

Table 4
Success in Experiment with Elasticity of Substitution $\sigma = 10$

Subsample	Observations	Var[log(y)]	Var[log(y_{KH})]	Var[log(y_{DKH})]	Success$_{KH}$	Success$_{DKH}$
All	79	1.26	.52	.63	.42	.50
OECD	22	.047	.030	.037	.65	.78
Non-OECD	57	.91	.36	.41	.39	.46
Africa	22	.75	.25	.25	.33	.34
Americas	23	.39	.20	.26	.50	.66
Asia and						
Oceania	18	.64	.32	.45	.49	.70
Europe	16	.16	.036	.044	.23	.29

Table 5
Success in Experiment with Minimum Size Equal to 10 Employees

Subsample	Observations	Var[log(y)]	Var[log(y_{KH})]	Var[log(y_{DKH})]	Success$_{KH}$	Success$_{DKH}$
All	79	1.26	.52	.78	.42	.62
OECD	22	.047	.030	.056	.65	1.2
Non-OECD	57	.91	.36	.49	.39	.54
Africa	22	.75	.25	.27	.33	.36
Americas	23	.39	.20	.34	.50	.86
Asia and						
Oceania	18	.64	.32	.63	.49	.98
Europe	16	.16	.036	.073	.23	.48

Of course, there is no clear indication of the correct threshold for
truncating the data. The fact that the official agency in some countries
(such as the Netherlands) collects data only for establishments with
more than 20 employees suggests that this is a reasonable threshold.
In any case, tables 5 and 6 address this question by reporting the experi-
ment results for the cases in which the threshold was 10 employees and
50 employees, respectively.

As expected, the results in tables 5 and 6 are qualitatively the same as
those in table 1. From a quantitative point of view, these experiments
indicate that the results are not very sensitive to the threshold. For a
relatively large change in the threshold, the main success measure varied
by less than 10%. Our view is that the benchmark results thus represent a
good compass reading.

D.   Distribution Information

To calibrate the model in our benchmark experiment we used the
plant size distribution, which has the advantage of fully exploiting

Table 6
Success in Experiment with Minimum Size Equal to 50 Employees

Subsample	Observations	Var[log(y)]	Var[log(y_{KH})]	Var[log(y_{DKH})]	Success$_{KH}$	Success$_{DKH}$
All	79	1.26	.52	.67	.42	.53
OECD	22	.047	.030	.049	.65	1.0
Non-OECD	57	.91	.36	.45	.39	.50
Africa	22	.75	.25	.26	.33	.35
Americas	23	.39	.20	.30	.50	.75
Asia and						
Oceania	18	.64	.32	.52	.49	.81
Europe	16	.16	.036	.058	.23	.37

the information contained in the data. A concern, however, is that potential noise in the data might yield misleading results. To address this issue, we redo our calculations using only the mean of each country distribution.

Note that the expression for the misallocation factor D (eq. [4]) can be split into two factors. The first has to do with the mean of the distribution, the other with its higher moments. In order to consider only the mean, we assume the second factor of the expression to be equal for all countries and redo the calculations using the first factor in the place of D:

$$
\begin{aligned}
D &\equiv \frac{\left[\sum_{i=1}^{N} M(1-\tau_i^*)^{\sigma-1}A_i^{\sigma-1}\right]^{\sigma/(\sigma-1)}}{\sum_{i=1}^{N} M(1-\tau_i^*)^{\sigma}A_i^{\sigma-1}} \\
&= M^{1/(\sigma-1)}\left\{\frac{\left[\sum_{i=1}^{N}(1-\tau_i^*)^{\sigma-1}A_i^{\sigma-1}\right]^{\sigma/(\sigma-1)}}{\sum_{i=1}^{N}(1-\tau_i^*)^{\sigma}A_i^{\sigma-1}}\right\}.
\end{aligned}
\tag{20}
$$

The results that use only the mean, reported in table 7, are similar to those for the benchmark case, except that the success rates are slightly smaller (56% rather than 58% for all countries). The exercise suggests that our findings are not driven by noise in the data set and that most of the relevant information is contained in the mean of the distribution.

E. Multiple Sectors

Our benchmark experiment assumes the economy to have only one sector. In this subsection, we redo our experiment under the assumption that the economy has multiple sectors, as in Hsieh and Klenow (2009). Specifically,

Table 7
Success in Experiment with Only Mean Information

Subsample	Observations	Var[log(y)]	Var[log(y_{KH})]	Var[log(y_{DKH})]	Success$_{KH}$	Success$_{DKH}$
All	79	1.26	.52	.71	.42	.56
OECD	22	.047	.030	.042	.65	.90
Non-OECD	57	.91	.36	.45	.39	.49
Africa	22	.75	.25	.26	.33	.34
Americas	23	.39	.20	.29	.50	.74
Asia and Oceania	18	.64	.32	.56	.49	.87
Europe	16	.16	.036	.054	.23	.35

we assume that the final good is produced by combining the output Y_s of S manufacturing industries, according to a Cobb-Douglas technology:

$$Y = \prod_{s=1}^{S} Y_s^{\theta_s}, \tag{21}$$

where $\sum_{s=1}^{S} \theta_s = 1$.

Expenditure minimization implies

$$P_s Y_s = \theta_s Y, \tag{22}$$

where P_s denotes the price of industry s and the final good price normalized to one. As before, each industry output is the aggregate of differentiated products:

$$Y_s = \left[\int_{\omega \in \Omega_s} y_i^{(\sigma-1)/\sigma} d\omega \right]^{\sigma/(\sigma-1)}, \tag{23}$$

where the measure of the set Ω_s represents the mass of available goods in sector s. There is a continuum of plants, each of which chooses to produce a different variety ω. These plants share the same Cobb-Douglas technology functional form but might differ in their productivity factors (as in eq. [3]) and maximize profits facing a plant-specific output distortion (as in eq. [4]).

To address misallocation distortion in this environment, we calculate the factor D as

$$D \equiv \prod_{s=1}^{S} \left\{ \frac{\left[\sum_{i=1}^{N_s} M_s (1 - \tau_i^*)^{\sigma-1} A_i^{\sigma-1} \right]^{\sigma/(\sigma-1)}}{\sum_{i=1}^{N_s} M_s (1 - \tau_i)^{\sigma} A_i^{\sigma-1}} \right\}^{\theta_s}. \tag{24}$$

As in Hsieh and Klenow (2009), we effectively make the misallocation factor equal to the weighted geometric average of the misallocation factor in each industry s. Note that this way of calculating the misallocation factor considers only misallocation that occurs *within* each sector. It does not consider the eventual misallocation of resources that makes sectors smaller or larger than their efficient size. A reason for calculating misallocation this way is that countries might specialize in sectors in which they have comparative advantage, in which case they could have sectors with sizes different from those of the U.S. economy (our benchmark) and nevertheless be efficient. In any case, it is noteworthy that calculating distortions this way yields a lower bound of misallocation.

To implement this exercise, we restrict our attention to countries for which we have at least 10 observations (i.e., 10 establishments) in at least 30 sectors, sector here referring to a two-digit SIC industry. This leaves a subsample of 32 countries. The choice of 30 sectors is arbitrary but turns out to be a reasonable compromise between a large number of sectors and a large number of countries.

As before, the calibration methodology consists of matching model distributions to actual histograms, but now this is done for each sector of each country. Again taking the U.S. economy as a benchmark, we find the distortion profiles for other countries. To obtain θ_s, the share of each sector in the economy, we use data on establishment revenues for the U.S. economy, also from the WorldBase data set and equation (22).

The results are presented in table 8, which compares the success measures for the one-sector and multiple-sector economies for the same subsample of countries. The success of the factor-only economy is the same in both cases since this model always contains only one sector. The success of the misallocation model is comparable in both specifications. Success is smaller in the multiple-sector model, but this is probably a consequence of the way it was formulated. This can be seen as another indication of the robustness of the results.

Another interesting observation with regard to the multiple-sector experiment is its relationship to Hsieh and Klenow's (2009) experiment. Their hypothetical "liberalizations" in China and India consider the elimination of various intrasector distortions such that capital and labor are hypothetically reallocated to equalize marginal products to the extent observed in the United States. Although they employ a different data set and calibration than we use, that the experiments share the same general framework invites comparison.

Hsieh and Klenow (2009) find the gains from reallocating resources to be on the order of 25%–40% in China and 50%–60% in India. According to our calculations, in the multiple-sector experiments, the gains for India and China are 31% and 35%, respectively. In the one-sector experiments, these gains are, again respectively, 41% and 62%. That our results

Table 8
Success in Experiment with Many Sectors

Experiment	Observations	Var[log(y)]	Var[log(y_{KH})]	Var[log(y_{DKH})]	Success$_{KH}$	Success$_{DKH}$
One sector	32	.41	.14	.23	.33	.55
Multiple sectors	32	.41	.14	.19	.33	.46

seem to be fairly consistent with theirs is a final reassuring sign of the robustness of our experiments.

VII. Conclusions

We calculated the implicit distortion needed to generate size distributions consistent with size histograms for a sample of 79 countries. We found the loss in output caused by these distortions to be quantitatively important. For our preferred calibration, the model explains 0.58 of the log variance of income per worker. This figure should be compared to the 0.42 success rate of the usual model. That is, when added to differences in resource accumulation (human and physical capital), differences in misallocation of resources add 0.16 (an increase of about 39%) to the explanatory power of our model of dispersion in cross-country income per worker. This result seems to be robust to changes in parameter calibrations and in the subsamples in which the data are more reliable (i.e., rich countries).

One potential improvement to our analysis would be to make use of richer data sets. In particular, a data set that also contained information on revenues and capital per establishment would support a richer model with many distortions, as in Hsieh and Klenow (2009). In the present case, it is possible that one distortion could cancel out or add to the effect of another distortion, thus affecting total misallocation. The work of Bartelsman et al. (2008) is a step in this direction.

There is also potential for improvement in the theoretical framework. Following the literature, we use the United States as an undistorted benchmark from which we derive other countries' distortions. An alternative approach would be to calibrate the distortions observed in the United States and obtain the characteristics of a truly undistorted economy to be used as a benchmark. This would require more modeling structure and assumptions.

Related to this last point, our exercise assumed that all countries share the same distribution of establishments' specific productivities as the United States, that is, that establishment productivities are not correlated with establishments' distortions. In contrast, richer models of plant dynamics such as that developed by Ericson and Pakes (1995) consider plants' development to be associated with "active learning." In such models, a plant's productivity tends to be connected to the distortion it faces. That is, distortions might lead a plant to invest more or less in R&D, which, in turn, would determine its productivity. As a

consequence, the effects of misallocation might be different from those calculated here.

Another modeling issue worth exploring concerns the amount of competition among establishments in an industry. Our model assumes that plants face a symmetric and constant elasticity of substitution that exogenously determines equilibrium markups. A richer specification could endogenously determine the distribution of markups and capture the impact of plants' entry costs on the degree of concentration in industries.

Finally, research should also shed light on how particular sources of inefficiency, such as credit market imperfections, macroeconomic volatility, defective bankruptcy procedures, or a malfunctioning regulatory environment, are driving cross-country differences in plant size distribution. This would be fruitful for drawing explicit policy implications. We leave this task to future work.

Endnotes

1. The consensus view in the development accounting literature is that two-thirds of these differences can be attributed to differences in efficiency or total factor productivity. Researchers have consequently attempted to explain why some countries are able to use their factors of production more efficiently and extract more output from them than others. In other words, the challenge of the growth literature is to reduce the role of measured productivity in explaining cross-country income. The traditional approach to tackling this puzzling question has been to explore the slow diffusion of technology from rich to poor countries (see Parente and Prescott 1994; Howitt 2000; Keller 2004).

2. The model implies a direct relation between productivity and size, and the model's size distribution relates to that of Lucas (1978). See Bartelsman et al. (2008) for a discussion of the limitations of this class of models and Sutton (1997) for a general discussion of firm size distribution.

3. See Davis and Henrekson (1999) and Hsieh and Klenow (2009) for a similar strategy.

4. Bartelsman, Haltiwanger, and Scarpetta (2009) review the measurement and analytical challenges of handling firm-level data and attempt to harmonize indicators of firm dynamics for a number of countries. Their harmonized data, however, are available for only a few (mostly industrialized) countries, and for many of these the data are confidential.

5. Dun & Bradstreet has been the leading U.S. source of commercial credit and marketing information since 1845. It operates in 205 countries and territories, either directly or through affiliates, agents, and associated business partners. It is also a government-approved source for assigning Standard Industrial Classification (SIC) codes to companies. Early uses of the D&B data include Caves' (1975) size and diversification pattern comparisons between Canadian and U.S. domestic plants as well as subsidiaries of U.S. multinationals in Canada and Lipsey's (1978) observations regarding the reliability of the U.S. data. More recently, Harrison, Love, and McMillian (2004) used D&B's cross-country foreign ownership information. Other research that has used D&B data includes Black and Strahan's (2002) study of

entrepreneurial activity in the United States and Acemoglu, Johnson, and Mitton's (2005) cross-country study of concentration and vertical integration.

6. Consistent capital data at the establishment level are available for few countries in our sample.

7. The countries in the sample are Algeria, Argentina, Australia, Austria, Bangladesh, Belgium, Benin, Bolivia, Brazil, Cameroon, Canada, Chile, China, Colombia, Congo, Costa Rica, Denmark, Dominican Republic, Ecuador, Egypt, El Salvador, Ethiopia, Finland, France, Gambia, Ghana, Greece, Guatemala, Haiti, Honduras, India, Indonesia, Iran, Ireland, Israel, Italy, Jamaica, Japan, Jordan, Kenya, Korea, Malawi, Malaysia, Mali, Mauritius, Mexico, Mozambique, Netherlands, New Zealand, Nicaragua, Niger, Norway, Pakistan, Panama, Papua New Guinea, Paraguay, Peru, Philippines, Portugal, Romania, Rwanda, Senegal, Sierra Leone, Singapore, Spain, Sweden, Switzerland, Syria, Thailand, Togo, Trinidad and Tobago, Tunisia, Uganda, United Kingdom, United States, Uruguay, Venezuela, Zambia, and Zimbabwe.

8. This is likely due to errors and differences in the classification of subsidiaries as U.S. or not.

9. We thank Daron Acemoglu for pointing out this problem to us and Kei-Mu Yi and Mark Bils for suggesting this solution. An alternative solution would be to take a stand on the distribution function and artificially complete the sample for poor countries. The approach of cutting out the plants below a threshold is nonparametric and probably minimizes distortions.

10. The "coworker mean" is defined as average firm size that weights each plant by its own size in calculating a mean; see Davis and Haltiwanger (1991).

11. For example, consider a country A that has two plants with 10 employees and one plant with 20 employees, and a country B that has three plants with 11 employees and one plant with 12 employees. If we choose $N = 12$, the "artificial histogram" of country A will have eight cells (or plants) with 10 employees and four cells with 20 employees. The artificial histogram of country B will have nine cells with 11 employees and three cells with 12 employees. In the case in which N is not a multiple of the number of plants, parts of different plants will be assigned to the same cell. But this is not quantitatively relevant given the large number of cells.

12. Continuing with our example, all eight cells with 10 employees of country A will be mapped to cells of 11 employees of country B. One of the cells of 20 employees of country A will be mapped to a cell of 11 employees of country B. And the remaining three cells of 20 employees of country A will be mapped to the three cells of 12 employees of country B.

13. We can think of the mapping between histograms as an identification problem that requires some assumption. Restuccia and Rogerson (2007) analyze various alternatives including the case in which there is no correlation between distortion and size distribution. The assumption adopted here has the property of minimizing distortion and thus underestimating the explanatory power of misallocation.

14. Although we chose $N = 100,000$ for our calculations, owing to graphical limitations, the figures depict the results for $N = 100$.

References

Acemoglu, D., S. Johnson, and T. Mitton. 2005. "Determinants of Vertical Integration: Finance, Contracts and Regulation." Working Paper no. 11424, NBER, Cambridge, MA.

Barro, R. J., and J. Lee. 2001. "International Data on Educational Attainment: Updates and Implications." *Oxford Economic Papers* 53:541–63.

Bartelsman, E., J. Haltiwanger, and S. Scarpetta. 2008. "Cross-Country Differences in Productivity: The Role of Allocative Efficiency." Unpublished manuscript, University of Maryland.

———. 2009. "Measuring and Analyzing Cross-Country Differences in Firm Dynamics." In *Producer Dynamics: New Evidence from Micro Data*, ed. T. Dunne,

J. B. Jensen, and M. J. Roberts. Chicago: University of Chicago Press (for NBER).

Basu, S., and J. G. Fernald. 1997. "Returns to Scale in U.S. Production: Estimates and Implications." *Journal of Political Economy* 105:249–83.

Bernard, A. B., J. Eaton, J. B. Jensen, and S. Kortum. 2003. "Plants and Productivity in International Trade." *American Economic Review* 93:1268–90.

Black, S. E., and P. E. Strahan. 2002. "Entrepreneurship and Bank Credit Availability." *Journal of Finance* 57:2807–33.

Caselli, F. 2005. "Accounting for Cross-Country Income Differences." In *Handbook of Economic Growth*, ed. P. Aghion and S. Durlauf. Amsterdam: Elsevier, North-Holland.

Caves, R. 1975. *Diversification, Foreign Investment, and Scale in North American Manufacturing Industries*. Ottawa: Economic Council of Canada.

Davis, S. J., and J. Haltiwanger. 1991. "Wage Dispersion between and within U.S. Manufacturing Plants, 1963–1986." *Brookings Papers on Economic Activity: Microeconomics*, 115–200.

Davis, S. J., and M. Henrekson. 1999. "Explaining National Differences in the Size and Industry Distribution of Employment." *Small Business Economics* 12:59–83.

Ericson, R., and A. Pakes. 1995. "Markov-Perfect Industry Dynamics: A Framework for Empirical Work." *Review of Economic Studies* 62:53–82.

Hall, B., and C. Jones. 1999. "Why Do Some Countries Produce So Much More Output per Worker than Others?" *Quarterly Journal of Economics* 114, no. 1:83–116.

Harrison, A., I. Love, and M. S. McMillian. 2004. "Global Capital Flows and Financing Constraints." *Journal of Development Economics* 75:269–301.

Howitt, P. 2000. "Endogenous Growth and Cross-Country Income Differences." *American Economic Review* 90:829–46.

Hsieh, C.-T., and P. J. Klenow. 2009. "Misallocation and Manufacturing TFP in China and India." *Quarterly Journal of Economics*, forthcoming.

Keller, W. 2004. "International Technology Diffusion." *Journal of Economic Literature* 42:752–82.

Klenow, P. J., and A. Rodriguez-Clare. 1997. "The Neoclassical Revival in Growth Economics: Has It Gone Too Far?" *NBER Macroeconomics Annual* 1997:73–102.

Lipsey, R. 1978. "The Creation of Microdata Sets for Enterprises and Establishments." *Annales de l'INSEE*, nos. 30–31:395–422.

Lucas, R. 1978. "On the Size Distribution of Business Firms." *Bell Journal of Economics* 9:508–23.

Melitz, M. J. 2003. "The Impact of Trade on Intra-industry Reallocations and Aggregate Industry Productivity." *Econometrica* 71:1695–1725.

Olley, G. S., and A. Pakes. 1996. "The Dynamics of Productivity in the Telecommunications Equipment Industry." *Econometrica* 64:1263–97.

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

Prescott, E. C. 1998. "Needed: A Theory of Total Factor Productivity." *International Economic Review* 39:525–51.

Restuccia, D., and R. Rogerson. 2007. "Policy Distortions and Aggregate Productivity with Heterogeneous Plants." Working Paper no. W13018, NBER, Cambridge, MA.

Rotemberg, J. J., and M. Woodford. 1992. "Oligopolistic Pricing and the Effects of Aggregate Demand on Economic Activity." *Journal of Political Economy* 100:1153–1207.

Schneider, F., and D. H. Enste. 2000. "Shadow Economies: Size, Causes and Consequences." *Journal of Economic Literature* 38:77–114.

Sutton, J. 1997. "Gibrat's Legacy." *Journal of Economic Literature* 35:50–59.