NBER WORKING PAPER SERIES

POLICY DISTORTIONS AND AGGREGATE PRODUCTIVITY WITH HETEROGENEOUS PLANTS

Diego Restuccia
Richard Rogerson

Working Paper 13018
http://www.nber.org/papers/w13018

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2007

We are grateful for comments of seminar participants at the 2003 NBER Summer Institute (Growth Workshop), University of Rochester, University of Iowa, 2005 Penn-Phily Fed Conference on Monetary and Macroeconomics, Federal Reserve Bank of Minneapolis, UCLA, Federal Reserve Bank of Chicago, New York University, 2004 Stanford Institute for Theoretical Economics (Heterogeneity and Aggregation in Macroeconomics Workshop), 2004 Meetings of the Society for Economic Dynamics (Florence), and HEC-Montreal. Rogerson thanks the NSF for financial support. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

Policy Distortions and Aggregate Productivity with Heterogeneous Plants
Diego Restuccia and Richard Rogerson
NBER Working Paper No. 13018
April 2007
JEL No. E2,O1

ABSTRACT

We formulate a version of the growth model in which production is carried out by heterogeneous plants and calibrate it to US data. In the context of this model we argue that differences in the allocation of resources across heterogeneous plants may be an important factor in accounting for cross-country differences in output per capita. In particular, we show that policies which create heterogeneity in the prices faced by individual producers can lead to sizeable decreases in output and measured TFP in the range of 30 to 50 percent. We show that these effects can result from policies that do not rely on aggregate capital accumulation or aggregate relative price differences. More generally, the model can be used to generate differences in capital accumulation, relative prices, and measured TFP.

Diego Restuccia
Finance Department
The Wharton School
University of Pennsylvania
3620 Locust Walk
Philadelphia, PA 19104-6367
USA
diego.restuccia@utoronto.ca

Richard Rogerson
Department of Economics
College of Business
Arizona State University
Tempe, AZ 85287
and NBER
richard.rogerson@asu.edu

1 Introduction

A large literature has emerged that attempts to use versions of the neoclassical growth model to understand cross-country differences in per capita incomes. A common assumption in much of this literature is a constant returns to scale aggregate production function that abstracts from heterogeneity in production units. Perhaps not surprisingly, therefore, much of this literature has been concerned with understanding the role of aggregate accumulation and how aggregate accumulation is affected by differences in (aggregate) relative prices.

Many important insights have emerged from this work. The thesis of this paper, however, is that the allocation of aggregate resources across uses may also be important in understanding cross-country differences in per capita incomes. That is, it is not only the level of factor accumulation that matters, but also how these factors are allocated across heterogeneous production units. And as a result, it is not only aggregate relative prices that matter but also the relative prices faced by different producers. Policies that leave aggregate relative prices unchanged but distort the prices faced by different producers will influence how resources are allocated across productive units and can potentially have substantial effects. Indeed there is substantial evidence of the importance of capital and labor allocation across plants as a determinant of aggregate productivity. For instance, Baily, Hulten, and Campbell (1992) document that about half of overall productivity growth in U.S. manufacturing in the 80's can be attributed to factor reallocation from low productivity to high productivity plants.[1]

We consider a version of the neoclassical growth model that incorporates heterogeneous production units as in Hopenhayn (1992) and Hopenhayn and Rogerson (1993). In the steady state of this model there is a non-degenerate distribution of plant-level productivity and the distribution of resources across these plants is a key element of the equilibrium resource allocation. In a calibrated version of the model we then study a class of distortions that lead to no changes in aggregate prices and no changes in aggregate factor accumulation. These distortions are to the prices faced by individual producers. Whereas in the competitive equi-

[1] This aspect of the growth process is also emphasized by Harberger (1998).

librium without distortions all producers face the same prices, we examine policy distortions whose direct effect is to create heterogeneity in the prices faced by individual producers. Because of this feature we refer to these distortions as *idiosyncratic distortions* to emphasize the fact that the distortion is (potentially) different for each producer. These idiosyncratic distortions lead to a reallocation of resources across plants. Although the policies we consider do not rely on changes in aggregate capital accumulation and in aggregate relative prices, we nonetheless find substantial effects of these policies on aggregate output and measured TFP. In our benchmark model we find that the reallocation of resources implied by such policies can lead to decreases in output and TFP in the range of 30 to 50 percent, even though the underlying range of available technologies across plants is the same in all policy configurations.

The policies that we consider are simple and abstract. In particular, we analyze policies that levy plant-level taxes or subsidies to output or the use of capital or labor. In reality, the list of policies that generate idiosyncratic distortions is both long and varied.[2] For instance, non-competitive banking systems may offer favorable interest rates on loans to select producers based on non-economic factors, leading to a misallocation of credit across establishments. Recent work by Peek and Rosengreen (2005) argues that such misallocation is highly prevalent in Japan. Banerjee and Minshi (2004) and Banerjee and Duflo (2005) present evidence that financial market imperfections lead to misallocation of credit across producers (see for instance Greenwood et al., 2007 for a model where the level of financial development affects the allocation of resources across productive uses). Governments may offer special tax deals and lucrative contracts to specific producers, all financed by taxes on other production activities. Public enterprises, which are usually associated with low productivity, may receive large subsidies from the government for their operation. Various product and labor market regulations may lead to distortions in the allocation of resources across establishments. Corruption may also lead to idiosyncratic distortions. The imposition

[2]Many policies or institutions can create heterogeneity in the costs or benefits of individual producers and affect the allocation of resources across producers. Our approach is to model the reallocation effects of these factors via taxes and subsidies. In the context of our model, the aggregate effects on output and productivity hinge on the reallocation of factors across plants with different productivity and not on whether the reallocation is caused by actual taxes and subsidies.

and enforcement of trade restrictions may also lead to distortions, and a substantial part of these may effectively be idiosyncratic. Each of these specific examples is of interest, and ultimately it is important to understand the quantitative significance of specific policies, regulations or institutions.

There is a growing literature studying the role of particular distortions on TFP and output. For example, Parente and Prescott (1999) have studied the role of monopoly-type arrangements in determining the use of inefficient technologies. Herrendorf and Texeira (2003) extend Parente and Prescott's model to allow for capital accumulation. If monopoly type arrangements are more prevalent in the investment sector, then these arrangements can lead to relative price and capital accumulation effects. Schmitz (2001) studies a similar channel, namely that low TFP in the investment sector leads to low capital accumulation, but in his model low TFP stems from a government policy supporting inefficient public enterprises. Low productivity in the investment sector seems to be at the core of low real investment rates in poor countries as argued by Hsieh and Klenow (2006). Lagos (2006) studies the effects of labor market institutions on aggregate TFP. Bergoeing et al. (2002) argue that bankruptcy laws are at the core of the fast recovery of Chile relative to Mexico in the wake of the debt crises in the early 80's (see also Bergoeing et al. 2004). Trade barriers and reforms are studied in the context of a general equilibrium model similar to ours in Chu (2002) (see also Melitz, 2003). However, we think it is valuable to begin with a generic representation of these types of policies in order to assess the overall quantitative significance of the potential effects as a complement to the studies that focus on specific channels.

More closely related to our paper are recent studies that emphasize the misallocation of resources across productive uses in aggregate productivity. Guner, Ventura, and Xu (2006) study policies that directly target the size of the establishment and find that these policies can have substantial effects in aggregate productivity. Bartelsmann, Haltiwanger, and Scarpetta (2006) study the effects of idiosyncratic distortions in the context of a model similar to ours using cross-country data on plants. Hsieh and Klenow (2006) study the impact of misallocation across establishments in explaining productivity in manufacturing in China and India. Alfaro et al. (2007) study income differences caused by the allocation of resources across heterogeneous plants using plant-level data for a sample of 80 countries.

Our model is implicitly a model of measured TFP. In addition to offering a theory to help account for differences in TFP, it can also potentially help shed light on observations about capital accumulation, relative prices, and TFP. For instance, in versions of the standard growth model, exogenous differences in TFP lead to lower capital accumulation. However, in the data there are several countries with high capital accumulation and low TFP. Our model offers a simple rationalization of this situation. If a country subsidizes the capital accumulation of low productivity units, then capital accumulation will increase but measured TFP will decline.

More generally, our model connects the literatures on capital accumulation and TFP. The literature on the role of capital accumulation emphasizes the impact of aggregate policy distortions on the return to capital investments, capital accumulation, and output, but TFP levels are exogenous and constant across countries.[3] Models of TFP, such as Parente and Prescott (2000), abstract from capital accumulation. How much of the cross country per capita income differences is accounted for by capital accumulation and other factors such as total factor productivity is a subject of great controversy.[4] Our theory of plant heterogeneity offers a link of these approaches to understanding per capita income differences across countries since idiosyncratic policy distortions can potentially lead to both capital accumulation and measured TFP differences.

The paper is organized as follows. In the next section we describe the model in detail, and in Section 3 we show how to construct the steady state equilibrium of the model. In Section 4 we calibrate our benchmark economy to data for the U.S. and in Section 5 we analyze the quantitative effects of idiosyncratic distortions in our calibrated model. Section 6 concludes.

[3]See for instance the work of Mankiw, Romer, and Weil (1992) and Chari, Kehoe, and McGrattan (1996).

[4]See for example, Klenow and Rodriguez-Clare (1997), Hall and Jones (1999), Prescott (1999), and Mankiw (1995).

2 Economic Environment

We consider a standard version of the neoclassical growth model augmented along the lines of Veracierto (2001) to allow for plant level heterogeneity as studied by Hopenahyn (1992) and Hopenhayn and Rogerson (1993). Plants have access to a decreasing returns to scale technology, pay a one-time fixed cost of entry, and a fixed cost of operation every period. Plants may die stochastically at an exogenous rate and hence, in steady state, there is ongoing entry and exit. We abstract from plant-level productivity dynamics by assuming that the productivity level of the plant remains constant over time. We study the competitive equilibrium of this model in which plants take the wage rate and the rental rate of capital as given and make zero expected profits. We then analyze policy distortions that affect output or factor prices faced by individual plants, the allocation of factors across plants, and therefore, aggregate measured TFP. In what follows we describe the environment in more detail.

2.1 Base Model

There is an infinitely-lived representative household with preferences over streams of consumption goods at each date described by the utility function,

$$\sum_{t=0}^{\infty} \beta^t u(C_t),$$

where C_t is consumption at date t and $0 < \beta < 1$ is the discount factor. Households are endowed with one unit of productive time each period and $K_0 > 0$ units of the capital stock at date 0.

Next we describe the technology. The unit of production is the plant. Each plant is described by a production function $f(s, k, n)$ that combines capital services k and labor services n to produce output. The function f is assumed to exhibit decreasing returns to scale in capital and labor jointly, and to satisfy the usual Inada conditions.[5] The parameter

[5]Our model is a single-good model in which a non-degenerate distribution of establishment sizes is sus-

s varies across plants and will capture the fact that technology varies across plants. Since our goal is to focus on the cross-sectional heterogeneity of plants we abstract from time series variation in s and hence assume that the value of s is constant over time for a given plant. In our quantitative work we assume

$$f(s,k,n) = sk^{\alpha}n^{\gamma}, \qquad \alpha,\gamma \in (0,1), \quad 0 < \gamma + \alpha < 1.$$

Note that in adopting such a specification we are implicitly assuming that the only difference across plants is the level of TFP. In particular, this functional form implies that capital to labor ratios are the same across plants in an equilibrium with no distortions. This assumption allows us to focus attention on the allocation of resources across units which differ along a single dimension, namely the level of TFP.

We also assume that there is a fixed cost of operation equal to c_f, measured in units of output. If the plant wants to remain in existence then it must pay the fixed cost. The net output produced by a plant that remains in existence is therefore given by $f(s,k,n) - c_f$. If a plant does not pay the fixed cost in any period then it ceases to exist.

Although plant-level TFP is assumed to be constant over time for a given plant, we assume that all plants face a probability of death. Specifically, we assume that in any given period after production takes place, each plant faces a constant probability of death equal to λ. It would be easy to allow this value to depend on the plant-level productivity parameter s, but we will assume it to be constant across types in the analysis carried out below.[6] Exogenous exit realizations are *iid* across plants and across time.

New plants can also be created, though it is costly. Specifically, in each period a new plant can be created by paying a cost of c_e measured in terms of output. After paying this cost a realization of the plant level productivity parameter s is drawn from the distribution

tained by decreasing returns at the establishment level. An alternative framework is to assume differentiated products and constant returns at the establishment level. In this alternative framework, the non-degenerate distribution of establishment sizes is sustained by curvature in preferences. Conceptually these frameworks are very similar. See Hsieh and Klenow (2006) and Alfaro et al. (2007) for empirical applications of the differentiated-products framework.

[6] As will be seen later, for our purposes what matters is the invariant distribution of plants across types, and whatever changes we introduce via λ would be undone by changes to the draws of s by new entrants.

with *cdf* $H(s)$. Draws from this distribution are *iid* across entrants. Let E_t denote the mass of entry in period t. We assume that there is an unlimited mass of potential entrants.

Feasibility in this model requires:

$$C_t + X_t + c_e E_t \leq Y_t - M_t c_f,$$

where C_t is aggregate consumption, X_t is aggregate investment, E_t is aggregate entry, Y_t is aggregate output, and M_t is the mass of producing firms. As it is standard, the aggregate law of motion for capital is given by:

$$K_{t+1} = (1 - \delta)K_t + X_t.$$

2.2 Policy Distortions

Our focus is on policies that create idiosyncratic distortions to plant-level decisions and hence cause a reallocation of resources across plants. As mentioned in the introduction, many different types of policies may generate such effects. While it is of interest to understand each such policy individually, the approach we take here is to analyze a generic family of distortions of this type. Specifically, we assume that each plant faces its own output tax or subsidy. In what follows we will simply refer to this distortion as the output tax, with the understanding that tax rates less than zero are possible and reflect subsidies. We will use τ to generically refer to the plant-level tax rate. To simplify our analysis we assume that τ can take on three values: a positive value reflecting that a plant is being taxed, a negative value reflecting that the plant is being subsidized, and zero reflecting no distortion for the plant. At the time of entry, the plant-level tax rate is not known, but its value is revealed once the plant draws its value of productivity s and before production takes place. Generically, we denote different specifications of policy by $\mathcal{P}(s, \tau)$ representing the probability that a plant with productivity s faces policy τ. We allow in $\mathcal{P}$ for the possibility that the value of the plant-level tax rate may be correlated with the draw of the plant-level productivity parameter, although this is not imposed in all our specifications. We also assume that the value of this tax rate remains fixed for the duration of the time for which the plant is in operation.

From the perspective of an entering plant, they face draws of s and τ, and what matters to them is the joint distribution over these pairs. For a given *cdf* H over the idiosyncratic draws of s, different specifications of policy will induce different joint distributions over pairs of (s, τ). We will represent this joint *cdf* by $G(s, \tau)$. (Alternatively, $G(s, \tau) = H(s) \times \mathcal{P}(s, \tau)$.)

A given distribution of plant-level tax and subsidies need not lead to a balanced budget for the government. We assume that budget balance is achieved on a per-period basis by either lump-sum taxation or redistribution to the representative consumer. We denote the lump-sum tax by T_t. Because our model does not have a labor/leisure decision lump-sum taxes have no effect on the model's equilibrium.

3 Equilibrium

We focus exclusively on the steady-state competitive equilibrium of the model. In a steady-state equilibrium the rental prices for labor and capital services will be constant, and we denote them by w and r respectively. The aggregate capital stock will be constant and there will also be a stationary distribution of plants across types. Before defining a steady-state equilibrium formally it is useful to first consider the decision problems of the agents in the model and to develop some notation. This discussion will also motivate an algorithm that can be used to recursively solve for the steady-state equilibrium. As we will see, the consumer problem will determine the steady state rental rate of capital. Given the rental rate of capital, the zero profit condition for entry of plants will determine the steady-state wage rate. Labor is supplied inelastically, and so in equilibrium total labor demand must equal unity. We show that this condition determines the amount of entry. We now go through the details.

3.1 Consumer's Problem

The consumer seeks to maximize lifetime utility subject to a budget constraint:

$$\sum_{t=0}^{\infty} p_t(C_t + K_{t+1} - (1-\delta)K_t) = \sum_{t=0}^{\infty} p_t(r_t K_t + w_t N_t + \Pi_t - T_t),$$

where p_t is the time zero price of period t consumption, w_t and r_t are the period t rental prices of labor and capital measured relative to period t output, Π_t is the total profit from the operations of all plants, and T_t is the lump-sum taxes levied by the government. N_t is total labor services supplied to the market, which will always be equal to one since the individual does not value leisure.

A standard argument using the first order conditions for this problem allows us to conclude that if there is a solution with r_t and C_t constant it must be that:

$$r = \frac{1}{\beta} - (1-\delta),$$

where r is the constant value of r_t. For future reference, the corresponding real interest rate, denoted by R, is given by

$$R = r - \delta = \frac{1}{\beta} - 1.$$

3.2 Incumbent Plant's Problem

The decision problem of a plant to hire capital and labor services is static since since there is no link between decisions made in different periods, that is, conditional upon remaining in operation a plant should simply hire labor and capital so as to maximize current period profits. And the decision of whether to remain in operation is equivalent to asking whether current period profits are non-negative since the plant's value of s does not change over time. Consider a plant with productivity level s and tax rate τ that faces (steady-state) input prices of r and w. Conditional upon producing, the maximum one period profit function $\pi(s,\tau)$ satisfies:

$$\pi(s,\tau) = \max_{n,k\geq 0} \left\{(1-\tau)sk^{\alpha}n^{\gamma} - wn - rk - c_f\right\}.$$

Conditional upon remaining in operation, optimal factor demands of this plant are thus given by:

$$\bar{k}(s,\tau) = \left(\frac{\alpha}{r}\right)^{\frac{1-\gamma}{1-\gamma-\alpha}} \left(\frac{\gamma}{w}\right)^{\frac{\gamma}{1-\alpha-\gamma}} (s(1-\tau))^{\frac{1}{1-\alpha-\gamma}},$$

$$\bar{n}(s,\tau) = \left(\frac{(1-\tau)s\gamma}{w}\right)^{\frac{1}{1-\gamma}} k^{\frac{\alpha}{1-\gamma}}.$$

Because both the plant-level productivity and tax rate are constant over time, the discounted present value of an incumbent plant is given by,

$$W(s,\tau) = \frac{\pi(s,\tau)}{1-\rho},$$

where $\rho = \frac{1-\lambda}{1+R}$ is the discount rate for the plant, R is the (steady-state) real interest rate, and λ is the exogenous exit rate.

3.3 Entering Plant's Problem

Potential entering plants make their entry decision knowing that they face a distribution over potential draws for the pair (s,τ). Letting W_e represent the present discounted value of a potential entrant, this value is given by:

$$W_e = \int_{(s,\tau)} \max_{\bar{x}\in\{0,1\}} \left[W(s,\tau), 0\right] dG(s,\tau) - c_e,$$

where the max inside the integral reflects the fact that the potential entrant will optimally decide whether to engage in production after observing their realized draw of (s,τ). We denote by $\bar{x}(s,\tau)$ the optimal entry decision with the convention that $\bar{x}=1$ means that the plant enters and remains in operation.

In an equilibrium with entry, W_e must be equal to zero since otherwise additional plants would enter. The condition $W_e = 0$ is thus referred to as the free-entry condition. Note, however, that the function $W(s,\tau)$ is completely determined by the values of endogenous variables w and r. Moreover, it is straightforward to see that this function is strictly decreasing in w and r. Since we have already argued that in steady state the value of r is determined by β and δ, it follows that there is at most one value of w for which $W_e = 0$. Hence, if there is an equilibrium with production then the free-entry condition will determine the wage rate.

3.4 Invariant Distribution of Plants

Let $\mu(s, \tau)$ denote the distribution of producing plants this period over plant level characteristics (s, τ). If the mass of entrants is E and the decision rule for production of entering plants is given by $\bar{x}(s, \tau)$ then next period's distribution of producers over (s, τ) pairs, denoted μ', satisfies:

$$\mu'(s, \tau) = (1 - \lambda)\mu(s, \tau) + \bar{x}(s, \tau)dG(s, \tau)E,$$

for all s and τ, where the first term represents the mass of incumbent plants that survive this period and the second term represents the mass of entering plants that remain in operation. In steady state the distribution μ will be constant over time, so we are interested in a fixed point of this mapping, or equivalently, an invariant distribution defined by this mapping. As long as death rates are bounded away from 0 this mapping will have a unique invariant distribution associated with it, and moreover, the invariant distribution will be linear in the mass of entry E. Letting $\hat{\mu}$ represent the invariant distribution associated with $E = 1$, it is easy to show that

$$\hat{\mu}(s, \tau) = \frac{\bar{x}(s, \tau)}{\lambda} dG(s, \tau), \qquad \forall s, \tau.$$

3.5 Labor Market Clearing

In the steady state, wage and capital rental rates determine the functions $\bar{k}(s, \tau)$, $\bar{n}(s, \tau)$, and $\bar{x}(s, \tau)$, and also the associated invariant distribution $\hat{\mu}$. Aggregate labor demand is then given by

$$N(r, w) = E \int_{(s,\tau)} \bar{n}(s, \tau) d\hat{\mu}(s, \tau).$$

Given values for w and r as determined as above, this equation can be used to determine the steady-state equilibrium level of entry. Recalling that labor supply is inelastic and equal to one, it follows that E satisfies:

$$E = \frac{1}{\int_{(s,\tau)} \bar{n}(s,\tau) d\hat{\mu}(s,\tau)}.$$

3.6 Definition of Equilibrium

We are now ready to formally define a steady-state competitive equilibrium for the economy.

A *steady state competitive equilibrium* with entry is a wage rate w, a rental rate r, a lump-sum tax T, an aggregate distribution of plants $\mu(s,\tau)$, a mass of entry E, value functions $W(s,\tau)$, $\pi(s,\tau)$, W_e, policy functions $\bar{x}(s,\tau)$, $\bar{k}(s,\tau)$, $\bar{n}(s,\tau)$ for individual plants, and aggregate levels consumption (C) and capital (K) such that:

(i) (Consumer optimization) $r = 1/\beta - (1-\delta)$,

(ii) (Plant optimization) Given prices (w, r), the functions π, W, and W_e solve incumbent and entering plant's problems and $\bar{k}$, $\bar{n}$, $\bar{x}$ are optimal policy functions,

(iii) (Free-entry) $W_e = 0$,

(iv) (Market clearing)

$$\begin{aligned} 1 &= \int_{s,\tau} \bar{n}(s,\tau) d\mu(s,\tau), \\ K &= \int_{s,\tau} \bar{k}(s,\tau) d\mu(s,\tau), \\ C + \delta K + c_e E &= \int_{s,\tau} (f(s,\bar{k},\bar{n}) - c_f) d\mu(s,\tau), \end{aligned}$$

(v) (Government budget balance)

$$T + \int_{s,\tau} \tau f(s,\bar{k},\bar{n}) d\mu(s,\tau) = 0,$$

(vi) (μ is an invariant distribution)

$$\mu(s,\tau) = E \frac{\bar{x}(s,\tau)}{\lambda} dG(s,\tau), \qquad \forall s,\tau.$$

4 Calibration

In this section we calibrate the model to data for the United States. In our calibration we treat the United States as an economy with no distortions. Several of the model's parameters are those of the growth model and we follow standard procedures for choosing those values. Relative to the growth model what is new are the parameters that determine the distribution of plants in equilibrium.

We let a period in the model correspond to one year in the data. We target a real rate of return of 4 percent, implying a value for β of 0.96. The extent of decreasing returns in the plant-level production parameter is an important parameter in our analysis. It is not straightforward how to pin down this parameter from cross-sectional data on establishments. Direct estimates of plant-level production functions across establishments and calibration studies point to a range of decreasing returns between 10 to 20 percent.[7] Our approach is to assume a conservative value for this parameter and to illustrate the quantitative implications of different values via sensitivity analysis. Our benchmark calibration assumes the extent of decreasing returns $(1 - \alpha - \gamma)$ to be 10 percent. We split the remaining share 1/3 to capital and 2/3 to labor, implying $\alpha = 0.3$ and $\gamma = 0.6$. We choose the depreciation rate of capital δ so that the investment to output ratio is equal to 20 percent. This choice implies $\delta = 0.08$. The implied capital to output ratio is 2.47, close to the observed capital to output ratio in the U.S. economy.

Another important component of the calibration is the range of values for plant-level productivity. Because we study policies that produce a reallocation of resources across plant types relative to the benchmark economy with no distortions, the range of plant-level productivity will determine the impact of factor allocation in aggregate productivity. In the benchmark economy there is a simple mapping between plant-level productivity and employment. As a result, the range of plant-level employment puts discipline on the range of plant-level productivity. In our model for the benchmark economy, the relative demand

[7]See for instance Pavcnik (2003), Atkeson, Khan, and Ohanian (1996), and Guner et al. (2006).

for labor between any two plants i and j is given by:

$$\frac{n_i}{n_j} = \left(\frac{s_i}{s_j}\right)^{\frac{1}{(1-\gamma-\alpha)}}.$$

From the U.S. Department of Commerce (1997), Census of Manufactures, the number of workers per-plant ranges from 1 to 7,090 workers. With our assumptions about α and β and normalizing the lowest level of plant productivity to 1, the above mapping implies a range of plant-level productivity from 1 to 2.43.[8]

We set the parameter $c_f = 0$, in our benchmark calibration, which implies that all plants that receive draws of s will produce output and remain in operation (since $s > 0$ for all plants). Since we focus on the steady-state implications of the model, endogenous entry and exit will affect the aggregate implications of distortions as long as they affect the invariant distribution of plants by productivity levels. We discuss the potential importance of this channel via sensitivity analysis in Section 5.4. The value of c_e is normalized to one. Effectively, any changes to this parameter can be undone by scaling the values of plant-level TFP.

The distribution H is chosen so that the invariant distribution of plant size across employment levels matches the data from the U.S. Census of Manufactures (1997). A key feature of the data is that a large number of small plants account for a small share of employment, whereas a small number of large plants account for a disproportionate large share of employment. For instance, plants with less than 10 workers are 51 percent of the plants but only 4 percent of the employment and 2.2 percent of the value added, while plants with 2,500 workers or more are 0.5 percent of the plants but 30 percent of the employment and 20 percent of the value added (see Figure 1). We approximate the distribution H on a grid with a large number of points. Because we have assumed that $c_f = 0$ and that λ is independent of s, the ratios of plant types in the invariant distribution μ are exactly the same

[8]Actually plants with more than 2,500 employees are top coded in the reported distribution of plants from the U.S. Department of Commerce (1997). We obtain a maximum of 7,090 workers by assuming that plants are uniformly distributed between 2,500 workers and this maximum to reproduce the average employment size of plants with more than 2,500 workers of 4,796.

as in the distribution H. We assume a grid of productivity values s with 100 points. We consider a log-spaced grid so we have more points at lower levels of productivity than at higher levels of productivity. From the data, we only observe the number of plants for a set of employment ranges. We assume that plants are uniformly distributed in each range so that the cumulative distribution function is a linear interpolation across the points for which we have data. In Figure 2 we document our approximated distribution H and the cumulative distribution across average plant sizes in the data. The distribution H (and therefore the invariant distribution μ) matches very well the size distribution of plants in the U.S. data.

Note that there is a close connection between the elasticity of the plant-level factor demand functions with respect to taxes and the elasticity of these functions with respect to plant-level TFP. Given our calibration procedure it follows that there is a close connection between the implied range of TFP values and the elasticity of plant-level factor demands with regard to taxes and subsidies. In particular, if the range of s values is large then these elasticities are small. We will return to this point later on in the paper when we discuss our results.

As noted earlier, we assume a constant exit rate λ across all plant productivity types and set this value to 10 percent. This generates an annual job destruction ratio of 10 percent which is roughly what Davis, Haltiwanger, and Schuh (1996) report for the U.S. manufacturing sector. Tybout (2000) reports annual exit rates for plants in developing countries that are roughly consistent with this value as well. We summarize parameter values and targets in Table 1.

Note that because we focus only on the steady state, there is no need to specify the utility function in order to solve for the equilibrium allocation. If we wanted to evaluate the welfare costs of distortions then we would need to specify the utility function, but since we will focus on quantifying the effects of various policies on TFP this will not be necessary.

It is of interest to look at some of the properties of the steady-state distributions in the benchmark economy. (See Figure 3 and Table 2.) First, although more than 50% of the plants have less than 10 workers, these plants represent a very small fraction of total value added and employment (less than 5 percent), and the same is true in the data (see again

Table 1: Benchmark Calibration to U.S. Data

Parameter	Value	Target
α	0.3	Capital income share
γ	0.6	Labor income share
β	0.96	Real rate of return
δ	0.08	Investment to output ratio
c_e	1.0	Normalization
c_f	0.0	Benchmark case
λ	0.1	Annual exit rate
s range	$[1, 2.46]$	Relative plant sizes
$H(s)$	see fig. 2	Size distribution of plants

Figure 1). Second, as commented earlier, because of the exponential functional form for the production function and the assumption that the exponents are independent of TFP, we see in Table 2 that output and labor shares are equalized, which implies that the distribution of labor and capital across plant types is the same as the distribution of output across plant sizes.

Table 2: Distribution Statistics of Benchmark Economy

	Plant Size by Employment		
	< 10	10 to 499	500 or more
Share of plants	0.51	0.47	0.02
Share of output	0.04	0.57	0.39
Share of labor	0.04	0.57	0.39
Share of capital	0.04	0.57	0.39
Average employment	4.2	64.8	1042.0

5 Quantitative Analysis of Distortions

In this section we study the quantitative impact of distortions to plant-level decision making. We report three main sets of results. Although our primary interest is in idiosyncratic distortions, we begin with an analysis of aggregate distortions. Specifically, we first analyze the consequences of an aggregate tax on output. This is of interest not only since it serves as a useful benchmark but also because the effects of this type of distortion are quantitatively different in a model with heterogeneous production units such as the one that we study than they are in a standard growth model. Although in our calibrated exercise the quantitative differences are not that large it is at least worth noting. We also study other forms of aggregate distortions that can be studied with entry and exit, namely differences in entry costs. According to Djankov et al. (2002) these costs differences are substantial across countries. Next we move on to consider idiosyncratic distortions. We first analyze the impact of idiosyncratic distortions when these distortions are uncorrelated with plant-level productivity s. Second, we analyze the impact of idiosyncratic distortions when these distortions are negatively (positively) correlated with plant level productivity, meaning that plants with low (high) values of s are subsidized and plants with high (low) levels of s are taxed.

The primary goal of these exercises is to assess the potential impact of reallocation on TFP and the cost of generating a given amount of reallocation. In general, policies that reallocate resources across plants will also have aggregate effects on capital accumulation. For example, a policy that subsidizes low productivity plants will cause a greater share of resources to be allocated to low productivity plants, as will a policy that taxes high productivity plants. But, whereas the subsidy will also cause capital accumulation to increase, the tax will cause capital accumulation to decrease. Because the effect of taxes on accumulation is relatively well-studied, in each case that we analyze we consider packages of idiosyncratic distortions such that there is no effect on aggregate capital accumulation. In this sense we focus on the TFP effects associated with reallocation and abstract from the capital accumulation effects.

5.1 Aggregate Distortions

In this subsection we report how the steady state is affected by a tax on output. For purposes of illustration we assess the consequences of a tax rate of 50 percent levied on all output. This tax rate results in a relative steady state output level of 0.62 between the distorted and undistorted economies. Note that in the growth model with a capital share equal to one half of the labor share (i.e., capital share equal to one third) this tax would generate a relative steady state level of output given by $0.5^{0.5} = 0.70$. The output effect is roughly 11 percent larger in our model than it is in the standard growth model. The reason for the difference is that in the present model there are decreasing returns to scale and when output is taxed there is also decreased entry of plants. This means that there is an increase in the number of workers per plant which decreases productivity. In fact, the drop in measured TFP in the model as a result of the tax on output is 11 percent and hence accounts for all the implied output difference between our model with plant heterogeneity and the standard growth model.

By explicitly considering a model with plant entry and exit, we can study a different form of aggregate distortion, namely the cost of entry for plants. Differences in entry costs are empirically relevant. Djankov et al. (2002) report differences in entry costs that are as large as a factor of 2 between poor and rich countries. We assess the impact of an increase in the cost of entry due to government regulation of 50 percent (specifically we increase c_e by 50 percent from the benchmark economy). This increase in entry costs reduces entry and therefore the amount of labor per plant. With decreasing returns at the plant level this entails an inefficiency. As a result, output and measured TFP fall by about 7 percent relative to the benchmark economy. Interestingly, the distortion to the entry cost induces no change in the capital to output ratio.

5.2 Uncorrelated Idiosyncratic Distortions

In this section we introduce idiosyncratic taxes and subsidies as discussed earlier. Here we assume that the distortions are uncorrelated with plant level productivity. In particular, we assume that half of the plants are taxed and half of the plants are subsidized. Such a configuration of distortions will cause resources to shift from the taxed plants to the subsidized plants. However, this will not entail a direct reallocation across productivity classes, since there is no correlation between plant-level TFP's and taxes.

We examine four different levels of this type of policy. We consider taxes of 10, 20, 30, and 40 percent. As described earlier, in each case we set the size of the subsidy so that the net effect on steady state capital accumulation is zero. This implies subsidies in the range of 5 to 7 percent.

It is interesting to note the apparent asymmetry of the size of the tax and subsidy rate. The reason for this asymmetry is that factor input demands from plants are very responsive to net factor costs in our calibration: a one percent increase in after tax returns leads to a ten percent increase in capital, holding factor prices constant. Hence, small differences in percent changes for equal size taxes and subsidies are greatly magnified, leading to the apparent asymmetry.

Table 3 summarizes the effects of the distortions on several variables of interest. The first row reports the level of output relative to the distortion free economy. Because aggregate inputs of labor and capital are the same in all cases, this is also the level of aggregate TFP relative to the distortion free economy. For completeness this is also reported in the second row. The third row reports the level of entry relative to the distortion free case. Since the total mass of plants operating is proportional to the mass of entry and the constant of proportionality is the same across all economies, this row also tells us the total mass of plants in operation relative to the distortion free economy. The final three rows report statistics related to the distortions. The variable Y_s/Y represents the output share of plants that are receiving a subsidy, the variable S/Y is the total subsidies paid out to plants receiving subsidies as a fraction of output, and the variable τ_s is the size of the subsidy required to generate a steady-state capital stock equal to that in the distortion-free economy.

Table 3: Effects of Idiosyncratic Distortions - Uncorrelated Case

	τ_t			
	0.10	0.20	0.30	0.40
Relative Y	0.98	0.95	0.94	0.94
Relative TFP	0.98	0.95	0.94	0.94
Relative E	1.00	1.00	1.00	1.00
Y_s/Y	0.80	0.93	0.98	0.99
S/Y	0.04	0.06	0.07	0.07
τ_s	0.05	0.07	0.07	0.07

We begin with the qualitative patterns. As expected, as the distortion increases so does the effect on output and TFP. Although not reported in the table, output shares across plant productivity types remain constant across all of these experiments. The source of the TFP differences is that subsidized plants become larger and taxed plants become smaller, so that whereas in the undistorted economy all plants with the same value of s are of the same size, in these economies there is a non-degenerate distribution of plant size within a plant level TFP class. With decreasing returns, this entails an efficiency loss. There is also potentially a change in the number of plants, but as the third row of the table indicates, this effect is zero, so that there is no change in the average level of capital or labor per plant.[9] As the distortion increases, the share of output accounted for by subsidized firms increases, as do the subsidy rate and the total payment of subsidies relative to output.

Next we turn to the quantitative magnitudes of these effects. Perhaps the most relevant result is that the overall magnitude of the effect on output and TFP is somewhat limited. As the table indicates, the maximum effect on TFP through this channel is 6 percent. Note that it takes a relatively small tax rate to generate the bulk of this effect. Even with a 10 percent tax rate the output share of subsidized firms is equal to 80 percent, and the maximum effect is virtually attained with a tax rate of 30 percent (see Figure 4). Although the maximum drop in TFP is relatively small, it is also interesting to note that few resources are required to finance this distortion. In particular, the total revenues needed to finance this maximum drop in TFP of 6 percent is only 7 percent of output. For the higher tax rates the values of S/Y and τ_s are virtually identical since the tax rate has decreased the tax base by so much that there is virtually no revenue generated.

It is not clear what the correct metric is to compare these distortions to aggregate distortions, but as one comparison, we ask what an aggregate proportional tax rate of 7 percent (i.e., a tax on the output of all plants) would do in this model. The answer is that output would fall by 5 percent. So, by this metric, although the maximum effect is relatively small, the size of the effect is roughly 20 percent higher than that generated by aggregate distortions.

[9]Since we focus on the quantitative importance of idiosyncratic distortions when aggregate capital accumulation is constant, most of our experiments feature a constant average employment size. More realistic reductions in average employment size across economies can be generated in the context of our model when aggregate capital accumulation is not kept constant.

It is of interest to note that the overall aggregate impact of idiosyncratic distortions depends on the fraction of plants that are taxed and subsidized. In our previous experiment we assumed that 50 percent of the plants were taxed and 50 percent subsidized. For the purpose of illustration, if 10 percent of the plants are subsidized and 90 percent taxed, the impact of a 40 percent tax would be a reduction in output and TFP of 20 percent (as opposed to 6 percent when 50 percent of the plants are subsidized). The subsidy rate required to produce this reallocation is 26 percent. Subsidizing fewer plants requires a larger subsidy to keep the aggregate capital stock constant and, as a result, a larger reallocation of factors and output across plants. This reallocation makes plants operate much farther away from their optimal size, ensuing the larger aggregate effects.

5.3 Correlated Idiosyncratic Distortions

The distortions considered in the last section were in some sense adding noise to the competitive market. Instead of all firms facing the same prices, each firm faces a different price, but there is nothing systematic about who faces what price. We found that the consequences of this were relatively minor. We now consider distortions which at least on the surface would seem to have the potential to do much more damage. In particular we consider the case where plants with low TFP receive a subsidy and plants with high TFP are taxed. In particular, we assume that 50 percent of the plants with low s receive a subsidy while the rest are taxed. Table 4 summarizes the results for this case (see also Figure 5).

Table 4: Effects of Idiosyncratic Distortions - Correlated Case

	τ_t			
Variable	0.1	0.2	0.3	0.4
Relative Y	0.87	0.78	0.73	0.72
Relative TFP	0.87	0.78	0.73	0.72
Relative E	1.00	1.00	1.00	1.00
Y_s/Y	0.57	0.83	0.95	0.99
S/Y	0.20	0.32	0.38	0.40
τ_s	0.35	0.39	0.40	0.40

Qualitatively the patterns are similar to those of the uncorrelated case: as distortions increase the drop in output and TFP increases and more resources are shifted toward subsidized plants. A key difference is that in this case the distortion is not to the size distribution of plants of a given productivity, but rather to the distribution of resources across plants of varying productivity. This distortion is much more important quantitatively. As the table shows, the maximum effect on TFP and output in this case is 28 percent, almost four times the effect in the uncorrelated case. The table also shows that this distortion is somewhat more costly to finance. To achieve the TFP reduction of 28 percent, subsidies totalling 40 percent of output are required, and since there are virtually zero revenues raised from taxation, this is the amount of resources that the government must raise in lump-sum taxes. Again, as a comparison we can ask what magnitude of effects would be generated by an aggregate output tax of 40 percent, and the answer is that this tax would reduce output by 40 percent. So, once again the message is that by this metric the effects due to idiosyncratic distortions seem to be comparable to those generated by aggregate distortions.

As it was the case for uncorrelated distortions, the magnitude of the effect on TFP gets amplified as fewer unproductive plants are subsidized. If only 10 percent of the lowest productivity plants are subsidized and the remaining taxed, there is a larger fall in output and TFP. When the tax rate is 40 percent, the drop in output and TFP is 45 percent. While the magnitude of aggregate effects is larger, the cost of this type of policy is also larger, about 83 percent of GDP.

While protecting and subsidizing low productivity plants is perversive in poor countries, large, presumably productive plants also get subsidized in some countries. The view that often motivates these policies is that larger and productive plants need to take on a bigger role in the development process. In the context of our model, policies that subsidize high productivity plants also have negative effects in output and TFP. These subsidies distort the optimal plant size even though subsidies entail a reallocation towards more productive plants. Overall the effect of these policies is a drop in measured TFP. For instance, in the context of our calibrated model, subsidizing 10 percent of the highest productivity plants and taxing the rest at 40 percent would imply a drop in output and TFP of 3 percent.

5.4 Discussion

Tax on Capital The above exercises assumed that the reallocation of resources was achieved through taxes and subsidies that were applied to plant-level output. We can also conduct the exercises assuming that either labor or capital input serves as the base. In Table 5 we present the results when capital serves as the base for two different levels of the tax rate. As before, we assume a subsidy that leaves total capital accumulation unchanged. Although the basic message is similar, there are a few differences from the case in which output is taxed/subsidized that bear mentioning. First, there is a more substantial reallocation of capital than there is of output, and this difference is particularly pronounced in the correlated case. Second, there is now also an effect on the mass of plants in operation, and this effect is of the same magnitude as the change in output and TFP. Third, the level of subsidies required to generate these changes are substantially smaller than those required in the case of output subsidies. Note that distortions levied through capital have an additional channel relative to distortions that work through output. In the case of output taxes, capital to labor ratios are unaffected by the distortions, but this is no longer true in the case of distortions that operate through factor prices. Finally, there is one result that seems somewhat perverse - namely that the subsidy rate required to maintain a constant aggregate capital stock decreases in both cases as the tax rate increases. The reason for this is that an increase in the tax rate also causes wages to decrease and this decrease in wages also affects the demand for capital.

Table 5: Idiosyncratic Distortions to Capital Rental Rates

	Uncorrelated		Correlated	
	$\tau_t = 0.50$	$\tau_t = 1.00$	$\tau_t = 0.50$	$\tau_t = 1.00$
Relative Y	0.97	0.95	0.89	0.82
Relative TFP	0.97	0.95	0.89	0.82
Relative E	0.97	0.95	0.89	0.82
Y_s/Y	0.82	0.91	0.41	0.59
K_s/K	0.89	0.96	0.65	0.83
S/Y	0.03	0.03	0.10	0.13
τ_s	0.10	0.09	0.44	0.42

Tax on Labor It turns out that our exact exercise cannot be carried out for the case of taxes and subsidies applied to labor. In particular it is generally not possible to distort the labor allocation across plants and also leave the aggregate capital stock unchanged by using only taxes and subsidies to plant-level labor. The reason for this is that labor is fixed, and any misallocation of labor necessarily affects the marginal product of capital. If more labor could be hired then this would raise the marginal product of capital and lead to increased capital accumulation, but our current formulation does not allow for this channel. One possible way to accommodate this is by adding a subsidy to capital accumulation or a subsidy to output. This leads to an extra degree of freedom and thus makes it somewhat difficult to compare the results. Given this issue, for the case of taxes and subsidies levied on labor we simply report the results for a case in which we set $\tau_t = \tau_s = 0.5$ and half of the plants are subsidized and the other half are taxed. This will necessarily result in a lower level of capital in the steady state, thus making the results not strictly comparable to those reported earlier. In view of this we direct our attention primarily to the effects on TFP rather than the effects on output. Results are reported in Table 6. The effect on TFP is somewhat larger here than in the case where taxes and subsidies are levied on output. In the uncorrelated case the drop in TFP is 8 percent and in the correlated case the drop is 30 percent, whereas with output as the base we obtained TFP decreases of 6 and 28 percent. As with the case of capital as a base, these subsidies also distort the capital to labor ratio at the plant level, thereby suggesting the possibility of somewhat larger effects. Also note that the level of subsidies required are much larger in the case of taxes on labor than in the cases of taxes on output and capital.

Sensitivity The final issue that we touch on here concerns the sensitivity of our results to changes in the underlying specification. In particular, because the TFP losses that we measure are due to moving resources across plants with different levels of TFP, one would presume that if we had larger dispersion in values of s across plants that the potential TFP effects would be much larger. In a mechanical sense this is certainly true. In this regard it is important to note that in our calibration the range of plant level TFP's is disciplined by

Table 6: Idiosyncratic Distortions to Wage Rates

	Uncorrelated	Correlated
Relative Y	0.89	0.60
Relative TFP	0.92	0.70
Relative K	0.89	0.60
Relative E	0.89	0.60
Y_s/Y	1.00	0.97
N_s/N	1.00	0.99
S/Y	0.60	0.58
τ_s	0.50	0.50

the data on relative plant sizes. Hence, the range of TFP values cannot be set arbitrarily.

More generally, however, while larger differences in the range of s will make it possible to obtain larger decreases in TFP associated with reallocation, the larger range of s will only arise if it requires a larger range of s values to reproduce a given relative plant size. This requires that the elasticity of employment with respect to s be smaller. But, if this is true, then it will require much larger subsidies to reallocate resources in such an economy. Hence, although there may be a greater potential for reallocation to produce larger decreases in TFP, it will also take much larger distortions to generate those losses. It is because of this that we feel it is useful in our current analysis to relate the TFP losses with the associated volume of subsidies. For example, if we change the extent of decreasing returns in our plant level production function then we can obtain much larger ranges of s but we also get less reallocation for a given tax and subsidy plan.

We have explored this issue via sensitivity analysis. In the interest of space here we simply report one such case. Following the choices of Veracierto (2001) we set $\alpha = 0.19$ and $\gamma = 0.64$, implying an overall returns to scale at the plant level of 0.83, a smaller value than our benchmark value of 90 percent. This reduces the elasticity of labor to changes in plant level TFP from 10 to approximately 6. Using the same calibration procedure we obtain an s range between 1 and 4.52 instead of 1 and 2.46 in our benchmark calibration. We find that in the case of correlated idiosyncratic distortions applied to output, the maximum effect on TFP is roughly 43 percent, which is 16 percentage points higher than the effect with the smaller decreasing returns to scale. However, to achieve this effect, it requires a tax rate of 60 percent and a subsidy rate of 78 percent. Previously we obtained drops in TFP of roughly 28 percent from a tax rate of 40 percent and a subsidy rate of 37 percent. So, although larger differences can be generated, it is more costly to generate them as well.

Nevertheless it remains of interest to examine how our findings would be affected by assuming production functions with different features such as fixed costs, capacity constraints or fixed proportions. We also note that larger TFP differences would also result if we assumed that c_f were greater than zero and there was some selection in terms of which entering plants choose to produce output. In this case, subsidies that are negatively correlated with plant productivity may reduce the productivity-entry threshold thereby bringing less productive technologies into the market. We have avoided this channel since by placing a lot of mass on plants with productivity below those being used in the distortion-free economy, it would seem to add an arbitrary element to the analysis. At the same time, it could be that policies in many countries do serve to allow plants to operate that would not operate at all in a market free of distortions, so this margin may be of practical importance. In fact, government policies such as trade protection and corporate bankruptcy laws are usually studied in the context of models with this margin (see for example Tybout 2000 and the references therein, and Bergoeing et al. 2002).[10]

6 Conclusion

We have analyzed distortions that lead to reallocation of resources across heterogeneous production units. Our results indicate that the impact of these distortions on aggregate output and TFP can be quite large. Based on one metric, we find that the effect of these distortions is roughly similar to the effects associated with distortions to aggregate relative prices. Given the pervasiveness of policies, regulations, and institutions that induce reallocations

[10]We have computed equilibrium reallocations for the case where there is selection in entry of plants (i.e., $c_f > 0$) and found that the amplification effect of this channel on the results is relatively small in our model. While selection on entering plants may result in unproductive plants operating in distorted economies, these plants still account for a small share of employment and output in the equilibrium of these economies. As a result, the overall impact on output and TFP of this effect is quantitatively small in our model.

of resources across productive units, it seems to us that this channel may be important in accounting for some of the cross-country patterns in output, capital accumulation and TFP.

References

[1] Alfaro, L., A. Charlton, and F. Kanczuk, 2007, "Firm-Size Distribution and Cross-Country Income Differences," manuscript, Harvard Business School.

[2] Atkeson, A., A. Khan, and L. Ohanian, 1996, "Are Data on Industry Evolution and Gross Job Turnover Relevant For Macroeconomics?" *Carnegie-Rochester Conference Series on Public Policy* 44: 216-250.

[3] Baily, M., C. Hulten, and D. Campbell, 1992, "Productivity Dynamics in Manufacturing Plants," *Brooking Papers on Economic Activity: Microeconomics* 187-267.

[4] Baneerjee, A. and K. Munshi, 2004, "How Efficiently is Capital Allocated? Evidence from the Knitted Garment Industry in Tipur," *Review of Economic Studies*, 71: 19-42.

[5] Banerjee, A. and E. Duflo, 2005, "Growth Theory Through the Lens of Development Economics," Handbook of Economic Growth, Vol 1A, Chapter 7.

[6] Bartelsman, E., J. Haltiwanger, and S. Scarpetta, 2006, "Cross-Country Differences in Productivity: The Role of Allocative Efficiency," manuscript, University of Maryland.

[7] Bergoeing, R., P. Kehoe, T. Kehoe, and R. Soto, 2002, "A Decade Lost and Found: Mexico and Chile in the 80's," *Review of Economic Dyanamics* 5(1): 166-205.

[8] Bergoeing, R., Loayza, and Repetto, 2004, "Slow Recoveries," *Journal of Development Economics* 75: 473-506.

[9] Chari, V., P. Kehoe, and E. McGrattan, 1996, "The Poverty of Nations: A Quantitative Exploration." Staff Report 204. Federal Reserve Bank of Minneapolis.

[10] Chu, T., 2002, "Exit Barriers and Productivity Stagnation," manuscript, University of Hawaii.

[11] Davis, S., J. Haltiwanger, and S. Schuh, 1996, *Job Creation and Destruction,* The MIT Press: Cambridge, Massachusetts.

[12] Djankov, S., R. La Porta, F. Lopez-de-Silanes, and A. Shleifer, 2002, "The Regulation of Entry," *Quarterly Journal of Economics*, 67(1): 1-37.

[13] Greenwood, J., J. Sanchez, and C. Wang, 2007, "Fianncing Development: The Role of Information Costs," manuscript, University of Pennsylvania.

[14] Gunner, N., G. Ventura, and Y. Xu, 2006, "Macroeconomic Implications of Size-Dependent Policies," manuscript, Penn State University.

[15] Harberger, A., 1998, "A Vision of the Growth Process," *American Economic Review,* 88(March): 1-32.

[16] Hall, R. and C. Jones, 1999, "Why Do Some Countries Produce so Much More Output per Worker Than Others? *Quarterly Journal of Economics*: 83-116.

[17] Herrendorf, B. and A. Texeira, 2003, "Monopoly Rights Reduce Income Big Time when there is Capital Accumulation," manuscript, Universidad Carlos III de Madrid.

[18] Hopenhayn, H., 1992, "Entry, Exit, and Firm Dynamics in Long Run Equilibrium," *Econometrica* 60: 1127-50.

[19] Hopenhayn, H. and R. Rogerson, 1993, "Job Turnover and Policy Evaluation: A General Equilibrium Analysis," *Journal of Political Economy,* 101(5): 915-38.

[20] Hsieh, C. and P. Klenow, 2006, "Missallocation and Manufacturing TFP in China and India," manuscript , Stanford University.

[21] Hsieh, C. and P. Klenow, 2006, "Relative Prices and Relative Prosperity," forthcoming, *American Economic Review.*

[22] Klenow, P. and A. Rodriguez-Clare, 1997, "The Neoclassical Revival in Growth Economics: Has it Gone Too Far?," NBER Macroeconomics Annual. University of Chicago Press.

[23] Lagos, R., 2006, "A Theory of TFP," forthcoming, *Review of Economic Studies.*

[24] Mankiw, G., D. Romer, and D. Weil. 1992. "A Contribution to the Empirics of Economic Growth." *Quarterly Journal of Economics* 107(2): 407-38.

[25] Mankiw, G., 1995, "The Growth of Nations," *Brookings Papers on Economic Activity: Macroeconomics,* 1: 275-326.

[26] Melitz, 2003, "The Impact of Intraindustry Trade Reallocations and Aggregate Industry Productivity," *Econometrica* 71(6): 1695-1725.

[27] Parente, S. and E. C. Prescott, 1999, "Monopoly Rights: A Barrier to Riches," *American Economic Review,* 89: 1216-33.

[28] Parente, S. and E. C. Prescott. 2000. *Barriers to Riches.* MIT Press. Cambridge, Massachusetts.

[29] Pavcnik, N., 2002, "Trade Liberalization, Exit, and Productivity Improvements: Evidence from Chilean Plants," *Review of Economic Studies*, 69: 245-76.

[30] Peek, J. and E. Rosengren, 2005, "Unnatural Selection: Perverse Incentives and the Misallocation of Credit in Japan," *American Economic Review,* 95(4).

[31] Prescott, E. C. 1998. "Needed: A Theory of Total Factor Productivity." *International Economic Review* 39: 525-52.

[32] Schmitz, J., 2001, "Government Production of Investment Goods and Aggregate Labor Productivity," *Journal of Monetary Economics* 47: 163-87.

[33] Tybout, J., 2000, "Manufacturing Firms in Developing Countries: How Well Do They Do, and Why?" *Journal of Economic Literature,* 38(1): 11-44.

[34] U.S. Department of Commerce (1997), Census of Manufactures.

[35] Veracierto, M., 2001, "Employment Flows, Capital Mobility, and Policy Analysis," *International Economic Review* 42(3): 571-95.

Figure 1: Distribution of Plants by Employment − U.S. Data

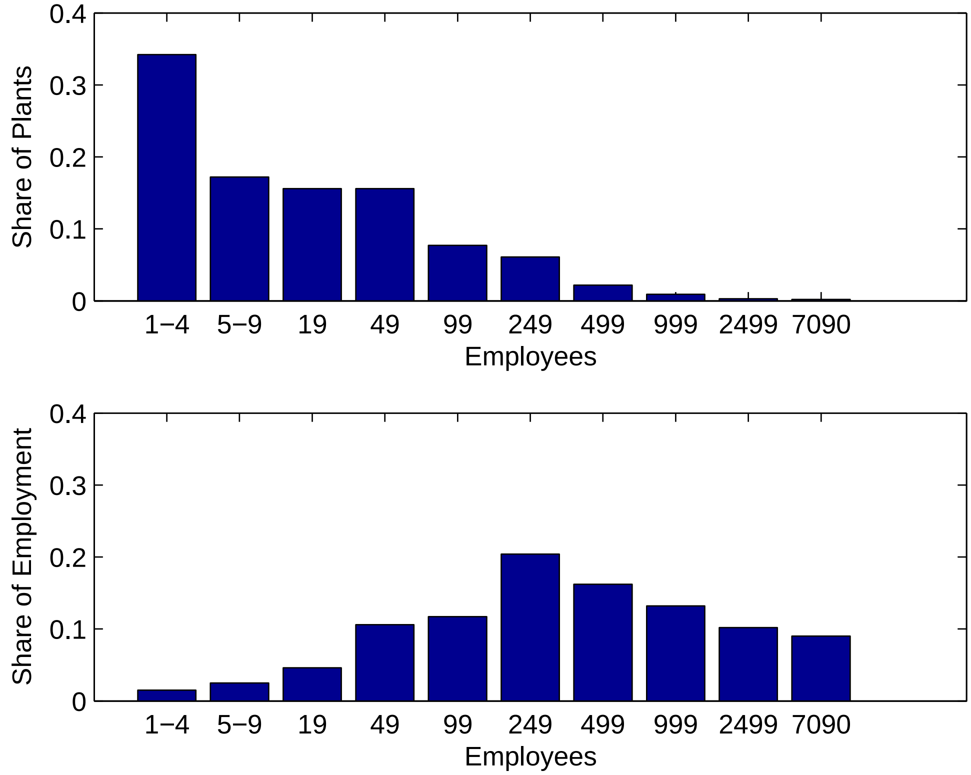

Figure 2: Distribution of Plants by Employment − Model vs. Data

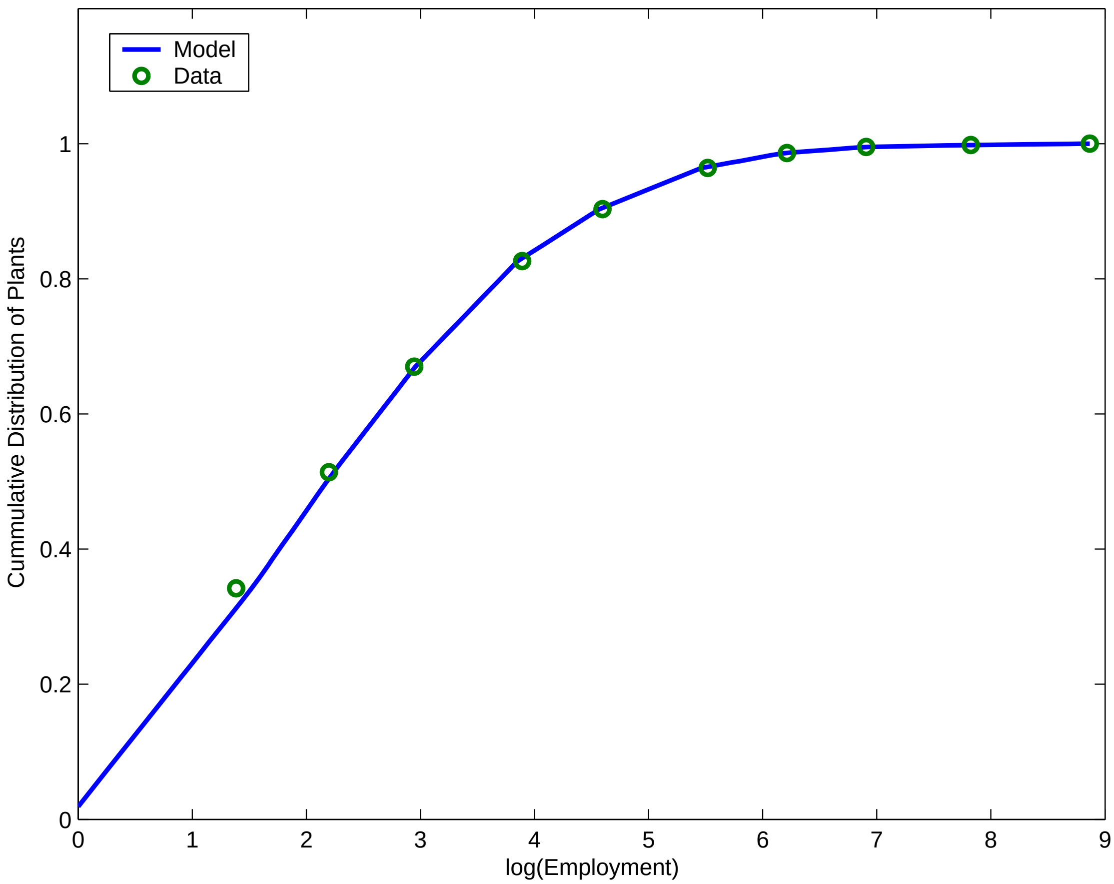

Figure 3: Share of Valued Added and Employment − Model vs. Data

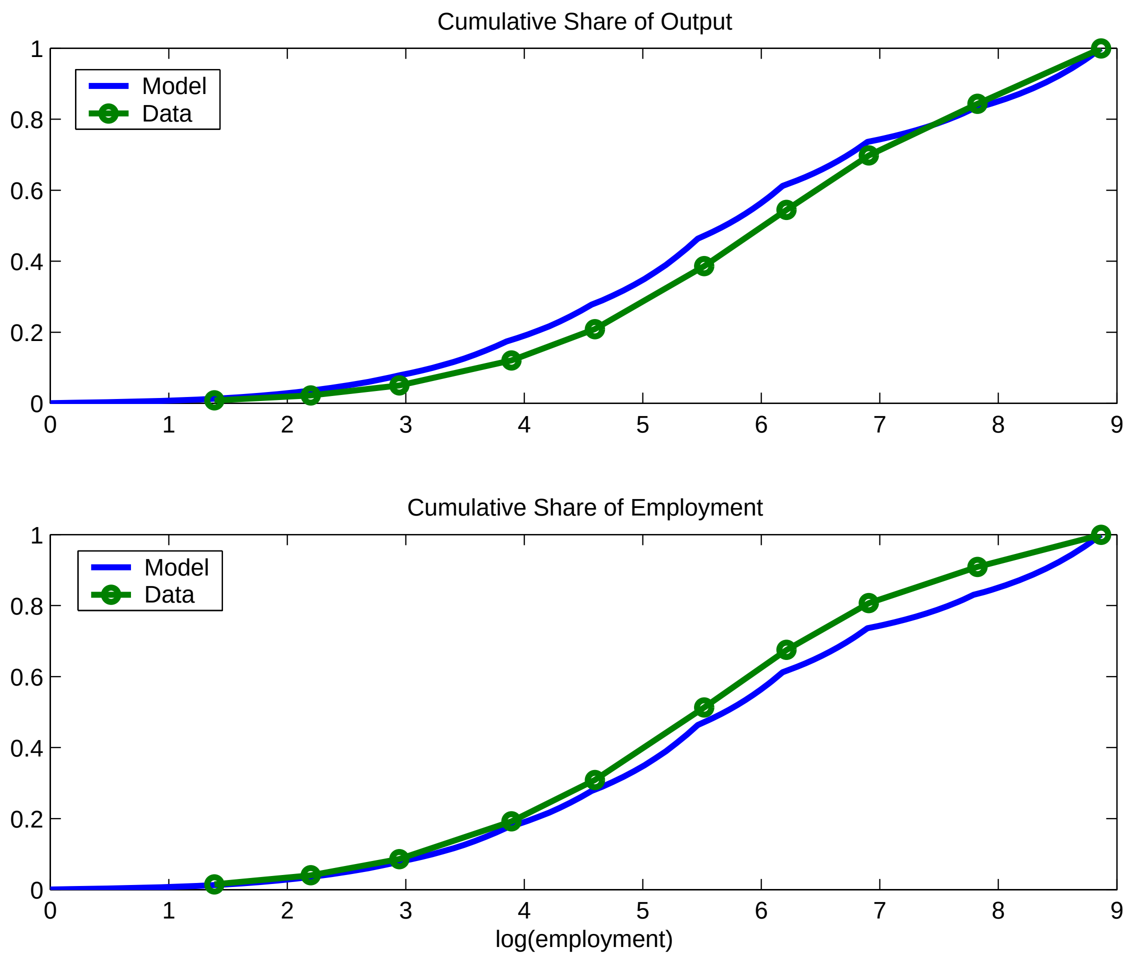

Figure 4: Share of Plants and Employment − Uncorrelated Experiments

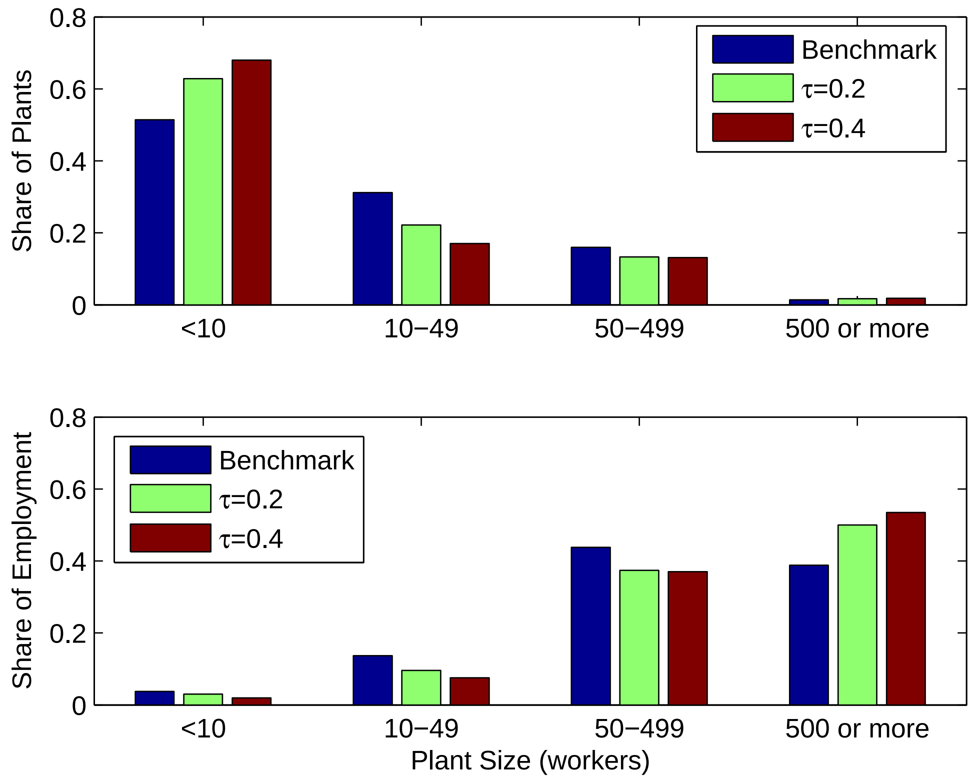

Figure 5: Share of Plants and Employment − Correlated 50%