Firm Size Distortions and the Productivity Distribution: Evidence from France[*]

Luis Garicano[†], Claire LeLarge[‡], John Van Reenen[§]

May 9, 2011

Abstract

Very Preliminary. A major empirical challenge in economics is to identify how regulations (such as labor protection) affect economic efficiency. Almost all countries have regulations that increase costs when firms cross a discrete size threshold. We show how these size-contingent regulations can be used to identify the equilibrium and welfare effects of regulation through combining a new model with the joint firm-level distribution of size and productivity. Our framework adapts the Lucas (1978) model to a world with size-contingent regulations and applies this to France where there are sharp increases in firing costs (which we model as a labor tax) when firms employ 50 or more workers. Using administrative data on the *universe* of firms 2002 through 2007, we show how this regulation has major effects on the distribution of firm size (a "broken power law") and productivity. We then econometrically recover the key parameters of the model in order to estimate the costs of regulation which appear to be substantial.

[†]London School of Economics and Centre for Economic Performance
[‡]CREST, Insee
[§]London School of Economics and Centre for Economic Performance

1 Introduction

A recent literature has documented empirically the large impact of misallocations of resources, due to distortions that raise the cost of labor or capital, on the distribution of productivity. As Restuccia and Rogerson (2008) argued[1] more efficient firms may have "too little" output or employment allocated to them due to various distortions in their economies. Hsieh and Klenow (2009) have shown that these misallocations go a long way towards explaining the gap in productivity between the US, China and India. In this paper, we focus on understanding the impact and the size of one specific distortion at one particular size on the French firm size distribution: a regulation increasing labor costs when firms reach 50 workers.

The idea that misallocations of resources may be partly behind the productivity gap is attractive in understanding the differences between the US and Europe. As Figure 1a shows, there appear to be far fewer French firms which are able to grow to the same scale as the productive US firms. Figure 1a shows two interesting patterns: a large bulge in the employment of firms around 50 workers in France, and a much larger share of very large firms in the US - the US has many more firms with over 2,500 employees[2]. This paper focuses on the first of those patterns, although we plan to examine the absence of very large French firms in later work.

[**Figures 1a and 1b about here**]

Labor legislation in France sharply increases firing costs when firms get to 50 employees. Specifically, firms with 50 or more employees formulate a "social plan," which is designed to facilitate reemployment, through training, etc. As a result, the costs of employing workers also rise (see Bertola and Bentolila, 1990) at that threshold. Figure 1b shows that indeed the legislation binds, so that there is a clear threshold effect at precisely 50 firms

What are the distortions in the size distribution, in the productivity distribution, and on aggregate productivity that result from those distortions? Our approach relies on revealed preference and on

---

[1] See also Bloom and Van Reenen (2010) and Petrin and Sivadasan (2010).A closely related literature is in development economics where some have pointed to the "missing middle", i.e. a preponderance of very small firms in poorer countries compared to richer countries (see Banerjee and Duflo, 2005). Many explanations have been put forward for this such as financial development, human capital, lack of competition in product markets, and social capital. One possibility, related to our approach, is size related labor regulations. Besley and Burgess (2004), for example, suggest that labor regulation is one of the reasons why the formal manufacturing sector is much smaller in some Indian states compared to others.

[2] Bartelsman et a (2009) examine misallocation using micro-data across many OECD countries and make a similar point. In particular, they find that the "Olley Pakes" (1996) covariance between size and productivity is much smaller in France (0.24 in their Table 1) and other European countries compared to the US (0.51 in their Table 1).

positive sorting effects. Some firms that would have been larger without the regulation choose to remain below the legal threshold to avoid these additional costs. We aim to identify these firms, calculate their counterfactual and use this observation to infer the cost of this legislation.[3]

There are different views on the underlying sources of heterogeneity in firm productivity. We follow Lucas (1978) in taking the stand that managerial talent is the primitive, and that the economy-wide observed resource distribution is as Manne (1965) felicitously put it, "a solution to the problem: allocate productive factors over managers of different ability so as to maximize output." Managers make discrete decisions or solve problems (Garicano, 2000). Making better decisions, or solving problems that others cannot solve, raises everyone's marginal product. This means that, in equilibrium, better managers must be allocated more resources. In fact, absent decreasing returns to managerial talent, the best manager must be allocated all resources. Given limits to managerial time or attention, the better managers are allocated more workers and more capital to manage. This results in a "scale-of-operations" effect whereby differences in talent are amplified by the resources allocated.[4] Lucas (1978) first explored these effects in an equilibrium setting.[5]

Consequently better managers, that is those that for whatever reason are able to generate more productivity, should be allocated (or equivalently, should choose) larger firm sizes. When managers are confronted with legislation that introduces a cost of acquiring a size that is beyond a certain threshold, they may choose to stay below the threshold and stay at an inefficiently small size. By studying the productivity of these marginal managers, we are able to estimate the cost of the legislation, the distortions in them, and thus the welfare cost of the legislation for the entire firm size distribution.

We start by setting up a simple model of the allocation of a single factor, labor to firms in a world where there are decreasing returns to managerial talent. We use it to study the effect of a step change in labor costs after a particular size and show that there are four main effects:

---

[3] In labor economics and macro, there has been an extensive discussion of the importance of Employment Protection Legislation (EPL) for unemployment and more recently productivity (e.g. Layard and Nickell, 1999; OECD, 2009). Our contribution is to combine a structural GE approach with detailed micro-economic data to quantify the costs of these regulations.

[4] In a model of this kind, the source of decreasing returns are on the production size, and are linked to limits to managerial time. For our purposes here, as Chiang and Klenow (2009) show, this source of decreasing returns is equivalent to having the decreasing returns come (as is more common in recent literature following Melitz (2003)) from the utility side.

[5] Such a scale of operations effects is at the heart of Rosen's (1982) theory of hierarchies, where efficiency units of labor controlled (and not just number of bodies) matter, and also in Garicano and Rossi-Hansberg (2006) where there is limited quantity-quality substitutability so that matching between workers and managers takes place. Empirically, this technology has been used to explain a wide-range of phenomena, most recently to calibrate the impact of scale of operations effects on CEO wages (Gabaix and Landier, 2008).

1. Equilibrium wages fall as a result of the reduction in the demand for workers (i.e. some of the tax incidence falls on workers)

2. Firm size increases for all firms below the threshold as a result of the general equilibrium effect on wages

3. Firm size reduces to precisely the regulatory threshold for a set of firms that are not productive enough to justify incurring the regulatory costs

4. Firm size reduces proportionally for all firms that are productive enough to incur the additional cost of regulation.

We use the model to guide our estimation of the impact of these costs. Specifically, we proceed in two ways. First, the theory tells us there is a deviation from the 'correct' firm size distribution as a result of the regulation. That is, we expect to see a departure from the usual power law firm size distribution[6] as firms bunch up below the threshold (50 workers). Given factors such as measurement error, the departure from the power law is not just at 49 but at a slightly smaller firm size. Then, at some point the distribution becomes again a power law, with a lower intercept. The key parameter of interest is this point where the power law 'recovers' again. Our econometric procedure follows the time series literature (notably Bai and Perron, 1998) to find these structural breaks in the power law.

Our second procedure is more direct, as it relies on the covariance of productivity and firm size. The idea for the procedure is to find the productivity of the marginal firm, that is to figure out which is indifferent between being at the regulatory threshold and the unrestricted size distribution. This can be done by calculating the maximum (or to avoid excessive noise, a high percentile) of the distribution of productivity at the size threshold and finding the equivalent productivity on the unrestricted size productivity relationship. This gives us a direct estimate of the productivity of the distorted firms and of the size they would have acquired without the threshold effect. Consequently, this procedure directly allows us to identify the cost of the regulation. We find a non-trivial cost of these regulations, such that firms at the threshold are reducing their desired size by up to 60%-70%[7].

---

[6]See Axtell (2001), Sutton (1997) and Gabaix, (2009). There is a large literature on the size and productivity distribution of firms in macro, trade, finance and IO. Appropriately, the first major study in this area was by Gibrat (1931) who studied French industrial firms, the main focus of the empirical part of our paper.

[7]Many empirical papers have shown that deregulation (e.g. Olley and Pakes, 1996), higher competition (e.g. Syverson, 2004) and trade liberalization (e.g. Pavcnik, 2002) have tended to improve reallocation by increasing the correlation between firm size and productivity.

The structure of the paper is as follows. Section 2 describes our theory and some extensions. Section 3 describes the institutional setting and data. Section 4 describes the empirical strategy of how we map the theory into the data. Section 5 contains the main results which come in two parts. First we show that the main empirical predictions of the model in terms of the size and productivity distribution are consistent with the data. Secondly, we estimate the parameters of the structural model and use this to show that the costs of the regulation are non-trivial. We present various extensions and robustness tests in Section 6 before drawing some conclusions in the final section.

2   Theory behind the estimation

We aim to estimate the distortions in the productivity distribution and the reallocation effect that result from an implicit tax on firm size that starts at a particular threshold. Our strategy relies on analyzing the choices of the firms that are trying to avoid the tax to estimate the cost of the tax. Having done that, we will be able to estimate the general equilibrium effects of the tax through the changes in firm size.

We study regulatory effects on the firm size distribution and on the productivity distribution in the simplest possible version of Lucas' model. There is only one input in production, labor, and a single sector. The primitive of the model is the distribution $\phi(c)$ of 'managerial ability' α (which we will measure as Total Factor Productivity, TFP), with cdf $\Phi : R^+ \rightarrow [0, 1]$. Ability is defined and measured by how much an agent can raise a team's output: a manager who has ability α and is allocated n workers produces $y = \alpha f(n)$. Larger teams produce more, $f' > 0$, but given e.g. limited managerial time, there are decreasing returns to the firm scale that a manager can manage, $f'' < 0$.

The key difference between our setting and the original Lucas model is that, in our application, there is a tax on firm size, which imposes a wedge between the wage the worker receives and the cost to the firm.[8] Since termination costs are generally denominated in years of salary, we assume this cost is a proportional increase in wage costs, taking the form of a labor tax. Moreover, this tax does not grow in a smooth way, but instead it begins hitting firms after they reach a given size N.

---

[8] In our application the 'tax' involves an extra marginal cost and also a fixed cost component. However, previous studies of this problem, such as particularly Abowd and Kramarz (2009) show that the fixed cost component are second order relative to the marginal cost component.

2.1 Individual Optimization

Let $\pi(\alpha)$ be the profits obtained by a manager with skill α when he manages a firm at the optimal size. These profits are then given by:

$$\pi(\alpha) = \max_n \alpha f(n) - w\bar{\tau}n \begin{cases} \bar{\tau} = 1 & \text{if } n < N \\ \bar{\tau} = \tau & \text{if } n \geq N \end{cases} \qquad (1)$$

where w is the worker's wage, n is the number of workers and τ is the tax, which only applies for firm over a minimum threshold of N (50) workers in our application).

Firm size at each side of the threshold is then determined by first order condition:

$$\alpha f'(n^*) - \bar{\tau}w = 0, \qquad \text{with } \tau = \bar{\tau} \text{ if } n \geq N \qquad (2)$$

so that $n^* = f'^{-1}(\frac{\tau w}{\alpha})$. Note that $n_\alpha > 0$, while n_τ and $n_w < 0$.

The size constraint is reached at size N and skill α_c given by:

$$\alpha_c = \frac{w}{f'(N)} \qquad (3)$$

Firms can legally avoid being hit by the regulation simply by choosing to remain small. The cost of this avoidance is increasing in the talent of the individual, and thus at a given ability level, given a choice between staying at $n = N$ and avoiding the tax, managers choose to pay the tax. The ability of the "marginal manager" that is unconstrained (α_u) is defined as:

$$\alpha_u f(N) - wN = \alpha_u f(n^*(\alpha_u)) - w\tau n^*(\alpha_u) \qquad (4)$$

where $n^*(\alpha_u)$ is the optimal firm size for an agent of skill α_u We call this threshold α_u, where u denotes the boundary of the unconstrained firms, and the firm size $n^*(\alpha_u)$, n_u.

2.2 Equilibrium

The most skilled agents choose to be manager-entrepreneurs, since they benefit from their higher ability in two ways. First, for a given firm size n, they earn more. Second, the most skilled individuals hire a larger team, $n(\alpha)$. We denote the ability threshold between managers and workers as $\alpha_{\min}$.

A competitive equilibrium is defined as follows:

Definition 1 *Given a distribution of managerial talent $\phi(\alpha)$, a per worker labor tax τ that binds all firms of size $n \geq N$, and a production function $\alpha f(n)$, a competitive equilibrium consists of:*

6

(*i*) *a wage level w paid to all workers*

(*ii*) *an allocation n(α) that assigns a firm of size n to a particular manager of skill α*

(*iii*) *a triple of cutoffs* $\{\alpha_{\min} \leq \alpha_c \leq \alpha_u\}$, *such that* $W = [0, \alpha_{\min}]$ *is the set of workers,* $M_1 = [\alpha_{\min}, \alpha_c]$ *is the set of unconstrained, untaxed managers,* $M_2 = [\alpha_c, \alpha_u]$ *is the set of constrained,* $n^* = N$, *but untaxed managers, and* $M_2 = [\alpha_u, \infty]$ *is the set of taxed managers such that:*

(1) No agent wishes to change occupation from manager to worker or to change from unconstrained to constrained.

(2) The choice of n(α) for each manager α is optimal given their skills, taxes τ and wages w;

(3) Supply of labor equals demand

Start with condition (1): an agent prefers to be a worker if $w > \alpha f(n) - wn$, or a manager if $w < \alpha f(n) - wn$, and thus we have:

$$\alpha_{\min} f(n) - wn = w \tag{5}$$

Equilibrium condition (2), from the first order conditon (2) implies that firm sizes are given by:

$$
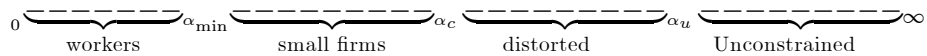\begin{aligned}
n(\alpha) &= 0 && \text{if } \alpha < \alpha_{\min} & (6)\\
n(\alpha) &= f'^{-1}(\frac{w}{\alpha}) && \text{if } \alpha_{\min} < \alpha < \alpha_c & (7)\\
n(\alpha) &= N && \text{if } \alpha_c < \alpha < \alpha_u & (8)\\
n(\alpha) &= f'^{-1}(\frac{\tau w}{\alpha}) && \text{if } \alpha_u < \alpha < \infty & (9)
\end{aligned}
$$

Thus we have four categories of workers as the following figure shows:

Equilibrium partition of individuals into workers and firm types by managerial ability

Finally, from condition (3), equilibrium requires that markets clear- that is the supply and demand of workers must be equalized. The supply of workers is $\Phi(\alpha_{\min})$, and the demand of workers by all available managers, $\int_{\alpha_{\min}}^{\infty} n(\alpha)d\Phi(\alpha)$, where $n(\alpha)$ is the continuous and piecewise differentiable function given as above, thus:

$$\Phi(\alpha_{\min}) = \int_{\alpha_{\min}}^{\infty} n(\alpha) d\Phi(\alpha). \tag{10}$$

Solving the model involves finding four parameters, the cutoff levels $\alpha_{\min}, \alpha_c, \alpha_u$, and the equilibrium wage w. For this we use the four equations (5), (10), (4) and (3)

The equilibrium is unique; the following proposition characterizes the comparative statics in the equilibrium:

Proposition 1 *The introduction of a variable cost of hiring workers starting at firm size N has the following effects:*

1. *Reduces equilibrium wages as a result of the reduction in the demand for workers*

2. *Increases firm size for all firms below the threshold, $[\alpha_1, \alpha_c]$, as a result of the general equilibrium effect on wages*

3. *Reduces firm size to the threshold N for all firms that are constrained, that is those in $[\alpha_c, \alpha_u]$*

4. *Reduces firm size for all firms that are taxed $[\alpha_c, \infty]$*

Example. Consider a power law with a slope similar, $\phi(\alpha) = \frac{0.6}{\alpha^{1.6}}$ and returns to scale parameter is $\theta = 0.9$. Figure 2 shows the firm size distribution for a firm size cut-off at 50 employees (needs label), and an employment tax of 1%. As in the distribution in the data, there is a spike at 50 employees that breaks the power law. Figure 3 reports the productivity distribution α as a function of firm size n. It shows that we should expect a spike in the productivity distribution at the point in which the regulation starts to bind. Essentially the maximum bar of this graph is the most productive firm that is affected by the regulation. We can trace the firm size simply by moving horizontally to the right in the graph, as in our empirical procedure.

2.3   Empirical Implications

The econometric work that follows aims to use the theory as a guide to estimate the welfare losses that result from this regulations.

As is well known, the firm size distribution generally follows a power law (see e.g. Axtell, 2001). Lucas (1978) shows that Gibrat's law implies that the managerial returns to scale function must be $f(n) = \alpha n^{\theta}$, and that for it to be consistent with a power law, the ability or productivity distribution must also be power, $\phi(\alpha) = k_{\alpha} \alpha^{-\beta_{\alpha}}$. As figure A2 shows, this is not a bad approximation; the

distribution of TFP is somewhere between log normal and power, but the fit for a power distribution is good for a large fraction of the data- the right hand side is linear in the $\log - \log$ plot.[9]

In this case, from the first order conditions, firm sizes are given, for a given wage, by:

$$n^* = \left(\frac{\alpha \theta}{w \overline{\tau}} \right)^{1/(1-\theta)}$$

A firm below the tax threshold N chooses a firm size $n^* = f'^{-1}(\frac{w}{\alpha})$; while a firm above the threshold chooses $n_\tau^* = f'^{-1}(\frac{\tau w}{\alpha})$.

If the distribution of $\phi(\alpha)$ follows a power law, the distribution of firm sizes $\chi(n)$ is also power (apart from the threshold), since by the change of variable formula, $\chi(n) = \phi(\alpha(n)) * n^{-\theta} w \overline{\tau} \frac{(1-\theta)}{\theta}$ (omitting the threshold). Thus the firm size distribution is given by

$$\chi(n) = \begin{cases} kn^{-\beta} & \text{if } n < n(\alpha_c) = N - 1 \\ m & \text{if } n = N - 1 \\ 0 & \text{if } N - 1 < n < n(\alpha_u) \\ k_\tau n^{-\beta} & \text{if } n > n(\alpha_u) \end{cases}$$

where $k = k_\alpha (1 - \theta) \left(\frac{w}{\theta} \right)^{\left(\frac{1-\beta}{1-\theta} \right)}$, $k_\tau = k_a (1 - \theta) \left(\frac{w\tau}{\theta} \right)^{\left(\frac{1-\beta}{1-\theta} \right)}$, the exponent $\beta = (\beta_\alpha (1 - \theta) + \theta)$, and $m = \int_{n_c}^{n_u} \phi(\alpha) d\alpha$.

In the empirical section, we estimate a power law on the undistorted segment roughly as follows (the exact specification is below:

$$\ln \chi(n) = \ln k - \beta \ln n + \delta (D_{n > n_u}) \tag{11}$$

where $D_{n > n_u}$ is is a dummy variable equal to unity for firms above the threshold n_u and zero otherwise. Thus, the coefficient on the dummy in the regression is:

$$\delta = \ln k_\tau - \ln k = \left(\frac{1 - \beta}{1 - \theta} \right) \ln \tau \tag{12}$$

which allows us to estimate the cost of the tax directly from the firm size distribution

Clearly, we can see two departures in Figure 1b from the predictions in the theory

1. The departure from the power law does not start at N, but earlier

---

[9]On the left hand side, part of the problem may be that we dropped firms having 0, 1 or 2 employees because, there, the coverage of FICUS is not great. This might explain why we have too few observations with low levels of TFP. At any rate, compared to a log-norml note that the tails are asymmetric, the right tail is way too thick, and the parameters of the fitted Gaussian (especially for scale) are almost degenerate.

2. The region immediately to the right of N does not have zero density, but rather there are some firms with positive employment levels just to the right of the regulatory cut-off, N.

Part of the problem is likely to be measurement error. As we will discuss later, the measurement of firm size that we have is not exactly the same one as the one used to determine whether a firm is subject to the regulation or not. From the perspective of the regulation, the relevant concept of employment is the number of workers at the precise date where the collective dismissal is announced. For example, it excludes apprentices and a few other categories. Our proxy for firm size is the arithmetic mean of the workforce at the end of the quarter of the fiscal year.[10]

Thus we assume we measure the firm size distribution with error, so that the observed firm size is $n = \widetilde{n} + \varepsilon$ w   here $\widetilde{n}$ is"true" employment and ε is random measurement error. Figure 2b shows the theoretical counterpart from Figure 2 when we add measurement error to the relationship. It shows that the figure has similar properties to the ones in the data. With measurement error, we have a set of firm sizes immediately above N where firms appear above the threshold but are mismeasured.[11] To a first approximation, we have in the firm size distribution probabilistic thresholds as follows:

Equilibrium distribution of individuals into workers and firm types by firm size
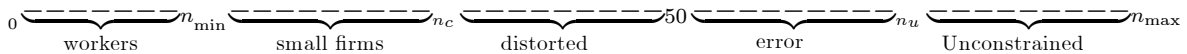
Our second empirical strategy consists in directly estimating the relationship

$$\alpha_u f(N) - wN = \alpha_u f(n^*(\alpha_u)) - w\tau n^*(\alpha_u) \tag{13}$$

By finding the marginal TFP, α_u on the constrained size distribution, we can obtain jointly the unconstrained firm size $(n^*(\alpha_u))$ and the implied tax τ.

---

[10]Fiscal definition, article 208-III-3 du Code General des Impots.

[11]In Section 6 we will discuss other potential sources of stochastic variation in the firm size distribution that would also justify some firms right above the threshold including selection issues such as heterogeneity in capital labor substitutability, dynamics and optimization errors. We show that those stories do not seem to account for other patterns in the data.

3   Empirical Setting and Data

3.1   Institutions: The French Labor market and Employment Costs

France is renowned for having a highly regulated labor market (see Abowd and Kramarz, 2003; Kramarz and Michaud, 2010). What is less well known is that most of these laws only bind on a firm when it reaches a particular employment size threshold. By far the most important size threshold is when a firm hits fifty employees - at this point of number of labour market regulations bind regarding the firm's ability to adjust its labor. Although there are some regulations that bind when a firm (or less often, a plant) reaches a lower threshold such as 10, 20 or 25 employees, 50 is generally agreed by labour lawyers and business people to be the critical threshold when costs rise significantly (see Appendix C)[12].

Perhaps the most important of these is a set of regulations introduced under a major piece of legislation in 1989. This required firms with 50 or more employees to formulate a "social plan" before laying off 10 or more workers (a "collective termination"). This social plan must place a limit on the total number of terminations and lay out plans to facilitate reemployment of terminated workers and will typically insist on an extensive retraining program. Union representatives or personnel delegates and the departmental director of the Ministry of Labor must also be informed of the plan. Two public meetings of the works council ("comité d'entreprise") must be organized with an interval between the meetings of 2–4 weeks depending upon the number of terminations proposed. The works council may require the firm to hire a consulting accountant (at the company's expense) to help the council with its analysis. During this period, the Ministry of Labor must be continuously informed of the proceedings, the plan, and the names of the proposed terminated workers. In addition to these firing costs in the 1989 law, there are some other pieces of regulation that bite at size 50 (see Appendix C).

How important are such provisions for firms? It is hard to know directly, as the opportunity cost of managerial time involved in preparing for such eventualities may be very great. Our framework is designed to recover the costs of such regulations. We treat such firing costs as an increase in the cost of labour. Firms face future shocks which will require them to adjust labor. Firms facing such a firing cost will effectively face a much higher cost in the eventuality that they face a negative shock. This affects the decision to hire and is (in expected value terms) very much like a labor tax.

---

[12]http://www.travail-emploi-sante.gouv.fr/informations-pratiques,89/fiches-pratiques,91/licenciement,121/le-plan-de-sauvegarde-de-l-emploi,1107.html

Since our analysis is fundamentally cross sectional we will model the firing cost as a labour tax.

There are other laws affecting French firms, so in one sense we are estimating a lower bound to the cost of regulation. But we are alert to the problem that some of the data is also affected by other laws which may also have a size-related threshold. Discussions with the labor ministry confirm that the threshold of 50 is the most important one in France, so it makes sense to begin our analysis here.

3.2 Data

Our main dataset is administrative data covering the universe of French firms between 2002 and 2007. These hold about 2.2m observations per year. These are the (mandatory) fiscal returns of all French firms ("FICUS") and are the appropriate level for analysis as it is on this administrative unit that the main laws pertains to. In addition to accurate information on employment (average number of workers in last quarter of the fiscal year), FICUS contains balance sheet information on labor, capital, investment, wage bills, materials, four digit industry affiliation, zipcode, etc. that are important in estimating productivity. More details of the dataset are given in the Appendix.

We take several approaches to estimating productivity. Our baseline results use the Levinsohn and Petrin (2003) method which extends the Olley and Pakes (1996) method of using a control function approach to deal with unobserved productivity shocks and selection when estimating production functions. Because we have a panel of firms we can implement this and estimate the production function coefficients separately by each four digit sector controlling for time dummies. The details of these regressions are reported in Appendix B. There are several issues with this approach (see Ackerberg et al, 2007) to estimating production functions so we also estimate TFP using a variety of other methods. For example, we consider a simple Solow-type approach where we calculate TFP as a residual using industry (and firm-specific) factor shares as weights. This led to qualitatively similar results to the ones presented here (see Appendix B for details).

In the current version of the paper we focus on firms whose main activity is in manufacturing (NACE2 class 15 to 35, 227 four digit industries), but we have similar results for other parts of the business sector (such as retail and business services - see Appendices).

4 Empirical Strategy

How costly is the employment protection legislation? We uncover this cost through revealed preference. Essentially, our approach is to identify the "constrained firms", those which legally avoid the

regulation by remaining too small, and identifying them. Once we have done this, we can calculate what they would have produced in the counterfactual world and thus we have an estimate of the cost of the regulation. In this section we explain how we apply our theoretical framework to the data we just reviewed.

4.1 Estimating the Broken Power Law from the firm size distribution

In equilibrium we have a matching correspondence $n(\alpha)$ that maps skill into firm size. In the regions where employment lies between $[n_c, 50]$ we find the constrained firms. These are firms that, given the choice between paying the labor cost τw^* and paying w^* but staying at size $n < 50$, prefer to stay below 50. The key econometric issue with these firms is that we do not have a simple invertible function that allows us to recover skills (or TFP) but instead have a mix of the firms that truly "belong" at these sizes, and the constrained ones, that is ability/productivity in both $\alpha < \alpha_c$ and in $[\alpha_c, \alpha_u]$. In $[50, n_u]$ we find a mass of firms that either are measured in error, as we assumed above, or belong there but chose not to cut their sizes. In the estimation procedure we assume that these are firms that stay at their first best size because they do not calculate the impact of Employment Protection Legislation on their decisions. Finally, we have unconstrained firms in $[n_u, n_{\max}]$ Once productivity exceeds a higher threshold α_u firms are sufficiently productive that they pay the tax in order to produce at a higher level.

Our objective is to estimate the thresholds n_u and n_c from the firm size distribution. Knowing these thresholds allows us, as we show later, to estimate the cost of the tax. Intuitively, knowing that a firm prefers having say 49 workers and not paying a cost, rather than having, say, 90 workers and incur the costs of EPL allows us to estimate the cost.

Let the observed distribution of shares be s_i, and the power law estimated from the unconstrained firms be $\widehat{k} n_i^{\widehat{\beta}}$. Then holding wages constant and the tax in place, the number of firms that are avoiding the regulation and choose to have a size smaller than N even though they would be expected to be above N is given by:

$$\sum_{i=n_c}^{N} s_i - \sum_{i=n_c}^{N} \widehat{k} n_i^{\widehat{\beta}}$$

That is, the bulge in the firm size distribution identifies the 'excess' of firms above what the power law would lead us to expect as. But we can also identify the firms that are constrained from the depression after the threshold:

$$\sum_{i=N}^{n_u} \widehat{kn_i^{\widehat{\beta}}} - \sum_{i=N}^{n_u} s$$

A constraint that relates the two thresholds is that the number of firms that are moved from one side to the other is the same either way, that is:

$$\sum_{i=n_c}^{N} s_i - \sum_{i=n_c}^{N} \widehat{kn_i^{\widehat{\beta}}} = \sum_{i=N}^{n_u} \widehat{kn_i^{\widehat{\beta}}} - \sum_{i=N}^{n_u} s_i \text{ or, equivalently} \tag{14}$$

$$\sum_{i=n_c}^{n_u} s_i = \sum_{i=n_c}^{n_u} \widehat{kn_i^{\widehat{\beta}}}$$

Thus our empirical problem is to estimate the power law distribution in (11) above (using the data on the undistorted part), subject to the constraint (14), that the number of firms in the estimated model is the same as under the power law.

Thus our estimating equation becomes:

$$\ln \chi(n) = \ln k - \beta \ln n + \delta(D_{n>n_u}) + \sum_{i=n_c}^{n_u} d_i \tag{15}$$

where all variables are defined as before (e.g. $D_{n>n_u}$ is a dummy variable that turns on to 1 for firms above the threshold n_u and is zero otherwise), but we have added d_i dummies that pick up the average number of firms in the distorted size categories, i.e. between the upper (n_u) and lower (n_c) employment thresholds. Equation (15) is estimated subject to the constraint (14).

Following Axtell (2001), we estimate equation (15) through OLS.[13], conditional on the 'structural breaks' at n_c and n_u. To find these structural break points, we follow Bai (1997) and Bai and Perron (1998) in their study of structural breaks in time series models. In our context, their result implies that for each partition $\{\{1, ...n_c\}, \{n_c...n_u\}, \{n_u, ...\}\}$, one obtains the OLS estimators of $\{k, \beta, \delta_1, \delta_2\}$ subject to constraint (14).[14] Letting the sum of squared errors generated by each of these partitions be $SSE(n_u, n_c)$, our estimates of the 'break points', n_u and n_c are:

$$(\widehat{n}_u, \widehat{n}_c) = \arg \min_{n_u, n_c} SSE(n_u, n_c) \tag{16}$$

---

[13] See Gabaix and Ibragimo (2008) for improvements in the OLS procedure using ranks, which is preferred for small samples and for the upper part of the distribution (not the middle, our focus).

[14] Perron and Qu (2005) show that the framework can accomodate linear restrictions on the parameter; and that the consistency and rate of convergence results hold and the limiting distribution is unaffected. However, our constraint is non-linear and no results exist on whether the results hold.

Bai and Perron (1998) show that, for a wide range of error specifications (including heteroskedastic like in our case) the break points are consistently estimated, and converge at rate $\acute{N}$, where $\acute{N}$ is the maximum firm size as long as $n_u - n_c > \varepsilon \acute{N}$, and $n_c < n_u$, (the break points are asymptotically distinct) which is true in our framework since we know $n_c < N < n_u$.

Armed with these parameter estimates we can the proceed to estimate τ using the results above in equation (12).

4.2 Using the relation between TFP and size to estimate the implicit tax from regulation

An alternative is to try to find the best constrained firms. Who are the best constrained manager? We calculate this by examining firm TFP in the 'bulge' of the distribution immediately before the employment cut off of 50. We take high percentiles of the empirical distribution to calculate the marginal manager who is "just constrained".

We know, from the theory, that:

$$
n(\alpha) = \left\{ \begin{array}{lll}
\left(\frac{\alpha\theta}{w}\right)^{1/(1-\theta)} & \text{if} & \alpha < \alpha_c \\
n(\alpha_c) = N - 1 & \text{if} & \alpha \in [\alpha_c; \alpha_u] \\
\left(\frac{\alpha\theta}{w.\tau}\right)^{1/(1-\theta)} & \text{if} & \alpha > \alpha_u
\end{array} \right.
$$

Empirically however, we rather focus on the "inverse" of this relation, because α is measured (estimated) with greater error than n:

$$
\alpha(n) \left\{ \begin{array}{lll}
= & \frac{w}{\theta}.n^{1-\theta} & \text{if} \quad n < n(\alpha_c) = N - 1 \\
\in & [\alpha_c; \alpha_u] & \text{if} \quad n = N - 1 \\
= & \frac{w.\tau}{\theta}.n^{1-\theta} & \text{if} \quad n \geq N
\end{array} \right.
$$

There are several things that limit this procedure. First, the threshold $n(\alpha_c) < 50$ from which we start observing "constrained firms" is difficult to identify with this approach. This is due to the fact that there is additional firm level heterogeneity (or estimation noise for TFP) which does not allow to identify $n(\alpha_c)$ accurately. We set $n(\alpha_c) = 47$, using our information from the FSD.

The threshold $n_\tau(\alpha_u) \geq 50$ is identified as the level of employment having a TFP similar to the highest quantile of TFP in the $[n(\alpha_c); 49]$ bin. More precisely, we use quantile regression to predict quantiles of TFP as a function of size on "undistorted" part of the firm size distribution, predict what should be the true quantiles on distorted parts, and identify $n_\tau(\alpha_u)$ from where the TFP distribution 'recovers.'

We identify neither the mass of "distorted firms" with this strategy, nor the proportion ρ of distorted firms (which is set at conventional values in the previous quantile regression approach); but on the other hand, conditional on ρ, we are able to compute the average TFP of constrained, and un-constrained firms in the $[n(\alpha_c); 49]$ bin.

The advantage of using the productivity distribution is that we can identify directly the τ parameter – either as (simply!) the difference in the intercepts in the two log-linear segments of the TFP/size relationship, or incorporating information about output and wages for marginal firms and using relation 4:

$$\alpha_u f(N - 1) - w.(N - 1) = \alpha_u f(n^*(\alpha_u)) - w\tau n^*(\alpha_u)$$

$$\iff \tau = \frac{w.(N - 1) + \alpha_u \left(f(n^*) - f(N - 1) \right)}{w.n^*}$$

To obtain the "just unconstrained" manager (and hence the employment threshold) we have to model the size-TFP relationship which we assume can be approximated by a nonparametric estimate of the relationship between TFP and size (we will use series and kernel estimates) with a break point around the cut-off. This allows us to estimate at what point such a marginal manager would rejoin the TFP distribution and hence the value of the employment cut-off, n_u.

In France, there are roughly 200,000 manufacturing firms per year. For TFP estimation, we use 1.2 million observations (2002 to 2007, with retrospective info for 2001).

4.3 Combining the two approaches of estimating the parameters

Since we have two methods of estimating the parameters they are implicitly "over-identified" and we could combine the estimates in various ways. Conditional on the model being correctly specified we could use a ML procedure to make more efficient inferences of the parameters. Alternatively we could take a method of moments approach and test some of the over-identifying assumptions.

We plan to do this in future versions, but for now we take the simpler approach of simply comparing what the implied parameter valued look like from the two alternative approached. |

5 Results

5.1 Qualitative analysis of the data

Before moving to the econometrics we first examine some qualitative features of the data to see whether they are consistent with our model. Many commentators have expressed skepticism about

the quantitative importance of employment regulations as it is sometimes hard to observe any clear change in the size distribution around important legal thresholds[15], so we first focus on this issue. Figure 5 presents the empirical distribution of firm size around the cut-off of 50 employees for two datasets. The dataset we use (FICUS) is the population dataset of the universe of French firms that forms the basis of our econometric work and is reproduced from Figure 1 in the top left corner (Panel 5.1). There is a sharp discontinuity in size precisely at 50 employees which is strong non-parametric evidence for the importance of the regulation. There are just over 300 firms with exactly 49 employees and then only about 130 with 50 employees. Importantly, the distribution which declines from 31 employees flattens after about 44 employees, just before the stacking up at 49 employees then dropping off a sharp cliff when size hits 50. Note that there are still some firms after 49 which would not be predicted by the pure version of the theory - we discuss this in detail later.

The other panels of Figure 5 compare FICUS with another dataset, DADs, that has been more typically used by labor economists. DADs is a worker-level dataset containing information on occupation (see Figure 13) and demographics. In Panel 5.2 we aggregate employment up to the appropriate level for each FICUS firm. We use employment dated on 31st December to more closely relevant thresholds for most regulations (see Appendix C). The discrete jump at 50 shows up here but not as clearly as the FICUS data. This suggests some greater measurement error or sampling issues in DADs. Panels 5.3 and 5.4 use an alternative definition of employment used by Insee, the French statistical agency, where the jump in firm size over the threshold is smoother still. Figures 5.3 uses Full-Time Equivalents which shows less of a jump than the straight count of employees in the previous panels (the main labor laws relate to the number of workers rather than Full-Time Equivalents, so this is expected). Panel 5.4 uses the average number of employees over the year, which completely fails to show any discontinuity at 50. We conclude from Figure 5 that measuring employment in a way that is more relevant for the laws, using population rather than sampled data and having very accurate employment measures is important in identifying a clear effect of labour market regulation.

Figure 6 shows the size distribution over a much wider range than Figures 1 and 5, using the power law approach of plotting firm size on a log scale on the x-axis and the share of firms in each employment class (also in logarithms) on the y-axis. The bulge in the proportion of firms with

---

[15]For example, Schivardi and Torrini (2008) on Italian data and Abidoye et al (2010) on Sri Lanka data. The authors find that there is slower growth just under the threshold consistent with the regulation slowing growth, but they find relatively little effect on the cross-sectional distribution.

around 45-50 employees is still discernible, although slightly harder to see given the wider scale of the graph. It is clear that the number of firms drops sharply at 50. Firm size seems to approximate a power law in the employment size distribution prior to the bulge around 50. After 50, there is the sharp fall in the number of firms and the line more flat than expected before resuming what looks like another power law. Broadly, outside a "distorted" region around 50 employees, one could describe this pattern a "broken power law" with the break at 50^{16}. The finding of the power-law for firm size in France is similar to that for many other countries and has been noted by other authors (e.g. Giovanni et al, 2010; Giovanni and Levchenko, 2010), but the finding of the break in the law precisely around the main labor market regulation is, we think, novel.

In Appendix A1 we look at a similar graph, but this time over an even larger size distribution covering all firms with between 1 and 1000 employees. As is well known the power law fits rather less well for the very small firms. Additionally, there does appear to be some break in the power law at firm size 10 and possibly as smaller one at firm size 20. This corresponds to the size thresholds from other pieces of labor and accounting regulations (see Appendix C). In order to avoid conflating these issues we focus our analysis on firms with 20 or more employees in Figure 6 and the rest of the paper. We can generalize the methods used here to other breaks in the Power Law which we will exploit in future versions[17].

The distribution of TFP is presented in Figure 7 broken down into five broad size classes. It is clear that larger firms have higher productivity as the mean and median of the distribution shifts to the right for larger firms. We would expect this from our basic model which, following Lucas, has the implication that more talented managers leverage their ability over a greater number of workers. Within each size band productivity looks approximately log normal (see Appendix for more descriptive statistics on overall productivity distribution).

Figure 8 continues the analysis of the covariance between productivity and firm size. In panel A we consider the mean level of TFP at each level of firm employment size. Panel A does this for firms of up to 100 employees and panel B for firms of up to 500 employees. In both panels productivity appears to rise monotonically with size, although there is more heteroskedacity for the larger firms as we would expect because there are fewer firms in each bin. Although we fit a fourth order polynomial in ln(size) to the TFP distribution, it is broadly log-linear. What is particularly

---

[16]See Howell (2002) for examples of how to estimate these types of distributions.More generally see Bauke (2007) for ways of consistently estimating power laws.

[17]Looking carefully at the line to the right of the 50 threshold in Figure 6 some more minor "bumps" are discernable (around 25 for example) which is suggestive of some (minor) effects of the other less important size regulations.

interesting for our purposes, however, is the "bulge" in productivity just before the 50 employee threshold. We mark these points in red. It looks very consistent with our model where some of the more productive firms who would have been just over 50 employees in the counterfactual world, choose to be below 50 employees to avoid the cost of the regulation. Firms just below the cut-off are a mixture of firms who would have had a similar employment level without the implicit tax and those firms whose size is distorted by the size-related regulation.

5.2 Econometric Implementation of the model

We are seeking to estimate the employment size cut-offs $n_c(\alpha_c)$, $n_u(\alpha_u)$ and the relevant parameters of the power law. As discussed above we use two strategies: the information from the firm size distribution and the joint distributions of TFP and firm size. These, together with our structural modelling assumptions enable us to identify the cost of the regulation, τ.

5.2.1 Estimation from the Firm Size Distribution

We turn first to the basic method of inferring what we need solely from the firm size distribution. The basic method can be seen by referring to Figure 9 which plots out the share of firms and firm employment size (similarly to Figure 6). Following equation (16) we fit the power law parameters and the employment cut-offs. We empirically estimate these upper and lower cut-offs as $n_c = 43$ and $n_u = 68$).

One intuitive way of seeing the procedure is as follows. Fix the lower employment threshold (say 43) and estimate the power law (conservatively) only on the part of the employment distribution below this and on the upper part of the size distribution that is undistorted (say under 42 and over 100)[18]. This procedure generates a mass of firms (entrepreneurs) displaced to the "bulge" in the distribution between n_c and N (i.e. 43 and 50) as shown in Figure 9. These firms are drawn from between N and n_u, and since we know the counterfactual slope of the power paw over this region, we can reallocate these firms so as to minimize the deviation from this counterfactual power law. n_u is estimated as the maximum employment bin which is attained in this procedure.

Rather than fixing n_c, the Bai and Perron (1998) procedure estimates this efficiently by minimizing a sum of squares criterion along with the other parameters in the model as in equation (16).

---

[18]We could in principle use all firms as small as one employee and up the largest firm in the economy. In practice the Power Law tends to be violated at these extremes of the distribution in all countries (e.g. Axtell (2001), so we follow that standard approach of trimming the upper and lower tails. We show that nothing is sensitive to these exact maximum and minimum employment thresholds as can be seen from the various figures.

This procedure gives us all the parameters necessary to estimate the implicit cost of the regulation which we calculate is equivalent to a labor tax of around 26% ($\tau = 1.26$).

5.2.2 Estimation from the relation between firm size and TFP

Our second method utilizes the estimates of productivity combined with firm size and is illustrated in Figure 10. We use this to calculate the indifference condition of the marginal manager as in Figure 3. The intuition behind Figure 10 is as follows. We estimate the relationship between TFP and firm size non-parametrically and fit a curve. Again we use data points that we have a strong prior that are unaffected by the regulation (e.g. as in the previous sub-section below 42 and above 100). As before, we initially fix the lower cut-off and observe TFP among the "distorted firms". We then find the highest TFP firm and find where this cuts the predicted TFP distribution. In Figure 10, the highest constrained TFP firm has a level of about 3.1 (see dashed horizontal line). This cuts the predicted TFP distribution at a point of a firm equal to 81 employees which is therefore our estimate of the upper threshold (i.e. $n_u = 81$). Using these estimated parameters, this implies a value of the implicit tax of 14% ($\tau = 1.14$). Compared to the previous method, the upper threshold is somewhat higher (81 vs. 68) and the implicit tax somewhat lower (14% vs. 26%), but both are in the same ballpark.

An empirical issue with this method is defining what is the "highest TFP". Although this is clear in theory, there are estimation errors in calculating firm-level TFP and the max will be heavily influenced by outliers. Consequently in Figure 10 we used the 95th percentile of the TFP distribution within each size bin. This is somewhat arbitrary of course, so Table 1 shows how our estimates of the upper employment threshold, n_u and implicit tax,τ , are affected by choices of what percentile to use (between the 99th and 50th). As can be seen from Table 1, there is a reasonable amount of stability whichever precise percentile is used in terms of the implied upper employment threshold: this ranges from 80 to 87. There is somewhat more sensitivity over the size of the implicit tax, however. In particular. using the 99th percentile produces very extreme estimates ($\tau = 3.7$) which suggests it is affected by outliers. The other estimates suggest that the baseline 95th percentile we use may be underestimating the cost of the regulation as the estimates range from 1.41 to 1.69.

6   Extensions and Robustness

In this section we consider several extensions to our framework and robustness tests of the results.

6.1 Estimates of GDP and welfare loss

[to follow]

6.2 Selection and other margins of adjustment to the regulation

The simplest version of the model focuses on the decision of the firm over how large to be based on employment. However, there are many other possible margins of adjustments that firms could potentially take to avoid the regulation. This would certainly add additional complexity to the mode. Broadly, the ability of these other margins of adjustment to alter our results depends on how easily workers can be substituted either for each other (across quality) or for other factors of production.

We begin with examining industries. There does not appear to be much of a break in the type of industries that firms are active in before and after the breaks in employment. Figure 11 illustrates this at the one and two digit levels, but we have also examined this at the three and four digit levels as well. It is also worth remembering that the TFP analysis is always within four digit sectors to avoid the problem of comparing apples and oranges.

A second way that firms could mitigate the cost of the regulation would be to substitute away from labor and into fixed capital. Figure 12 examines investment flows and capital intensity by firm size. There is an increase in capital intensity as firms get larger, but this is rather concave, falling off after firms reach around 100 employees. The relationship is noisy, but there is no clear evidence of any discontinuity around the thresholds. This may suggest that capital-labor substitution is not so easy over the margins we are looking at.

A third way of adjusting would be to use temporary and outsourced workers who are not covered by the regulation. Again, there was little evidence of this.

A fourth method would be by substituting across workers of different occupational types - here we do have some suggestive evidence suggesting firms are using this way of adjusting. Figure 13 shows that change in the share of the main three skill groups in French firms across firm size (managers - the most skilled group, manual workers - the least skilled group and clerical workers - the main middle group). Panel A shows the share of managers (excluding the CEO). This share seems to rise with firm size, but there is a clear change in the pattern around the threshold with firms choosing to increase their proportion of managers just after the regulatory threshold. Panel B shows almost a mirror image for manual workers - firms seem to reduce their reliance on less skilled workers

21

around the threshold. The middle group of workers in Panel C is relatively unaffected (the smaller residual groups look broadly like Panel C). this indicates a pattern whereby instead of expanding the quantity of workers a as it nears the threshold, firms will increase the quality of employees by substituting away from low skilled manuals to more skilled managers. This enables them to increase output without necessarily increasing employment and paying the extra regulatory cost.

6.3   Growth Analysis near the threshold

[Showing that growth is slower for firms just under the 50 threshold]

7   Conclusions

How costly is labor market regulation? This is a long-debated subject in policy circles and economics. We have tried to shed light on this issue by introducing a structural methodology that combines a simple theoretical general equilibrium approach based on the well known Lucas (1978) model of the size and productivity distribution of firms. We introduce size-specific regulations into this model, exploiting the fact that in most countries EPL only bites when firms cross specific size thresholds. We show how such a model generates predictions about the changes in the size and productivity distribution and moreover, can be used to generate an estimate of the implicit tax of the regulation. Intuitively, firms will optimally choose to remain small to avoid the regulation, so the size distribution becomes distorted with "too many" firms just below the size threshold and "too few" firms just above it. Furthermore, the distribution of productivity is also distorted: some of those firms just below the cut-off are "too productive" as they have been prevented from growing to their optimal size by the regulation. We show how the regulation creates welfare losses by (i) allocating too little employment to more productive firms who choose to be just below the regulatory threshold, (ii) allocating too little employment to more productive firms because they bear the implicit labor tax (whereas small firms do not) and (iii) through reducing equilibrium wages (due to some tax incidence falling on workers) this encourages too many individuals to become small entrepreneurs rather than working as employees for more productive entrepreneurs.

We implement this model on the universe of firms in the French private economy. France has onerous labor laws which bite when a firm has 50 employees, so is ideally suited to our framework. We find that the qualitative predictions of the model fit very well: (i) there is a sharp fall off in the firm size distribution precisely at 50 employees resembling a "broken power law" and (ii) there is a

bulge in productivity just to the left of the size threshold. Having good employment measures over the population of firms helps a lot.

We then estimate the key parameters of the theoretical model using two approaches based on the size distribution and the joint distribution of TFP and size. Both methods suggest substantial costs of the employment regulation which seems to place an additional cost on labor in the range of 20-30% of the wage.

This is just the preliminary sketch of our research program. We need to do a lot more testing of the results and extensions to the greater institutional complexity of the labor market. We believe that our approach is a simple, powerful and potentially fruitful way to tackle the vexed problem of the impact of regulation on modern economies.

REFERENCES

Abidoye, Babatunde, Orazem, Peter and Milan Vodopivec (2009) "Firing Cost and Firm Size: A study of Sri Lanka's Severance Pay system" World Bank Social Protection Discussion Paper No. 2009.

Abowd, John M. and Kramarz, Francis (2003) "The costs of hiring and separations" *Labour Economics*, 10(5), 499-530.

Ackerberg, Daniel, Larnier Benkard, Stephen Berry and Ariel Pakes (2007) "Econometric Tools for Analyzing Market Outcomes" *Handbook of Econometrics 6A*, 63, 4172-4276, Amsterdam: Elsevier.

Axtell, Robert (2001) "Zipf Distribution of US firm sizes" *Science*, 293 (5536), 1818-20.

Bai, Jushan (1997), "Estimating Multiple Breaks Once at a Time," Econometric Theory, 13: 315-352.

Bai, Jushan and Pierre Perron (1998), "Estimating and Testing Linear Models with Multiple Structural Changes," Econometrica, 66(1): 47-78.

Banerjee, Abhijit and Esther Duflo (2005) "Growth Theory through the lens of development" Chapter 7 474-544 in Philippe Aghion and Steven Durlauf (eds) *Handbook of Economic Growth*, Amsterdam: Elsevier

Bartelsman, Erik, John Haltiwanger and Stefano Scarpetta (2009) "Cross Country Differences in Productivity: The Role of Allocation and Selection" NBER Working Paper 15490

Bauke, Keiko "Parameter Estimation for power law distributions by maximum likelihood methods", mimeo, Rudolf Peirtls Centre for Theoretical Physics, Oxford

Bentolila, Samuel and Bertola, Giuseppe (1990), "Firing Costs and Labour Demand: How Bad Is Eurosclerosis?," Review of Economic Studies, 57(3): 381-402.

Besley, Timothy and Robin Burgess (2004), "Can Labor Regulation Hinder Economic Performance? Evidence from India," Quarterly Journal of Economics, 119(1).

Ceci-Renaud, Nila and Paul-Antoine Chevalier (2011), "L'impact des seuils de 10, 20 et 50 salariés sur la taille des entreprises françaises", Economie et Statistique, vol. 437,
     pp. 29- 45.

Gabaix, Xavier and Rustam Ibragimov, "Rank-1/2: A Simple Way to Improve the OLS Estimation of Tail Exponents," Journal of Business Economics and Statistics, customer. 29(1), 2011, p. 24–39.

Gabaix, Xavier and Augustin Landier (2008) "Why has CEO pay increased so much?" *Quarterly Journal of Economics* 123, 49-100.

Gabaix, Xavier (2009) "Power laws in economics and finance" *Annual Review of Economics* 1, 255-293.

Garicano, Luis (2000), "Hierarchies and the Organization of Knowledge in Production," Journal of Political Economy, 108: 5, pp. 874-904.

Garicano, Luis and Rossi-Hansberg, E. (2006), "Organization and Inequality in a Knowledge Economy," Quarterly Journal of Economics, CXXI (2006), 1383-1435.

Gibrat, Robert (1931) *Les Inégalités économiques*, Paris: Librarie du Recueil Sirey.

Giovanni, Julian, Andrei Levchenko and Romain Ranciere (2010) "Power Laws in Firm Size and Openness to Trade: Measurement and Implication" IMF Working Paper 10/109.

Giovanni, Julian, and Andrei Levchenko (2010) "Firm entry, Trade and Welfare in a Zipf's World", University of Michigan, mimeo.

Howell, Leonard W. (2002) "Maximum likelihood estimation of the broken power law spectral parameters with detector design applications" *Nuclear Instruments and Methods in Physics Research A*, 489, 422-438.

Kramarz, Francis and Michaud, Marie-Laure (2010) "The shape of hiring and separation costs in France" *Labour Economics*, 17(1), 27-37.

Levinsohn, James and Amil Petrin (2003) "Estimating Production Functions using inputs to control for unobservable" *Review of Economic Studies*, 70(2), 317-341.

Lucas, Robert E. (1978), "On the Size distribution of business firms" *Bell Journal of Economics*, 9, 2, 508-523.

Manne, H. G. (1965), "Mergers and the Market for Corporate Control," Journal of Political Economy, customer. 73, No. 2: 110-120.

Olley, Stephen and Pakes, Ariel (1996) "The Dynamics of Productivity in the Telecommunications Industry" *Econometrica*, 6, 1263-1297.

Pavcnik, Nina (2002) "Trade Liberalization, Exit, and Productivity Improvements: Evidence from Chilean Establishments", The Review of Economic Studies, 69(1): 245-276.

Perron and Qu (2005)

Petrin, Amil and Jagadeesh Sivadasan (2010) "Estimating Lost Output from Allocative Inefficiency with Application to Chile and Firing Costs", University of Minnesota, mimeo.

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

Rosen, Sherwin (1982), "Authority, Control and the Distribution of Earnings," Bell Journal of Economics, XIII, 311-323.

Syverson, Chad (2004a) "Market Structure and Productivity: A concrete example" *Journal of Political Economy*, 112, 6, 1181-1222.

Schivardi, Fabiano and Roberto Torrini (2008) "Identifying the effects of firing restrictions through size-contingent differences in regulation," *Labour Economics*, 15(3), 482-511.

Sutton, John (1997) "Gibrats legacy" *Journal of Economic Literature*, XXXV (1) 40-59.

A   Appendix A: Omitted Proofs

[To add]

B   Appendix B: Details on TFP Estimation and Results

There is no one settled way of best estimating TFP and there are many approaches suggested in the literature. Fortunately, at least at the micro-level, different methods tend to produce results where the correlation of TFP estimated by different methods is usually high (see Syverson, 2010).

In the baseline result we follow the method of Levinsohn and Petrin (2003) who propose extending the Olley and Pakes (1996) control function method to allow for endogeneity and selection. Olley and Pakes proposed inverting the investment rule to control for the unobserved productivity shock (observed to firm but unobserved to econometrician) that affects the firm's decision over hiring (and whether to stay in business). Because of the problem of zero investment regimes (common especially

among smaller firms that we use in our dataset) Levinsohn and Petrin (2003) recommended using materials as an alternative proxy variable that (almost) always takes an observed positive value.

We use this estimator to estimate firm-level production functions on French panel data 2002-2007 (using the unbalanced panel) by each of the four-digit manufacturing industries in our dataset. We also did the same for the retail sector and the business services sector. The production functions take the form (in each industry):

$$\ln y_{it} = \beta_n \ln n_{it} + \beta_k \ln k_{it} + \beta_m \ln m_{it} + \omega_{it} + \tau_t + \eta_{it} \tag{17}$$

where $y =$ output, $n =$ labour, $k =$ capital, $m =$ materials, ω is the unobserved productivity shock, τ_t is a set of time dummies and η is the idiosyncratic error of firm i in year t. From estimating the parameters of the production function we can then recover our estimate of the persistent component of TFP. Note that TFP is always normalized within industry and year.

There are of course many problems with these estimation techniques. For example, Ackerberg et al (2006) focus on the problem of exact multicollinearity of the variable factors conditional on the quasi-fixed factors given the assumption that input prices are assumed to be common across firms. Ackerberg et al (2007) suggest various solutions to this issue.

We consider alternative ways to estimate TFP including the more standard Solow approach. Here we assume that we can estimate the factor coefficients in equation (17) by using the observed factor shares in revenues. We do this assuming constant returns to scale, so $\beta_n = \frac{wn}{py}; \beta_m = \frac{cm}{py}$ and $\beta_k = 1 - \frac{wn}{py} - \frac{cm}{py}$ where c = the price of materials. We used the four digit industry factor shares averaged over our sample period for the baseline but also experimented with some firm-specific (time invariant) factor shares. As usual these alternative measures led to similar results.

A problem with both of these methods is that we do not observe firm-specific prices so the estimates of TFP as we only control for four digit industry prices. Consequently, the results we obtain could be regarded as only revenue-based TFPR instead of quantity-based TFPQ (see Hsieh and Klenow, 2009). TFPQ is closer to what we want to theoretically obtain as our estimate of α.In practice, there is a high correlation between these two measures as shown by Foster et al (2008) who have actual data on plant level input and output prices. So it is unclear whether this would make too much of a practical difference to our results.

An alternative approach would be to follow de Loecker (2010) and put more structure on the product market. For example, assuming that the product market is monopolistically competition enables the econometrician in principle to estimate the elasticity of demand and correct for the mark-up implicit in TFPR to obtain TFPQ. We will pursue this in future work.

C Appendix C: More Details of some Size-Relation Labor Market Regulations in France

The main bite of labor (and some accounting) regulations comes when the firm reaches 50 employees. But there are also some other size-related thresholds at other levels. The main other ones comes at 10-11 employees. For this reason we generally trim the analysis below 12 employees to mitigate nay bias induced in estimation from these other thresholds. For more details on French regulation see inter alia Abowd and Kramarz (2003) and Kramarz and Michaud (2010).

C.1 Labor Regulations

From fifty employees:

- Obligation to use a complex redundancy plan with oversight, approval and monitoring from Ministry of Labor in case of a collective redundancy for 9 or more employees (threshold based on total employment at time of redundancy). See text.

- Appointing a shop steward if demanded by workers (threshold exceeded for 12 consecutive months during the last three years);

- Obligation to establish a committee on health, safety and working conditions (HSC) and train its members (threshold exceeded for 12 months during the last three years)

- Obligation to establish a profit sharing (threshold exceeded for six months during the accounting year within one year after the year end to reach an agreement);

- Obligation to establish a staff committee with business meeting at least every two months (plant level: threshold exceeded for 12 months during the last three years)

From twenty-five employees:

- Duty to supply a refectory if requested by all employees;

- Electoral colleges for electing representatives. Increased number of delegates from 26 employees.

From twenty employees:

- Contribution to the National Fund for Housing Assistance;

- Increase the contribution rate for continuing vocational training of 1.05% to 1.60%

- Compensatory rest of 50% for mandatory overtime beyond 41 hours per week

From eleven employees:

- Allowance of at least six months salary if terminated without cause or serious;

- Obligation to conduct the election of staff representatives(threshold exceeded for 12 consecutive months over the last three years).

- **From ten employees:**

- Monthly payment of social security contributions, instead of a quarterly payment (according to the actual last day of previous quarter);

- Obligation for payment of transport subsidies (Article L. 2333-64 of the General Code local authorities);

- Increase the contribution rate for continuing vocational training of 0.55% to 1.05% (threshold exceeded on average 12 months).

C.2   Accounting rules

The additional requirements depending on the number of employees of enterprises, but also limits on turnover and total assets are as follows:

From fifty employees:

- loss of the possibility of a simplified presentation of Schedule 2 to the accounts (also if the balance sheet total exceeds 2 million or if the CA exceeds 4 million);

- requirement for LLCs, the CNS, limited partnerships and legal persons of private law to designate an auditor (also if the balance sheet total exceeds 1.55 million euros or if the CA is more than 3.1 million euros, applicable rules of the current year).

From ten employees:

- loss of the possibility of a simplified balance sheet and income statement (also if the CA exceeds 534 000 euro or if the balance sheet total exceeds 267 000 euro, applicable rule in case of exceeding the threshold for two consecutive years).

D   Appendix D: Data

Figure 1a: The Firm size distribution in the US and France compared
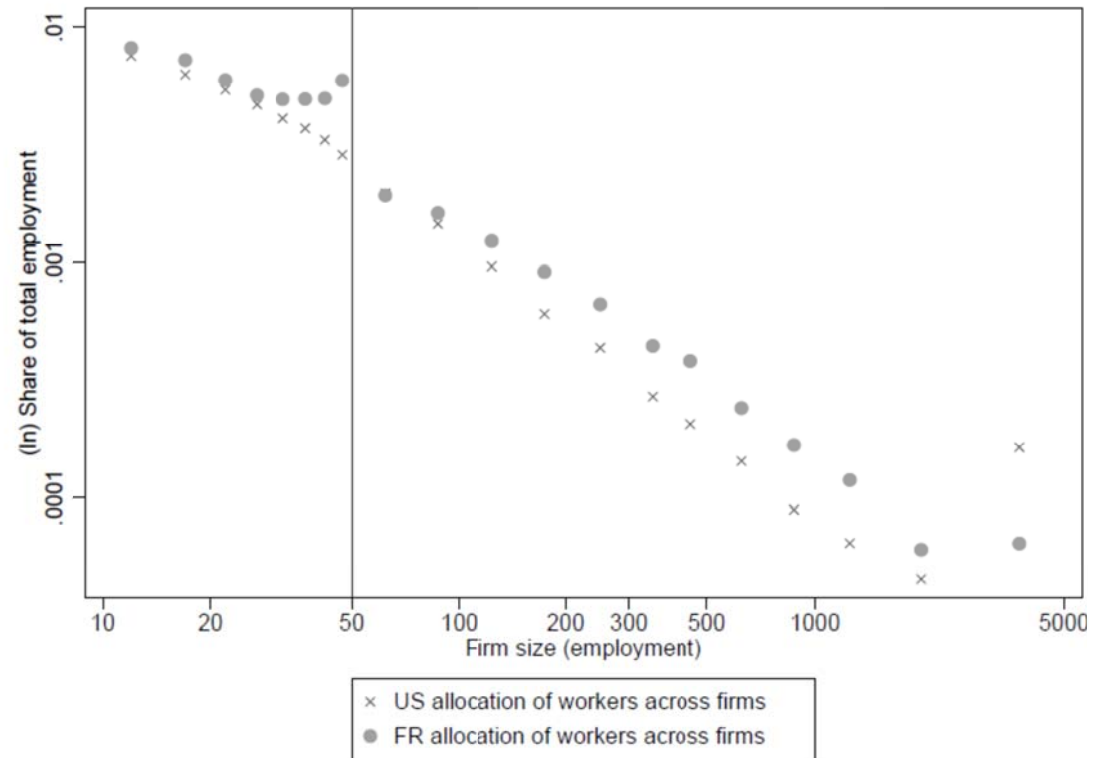
Source: FICUS for France and Census for the US. Population databases of all firms.

Notes: This is the distribution of firms (not plants). Authors' calculations

Figure 1b: Number of Firms by employment size in France
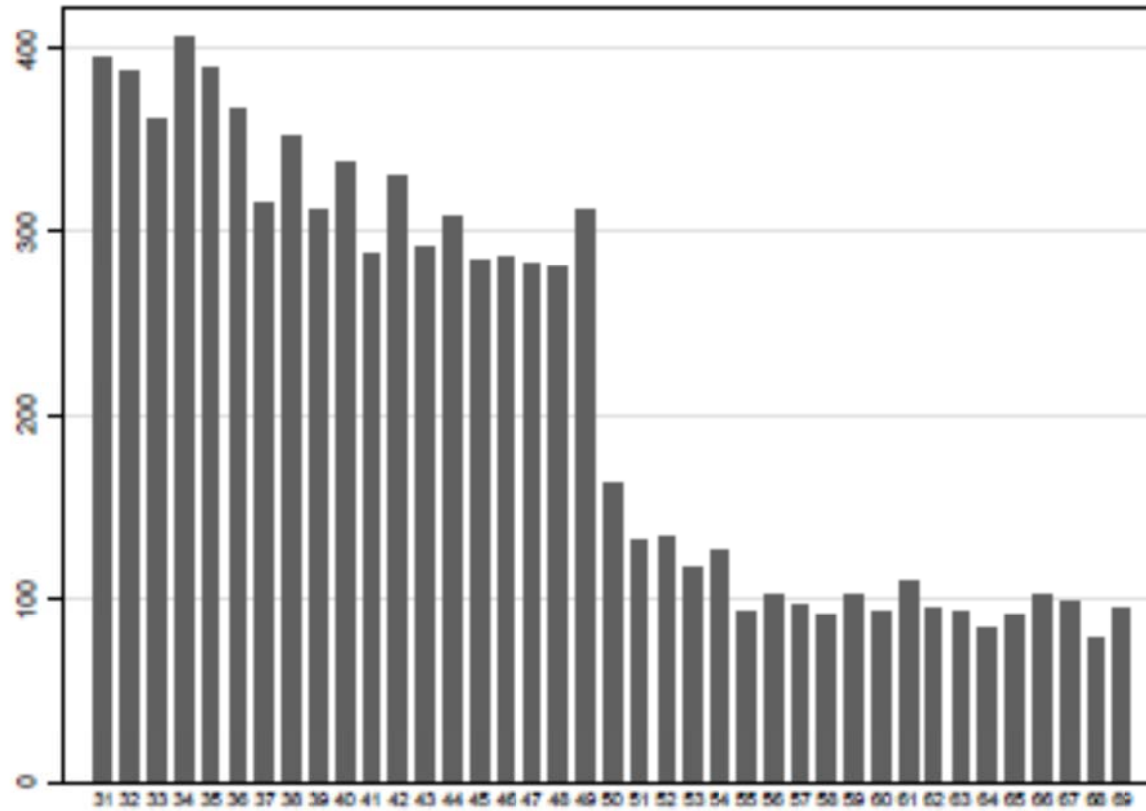
Source: FICUS

Notes: This is the population of manufacturing firms in France with between 31 and 69 employees. This plots the number of firms in each exact size category (i.e. raw data, no binning). There is a clear drop when the employment regulation begins for firms with 50 or more employees.

Figure 2: Theoretical Firm size distribution with regulatory constraint

FirmSizeDistribution
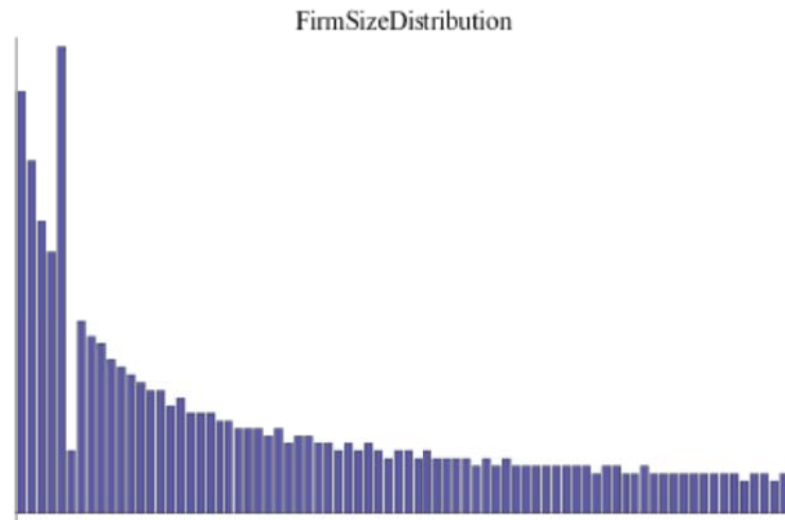
Notes: This figure shows the theoretical firm size distribution. The tallest bar represents the point at which the size constraint binds.

Figure 2.b: Theoretical Firm size distribution with regulatory constraint and measurement error
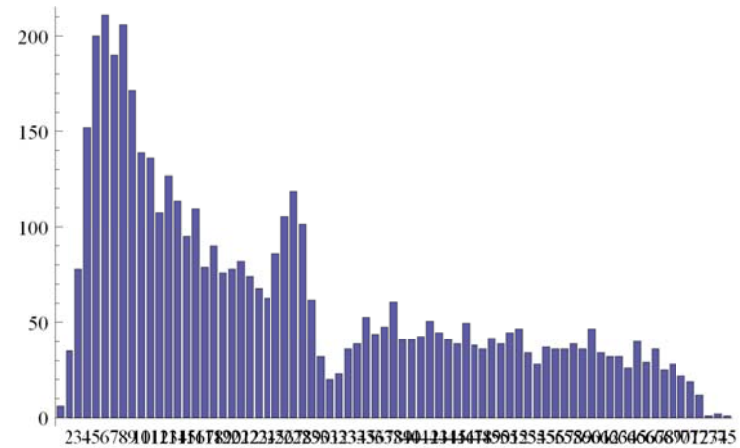
Notes: This figure shows the theoretical firm size distribution with additive N(0,3) measurement error. The tallest bar represents the point at which the size constraint bins.

Figure 3: Theoretical Relationship between TFP (managerial talent) and firm size
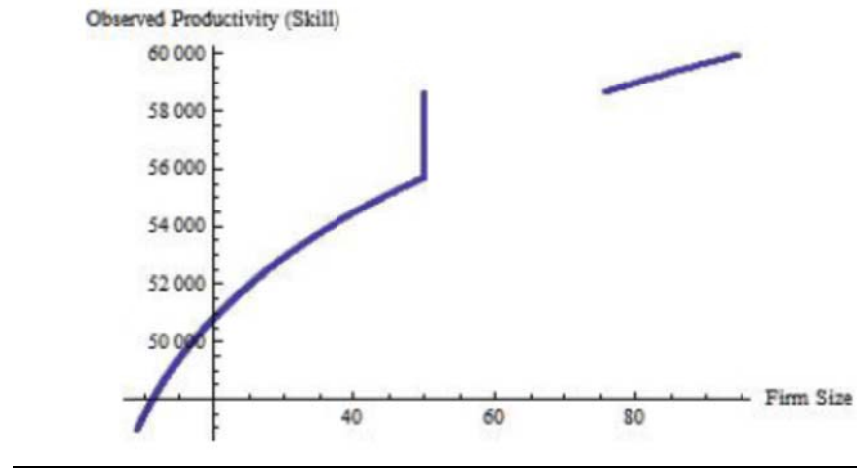
Notes: This figure shows the theoretical relationship between TFP and firm size. There is a mass of firms at employment size=50 where the regulatory constraint binds.

Figure 4: Definitions of regimes in terms of employment
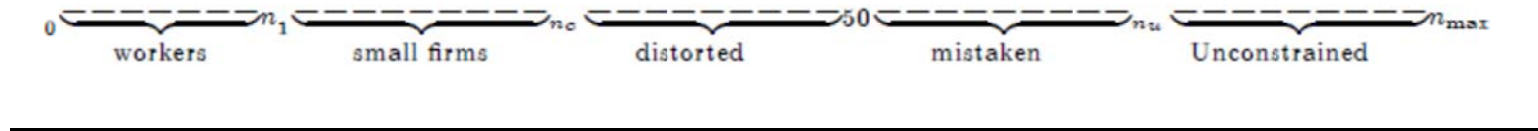
Notes: This figure shows the definitions of different regimes in our model. Below n_1 individuals choose to be workers rather than managers. Between n_1 and n_c there are "small firms". Between n_c and 50 are some firms who are affected by the regulatory constraint by choosing to be smaller than they otherwise would have been: we call this the "distorted" regime. Between n_u and the largest firm n_max are firms who are choosing to pay the tax rather than keep themselves small. In the pure theory model there should be no firms between 50 and n_u, but we discuss in the text reasons why there are empirically some firms located here (essentially measurement error and "mistakes").

Figure 5:

Number of Firms by employment size: alternative datasets and definitions
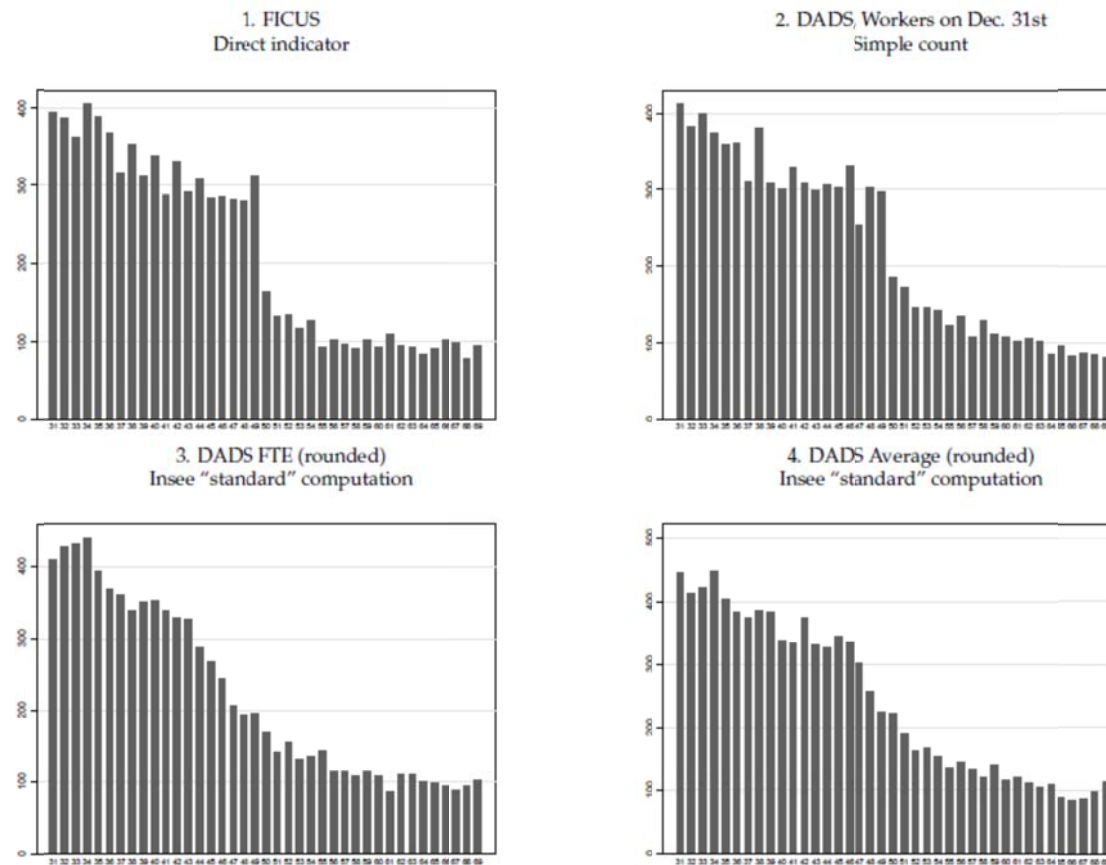
Figure 6:

Share of Firms by employment size: Distribution over a wider range of firms
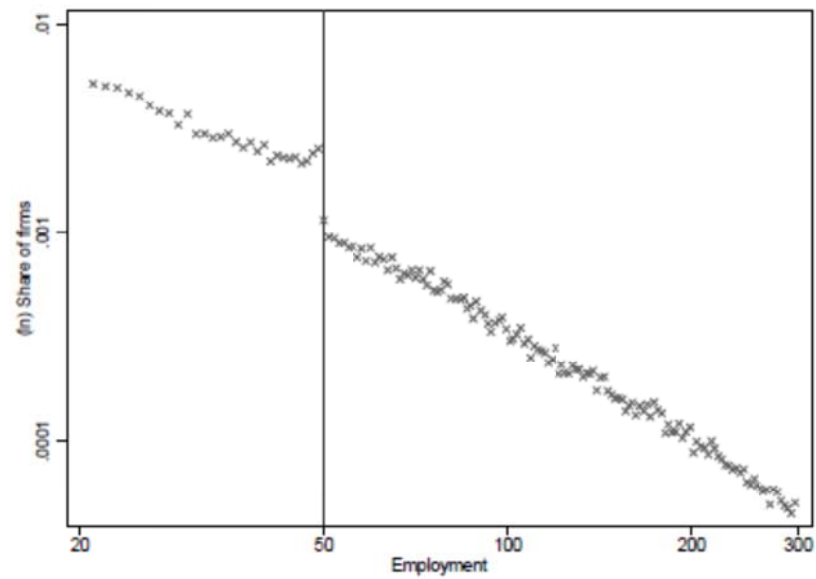
Notes: The raw data is in "power law" form, where the share of firms in the population is presented on the y-axis and the employment size of each firm is presented on the presented on the x-axis (on a log scale). The vertical line at 50 employees indicates where the regulation binds.

Figure 7: TFP Distribution by Firm Size
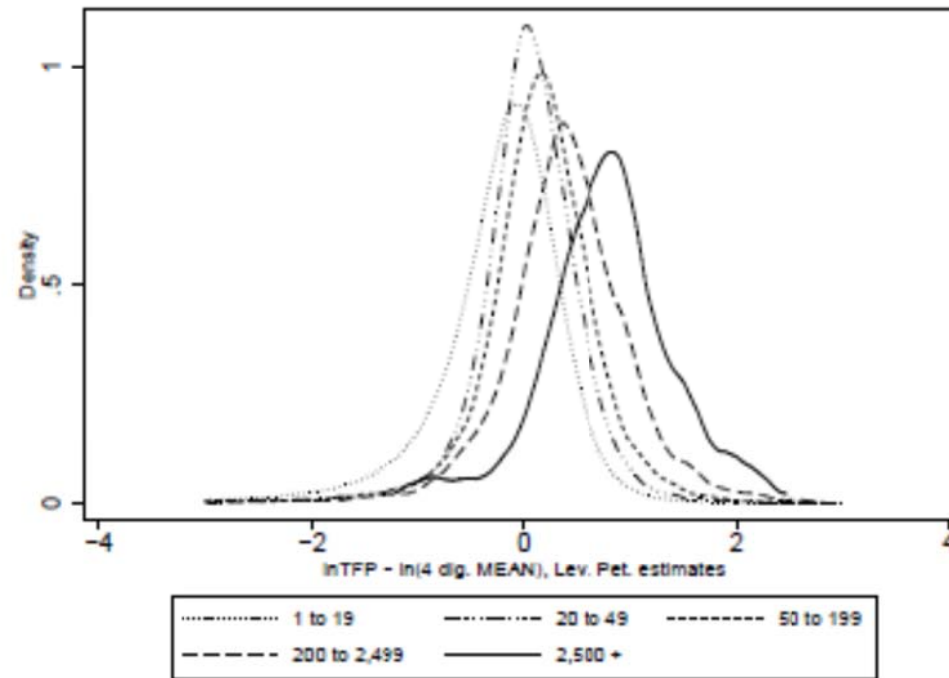
Notes: This figure plots the (kernel density smoothed) distribution of TFP (estimated by the Levinsohn Petrin method) across four size classes. TFP is relative to the four digit industry by year average.

Figure 8: TFP Distribution around the regulatory threshold of 50 employees

Panel A: Short Employment span
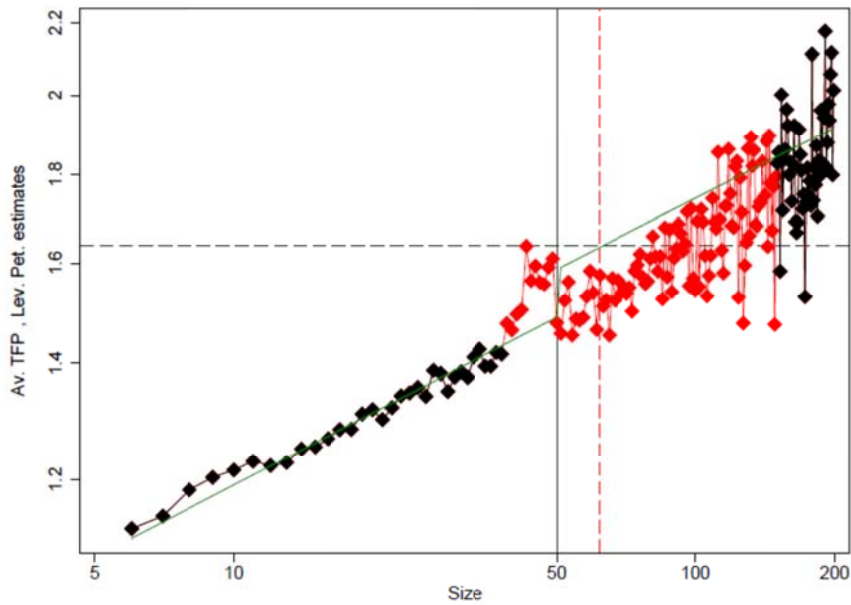
Panel B: Longer Employment span
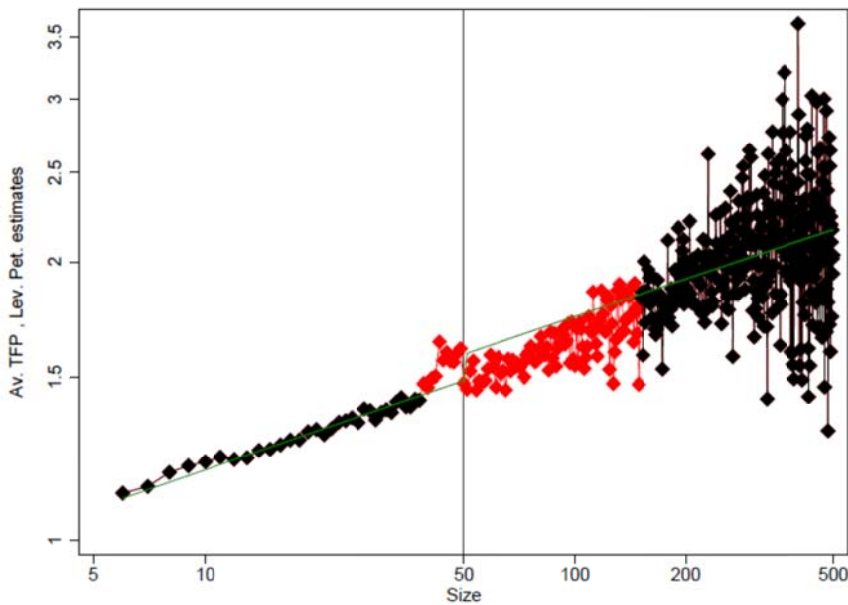
Notes: This figures plots the mean level of TFP by firm employment size using an upper support of 100 (Panel A) or 500 (Panel B). A third order polynomial is displayed in both panels using only data from the "undistorted" points (shown in red).

Figure 9: Estimating parameters using the Firm Size distribution

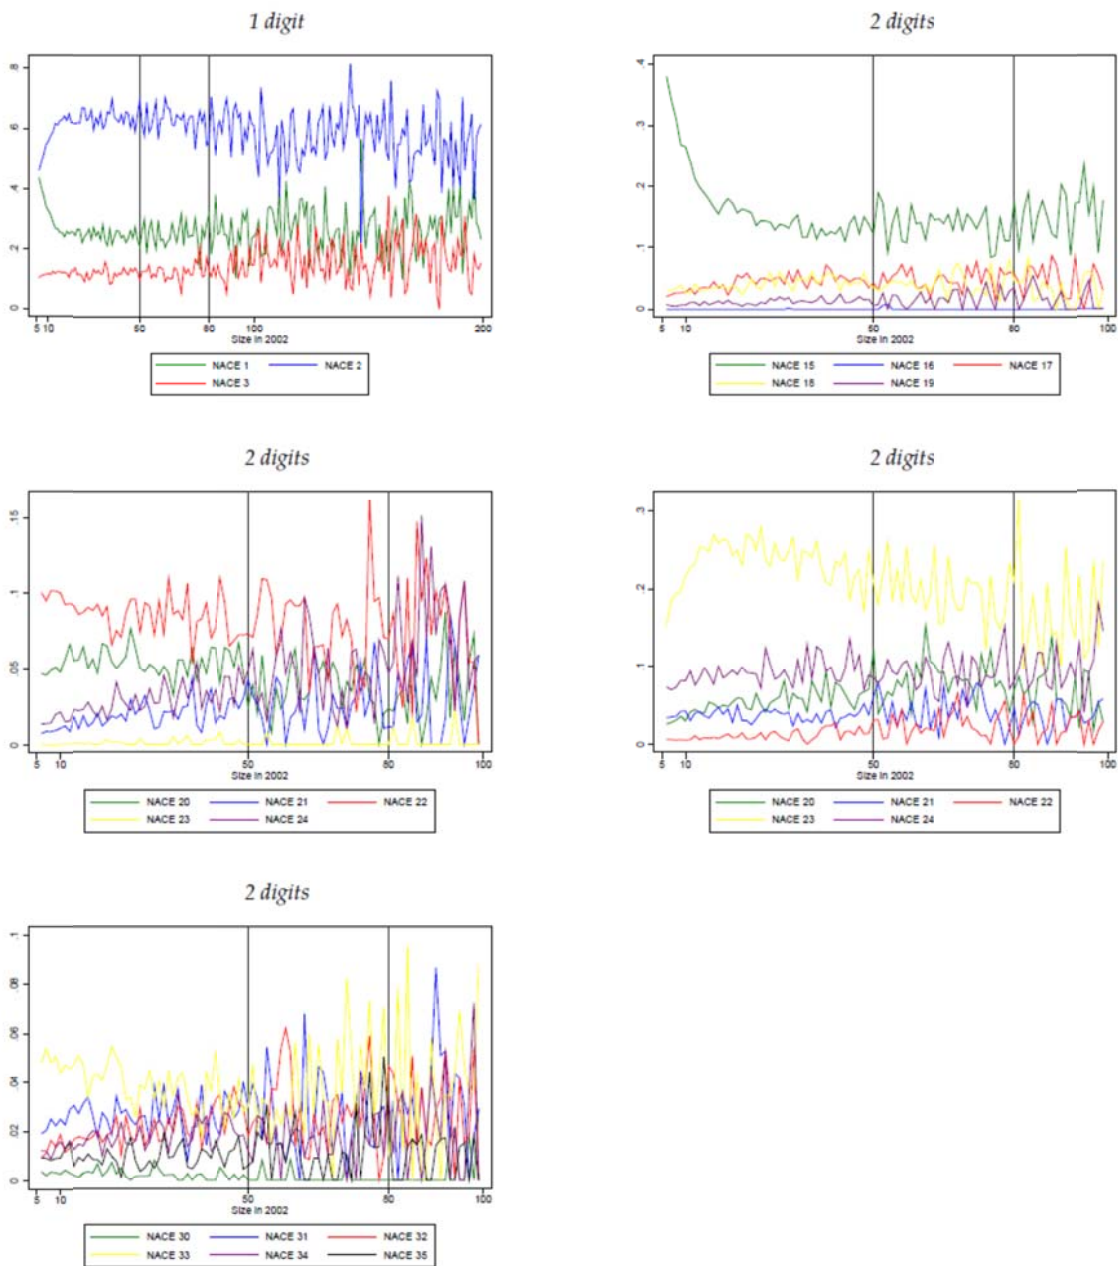## (implication is that τ =1.265 and n_u =68)
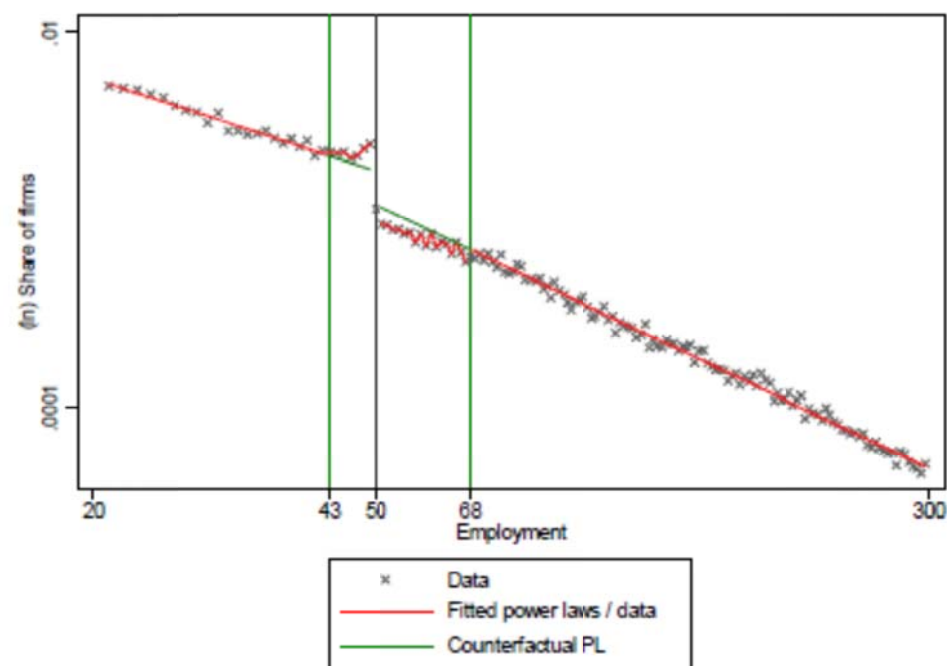
Notes: This figure illustrates how we calculate the parameters of the model from the size distribution. We estimate separate power laws between employment and the share of firms in two regimes: (i) employment size 20 to 42 and (ii) (employment size 100 to 300). We use these as the counterfactual slopes in the counterfactual world but for the regulation (these are projected over the potentially distorted regions between employment size 43 to 100) . The adding up constraint for total employment implies the values of the two unknown parameters, τ and n_u (see text). The vertical lines are the lower threshold (n_c =43), the regulatory constraint (n=50) and the estimated upper employment threshold (n_u =68).

Figure 10: Estimating parameters using the TFP Distribution

(implication is that $\tau = 1.14$, $n_u = 76$)
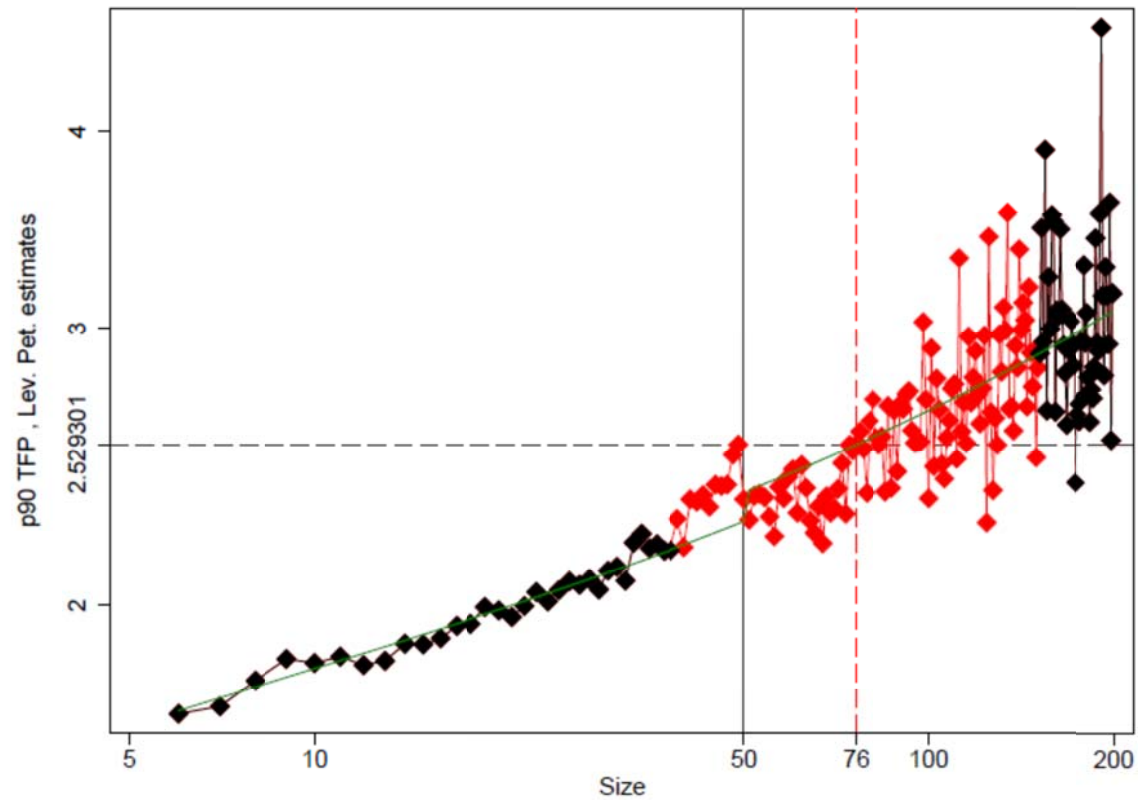
Notes: This figure illustrates how we calculate the parameters of the model from the TFP and size distribution. We predict the 90th percentile of TFP distribution from a fourth order polynomial in size using only data on the part of the distribution which is away from the thresholds (the black points). The vertical lines show the regulatory constraint (n=50) and the estimated upper employment threshold ($n_u = 76$). See text for more details.

Figure 11: Industry Composition around the threshold
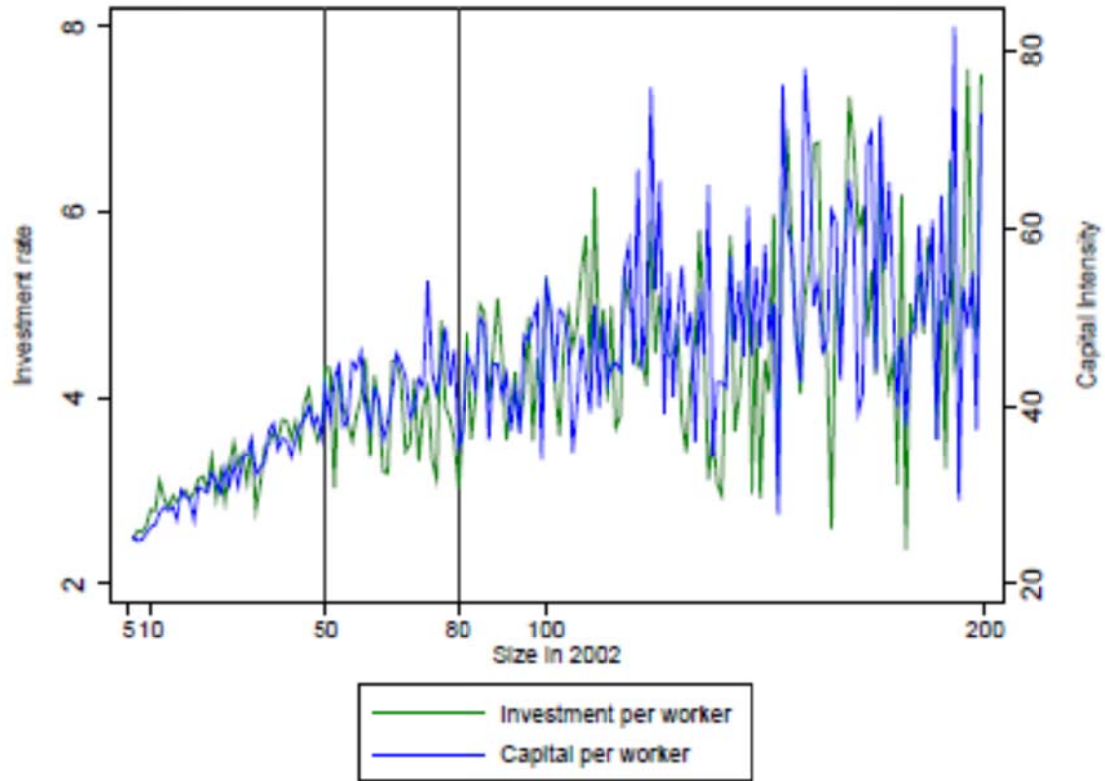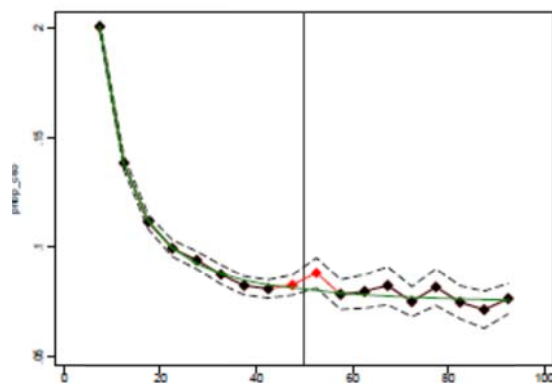
Figure 13: Adjustment in Types of Labor?

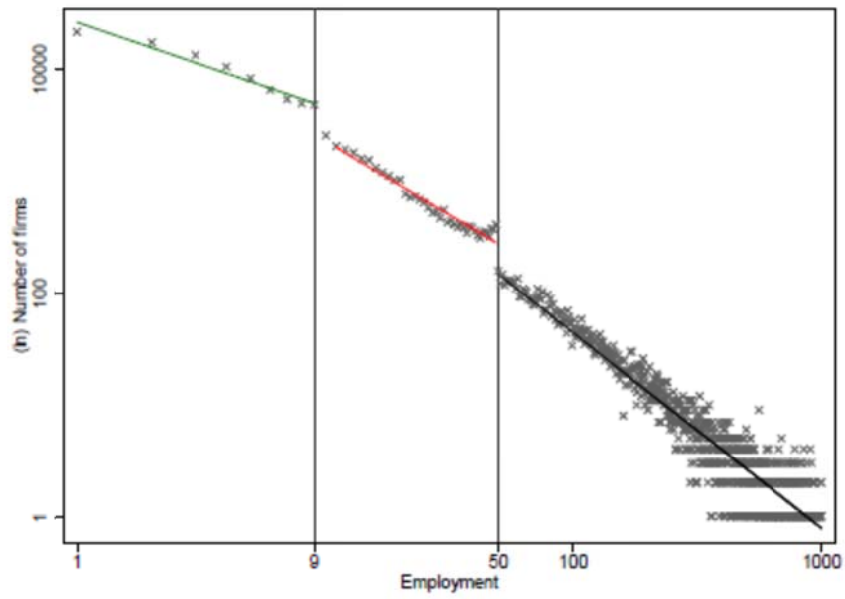## A. Share of Managers
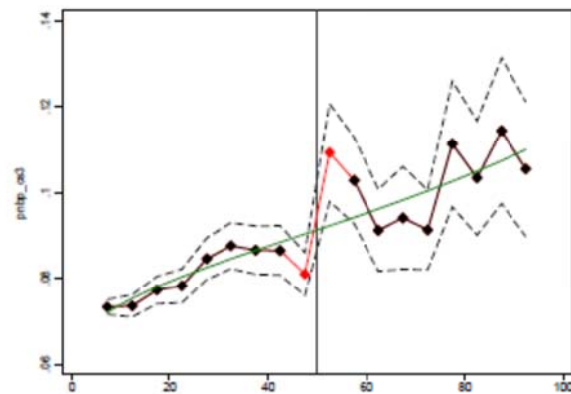
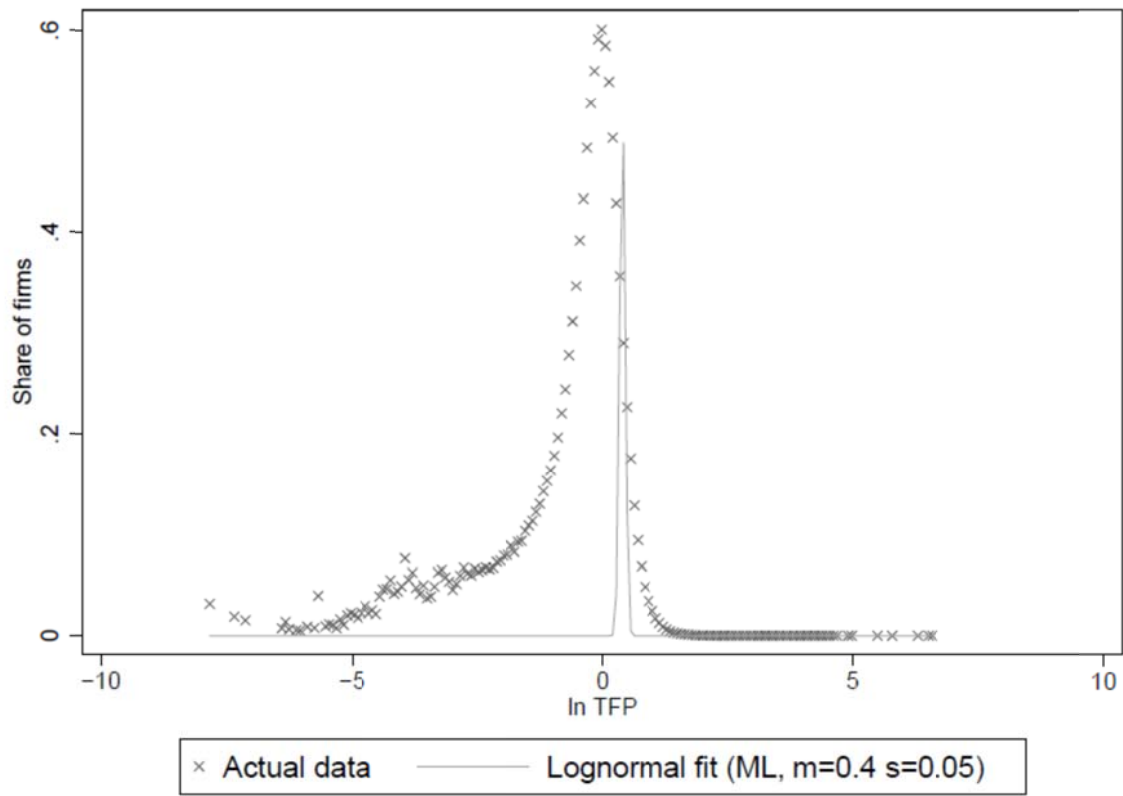## B. Share of manual workers
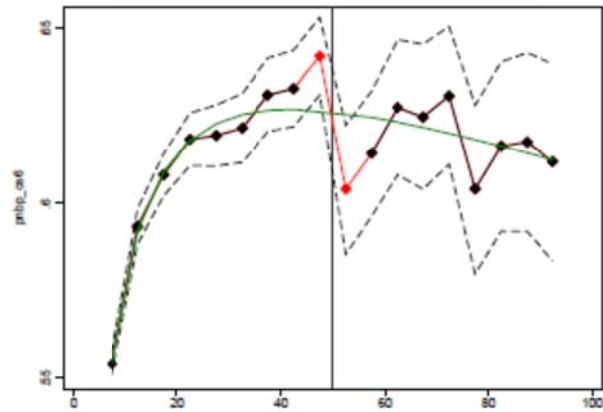
C. Share of Clerical Workers

Notes: This looks at the share of the main three occupational groups using combined FICUS and DADs data.
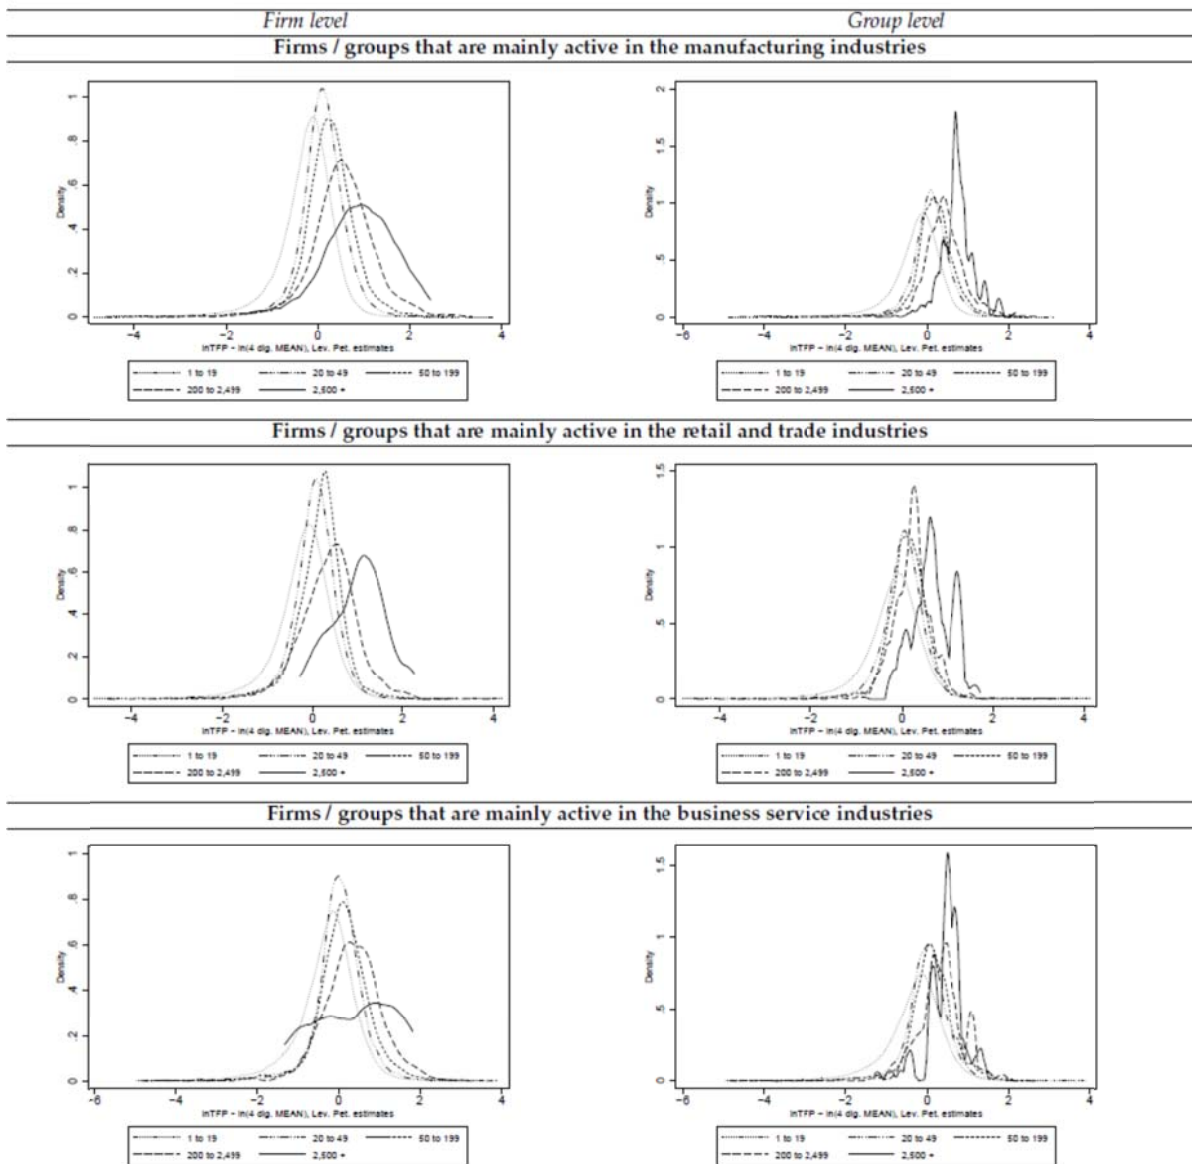
Appendices

Figure A1: Size Distribution over a larger range (up to 1000 employees)

Figure A2: TFP Distributions

Legend: × Actual data — Lognormal fit (ML, m=0.4 s=0.05)

Figure A3: TFP Distributions by Size by broad industrial Sector (broken down into firms (on left) and enterprise groups (on right)

Figure A4: Transition Probabilities

After 1 year

After 2 years

After 5 years

Entry