NBER WORKING PAPER SERIES

SIZE-DEPENDENT REGULATIONS, FIRM SIZE DISTRIBUTION, AND REALLOCATION

François Gourio
Nicolas A. Roys

Working Paper 18657
http://www.nber.org/papers/w18657

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2012

We thank Nick Bloom, Jeff Campbell, Theodore Papageorgiou, Marcelo Veracierto, and conference and seminar participants for useful discussions and comments. The views expressed here are those of the authors and do not necessarily represent those of the Federal Reserve Bank of Chicago or the Federal Reserve System or the National Bureau of Economic Research.

Size-Dependent Regulations, Firm Size Distribution, and Reallocation
François Gourio and Nicolas A. Roys
NBER Working Paper No. 18657
December 2012
JEL No. E23,L11,L25,O1

ABSTRACT

In France, firms with 50 employees or more face substantially more regulation than firms with less than 50. As a result, the size distribution of firms is visibly distorted: there are many firms with exactly 49 employees. We model the regulation as a sunk cost that must be paid the first time the firm reaches 50 employees, and we estimate the model using indirect inference by fitting these salient features of the size distribution. The key finding is that the legislation acts like a sunk cost equivalent to approximately one year of an average employee salary. Removing the regulation improves labor allocation across firms, leading to a productivity gain of around 0.3%, holding the number of firms fixed. However, if firm entry is elastic, the steady-state gains are significantly smaller.

François Gourio
Department of Economics
Boston University
270 Bay State Road
Boston, MA 02215
and CEPREMAP
and also NBER
fgourio@bu.edu

Nicolas A. Roys
University of Wisconsin-Madison
Department of Economics
1180 Observatory Drive
Madison 53703 WI
nroys@ssc.wisc.edu

1 Introduction

In many countries, small firms face lighter regulation than large firms. Regulation, broadly defined, takes many forms, from hygiene and safety rules, to mandatory elections of employee representatives, to larger payroll taxes. The rationale for exempting small firms from some regulations is that the compliance cost is too high relative to their sales. A necessary consequence, however, is that regulations are phased in as the firm grows, generating an implicit marginal tax. Because regulations are typically phased in at a few finite points, they are sometimes referred to as "threshold effects": for instance, in the case of France, a first important set of regulations applies to firms with more than 10 employees, and a second important set of regulations applies to firms with more than 50 employees. As a result, the firm size distribution is distorted, with few firms with exactly 10 (or 50) employees and a large number of firms with 9 (or 49) employees. Figure 1 plots the firm size distribution in our French data, illustrating this well-known pattern.

These distortions have generated a large interest from public policy circles; for instance numerous commissions drew attention to this issue (see for instance Cahuc and Kramarz (2004)). In spite of this interest, there is little work formally modeling these policies to understand and evaluate their effects.

On the positive side, a structural model is needed to understand the exact sources of distortion. It is not obvious how the regulations should be modeled, given their scope and complexity (which we discuss in detail in section 2). Are regulations equivalent to higher fixed costs, higher proportional taxes on labor, or to a sunk cost? The puzzle that quickly emerges is, why are there *any* firms at all with exactly 50 employees given the higher fixed costs? Our intuition is that many of these regulations might be better approximated as a sunk cost (i.e. a one-time investment), since a large fraction of the cost is learning the regulation. The presence of the sunk cost also helps explain why there are some firms that have exactly 50 employees: firms are reluctant to have more than 50 employees the first time that they reach that limit, but they do not care about the limit in subsequent periods, since the cost is already paid.

On the normative side, what are the potential benefits of removing, or smoothing, the regulation thresholds? The visibly distorted firm distribution suggests that productivity could be increased if firms close to the threshold grow, as labor would be reallocated towards more productive firms.

To address these questions, we introduce and estimate a simple structural model that takes into account the phase-in of the regulation. Our model incorporates both a sunk cost of complying with the regulation (which captures both the cost of learning the regulation and the cost of any one-time investment that it requires), a higher per-period payroll tax and higher fixed costs. We show that the later two have similar implications for employment, hence we concentrate on the case of the payroll tax.

Our model can be solved using standard stochastic dynamic optimization techniques (Dixit and Pindyck (1994), Stokey (2008)), and we obtain the cross-sectional distribution in closed form. This is useful when we turn to the estimation because simulating accurately the highly skewed cross-sectional distribution of firms is challenging.

For clarity, we fit two polar cases of our model to the data: first, the case where the regulation involves only a sunk cost, and second, the case where the regulation involves only a per-period higher payroll tax.

In the first case, we find that the regulation is equivalent to a sunk cost of about one year of an average employee wage. We next use our model estimates to infer the social cost of the regulation. Holding the number of firms constant, we find a productivity loss of 0.3% due to misallocation of labor across firms. However when we allow the number of firms to adjust, we find a much smaller effect, around 0.04%. This suggests that these regulations may not have large aggregative effects. In the second case, where the regulation is modeled as a per period tax, this tax is estimated to be fairly small, about 0.26%, and its aggregate effects are even smaller than in the sunk cost model. Finally, we also provide some evidence that the sunk cost model fits the data better.

The rest of the paper is organized as follows. We first discuss the related literature. Section 2 presents the data and some reduced-form evidence that motivates our analysis. Section 3 discusses the model. Section 4 covers our estimation method and presents the empirical results. Section 5 uses these estimates to conduct some policy experiments. Section 6 concludes.

Related Literature Our paper is related to a recent growing literature which studies the effect of misallocation on aggregate productivity and welfare. Building on Hopenhayn and Rogerson (1993), Restuccia and Rogerson (2008) and Buera et al. (2011) suggest that misallocation is an important determinant of total factor productivity (TFP). Hsieh and Klenow (2009) and Bartelsman et al. (2009) present empirical evidence consistent with higher misallocation in poorer countries with lower TFP. Closely related to our paper, Guner et al. (2008) suggest that for the same reason, size-dependent policies have a large negative impact on total factor productivity.

In the studies of Restuccia and Rogerson (2008) and Guner et al. (2008), distortions arise due to implicit "taxes". However these taxes are not directly measured. The regulations that we discuss are a prime example of these distortions, and they very clearly affect the firm distribution, consistent with these studies. While our aim is more modest than the macro studies, since we focus on one particular distortion, we believe that our focus allows a credible identification of the effect of government regulation on firms outcomes. In particular, we evaluate whether it is feasible to match the distortion in the firm size distribution, which is the *prima facie* evidence that the regulation matters.

There are several existing studies documenting the distortion in size distribution in France (for instance Cahuc and Kramarz (2004) or Ceci-Renaud and Chevalier (2011)), but we are not aware of any structural modeling that tries to apprehend the costs of the distortion. While finishing this paper, we became aware of a very recent working paper Garicano et al. (2012) that shares some of our goals and approach. The key differences between our papers are that we focus on the sunk cost element of the regulation and aim to fit the distribution around the threshold, whereas they focus on the labor tax element and aim at the entire firm distribution. Hence, while we use similar data, we have different models, estimation methods and targets. Overall, our results are complementary. We compare our results in more detail in sections 4 and 5.

2 Motivating Evidence

We first describe briefly the institutional background, then we present our data sources, and finally we show some simple reduced-form evidence of the threshold effects.

2.1 Institutional Background

This section draws heavily from Ceci-Renaud and Chevalier (2011). Labor laws in France as well as various accounting and legal rules make special provisions for firms with more than 10, 11, 20, or 50 employees.

These regulations are not all based on the same definition of “employee”. Labor laws, which are likely the most important, are based on the full-time equivalent workforce, computed as an average over the last 12 months. The full-time equivalent workforce takes into account part-time workers, as well a temporary workers, but not trainees or subsidized employment (*contrats aidés*). On the other hand, several rules are based on sales as well as employment.

The main additional regulations as the firm reaches 50 employees are:

- possibly mandatory designation of an employee representative;

- a committee for hygiene, safety and work conditions must be formed and trained;

- a *comité d'entreprise* must be formed, that must meet at least every other month; this committee, that must have some office space and receives a subsidy equal to 0.2% of the total payroll, has both social objectives (e.g., organizing cultural or sports activities for employees) and an economic role (mostly on an advisory basis);

- higher payroll tax subsidizing training which goes from 0.9% to 1.5% (*formation professionelle*);

- in case of firing of more than 9 workers for "economic reasons", a special legal process must be followed (*plan social*). This legal process implies potentially a larger cost and higher uncertainty for the firm.

We emphasize that these are just a subset of the regulations which apply. This is enough to give a glimpse of why one may expect them to be important, and also the difficulty of modeling these rules in a simple model: while some are simply monetary rules (e.g., higher taxes), many add an element of uncertainty, and many require the firm to do some organizational work.

2.2 Data

We use a panel data of firms assembled by the French National Statistical Institute (INSEE), that covers the 1994-2000 period. This panel, known as BRN (Bénéfices Réels Normaux), contains employment as well as standard accounting information on total compensation costs, value added, current operating surplus, gross productive assets, etc. The BRN data are exhaustive of all private companies with a sales turnover of more than 3.5 million Francs (around 530,000 Euros) and liable to corporate taxes under the standard regime, and include some other smaller firms. For our purpose, the 3.5 million threshold implies that we have all firms with more than 30 employees or so. Hence we focus on the threshold at 50 employees, for which our data is essentially exhaustive.

We removed from the sample firms with strictly less than 20 employees when we estimate the model. This generated a sample of 44,1890 firms that we follow for 7 years, or 309,323 firm-year observations.

2.3 Preliminary data analysis

Figure 1 plots the distribution of employment for the entire period (1994-2000) and is truncated at 100 employees. Figure 2 zooms on the distribution between 40 and 60 employees. There are clearly large discontinuities around the thresholds of 10 and 50 employees. On the other hand, the threshold for 20 employees appears less significant. Many surveys reveal "rounding" of employment, but this figure shows the opposite pattern.

Table 1 reports the number of firms by number of workers over the range $40 - 60$ normalized by the fraction of firms between 40 and 60. There is a clear drop in the number of firms after 49 employees. For example, there more than three times as many firms with 49 employees as firms with 51 employees.

A useful way to summarize the break in this distribution is to approximate it with a power law distribution. The power law assumption states that the probability that firm size is greater than x is proportional to $x^{-\xi}$. Formally, $P(\text{Size} > x) = Cx^{-\xi}$ where C and ξ are constants.

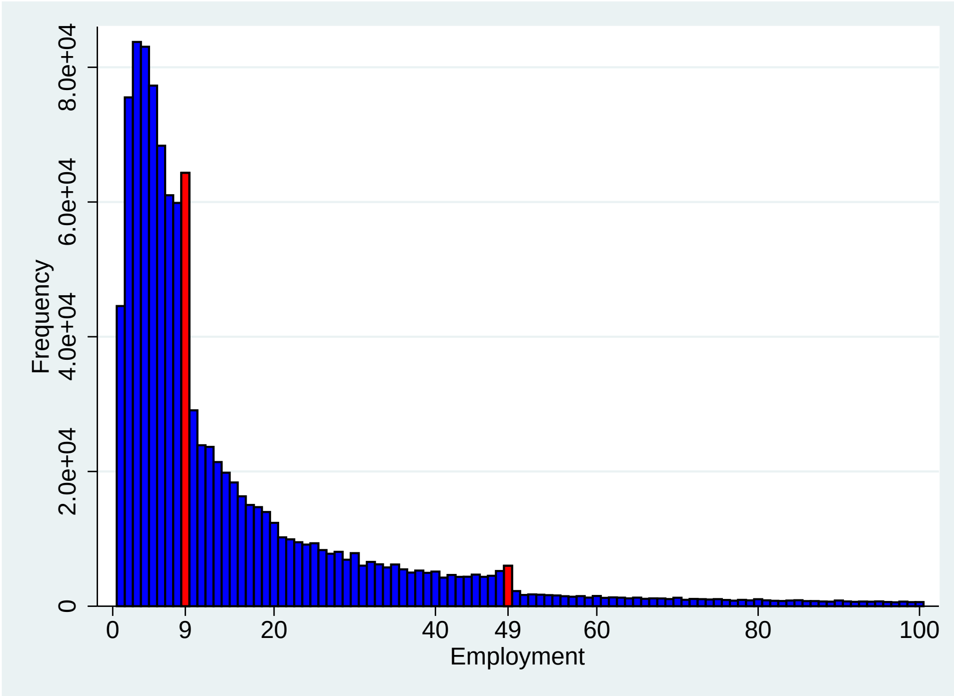

Figure 1: Firm size distribution by employment between 1 and 100 employees.

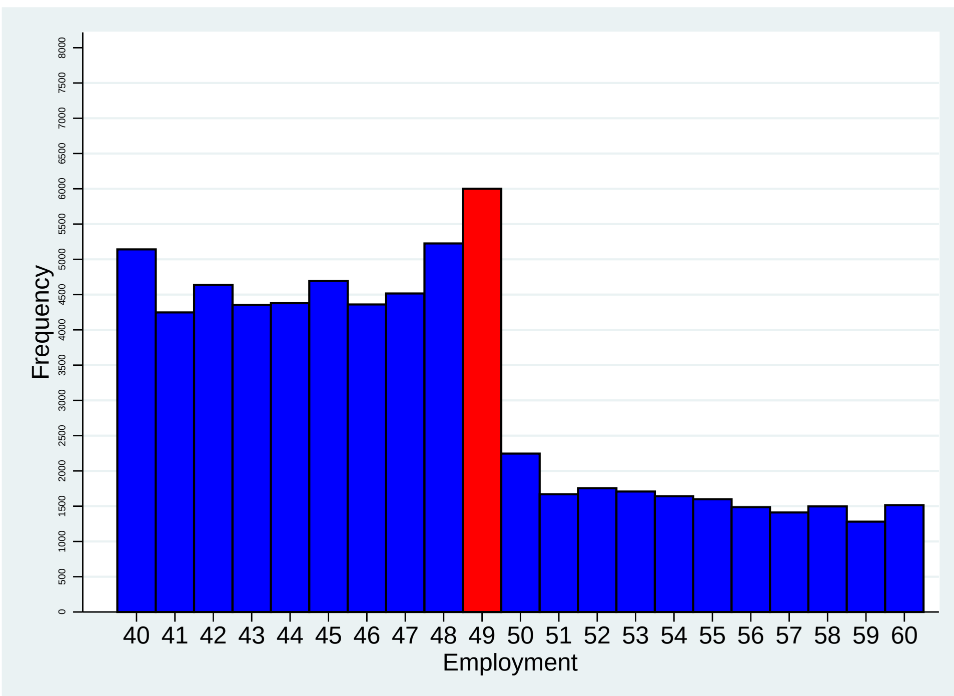

Figure 2: Firm size distribution by employment between 40 and 60 employees.

	Fraction	Std. Dev.	# Firms		Fraction	Std. Dev.	# Firms
40	8.05	0.1089	5141	50	3.51	0.0721	2246
41	6.65	0.1004	4247	51	2.61	0.0651	1669
42	7.26	0.1035	4637	52	2.75	0.0655	1755
43	6.82	0.1023	4355	53	2.67	0.0648	1708
44	6.85	0.1001	4378	54	2.57	0.0631	1641
45	7.34	0.1049	4692	55	2.50	0.0617	1598
46	6.82	0.1032	4360	56	2.34	0.0564	1486
47	7.07	0.0998	4516	57	2.21	0.0587	1411
48	8.18	0.1129	5225	58	2.34	0.0583	1497
49	9.39	0.1167	6001	59	2.00	0.0564	1281

Table 1: Fraction is the number of firms for each employment size over the range $40 - 60$, normalized by the total number of firms between 40 and 60; Std. Dev is the standard error of the fraction of firms for each employment level; and #Firms is the raw number of firms in each bin.

Figure 3 displays the results of two estimations. First, we estimate the parameters C and ξ of the power law for firms with more than 100 employees. The power law seems to approximate well the firm size distribution for all but the largest firms. This is a well-known result (See Axtell (2001) and Di Giovanni et al. (2011) among others). Second, we run a regression of the log frequency on log size, with or without a structural break at size 50. The presence of a structural break is clearly visible from this second figure. In the Appendix, Figure 13 shows that the same pattern hold across sectors.

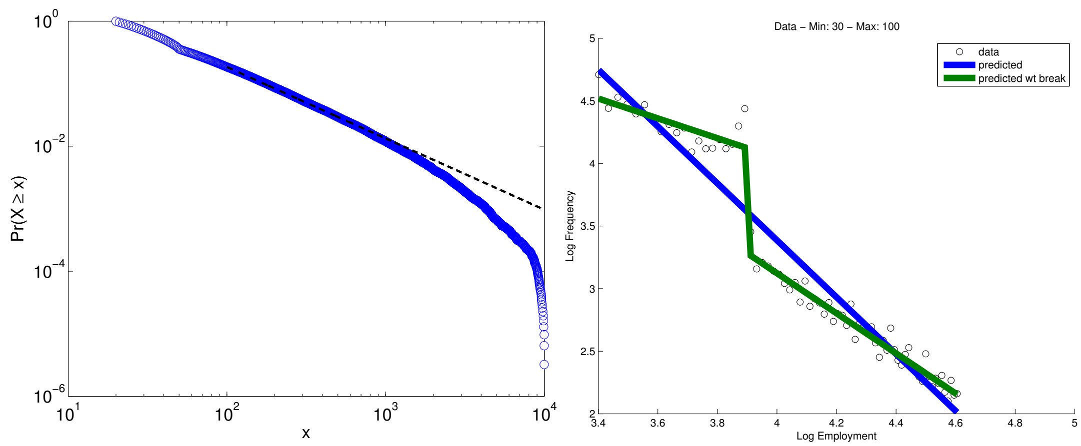

Figure 3: Power Law Estimation: (a) Estimation by Maximum Likelihood for all the firms with a number of employees greater than 100; (b) Regression of the logged number of firms on the logged number of employees with and without a structural break at 50 for firms with employment level between 30 and 100.

The dynamics of firms around the threshold are also affected. Figure 4 reports the probability that a firm has an employment level in particular bin in period t conditional on having employment in the same bin at time $t-1$. Each bin has a width of 5 employees. Overall, this probability declines with firm size, as inaction

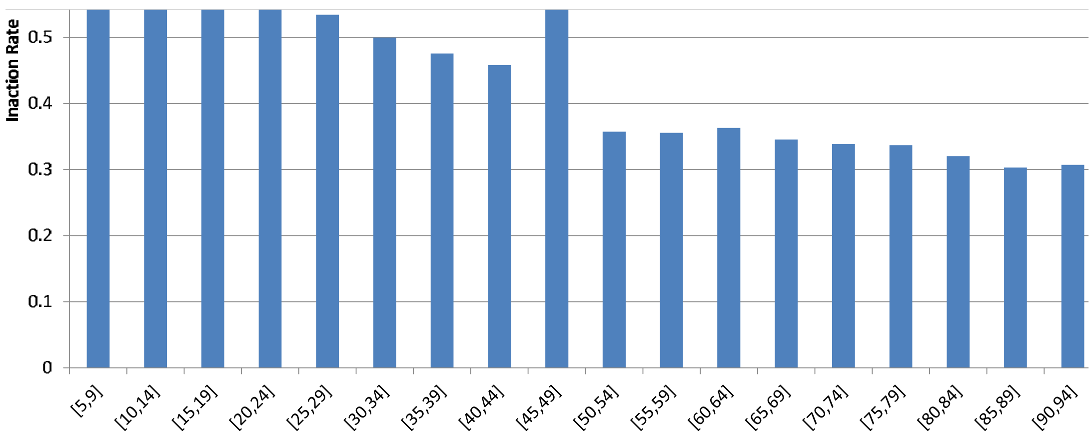

Figure 4: $P(n_t \in \text{bin}|n_{t-1} \in \text{bin})$: probability that a firm has an employment level in particular bin in period t conditional on having employment in the same bin at time $t-1$. Each bin has a width of 5 employees.

is more likely for small firms. Yet, this probability increases right before the threshold. Firms in the bin 45-49 are significantly more likely to remain in that bin next year (57% compared to 47% for the bin 35-39 and 35% for the bin 55-59). This suggests that the presence of the threshold leads to inaction and hence slows down the growth of employment. The same patterns hold if we compute the inaction rate for each level of employment as represented in Figure 5. The probability of keeping employment constant between two consecutive years is 34% at 49 employees which is much higher than this statistics at 40 employees (17%) or 59 employees (11%).

To assess the statistical significance of this result, we estimate a probit characterizing the probability of not adjusting employment. Explanatory variables are a set of dummies variables indicating whether or not last period employment was $45, ..., 55$, the growth rate of production, last period employment, and a set of time dummies capturing aggregate shocks. The estimation uses firms with at least 5 employees over the period 1994-2000. Table 2 reports the coefficients. The probability of inaction increases for firm with a number of employees between 45 and 49. The largest increase is observed for firms of size 49.

Finally, Figure 6 plots labor productivity by size. There are two patterns in this picture: first, labor productivity is higher for large firms, as is well known. Second, while there is substantial noise in this figure, a peak of labor productivity is obtained for 49 employees. This is also a natural implication of the regulation: because firms are reluctant to go over the threshold, they hire less labor than they would, generating larger output per worker.

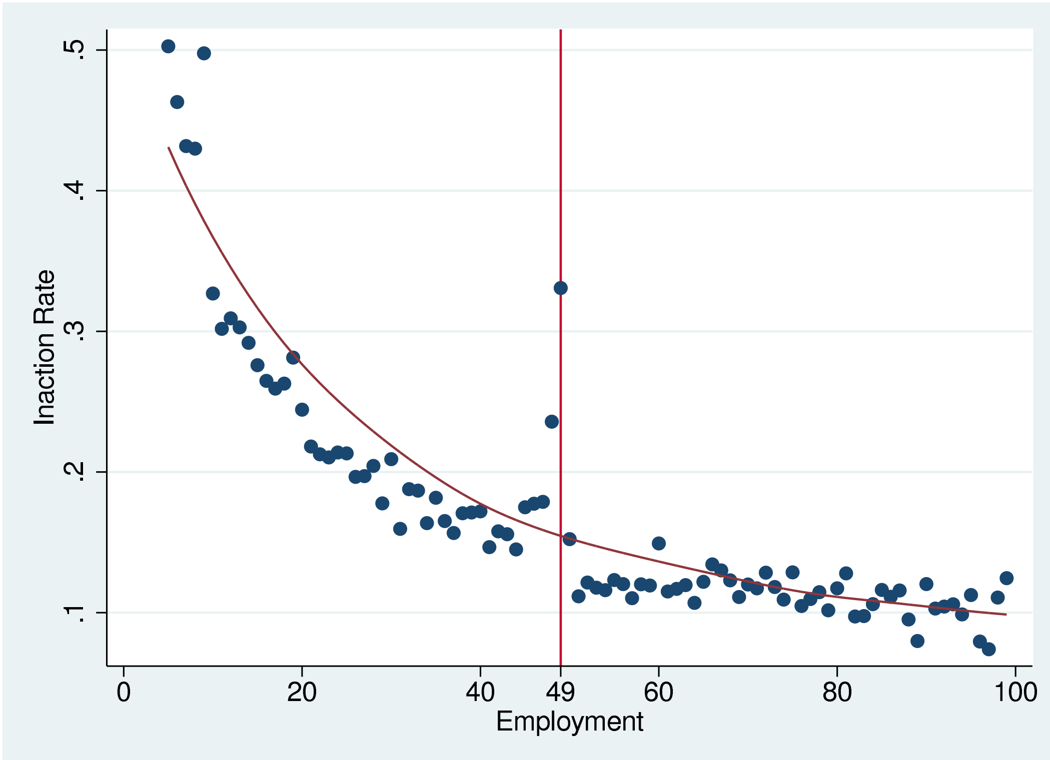

Figure 5: Inaction rate: probability that employment stays constant between two years for each employment level between 5 and 100. Each dot represents a particular employment level. The solid line is a locally weighted regression of the inaction rate on the employment level with bandwidth 0.8. The vertical line represents a level of employment of 49.

Variable	**Coefficient**	**(Std. Err.)**
Production Growth Rate	-0.2136	(0.0060)
Log of Previous Period Employment	-0.4680	(0.0020)
Size 45	0.1037	(0.0235)
Size 46	0.1182	(0.0243)
Size 47	0.1398	(0.0239)
Size 48	0.3561	(0.0209)
Size 49	0.6619	(0.0185)
Size 50	0.0715	(0.0350)
Size 51	-0.1285	(0.0442)
Size 52	-0.0574	(0.0421)
Size 53	-0.0678	(0.0431)
Size 54	-0.0677	(0.0441)
Size 55	-0.0292	(0.0438)

Table 2: Estimation of a probit characterizing the probability of not adjusting employment. Dependent variable is the inaction rate. Explanatory variables are a set of dummies variables indicating whether or not last period employment was $45, ..., 55$, the growth rate of production, last period logged employment, and a set of time and sectoral dummies. (Standard Errors in Parentheses).

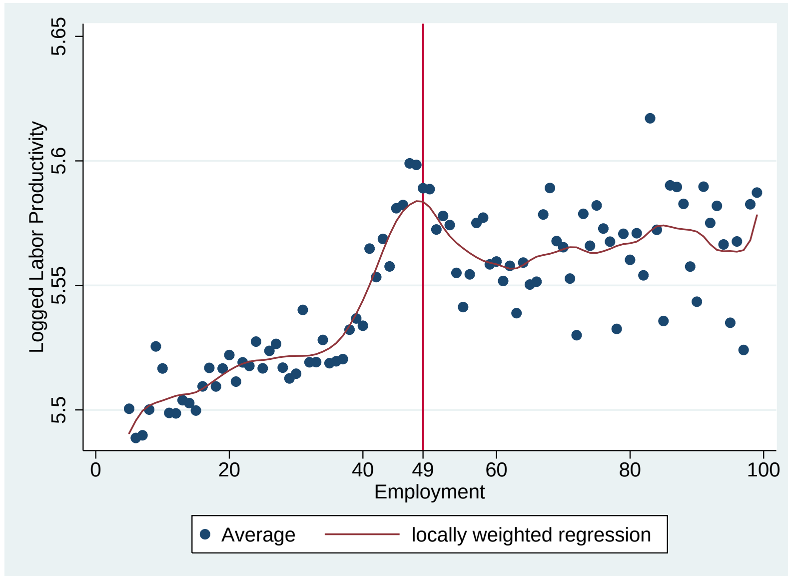

Figure 6: Mean logged labor productivity calculated as the ratio of value-added and employment for each employment level between 5 and 100. Each dot represents an employment level. The solid line is a locally weighted regression of logged labor productivity on the employment level with bandwidth 0.18. The vertical line represents a level of employment of 49.

3 Model

In this section, we introduce and solve a simple dynamic model of production and employment, based on Lucas (1978). For simplicity we assume that there is only one threshold. Firms face a regulation which requires them to pay a sunk cost the first time that their employment exceeds the threshold $\underline{n}$, and firms face higher per-period costs if they currently have more than $\underline{n}$ employees. Hence, our model incorporates both types of costs. When we estimate the model, we consider the two polar cases of zero per-period costs and zero sunk cost.

We start with a partial-equilibrium model, which is the basis of our estimation strategy. Section 5 embeds our model of the firm in a general equilibrium framework to perform some policy experiments.

3.1 Model assumptions

Time is continuous and there is no aggregate uncertainty. There is a continuum of firms, which are ex-ante homogeneous but differ in their realization of idiosyncratic shocks. Each firm operates a decreasing-return to scale, labor-only production:

$$y = e^z n^\alpha,$$

where $\alpha \in (0, 1)$ and e^z is the exogenous productivity (e denotes the exponential function). For simplicity we assume that exit is exogenous and occurs at rate λ. Note that we abstract from fixed costs in this problem; given that we assume exogenous exit, this is without loss of generality. Fixed costs do not affect the employment decision, and we do not use profits data in our estimation.

We assume that log productivity z follows a Brownian motion,

$$dz = \mu dt + \sigma dW_t.$$

This specification is attractive not only because of its tractability, but because it is consistent with two robust features of the data: (i) firm-level shocks are highly persistent, if not permanent; (ii) the firm size distribution follows a Pareto distribution. As we show below (and as is well known), the geometric Brownian motion dynamics generate a stationary distribution that is Pareto.

We also assume that all firms enter with the same productivity z_0. This simplification has little impact on our results since we do not focus our estimation on small firms (which is where the entrants start).

Employment n can be costlessly adjusted, and the wage is w. For simplicity, we assume that n is a continuous choice (i.e., we do not impose indivisibility). If n is greater than $\underline{n}$, a proportional tax τ applies to the wage rate and a fixed cost c_f has to be paid. We assume that the proportional tax applies to all employment, including that below $\underline{n}$, but this is without loss of generality, since we allow the fixed cost c_f to be negative (i.e., the tax could apply only to employment in excess of $\underline{n}$). The first time a firm crosses the threshold $\underline{n}$, it has to pay a sunk cost F. This cost captures the investment necessary to comply with the regulation, including the physical cost of buying an equipment, but also the informational costs such as learning about the regulation and perhaps consulting with lawyers or accountants. These informational costs may also reflect wasted managerial time.

The presence of the sunk cost makes this a dynamic optimization problem. Let $s \in \{0, 1\}$ denote whether a firm has already paid the sunk cost in the past. The state of the firm is summarized by (z, s).

3.2 Static subproblem

We first study the static problem, to determine the firm profit function which will enter the dynamic optimization.[1] To find the optimal labor demand and profit of the firm, we first solve the firm's problem conditional on operating below the threshold, then we find the solution conditional on operating above the

[1]This section thus does not depend on assumption that z is a Brownian motion.

threshold, and finally we find the overall solution by combining these results.

The current-period profit function for a firm which operates below the threshold is:

$$\pi^b(z) = \max_{0 \leq n < \underline{n}} \{e^z n^\alpha - wn\}. \tag{1}$$

The superscript b stands for "below the threshold". Optimal employment is given by:

$$n^b(z) \quad = \begin{cases} \left(\frac{\alpha}{w}\right)^{\frac{1}{1-\alpha}} e^{\frac{z}{1-\alpha}} & , \text{ if } z < \underline{z} \\ \underline{n}^- & , \text{ if } z \geq \underline{z}. \end{cases}$$

where $\underline{z} = \log\left(\underline{n}^{1-\alpha}\frac{w}{\alpha}\right)$ and $\underline{n}^-$ indicates a value just below $\underline{n}$. Profits are given by the formula

$$\begin{aligned} \pi^b(z) &= e^{\frac{z}{1-\alpha}} \left(\frac{\alpha}{w}\right)^{\frac{\alpha}{1-\alpha}} (1-\alpha), \text{ if } z < \underline{z} \\ &= e^z \underline{n}^\alpha - w\underline{n}, \text{ if } z \geq \underline{z}. \end{aligned}$$

The current-period profit function for a firm that decides to operate above the threshold, and hence to face the regulation, is:

$$\pi^a(z) = \max_{n \geq \underline{n}} \{e^z n^\alpha - w(1+\tau)n - c_f\}. \tag{2}$$

where the superscript a stands for "above the threshold". The firm operates above the threshold if z is greater than a cutoff value $\overline{z}$, defined as the solution to

$$e^{\frac{z}{1-\alpha}} \left(\frac{\alpha}{w(1+\tau)}\right)^{\frac{\alpha}{1-\alpha}} (1-\alpha) - c_f = e^z \underline{n}^\alpha - w\underline{n}.$$

It is easy to see that $\overline{z} > \underline{z}$, provided that there is a cost of operating above the threshold: $\tau\overline{n} + c_f > 0$. We will maintain this realistic assumption throughout the paper.

Summarizing, optimal employment if the firm decides to operate above the threshold is

$$n^a(z) \quad = \begin{cases} \underline{n}^+ & \text{if } z < \overline{z}, \\ \left(\frac{\alpha}{w(1+\tau)}\right)^{\frac{1}{1-\alpha}} e^{\frac{z}{1-\alpha}} & , \text{ if } z \geq \overline{z}, \end{cases} \tag{3}$$

This leads to profits

$$
\begin{aligned}
\pi^a(z) &= e^{\frac{z}{1-\alpha}} \left(\frac{\alpha}{w(1+\tau)}\right)^{\frac{\alpha}{1-\alpha}} (1-\alpha) - c_f, \text{ if } z \geq \overline{z}, \\
\pi^a(z) &= e^z \underline{n}^\alpha - w(1+\tau)\underline{n} - c_f, \text{ if } z < \overline{z}.
\end{aligned}
$$

Combining our results, we can now write the firm profit, as a function of the current productivity and state $s \in \{0,1\}$. Recall that $s = 0$ means that the firm has not paid the sunk cost and hence is forced to operate below the threshold, whereas a firm with $s = 1$ can choose to operate either below or above the threshold. Mathematically,

$$
\pi(z,0) = \pi^b(z),
$$

$$
\pi(z,1) = \max\left\{\pi^a(z), \pi^b(z)\right\}.
$$

We can obtain a formula for $\pi(z,1)$ by noting the following: (i) if $z < \underline{z}$, $\pi^b(z) > \pi^a(z)$, since the firm pays lower wages and fixed costs; (ii) for $z > \underline{z}$, the firm will decide to operate above the threshold; (iii) if $z \in (\underline{z},\overline{z})$, it is optimal to remain just below the threshold. Hence,

$$
\begin{aligned}
\pi(z,1) &= e^{\frac{z}{1-\alpha}} \left(\frac{\alpha}{w}\right)^{\frac{\alpha}{1-\alpha}} (1-\alpha) \text{ for } z < \underline{z}, \\
&= e^z \underline{n}^\alpha - w\underline{n} \text{ for } \underline{z} \leq z \leq \overline{z}, \\
&= e^{\frac{z}{1-\alpha}} \left(\frac{\alpha}{w(1+\tau)}\right)^{\frac{\alpha}{1-\alpha}} (1-\alpha) - c_f \text{ for } z > \overline{z}.
\end{aligned}
$$

For completeness, we also state the profit function:

$$
\begin{aligned}
\pi(z,0) &= e^{\frac{z}{1-\alpha}} \left(\frac{\alpha}{w}\right)^{\frac{\alpha}{1-\alpha}} (1-\alpha) \text{ for } z < \underline{z}, \\
&= e^z \underline{n}^\alpha - w\underline{n} \text{ for } z \geq \underline{z}.
\end{aligned}
$$

and the employment demand:

$$
\begin{aligned}
n(z,0) &= \left(\frac{\alpha}{w}\right)^{\frac{1}{1-\alpha}} e^{\frac{z}{1-\alpha}} \text{ for } z < \underline{z}, \\
&= \underline{n}^- \text{ for } z > \underline{z}.
\end{aligned}
$$

$$
\begin{aligned}
n(z,1) &= \left(\frac{\alpha}{w}\right)^{\frac{1}{1-\alpha}} e^{\frac{z}{1-\alpha}} \text{ for } z < \underline{z}, \\
&= \underline{n}^{-} \text{ for } \underline{z} \leq z \leq \overline{z}, \\
&= \left(\frac{\alpha}{w(1+\tau)}\right)^{\frac{1}{1-\alpha}} e^{\frac{z}{1-\alpha}} \text{ for } z \geq \overline{z}.
\end{aligned}
$$

To illustrate the logic, figure 7 plots profits as a function of employment for a high, medium and low value of productivity. The left column is the case of a proportional wage tax while the right column represents a fixed cost. This figure illustrates that firms with low productivity decide to operate below the threshold, since it is where their profits are highest. The high productivity firms operate above the threshold. The medium productivity firms operate exactly at (i.e. just below) the threshold. This figure also shows that the two types of per-period costs (fixed cost or wage tax) lead to the same implications for employment. Unless one uses data on productivity or profits, it is indeed impossible to distinguish the two. In our empirical work we focus on the case of a wage tax, because one provision of the law explicitly implies higher payroll taxes.

3.3 Dynamic optimization

Given the process for z, and the probability of exit λ, the firm's value maximization problem can be written formally as choosing a stopping time T to cross the threshold. Formally, for a firm that has productivity z today:

$$
V(z,0) = \sup_{T \geq 0} E\left[\int_0^T e^{-(r+\lambda)t}\pi(z_t,0)dt + \left(\int_T^{\infty} e^{-(r+\lambda)t}\pi(z_t,1)dt - Fe^{-(r+\lambda)T}\right)\right]. \tag{4}
$$

(Note that in writing this expression, we normalized the exit value to zero; since exit is exogenous, this is without loss of generality.) Intuitively, the firm will make the switch if its productivity becomes large enough; denote by z^* the cutoff that triggers the firm to pay the sunk cost. A standard option value argument implies that z^* will be greater than $\overline{z}$: given that the evolution of productivity z is uncertain, the firm will delay paying the sunk cost rather than invest as soon as it expects the investment to be just profitable in the present discount value sense.

This section presents the solution of the model using directly some results in Stokey (2008) for a general option exercise problem.[2] First, we rewrite the problem explicitly as choosing a cutoff z^*, given the current

[2] An alternative solution method, using the more intuitive Hamilton-Jacobi-Bellman equations and smooth pasting conditions, is presented in the appendix.

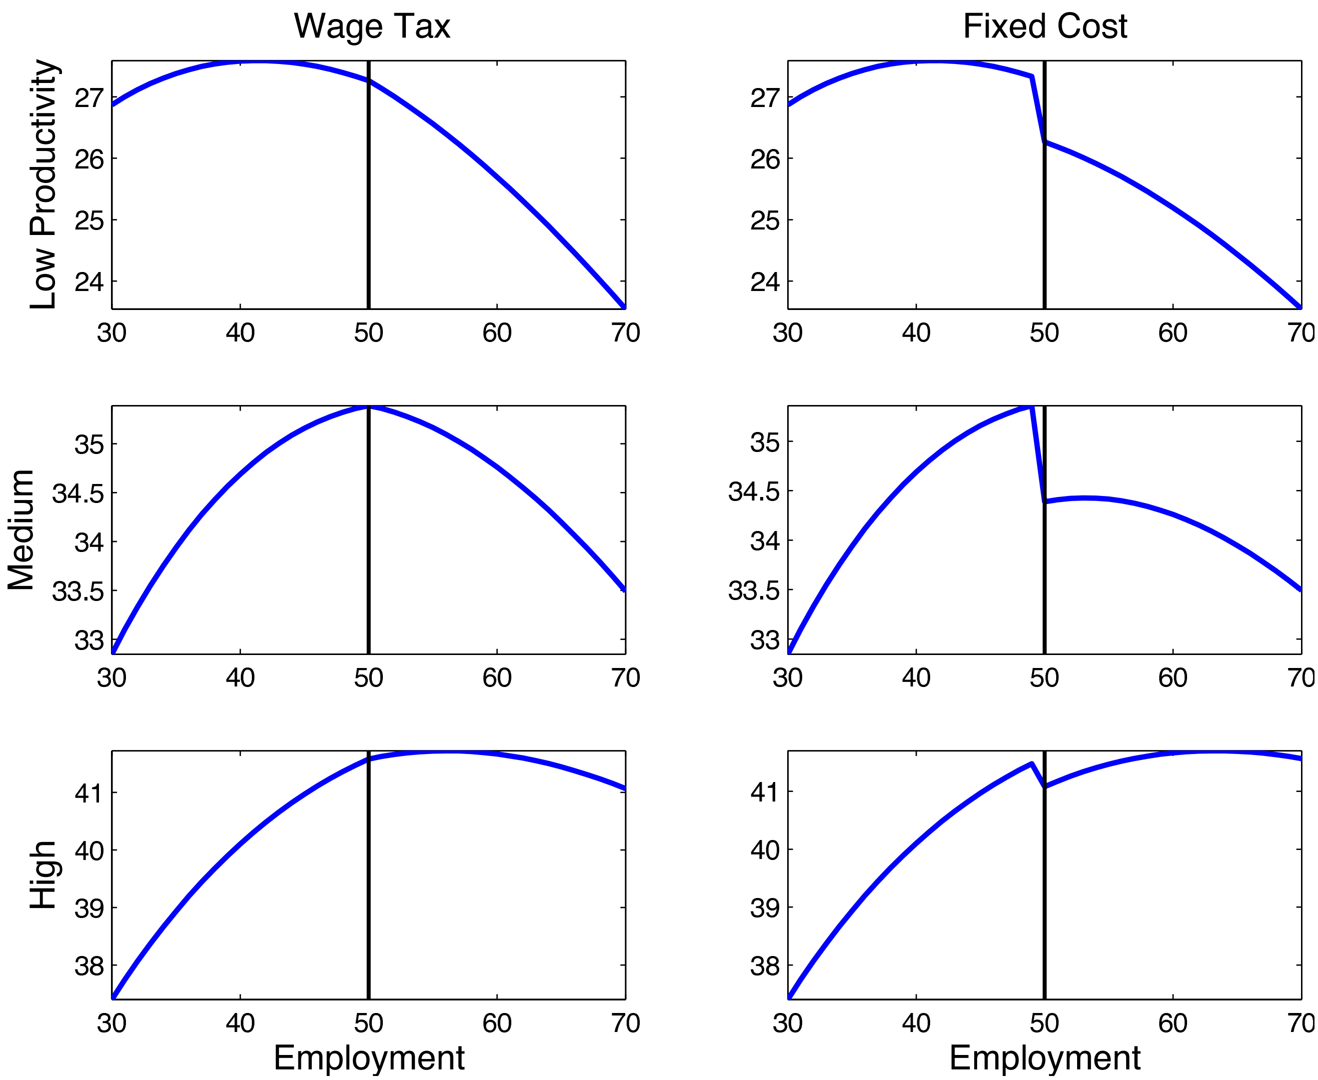

Figure 7: Profits, as a function of employment, for three different values of productivity (top, middle and bottom panels). Left panel: Wage Tax if operate with more than 50 employees; Right panel: Fixed cost if operate with more than 50 employees. This figure is a numerical illustration and not based on estimated parameters.

value z:

$$V(z,0)=\sup_{z^*\geq z} E_z\left[\int_0^{T(z^*)} e^{-(r+\lambda)t}\pi(z_t,0)dt+e^{-(r+\lambda)T(z^*)}\left(V(z^*,1)-F\right)\right], \tag{5}$$

with

$$V(z^*,1)\equiv E_{z^*}\left[\int_0^{\infty} e^{-(r+\lambda)t}\pi(z_t,1)dt\right],$$

and with R_1 and R_2 the roots of the quadratic $\frac{\sigma^2}{2}R^2+\mu R-(\lambda+r)=0$, i.e. with $J=\sqrt{\mu^2+2(r+\lambda)\sigma^2}$, we have $R_1=\frac{-\mu-J}{\sigma^2}<0$, and $R_2=\frac{-\mu+J}{\sigma^2}>0$.

The next proposition derives the optimal policy. In the language of Stokey (2008), R_1 discounts the time the process z will spend between $\overline{z}$ and z^*.

Proposition. *The solution to the firm problem (equation (5)) is z^*, the unique value satisfying:*

$$-R_1\int_{\overline{z}}^{z^*} e^{R_1(z^*-z)}\left[\pi^a(z)-\pi^b(z)\right]dz=(r+\lambda)F. \tag{6}$$

Proof. See appendix. □

For given structural parameters $\{\alpha,\overline{n},\mu,\sigma,\tau,c_f,F,r,\lambda\}$, this equation allows us to find z^* numerically easily. We conclude this subsection by noting some intuitive comparative statics: higher uncertainty, higher sunk costs, or higher fixed costs, all make it optimal to wait longer before crossing the threshold. This is the standard real option effect.

Corollary. *z^* is increasing in σ^2,F,τ_w,τ_f and $\underline{n}$.*

Proof. Differentiation of equation (6) gives the results. □

3.4 Stationary Distribution

Given our interest in the size distribution, we derive the joint cross-sectional distribution over (z,s) in closed form. Denote the probability density function as $f(z,s)$. Recall that firms enter with $z=z_0$, and z then evolves according to a Brownian motion with parameters (μ,σ). Firms switch from $s=0$ to $s=1$ as soon as z reaches z^*, and exit upon the realization of a Poisson process with parameter λ. We can write the Kolmogorov Forward equation, which reflects the conservation of the total number of firms, net of exit:

$$-\mu\frac{\partial f(z,0)}{\partial z}+\frac{\sigma^2}{2}\frac{\partial^2 f(z,0)}{\partial z^2}=\lambda f(z,0), \tag{7}$$

which holds for all $z < z_0$ and all $z \in (z_0, z^*)$. (See Dixit and Pindyck (1994), appendix of chapter 3, for a heuristic derivation, and chapter 8 for an application similar to our case.) The equation needs not hold for $z = z_0$, since there is entry of new firms.

The same equation applies to firms which have made the switch:

$$-\mu \frac{\partial f(z,1)}{\partial z} + \frac{\sigma^2}{2} \frac{\partial^2 f(z,1)}{\partial z^2} = \lambda f(z,1), \tag{8}$$

which holds for all $z \in (-\infty, z^*)$ and for all $z \in (z^*, +\infty)$.

Last, we need to state the boundary conditions. The first one is simply the requirement that f is a density, i.e.

$$\int_{-\infty}^{+\infty} f(s,1)ds + \int_{-\infty}^{+\infty} f(s,0)ds = 1.$$

To derive the other boundary conditions, the easiest approach is to approximate the Brownian motion with a discrete random walk, as in Dixit and Pindyck (1994). This yields the conditions

$$f(z^*, 0) = 0,$$

and $f(.,0)$ must be continuous at z_0, while $f(.,1)$ must be continuous at z^* :

$$\lim_{s \to z_0^-} f(s,0) = \lim_{s \to z_0^+} f(s,0),$$

$$\lim_{s \to z_-^*} f(s,1) = \lim_{s \to z_+^*} f(s,1).$$

Finally, a balance condition holds for $z = z^*$, reflecting that the number of firms which reach z^* and have $s = 0$ is equal to the number of firms which enter at $s = 1$ with $z = z^*$, and is equal to the number of firms with $s = 1$ which exit in any time period: this leads to

$$-\frac{\sigma^2}{2} f'(z^*, 0) = \lambda \int_{-\infty}^{\infty} f(s,1)ds.$$

Given these boundary equations, solving for the cross-sectional distribution involves some simple algebra, which is relegated to the appendix. The result is:

$$
\begin{aligned}
f(z,0) &= \frac{\beta_1\beta_2}{\beta_1-\beta_2}\left(e^{\beta_2(z-z_0)} - e^{\beta_1(\hat{z}-z_0)}e^{\beta_2(z-\hat{z})}\right), \text{ for } z < z_0, \\
&= \frac{\beta_1\beta_2}{\beta_1-\beta_2}\left(e^{\beta_1(z-z_0)} - e^{\beta_1(\hat{z}-z_0)}e^{\beta_2(z-\hat{z})}\right), \text{ for } z^* > z > z_0,
\end{aligned}
$$

and

$$
\begin{aligned}
f(z,1) &= \frac{\beta_1\beta_2}{\beta_1-\beta_2}e^{\beta_1(\hat{z}-z_0)}e^{\beta_2(z-\hat{z})}, \text{ for } z < z^*, \\
&= \frac{\beta_1\beta_2}{\beta_1-\beta_2}e^{\beta_1(z-z_0)}, \text{ for } z > z^*.
\end{aligned}
$$

This expression implies that z has an exponential distribution in the upper tail. Since log employment and log sales are both proportional to z, employment and sales follow Pareto distributions, and the c.d.f. of employment is proportional to n to the power $\beta_1(1-\alpha)$.[3]

Figure 8 below illustrates some properties of our model. (This figure is drawn using our parameter estimates.) For now focus on the left column. The second panel from the top shows the distribution in the absence of regulation - it is is Pareto. The bottom panel depicts the distribution with a per-period wage tax (here too, the results would be similar with a fixed cost). There is a substantial "hole" in the distribution with no firms whatsoever between 50 and 54 employees. This figure presents an empirical challenge, because in the data there are many firms with an employment level slightly greater than 49. It would be incredible to attribute the presence of all these firms to measurement error. Last, the third row shows the impact of a sunk cost on the firm size distribution. The sunk cost model does not suffer from the same deficiency as the fixed cost model: there are no holes in the distribution, and in particular some firms have exactly 50 employees. These are firms that crossed the threshold in the past and that were subsequently hit by negative productivity shocks. Finally, the right hand column adds some classical measurement error to employment, which obviously helps smoothing out the distribution.

To establish the economic relevance of these regulations, we now turn back to the data and propose a simple structural estimation of our model.

[3]Note that this implies some restrictions on β_1 to ensure that employment be finite. This in turn restricts the parameters μ, λ, σ^2. Our estimated parameters satisfy these restrictions, so we do not need to impose them in practice.

4 Estimation

This section proposes a simple estimation of our model using indirect inference. We take advantage of our closed form solutions which make calculating model moments computationally easy.

As discussed below, we incorporate classical measurement error in (log) employment; the standard deviation of measurement error is σ_{mrn}. Table 3 lists our parameters. The full set of structural parameters is the vector $\theta = (r, w, \alpha, z_0, \lambda, \mu, \sigma, \tau, F, \sigma_{\text{mrn}})$. We partition this vector into two vectors, i.e. $\theta = (\theta_p, \theta_e)$ where $\theta_p = (r, w, \alpha, z_0)$ includes parameters that are set a priori, and $\theta_e = (\lambda, \mu, \sigma, \tau, F, \sigma_{\text{mrn}})$ is the vector of estimated parameters.

Like calibration, indirect inference works by selecting a set of statistics of interest, which the model is asked to reproduce.[4] These statistics are called sample auxiliary parameters $\hat{\Psi}$ (or target moments). For an arbitrary value of θ_e, we use the structural model to generate S statistically independent simulated data set and compute simulated auxiliary parameters $\Psi^s(\theta_e)$. The parameter estimate $\hat{\theta}_e$ is then derived by searching over the parameter space to find the parameter vector which minimizes the criterion function:

$$\hat{\theta}_e = \arg \min_{\theta_e \in \Theta_e} \left(\hat{\Psi} - \frac{1}{S} \Psi^s(\theta_e) \right)' W \left(\hat{\Psi} - \frac{1}{S} \Psi^s(\theta_e) \right)$$

where W is a weighting matrix and Θ_e the estimated parameters space. We choose the identity matrix as using the optimal weighting matrix is fraught with a well-known small sample bias problem (Altonji and Segal (1996)). This procedure generates a consistent estimate of θ_e. The minimization is performed using Nelder-Mead simplex algorithm. We used different starting values to find the global minima. To simulate the model, we draw from the stationary distribution derived in the previous section.

The standard errors are obtained using 500 bootstrap repetitions. In each bootstrap repetition, a new set of data auxiliary parameters is produced using Block-Bootstrap (Hall and Horowitz (1996)). An estimator $\hat{\theta}_e^b$ is found by minimizing the weighted distance between the recentered bootstrap auxiliary parameters and the recentered simulated auxiliary parameters.

4.1 Predefined Parameters

Some parameters are not estimated because they are either normalization or are fairly standard. We set the real interest rate r to 5 percent. The wage rate is normalized to 1. We assume that α equals 0.66, as in Cooper et al. (2007). This parameter is a reduced form for the labor share, decreasing returns to scale and

[4] See Gourieroux et al. (1993) for a general discussion of indirect inference.

Parameters	Definition	
r	interest rate	fixed
α	curvature profit function	fixed
z_0	Entry TFP level	fixed
w	wages	normalized
λ	death probability	estimated
μ	drift	estimated
σ	std dev shocks	estimated
τ	proportional tax on wages	estimated
F	sunk cost	estimated
σ_{mrn}	measurement error	estimated

Table 3: Economic Parameters

the elasticity of demand.[5] Finally, the parameter z_0 is irrelevant for the statistics that we target given the Pareto distribution implied by our model.

4.2 Measurement Error

There is likely to be some measurement error in our employment variable, which is the arithmetic average of the number of employees at the end of each quarter. Further, it is the relevant measure of employment for some but not all of the regulations. For instance, some regulations are based on employment measured in full-time equivalent and some other regulations apply if there is more than 50 employees in the firm for more than 12 months. Finally, measurement error also more broadly captures time aggregation problems as well as adjustment cost or search frictions which lead to an imperfect control of the size of the workforce. However, since our data is based on administrative sources, it has a relatively high quality, and we think measurement error is limited.

We explicitly introduce measurement error into the simulated moments to mimic the bias these impute into the actual data moments. We do so by multiplying employment by mrn_{it} that is i.i.d over firm and time and follow a log-normal distribution with mean $-\frac{1}{2}\sigma_{\text{mrn}}^2$ and standard deviation σ_{mrn}.[6]

4.3 Auxiliary Parameters and Identification

Table 5 lists the auxiliary parameters (target moments) that we picked. We set to match bins of the distribution of employment around the thresholds, the average and volatility of growth in employment, the

[5]There is little agreement on this parameter. We have experimented with different values. While the precise value of α does matter for some of our parameter estimates, it has a small effect on our policy experiments. This is because a higher α implies a lower volatility of fundamental shocks to fit the size distribution, and these also have offsetting effects on the benefits to reallocation.

[6]We also estimated the model with measurement error represented as the difference of two Poisson distributed random variables and we have found very similar results to those we report here.

	Sunk Cost Model		Proportional Tax Model	
λ	0.0499	(0.0016)	0.0205	(0.0003)
μ	0.0029	(0.0002)	-0.0035	(0.0001)
σ	0.0842	(0.0015)	0.0751	(0.0008)
σ_{mrn}	0.0107	(0.0005)	0.0377	(0.0010)
F	1.0810	(0.0672)	0	
τ	0		0.0026	(0.0001)

Table 4: Parameter Estimates (Standard Errors in Parentheses)

slope of the power law. The rationale for the bins is that we want to reproduce well the firm size distribution jump that is the visual evidence that the regulation matters. The rational for the last three moments is that we want the model to be consistent with key features of firm dynamics.

Identification of the models's parameters is achieved by a combination of functional form and distributional assumptions, and is difficult to prove, but the intuition is straightforward. Heuristically, the median net employment growth is informative about the drift μ. The bins between 40 and 60 are informative regarding the frictions parameters τ and/or F and the variance of measurement error σ_{mrn}. The variance of employment growth is informative about the variance of productivity shocks σ and the variance of measurement error σ_{mrn}. The slope of the power law is informative regarding the variance of productivity shocks σ, the drift μ and the exit rate λ.

4.4 Estimation Results

We do not attempt to estimate a model with both sunk cost and wage tax since identification is delicate. Rather we compare the model with sunk cost and the model with per-period wage tax.

Table 4 reports the structural parameters estimates together with estimated standard errors. The first column estimates the sunk cost model with $\tau = 0$ and the second column estimates the proportional tax model with $F = 0$. In the appendix, we present and discuss the estimation results when splitting our sample into four sectors (manufacturing, retail, construction and services).

The data are consistent with a regulation that acts like a sunk cost of about one year of a worker wages or a small proportional tax on wages of 0.26%. The estimates for the drift and the variance are not very sensitive to the specification of the regulation. Shocks to total factor productivity are estimated to be 8% per year, which is in line with standard estimates.

Measurement error is much larger for the model with a proportional tax than for the model with a sunk cost. The main reason is that the model with a proportional tax implies that firms do not want to operate on the right side of the threshold. Even with a small estimated tax, the model predicts that when it is

	Data	Sunk Cost Model	Proportional Tax Model
# firms			
40-45	0.3565	0.3392	0.3395
45-50	0.3883	0.3852	0.3847
50-55	0.1413	0.1517	0.1504
55-60	0.1140	0.1238	0.1253
Median $\Delta \log n$	0	0.0054	-0.0096
$V(\Delta \log n)$	0.0485	0.0607	0.0516
Power Law	1.1417	1.1471	1.1503

Table 5: Auxiliary Parameters

	Data	Proportional Tax	Sunk Cost
Fraction of firms			
>200	0.4711	0.4505	0.4515
>500	0.1615	0.1570	0.1578
>1000	0.0674	0.0707	0.0713
>5000	0.0059	0.0111	0.0112

Table 6: Firm Size Distribution - share of firms among firms with more than 100 employees, in the data, in the model with tax cost, and in the model with sunk cost.

optimal to cross the threshold, firms operate at a size of at least 55 employees. Without measurement error, there are no firms with an employment level between 50 and 55. This is illustrated in the bottom left panel of Figure 8. The model with a sunk cost does not suffer from this feature. Some firms operate naturally to the right of the threshold: firms that have crossed the threshold in the past and that were subsequently hit by negative productivity shocks. As a result, the amount of measurement error is much lower in the sunk model, with an estimated standard deviation of 1% compared to the model with a proportional tax that implies a measurement error close to 4%.

The exit rate is estimated to be 5% in the sunk cost model and 2% in the proportional tax model. This is low, and reflects that our model do not allow for endogenous exit. Hence the mode underpredicts the exit rate of small firms and is better description of the exit behavior of larger firms.

Table 5 reports the fit of the models. They are both able to fit the moments well, and in particular the jump in the bins distribution between 45-49 and 50-54. Figure 9 reproduces the log frequency - log size regression using simulated data from the model with a sunk cost. Both model reproduces the discontinuity at 50 and the change in the constant of the power law distribution observed in the data.

We finally examine the ability of the model to account for the large firms' size distribution. Table 6 reports the fraction of firms above 200, 500, 1000 and 10000 employees normalized by the fraction of firm of more than 100 employees. Although these moments are not directly targeted in the estimation, both models do a reasonable job.

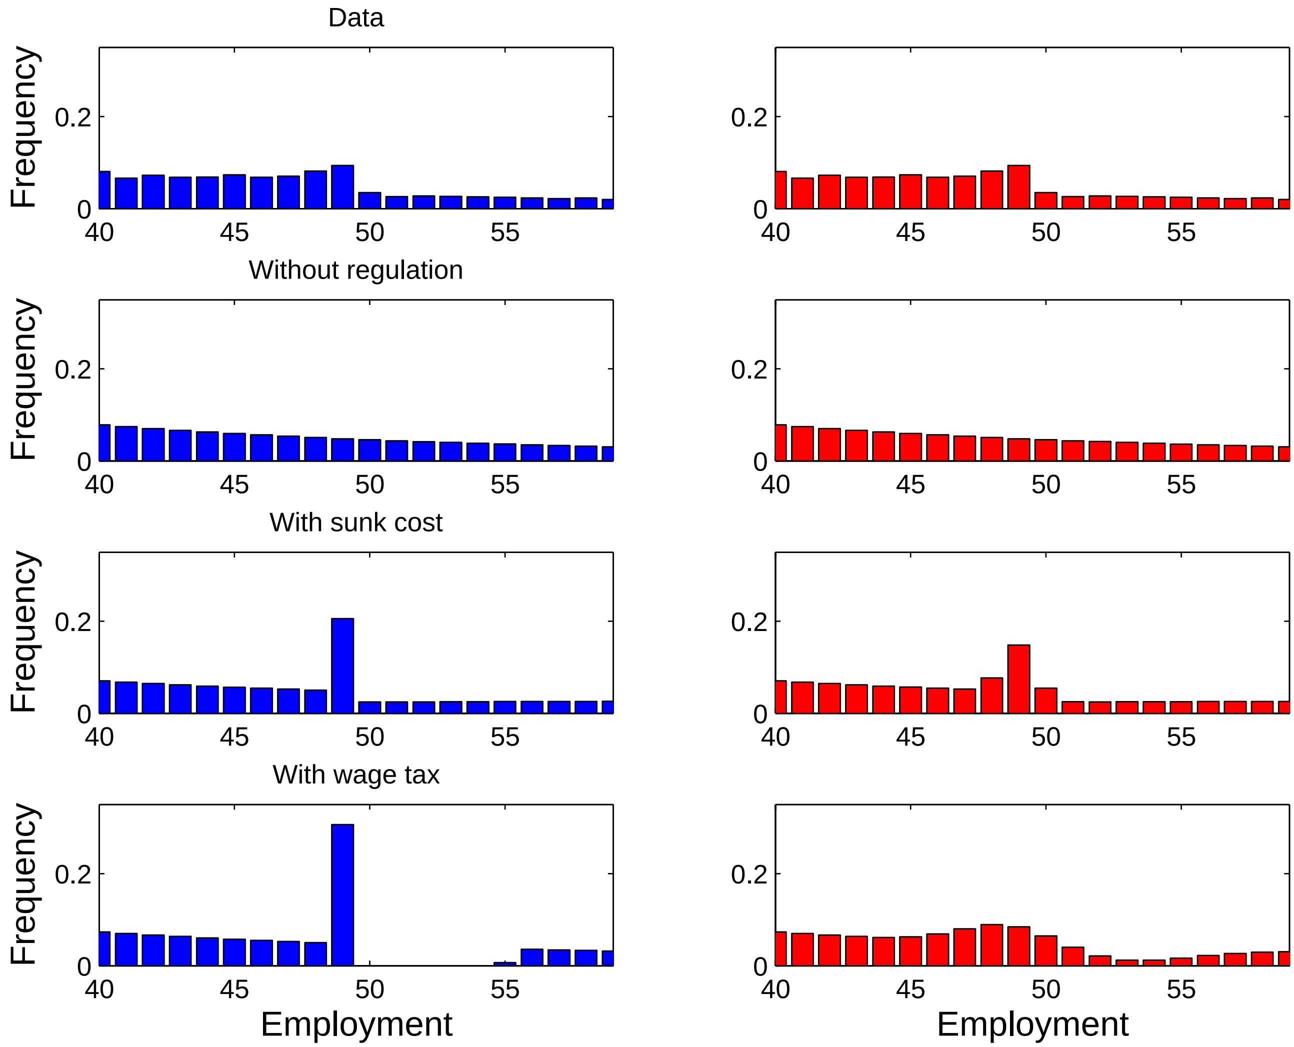

Figure 8: Distribution of firm employment (between 40 and 59 employees), in the data and in the model. The distribution is normalized by the total number of firms between 40 and 59 employees.

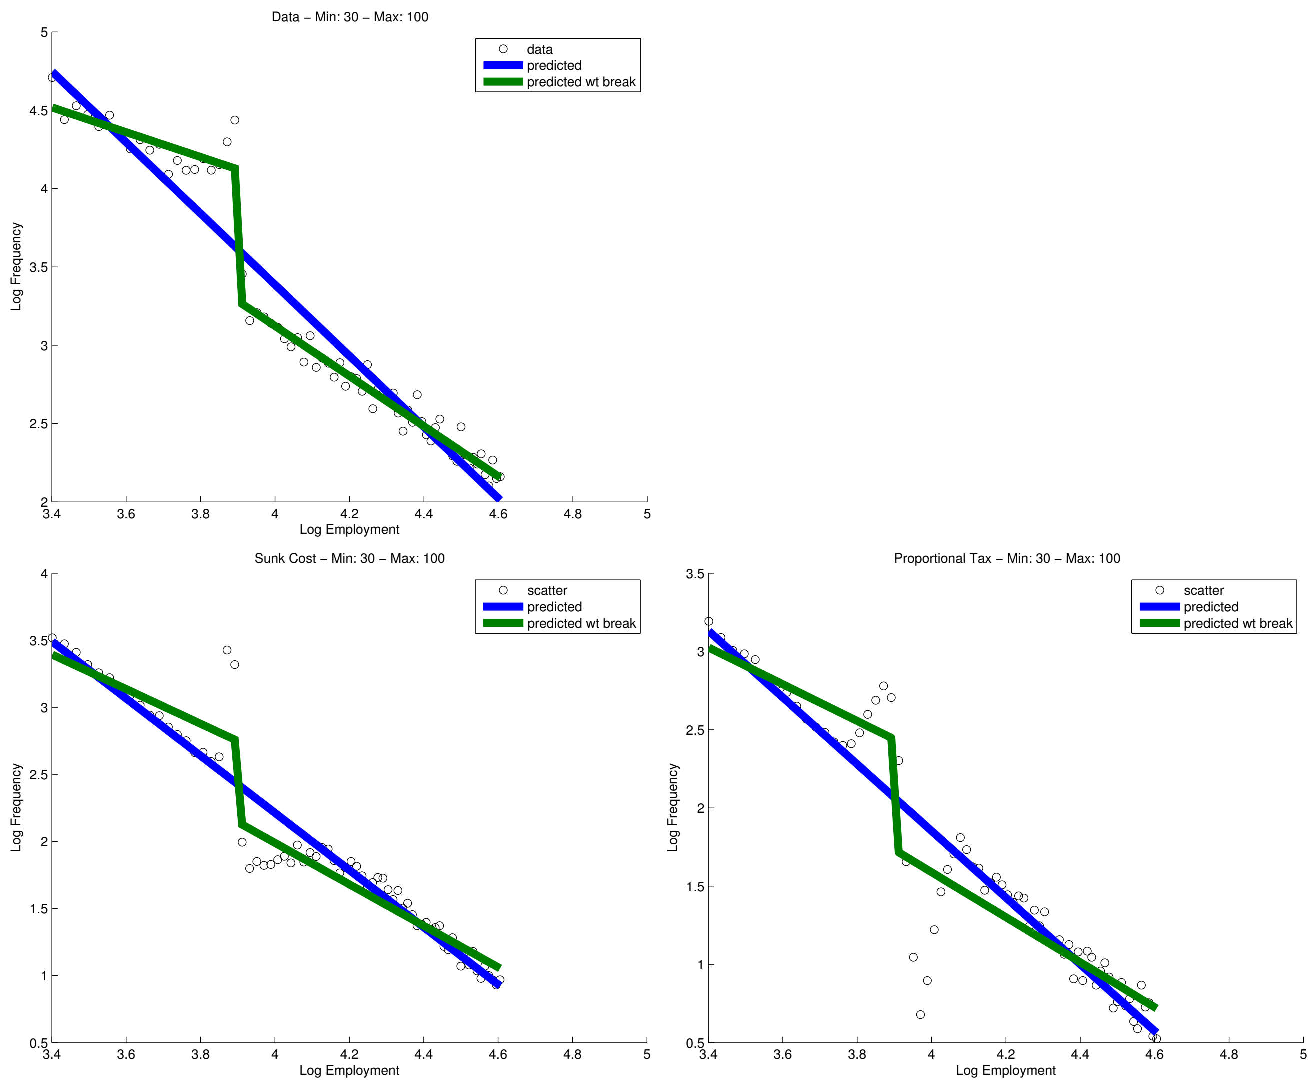

Figure 9: Broken Power Law: Regression of the logged number of firms on the logged number of employees with and without a structural break at 50 for firms with employment level between 30 and 100 using (1) the data, (3) simulated data from the sunk cost model and (3) simulated data from the proportional tax model.

We close by asking, is there any additional means of distinguishing the sunk cost and wage tax model? First, as figure 8 shows, the wage tax model implies that the true distribution of employment (without measurement error) has a very large spike at 49, about 50% larger than the sunk cost model, and in addition there are, in reality, no firms with employment just above 50. With measurement error, the model also mechanically underpredicts the distribution to the right of the threshold since the spike is smoothed out by measurement error. We are somewhat skeptical that measurement error is so large in our data which is based on administrative sources.[7]

A more direct test of the sunk cost model is to calculate the firm size distribution, conditional on having been above 50 in the past.[8] In the sunk cost model, this conditional distribution should not exhibit a spike at 49, so this is a stringent test of our theory. First, figure 10 compares the conditional and unconditional distribution in the data. While there is still a spike at 50, its size is dramatically reduced by conditioning. Whereas in the unconditional distribution, there are 2.7 times more firms with 49 employees than with 50, in the conditional distribution, this ratio is only 1.5. To go further and see how measurement error affects this statistic, figure 11 presents the results for the data and for the models at estimated parameter values. The second panel on the left confirms that without measurement error, there is no spike altogether. With our small measurement error, there is a small spike, similar to the data. The intuition is that the conditioning (on employment being greater than 50 in the past) is slightly noisy since employment is not perfectly measured; hence we capture some firms that are still below the threshold, and remain there. In contrast, the wage tax model produces a large spike that is smoothed out by measurement error. Overall, the sunk cost model appears closer to the data.

5 Policy Experiments

In the previous section, we estimated the regulatory cost as perceived by firms. In this section, we use our estimates to infer the aggregate effect of the regulation on output, employment and productivity.

From the point of view of a social planner, the regulation misallocates labor across firms and hence reduces total factor productivity. Moreover, the regulation affects the incentives of firms to enter. To demonstrate this, we consider three experiments, which differ in the set of equilibrium feedback that they allow. We first discuss the conceptual framework for our experiments, then we present and discuss the results.

[7] Garicano et al. (2012) using a similar model also estimate large measurement error.

[8] We thank Theodore Papageorgiou for this suggestion.

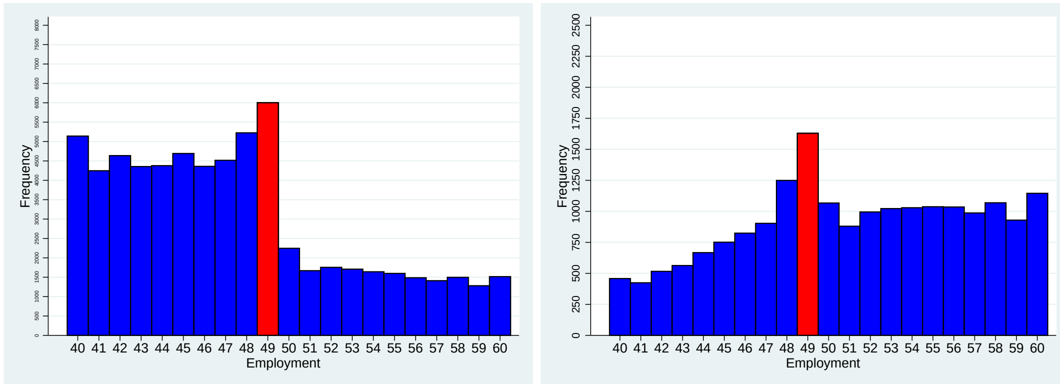

Figure 10: Distribution of firm employment (between 40 and 59 employees). Left Panel: unconditional. Right Panel: conditional on having had more than 50 employees in the past.

5.1 Three Experiments

For the purpose of estimation, we do not need to take a stand on the determination of the number of firms or of market prices: given the observed number of firms and factor prices, we use cross-sectional information to identify our parameters. However, for our policy experiments, it can matter whether the number of firms, or prices adjust in response to a change in the regulation. Our three experiments differ in their assumptions about these equilibrium feedbacks.

Our first experiment abstracts from all feedbacks and considers the effect of removing entirely the regulation, holding the wage, the interest rate and the number of firms fixed. Concretely, we first solve the firm problem with our estimated regulation, and obtain firm's optimal decisions, $n(z;w,\theta)$, $y(z;w,\theta)$ and $z^*(w,\theta)$ where for clarity we now index all policy functions by the wage w as well as θ, the vector of parameters (which includes the regulation). We then calculate aggregate employment and output using the cross-sectional distribution $f(z;\theta)$:

$$N(w,\theta) = \int_{-\infty}^{\infty} n(z;w,\theta) f(z;\theta) dz,$$

$$Y(w,\theta) = \int_{-\infty}^{\infty} y(z;w,\theta) f(z;\theta) dz.$$

We can compute the policy rules, cross-sectional distributions, and aggregate labor and output for different regulations, corresponding to different vectors θ; in particular if we remove the regulation entirely, we set $\theta = \theta_0$ and employment and output are $N(w,\theta_0)$ and $Y(w,\theta_0)$. We then calculate the percentage change in N and Y from θ to θ_0. This experiment implicitly assumes a perfectly elastic supply of labor, and an

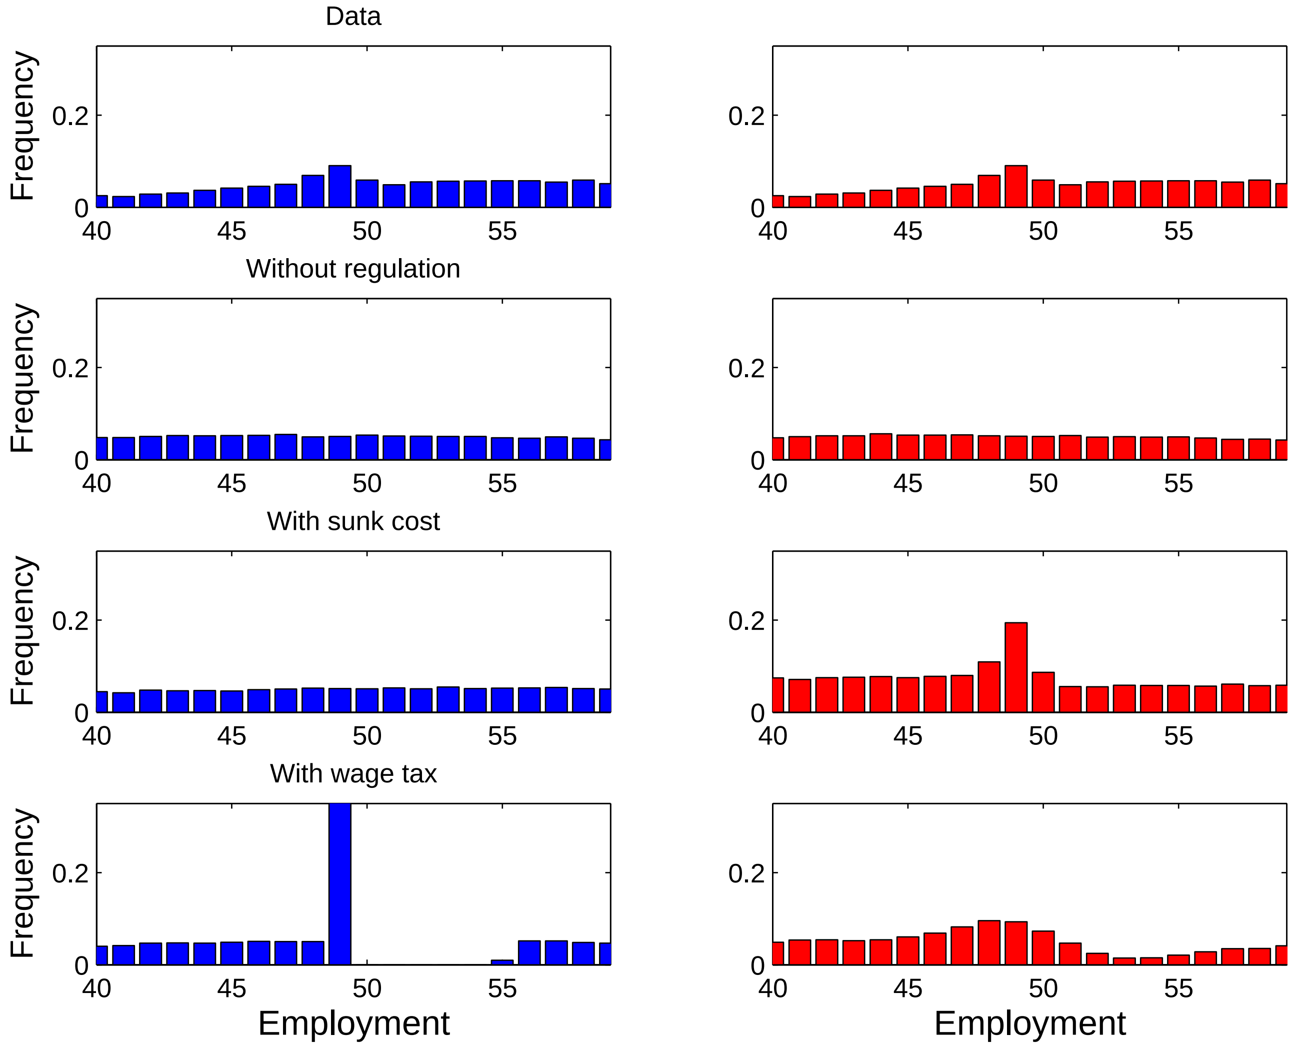

Figure 11: Distribution of firm employment (between 40 and 59 employees), conditional on having had more than 50 employees in the past, in the data and in the model. The distribution is normalized by the total number of firms between 40 and 59 employees.

inelastic supply of firms. We also discuss briefly the effect of applying the regulation to all firms, and to move the threshold to 75 employees rather than 50 employees; these all correspond to different vectors θ.

Our second experiment asks, how much of an increase in output can we obtain, holding total employment constant, simply by reallocating labor across firms. That is, suppose we solve the allocation problem:

$$Y(N;\theta) = \max_{\{n(z)\}_{z=-\infty}^{\infty}} \int_{-\infty}^{\infty} e^z n(z)^\alpha f(z;\theta) dz$$
$$s.t. \quad : \quad \int_{-\infty}^{\infty} n(z) f(z;\theta) dz \leq N,$$

where $n(z)$ is the employment of firms with productivity z. This is the solution of the aggregate production function (expressing maximum possible aggregate output as a function of available aggregate labor). We report the percentage change in $Y(N;\theta)$ as θ is varied. An alternative way to think about this experiment is that labor supply is fully inelastic and the wage adjusts to keep the same total employment.

Our third experiment adds endogenous entry and labor supply to the model by embedding our firm dynamics in a general equilibrium framework as in Hopenhayn and Rogerson (1993). Since this model is well know, we describe it only briefly here. First, there is a representative agent with utility function

$$\int_0^{\infty} e^{-\rho t} \left(\log(C_t) - B \frac{N_t^{1+\phi}}{1+\phi} \right) dt.$$

This agent supplies work to the market, at the wage w_t, and buys or sells assets at the interest rate r_t. In equilibrium, the only assets are the firms. We consider a steady-state stationary equilibrium: there is no aggregate variation, since a law of large numbers apply, and macroeconomic aggregates are constant. As a result, the interest rate is constant, $r_t = r = \rho$.

For a given wage, we can solve the value function $V(z, s; w, \theta)$, policy functions $n(z, s; w, \theta)$ and $z^*(w, \theta)$, and stationary distribution $f(z, s; w, \theta)$ as in section 3. We have added the wage as an explicit argument to these functions to emphasize the dependence. Since all firms enter with a productivity z_0, the free entry condition reads,

$$k = V(z_0, 0; w, \theta). \tag{9}$$

Denote the flow of firms entering per unit of time E, and denote $Mf(z, s; w)$ the stationary distribution of firms. With exogenous exit at rate λ, the flow of entrants per unit of time E must equal λM in a stationary equilibrium.

Total output is then given by

$$Y(w;\theta) = M \int_{-\infty}^{\infty} e^{z} n(z;w,\theta)^{\alpha} f(z;\theta) dz, \tag{10}$$

and total labor is

$$N(w;\theta) = M \int_{-\infty}^{\infty} n(z;w,\theta) f(z;\theta) dz. \tag{11}$$

Labor supply satisfies the first order condition

$$B.C_t.N_t^{\phi} = w_t, \tag{12}$$

and the goods market constraint is

$$C_t + E_t k = Y_t. \tag{13}$$

A stationary equilibrium is then given by $\{Y, C, E, M, w, N\}$ such that $E = \lambda M$ and the equations (9)-(13) are satisfied. In this model, the free entry condition pins down the equilibrium wage. Given this wage, the number of firms adjusts the scale of the economy so that labor demand equals labor supply; that is, there is a perfectly elastic supply of firms.

This third experiment is itself divided in two different cases, labeled (a) and (b): we first consider the case of perfectly inelastic labor supply ($B = 0$ and $N = \overline{N}$), and then we consider the case of an elastic labor supply ($B > 0$).

We close by mentioning three issues that affect all experiments. First, we need to take a stand on whether the regulation cost is a real resource cost (that must be deducted from the resource constraint) or is a transfer (which is rebated lump-sum to households). In reality it is likely that both components are present. Hence we will present the results for the two possible assumptions. Second, the calculations above focus on steady-state effects and abstract from transitional dynamics. We believe this is appropriate to examine the long-run productive effects of the regulation, but of course this makes the welfare comparison inaccurate. Last, our calculations have little to say on the desirability of the regulations themselves since we do not model the benefits of the regulation.

5.2 Results

Table 7 and 8 present the results of our three experiments for the sunk model and for the proportional tax model. In all cases, we adjust z_0 so that the wage is one, as assumed in our estimation. For experiment 3a, we

Experiment	Y	N	w	M	C
1:Partial Equilibrium	0.866	0.885			
2:Pure Reallocation	0.288				
3a:GE, Inelastic Labor	-0.009		0.0097	-0.849	0.046
3b:GE, Elastic Labor	-0.027	-0.018	0.0097	-0.867	0.028

Table 7: Policy Experiments for the model estimated with sunk cost

Experiment	Y	N	w	M	C
1:Partial Equilibrium	0.337	0.510			
2:Pure Reallocation	0.099				
3a:GE, Inelastic Labor	-0.162		0.011	-0.477	-0.149
3b:GE, Elastic Labor	-0.082	0.080	0.011	-0.397	-0.069

Table 8: Policy Experiments for the model estimated with wage tax

calibrate the entry cost k to replicate the average firm size (7.5 employees per firm). For experiment 3b, we further need to calibrate labor supply preferences. We set an elasticity of labor $\phi = 1$ (see Chetty (2012) for a discussion), and B such that total employment is 0.25. These are standard values in the macroeconomics literature.

The first experiment shows that removing the sunk cost regulation leads to a significant increase in output and employment, close to one percent, as many medium-sized firms grow by going over the threshold and hence increase labor demand. Average labor productivity falls sightly as many firms that were previously constrained in their employment are now able to increase it. Interestingly, and this holds for all our experiments, the output (and employment) gains from extending the threshold to 75 employees rather than 50 are only 0.12 percent (not reported in the table), much smaller than the gains from entirely eliminating the threshold.[9]

Our second experiment shows what happens if we force total employment to remain constant. This is equivalent to taking the results of the first experiment and increasing the wage to make employment return to its initial value. In this case, the output gain is more modest. Very large firms and very small firms contract because of the higher wage. But intermediate firms grow as they now go over the threshold. The productivity gains from the reform are significant. We note that this result goes some way towards addressing the observation that France has relatively less medium-sized firms than comparable countries (See Bartelsman et al. (2009) or Bartelsman et al. (2013) among others).

Our third experiment, that adds endogenous entry, yields quite different results. Allowing the number of firms to adjust reduces dramatically the steady-state output gains. Since firms close to the threshold

[9] Of course, if the threshold is pushed sufficiently high, the gains converge to those obtained by fully eliminating the thresholds; but this convergence is slow.

can grow, the economy needs fewer firms, which economizes on entry costs. Overall, output actually falls slightly in the new steady-state, but the reduced entry costs imply that consumption rises.[10] If labor supply is elastic, the wealth gains from removing the threshold further reduce labor supply and output. However, this effect is fairly small. This points to another of our results: whether we model the regulation as a tax or as a real resource cost has very little effect on these experiments. For instance, if in experiment 3b the regulation was a transfer instead of a real resource cost, the decline of output would be -0.024 percent instead of -0.027 percent, and employment would contract by -0.022 percent instead of -0.018 percent (unreported in the table).

The same kind of intuition for the policy experiments applies if the regulation is a wage tax rather than a sunk cost. The effects are smaller because we estimate a fairly small tax. The main difference is that the wage tax directly affects the demand for labor, leading to a relatively larger decline of employment.

Finally, the motivation for the phase-in of the regulation at 50 employees is that it is too costly to impose the compliance cost on small firms. We can evaluate this argument by considering the counterfactual, what would happen if all firms were subject to the regulation? With free entry, this would have dramatic effects on the number of firms. For instance, in experiment 3b, the effect of imposing the sunk cost on everyone is to reduce output by 3.80 percent, with the number of firms declining by a whooping 11.45 percent. It is safe to say, then, that applying the regulation to all firms would be quite costly, which suggests that the phase-in is perhaps not such a bad policy.[11]

One criticism of these experiments is that the free entry assumption is too extreme. In this spirit, figure 12 presents the results where we vary the elasticity of supply of firms. To do so, we extend this model by relaxing the assumption that entry is perfectly elastic at cost k. To generate an upward-sloping supply of entrants to the economy, we suppose that in each period there is a pool N of potential entrants, which differ in their entry cost. The entry cost is distributed according to the cumulative distribution function H. In a given period, only potential entrants with an entry cost below $V(z_0, 0; w)$ will enter. Denote k^* the threshold value for k. The flow of entrants E will equal $NH(k^*)$ and the free entry condition is $V(z_0, 0; w) = k^*$.

We parametrize the c.d.f. H as a log-normal distribution with standard deviation σ_v. This parameter captures the heterogeneity of entry costs and hence the (inverse) elasticity of supply of entrants. For each value of σ_v, we recalibrate the model and run the policy experiments. Figure 12 shows that as we reduce σ_v, the results approach experiment 3, where entry is perfectly elastic: there is a large decline in the number of

[10] Our results echo those of Jaef (2012), who shows that incorporating entry and exit reduces the gains to reallocation.

[11] The study of Garicano et al. (2012) obtains large welfare gains in part because they estimate a large tax and large measurement error, and in part because of different assumptions about the equilibrium feedback.

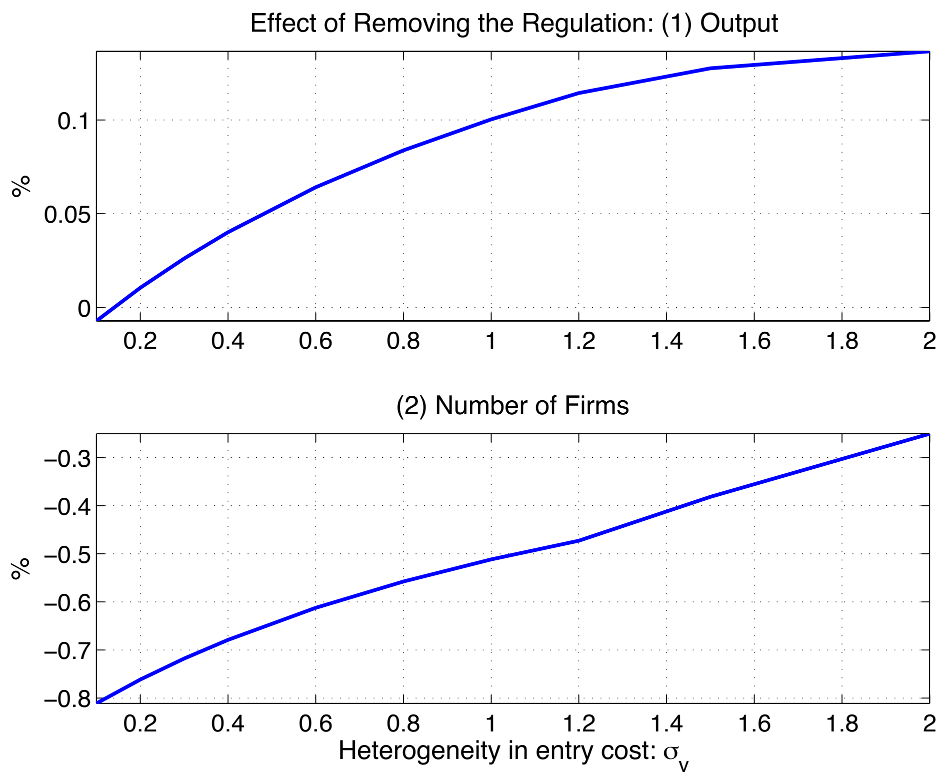

Figure 12: Comparative Statics of Policy Experiments: change in the number of firms and in total output when the sunk cost is removed, as a function of the standard deviation of entry costs.

firms M and a smaller increase, or even a decrease, in output Y. As we increase heterogeneity in entry costs σ_v and hence reduce the elasticity of firms, we see a smaller reaction in the number of firms and a larger increase in output. It is however difficult to pin down a realistic value for σ_v from cross-section data alone.

6 Conclusion

Our paper studies a particular regulation which clearly distorts the firm size distribution, leading to an obvious misallocation of labor - a channel that has been emphasized in the recent literature. Our results provide a "case study" that is complementary to broader macro approaches (Hopenhayn and Rogerson (1993), Restuccia and Rogerson (2008) and Buera et al. (2011)).

We obtain plausible estimates of the costs of the regulation and find that their aggregate effects are significant if firm entry is inelastic. However these effects are limited if firm entry is elastic enough.

There are several interesting extensions. First, incorporating labor adjustment costs or search frictions would be useful to take into account the imperfect, and costly control over the size of the labor force. Second, introducing in the model other inputs, such as capital, would generate some factor substitution close to the threshold: if firms do not want to increase employment, they may react by using other factors. Finally, given the limitations of our data for small firms, we have abstracted from the existence of other thresholds (at 10 and 20 employees), but incorporating them would be useful to quantify the total effect of these regulations on a firm's life-cycle.

References

Altonji, J. and Segal, L. (1996). Small-Sample Bias in GMM Estimation of Covariances Structures. *Journal of Business & Economic Statistics*, 14(3):353–366.

Axtell, R. (2001). Zipf distribution of us firm sizes. *Science*, 293(5536):1818–1820.

Bartelsman, E., Haltiwanger, J., and Scarpetta, S. (2004). *Microeconomic evidence of creative destruction in industrial and developing countries*. World Bank.

Bartelsman, E., Haltiwanger, J., and Scarpetta, S. (2009). Cross-country differences in productivity: the role of allocation and selection. NBER Working Papers 15490, National Bureau of Economic Research.

Bartelsman, E., Scarpetta, S., and Schivardi, F. (2013). Comparative analysis of firm demographics and survival: evidence from micro-level sources in oecd countries. *Industrial and Corporate Change*, 14(3):365–391.

Buera, F., Kaboski, J., and Shin, Y. (2011). Finance and development: A tale of two sectors. *The American Economic Review*, 101(5):1964–2002.

Cahuc, P. and Kramarz, F. (2004). *De la précarité à la mobilité: vers une sécurité sociale professionnelle*. La documentation française.

Ceci-Renaud, N. and Chevalier, P. (2011). Les seuils de 10, 20 et 50 salariés: impact sur la taille des entreprises françaises. *Economie et Statistique*, (437):29–45.

Chetty, R. (2012). Bounds on elasticities with optimization frictions: A synthesis of micro and macro evidence on labor supply. *Econometrica*, 80(3):969–1018.

Cooper, R., Haltiwanger, J., and Willis, J. (2007). Search frictions: Matching aggregate and establishment observations. *Journal of Monetary Economics*, 54:56–78.

Di Giovanni, J., Levchenko, A., and Ranciere, R. (2011). Power laws in firm size and openness to trade: Measurement and implications. *Journal of International Economics*, 85(1):42–52.

Dixit, A. K. and Pindyck, R. S. (1994). *Investment under uncertainty*. Princeton University Press.

García-Santana, M. and Pijoan-Mas, J. (2011). Small scale reservation laws and the misallocation of talent. CEPR Working Papers 8242, Centre for Economic Policy Research.

Garcia-Santana, M. and Ramos, R. (2012). Dissecting the size distribution of establishments across countries. mimeo, CEMFI.

Garicano, L., Lelarge, C., and Van Reenen, J. (2012). Firm size distortions and the productivity distribution: Evidence from france. mimeo, London School of Economics.

Gourieroux, C., Monfort, A., and Renault, E. (1993). Indirect inference. *Journal of applied econometrics*, 8(S1):S85–S118.

Guner, N., Ventura, G., and Xu, Y. (2008). Macroeconomic implications of size-dependent policies. *Review of Economic Dynamics*, 11(4):721–744.

Hall, P. and Horowitz, J. (1996). Bootstrap critical values for tests based on generalized-method-of-moments estimators. *Econometrica*, pages 891–916.

Hobijn, B. and Şahin, A. (2013). Firms and flexibility. *Economic Inquiry*, 51(1):922–940.

Hopenhayn, H. A. and Rogerson, R. (1993). Job Turnover and Policy Evaluation: A General Equilibrium Analysis. *The Journal of Political Economy*, 101(5):915–938.

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

Jaef, F. (2012). Entry, exit and misallocation frictions. mimeo, International Monetary Fund.

Lucas, R. E. (1978). On the Size Distribution of Business Firms. *The Bell Journal of Economics*, 9(2):508.

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

Stokey, N. L. (2008). *The Economics of Inaction: Stochastic Control Models with Fixed Costs*. Princeton University Press.

Appendix

This appendix is not intended for publication. The first section presents estimation results for each of four broad sectors of the economy (manufacturing, construction, retail and services). The second section presents the proofs of some model results and formulas.

A Sectoral Results

Tables 9 and 10 report the structural parameters estimates (with standard deviation in parenthesis) of the sunk cost model and the proportional tax model respectively, for each major sector of the economy. Tables 11 and 12 present the model fit for each sector.

There is some heterogeneity across sectors. Most notably, firms in the service sector have more volatile employment growth. The model interprets this fact as a higher volatility of productivity/demand shock z. Given that the Pareto power law exponent is not too different across sector, the model requires a higher exit rate of the service sector. However, the firm size distortion around 50 is comparable in all sectors, as shown in figure 13 and table 11. As a result, the estimated regulation costs are fairly similar. The estimated sunk cost varies between 90% and 115% of a worker annual wages. Similarly, the estimated proportional tax on wages varies between 0.25% and 0.29%. Hence, overall our results are robust across these subsamples.

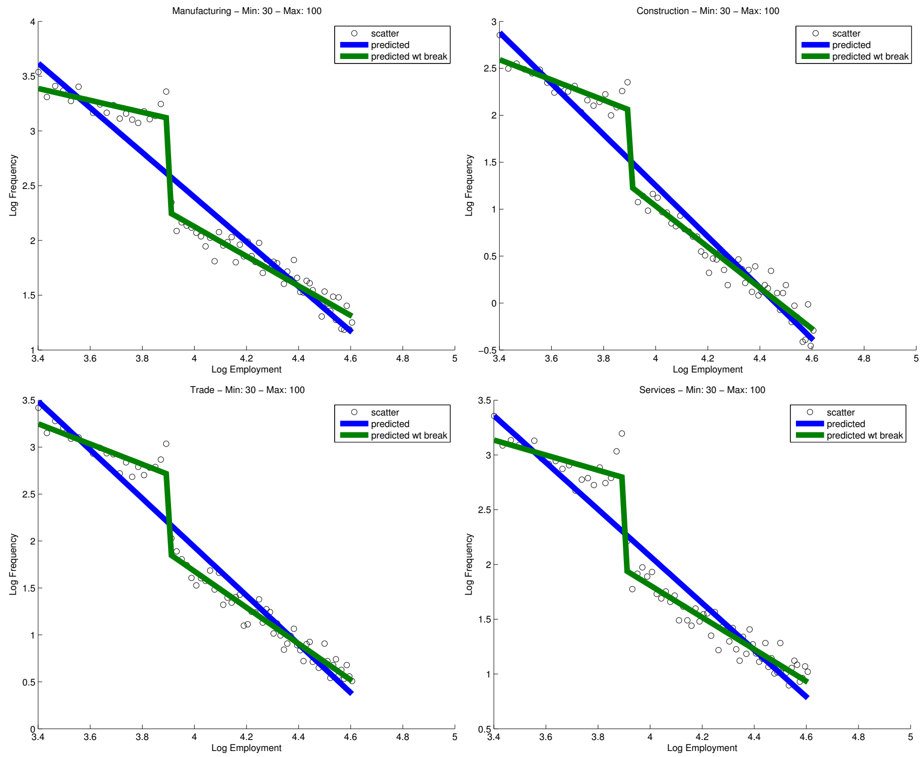

Figure 13: Broken Power Law:Regression of the logged number of firms on the logged number of employees with and without a structural break at 50 for firms with employment level between 30 and 100.

	All	Manufacturing	Construction	Trade	Services
λ	0.0499	0.0401	0.0499	0.0510	0.0798
	(0.0016)	(0.0004)	(0.0008)	(0.0016)	(0.0056)
μ	0.0029	0.0077	0.0071	0.0046	0.0068
	(0.0002)	(0.0001)	(0.0002)	(0.0005)	(0.0016)
σ	0.0842	0.0575	0.0535	0.0700	0.0985
	(0.0015)	(0.0003)	(0.0009)	(0.0012)	(0.0019)
σ_{mrn}	0.0107	0.0132	0.0164	0.0134	0.0113
	(0.0005)	(0.0009)	(0.0015)	(0.0012)	(0.0009)
F	1.0810	1.1486	1.1011	1.0666	0.9027
	(0.0672)	(0.0612)	(0.0881)	(0.0858)	(0.0789)

Table 9: Sunk Cost Model - Parameter Estimates (Standard Error in Parentheses)

	All	Manufacturing	Construction	Trade	Services
λ	0.0205	0.0246	0.0387	0.0303	0.0581
	(0.0003)	(0.0004)	(0.0007)	(0.0007)	(0.0013)
μ	-0.0035	0.0024	0.0044	-0.0009	0.0011
	(0.0001)	(0.0001)	(0.0003)	(0.0002)	(0.0003)
σ	0.0751	0.0593	0.0525	0.0705	0.0968
	(0.0008)	(0.0012)	(0.0010)	(0.0007)	(0.0007)
σ_{mrn}	0.0377	0.0399	0.0390	0.0409	0.0350
	(0.0010)	(0.0016)	(0.0022)	(0.0017)	(0.0013)
τ	0.0026	0.0026	0.0025	0.0026	0.0029
	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)

Table 10: Proportional Tax Model - Parameter Estimates (Standard Error in Parentheses)

	All		Manufacturing		Construction		Trade		Services	
	Data	Model	Data	Model	Data	Model	Data	Model	Data	Model
# firms										
40-45	0.3565	0.3392	0.3609	0.3369	0.3682	0.3496	0.3672	0.3458	0.3349	0.3359
45-50	0.3883	0.3852	0.3877	0.3791	0.3819	0.3769	0.3849	0.3802	0.3956	0.3954
50-55	0.1413	0.1517	0.1378	0.1497	0.1396	0.1522	0.1393	0.1523	0.1486	0.1483
55-60	0.1140	0.1238	0.1136	0.1341	0.1104	0.1214	0.1087	0.1216	0.1209	0.12033
Median $\Delta \log n$	0	0.0054	0	0.0194	0	0.0179	0	0.0108	0.0080	0.0163
$V(\Delta \log n)$	0.0485	0.0607	0.0319	0.0286	0.0306	0.0250	0.0437	0.0422	0.0839	0.0829
Power Law	1.1417	1.1471	1.0600	1.0628	1.3300	1.3331	1.2600	1.2632	1.1600	1.1603

Table 11: Auxiliary Parameters - Sunk Cost Model

	All		Manufacturing		Construction		Trade		Services	
	Data	Model	Data	Model	Data	Model	Data	Model	Data	Model
# firms										
40-45	0.3565	0.3395	0.3609	0.3343	0.3682	0.3494	0.3672	0.3461	0.3349	0.3367
45-50	0.3883	0.3847	0.3877	0.3815	0.3819	0.3752	0.3849	0.3761	0.3956	0.3963
50-55	0.1413	0.1504	0.1378	0.1580	0.1396	0.1529	0.1393	0.1552	0.1486	0.1480
55-60	0.1140	0.1253	0.1136	0.1261	0.1104	0.1225	0.1087	0.1226	0.1209	0.1190
Median $\Delta \log n$	0	-0.0096	0	0.0052	0	0.0110	0	-0.0020	0.0080	0.0010
$V(\Delta \log n)$	0.0485	0.0516	0.0319	0.0337	0.0306	0.0269	0.0437	0.0463	0.0839	0.0832
Power Law	1.1417	1.1503	1.0600	1.0568	1.3300	1.3380	1.2600	1.2557	1.1600	1.1583

Table 12: Auxiliary Parameters - Proportional Tax Model

B Proofs

B.1 Proof of Proposition 1

First, note that the function $V(.,1)$ is twice continuously differentiable (see Stokey (2008) Chapter 5.6 for a proof). Using the previously computed $\pi(z,1)$ gives:

$$\begin{aligned} V(z^*,1) &= \frac{1}{J}\left[\int_{z^*}^{\infty} e^{R_2(z^*-z)}\pi(z,1)dz + \int_{-\infty}^{z^*} e^{R_1(z^*-z)}\pi(z,1)dz\right], \\ &= \frac{1}{J}\Bigg[\int_{z^*}^{\infty} e^{R_2(z^*-z)}\pi^a(z)dz + \int_{\overline{z}}^{z^*} e^{R_1(z^*-z)}\pi^a(z)dz \\ &\quad + \int_{\underline{z}}^{\overline{z}} e^{R_1(z^*-z)}\pi^b(z)dz + \int_{-\infty}^{\underline{z}} e^{R_1(z^*-z)}\pi^b(z)dz\Bigg]. \end{aligned}$$

Define, for all $x \leq z^*$,

$$\begin{aligned} H(x, z^*) &\equiv E_x \left[\int_0^{T(z^*)} e^{-(r+\lambda)t} \pi(z_t, 0) dt + e^{-(r+\lambda)T(z^*)} \left(V(z^*, 1) - F \right) \right], \\ &= \frac{1}{J} \left[\int_x^{z^*} e^{R_2(x-z)} \pi(z, 0) dz + \int_{-\infty}^{x} e^{R_1(x-z)} \pi(z, 0) dz - e^{R_2(x-z^*)} \int_{-\infty}^{z^*} e^{R_1(z^*-z)} \pi(z, 0) dz \right] \\ &\quad + e^{R_2(x-z^*)} \left(V(z^*, 1) - F \right). \end{aligned}$$

Then, $V(x, 0) = \sup_{z^* \geq x} H(x, z^*)$. Note that $H(x, z^*)$ is twice continuously differentiable. The FOC for a maximum at $z^* \geq \overline{z}$ is

$$\begin{aligned} 0 &\leq f_{z^*}(x, z^*) \\ &= \frac{1}{J} \left[e^{R_2(x-z^*)} \pi(z^*, 0) + R_2 e^{R_2(x-z^*)} \int_{-\infty}^{z^*} e^{R_1(z^*-z)} \pi(z, 0) dz \right] \\ &\quad + \frac{1}{J} \left[-e^{R_2(x-z^*)} \pi(z^*, 0) - R_1 e^{R_2(x-z^*)} \int_{-\infty}^{z^*} e^{R_1(z^*-z)} \pi(z, 0) dz \right] \\ &\quad - R_2 e^{R_2(x-z^*)} \left(V(z^*, 1) - F \right) + e^{R_2(x-z^*)} V_{z^*}(z^*, 1) \\ &= e^{R_2(x-z^*)} \left[\frac{R_2 - R_1}{J} \int_{-\infty}^{z^*} e^{R_1(z^*-z)} \pi(z, 0) dz - R_2 \left(V(z^*, 1) - F \right) + V_{z^*}(z^*, 1) \right], \end{aligned}$$

with equality if $z^* > \overline{z}$. Hence,

$$\begin{aligned} & V(z^*, 1) \\ = \; & \frac{1}{J} \left[\int_{z^*}^{\infty} e^{R_2(z^*-z)} \pi^a(z) dz + \int_{\overline{z}}^{z^*} e^{R_1(z^*-z)} \pi^a(z) dz + \int_{\underline{z}}^{\overline{z}} e^{R_1(z^*-z)} \pi^b(z) dz + \int_{-\infty}^{\underline{z}} e^{R_1(z^*-z)} \pi^b(z) dz \right] \\ & V_z(z^*, 1) \\ = \; & \frac{R_2}{J} \int_{z^*}^{\infty} e^{R_2(z^*-z)} \pi(z, 1) dz + \frac{R_1}{J} \int_{-\infty}^{z^*} e^{R_1(z^*-z)} \pi(z, 1) dz \\ = \; & \frac{R_2}{J} \int_{z^*}^{\infty} e^{R_2(z^*-z)} \pi^a(z) dz \\ & + \frac{R_1}{J} \left[\int_{\overline{z}}^{z^*} e^{R_1(z^*-z)} \pi^a(z) dz + \int_{\underline{z}}^{\overline{z}} e^{R_1(z^*-z)} \pi^b(z) dz + \int_{-\infty}^{\underline{z}} e^{R_1(z^*-z)} \pi^b(z) dz \right]. \end{aligned}$$

Plugging ing the FOC gives

$$(R_1 - R_2) \int_{\overline{z}}^{z^*} e^{R_1(z^*-z)} \left[\pi^a(z) - \pi^b(z) \right] dz + R_2 J F = 0,$$

which simplifies to

$$R_1 \int_{\overline{z}}^{z^*} e^{R_1(z^*-z)} \left[\pi^a(z) - \pi^b(z)\right] dz + (r+\lambda)F = 0.$$

It is easy to see that there exists a unique value of z^* that satisfies the preceding equality. Moreover, one can compute these integrals easily given our formulas for $\pi^a(z)$ and $\pi^b(z)$.

B.2 Alternative Derivation of Optimal Policy using Dynamic Programming

We start by writing the Hamilton-Jacobi-Bellman equation satisfied by V :

$$(r+\lambda)V(z,1) = \pi(z,1) + \mu V_z(z,1) + \frac{\sigma^2}{2} V_{zz}(z,1), \tag{14}$$

for any z, and

$$(r+\lambda)V(z,0) = \pi(z,0) + \mu V_z(z,0) + \frac{\sigma^2}{2} V_{zz}(z,0), \tag{15}$$

for $z < z^*$. Note that $\pi(z,0)$ and $\pi(z,1)$ are only C^1 (continuous differentiable): the second derivative is discontinuous at $z = \underline{z}$ for $\pi(.,0)$ and $\pi(.,1)$, and at $z = \overline{z}$ for $\pi(.,1)$.

The boundary conditions given by value matching:

$$V(z^*,1) = V(z^*,0) - F, \tag{16}$$

and by the smooth pasting condition:

$$V_z(z^*,1) = V_z(z^*,0). \tag{17}$$

The general solution of the associated homogeneous ODE (i.e., without the term π) is $A_1 e^{R_2 z} + A_2 e^{R_1 z}$, where R_1 and R_2 are the roots of the quadratic

$$\frac{\sigma^2}{2} X^2 + \mu X - (r+\lambda) = 0, \tag{18}$$

i.e.$R_2 = \frac{-\mu + \sqrt{\mu^2 + 2(r+\lambda)\sigma^2}}{\sigma^2} > 0$ and $R_1 = \frac{-\mu - \sqrt{\mu^2 + 2(r+\lambda)\sigma^2}}{\sigma^2} < 0$.

The specific forms of $\pi(z,0)$ and $\pi(z,1)$ make it possible to find particular solutions. Starting with the first equation, we guess that

$$\begin{aligned}\widetilde{V}(z,0) &= b_0 e^{\frac{z}{1-\alpha}}, \text{ for } z < \underline{z}, \\ &= b_1 e^z + b_2, \text{ for } z > \underline{z},\end{aligned}$$

is a solution of 15, for constants b_0, b_1, b_2 to be determined.

$\widetilde{V}$ satisfies the ODE for $z < \underline{z}$, provided that b_0 solves:

$$(r+\lambda)b_0 = \left(\frac{\alpha}{w}\right)^{\frac{\alpha}{1-\alpha}}(1-\alpha) + \mu\frac{b_0}{1-\alpha} + \frac{\sigma^2}{2}\frac{b_0}{(1-\alpha)^2},$$

or

$$b_0 = \frac{\left(\frac{\alpha}{w}\right)^{\frac{\alpha}{1-\alpha}}(1-\alpha)}{r+\lambda-\frac{\mu}{1-\alpha}-\frac{\sigma^2}{2(1-\alpha)^2}}.$$

For $z > \underline{z}$, we require that

$$(r+\lambda)\left(b_1 e^z + b_2\right) = e^z \underline{n}^\alpha - w\underline{n} + \mu b_1 e^z + \frac{\sigma^2}{2} b_1 e^z,$$

i.e.

$$b_2 = -\frac{w\underline{n}}{r+\lambda},$$

$$b_1 = \frac{\underline{n}^\alpha}{r+\lambda-\mu-\frac{\sigma^2}{2}}.$$

The general solution of the first equation is thus

$$\begin{aligned}V(z,0) &= \widetilde{V}(z,0) + A_1 e^{R_2 z} + A_2 e^{R_1 z} \\ &= \frac{\left(\frac{\alpha}{w}\right)^{\frac{\alpha}{1-\alpha}}(1-\alpha)}{r+\lambda-\frac{\mu}{1-\alpha}-\frac{\sigma^2}{2(1-\alpha)^2}} e^{\frac{z}{1-\alpha}} + A_1 e^{R_2 z} + A_2 e^{R_1 z}, \text{ for } z < \underline{z} \\ &= \frac{\underline{n}^\alpha}{r+\lambda-\mu-\frac{\sigma^2}{2}} e^z - \frac{w\underline{n}}{r+\lambda} + A_1 e^{R_2 z} + A_2 e^{R_1 z}, \text{ for } z \geq \underline{z}.\end{aligned}$$

Turning to the second equation, we again look for one solution, which we guess as

$$\begin{aligned}\widetilde{V}(z,1) &= e^{\frac{z}{1-\alpha}}b_3,\ \text{for } z<\underline{z},\\ &= e^{z}b_4+b_5,\ \text{for } \overline{z}>z>\underline{z},\\ &= e^{\frac{z}{1-\alpha}}b_6+b_7,\ \text{for } z>\overline{z}.\end{aligned}$$

The scalars b_3, b_4, b_5, b_6, b_7 must satisfy:

$$\begin{aligned}b_3 &= \frac{\left(\frac{\alpha}{w}\right)^{\frac{\alpha}{1-\alpha}}(1-\alpha)}{r+\lambda-\frac{\mu}{1-\alpha}-\frac{\sigma^2}{2(1-\alpha)^2}} = b_0,\\ b_4 &= \frac{\underline{n}^{\alpha}}{r+\lambda-\mu-\frac{\sigma^2}{2}} = b_1,\\ b_5 &= -\frac{w\underline{n}}{r+\lambda} = b_2,\\ b_6 &= \frac{\left(\frac{\alpha}{w(1+\tau)}\right)^{\frac{\alpha}{1-\alpha}}(1-\alpha)}{r+\lambda-\frac{\mu}{1-\alpha}-\frac{\sigma^2}{2(1-\alpha)^2}} = \frac{b_0}{(1+\tau)^{\frac{\alpha}{1-\alpha}}},\\ b_7 &= -\frac{c_f}{r+\lambda},\end{aligned}$$

and the general solution is

$$V(z,1)=\widetilde{V}(z,1)+A_3e^{R_2z}+A_4e^{R_1z}$$

Finally we need to determine A_1, A_2, A_3, A_4 and z^*. A standard argument implies that $A_3=0$ (the investment option values goes to 0 if $z\to\infty$). Moreover, $A_4=0$ since as $z\to-\infty$ the firm value remains finite. Last, $A_2=0$ for the same reason. The two scalars A_1 and z^* are thus determined by the following system of two equations in two unknowns:

$$\widetilde{V}(z^*,1)=\widetilde{V}(z^*,0)+A_1e^{R_2z^*}-F,$$

$$\widetilde{V}_z(z^*,1)=\widetilde{V}_z(z^*,0)+A_1R_2e^{R_2z^*}.$$

Given the formulas for $\widetilde{V}$ and that $z^*>\overline{z}>\underline{z}$, this can be rewritten as:

$$e^{\frac{z^*}{1-\alpha}}b_6+b_7=\frac{\underline{n}^{\alpha}}{r+\lambda-\mu-\frac{\sigma^2}{2}}e^{z^*}-\frac{w\underline{n}}{r+\lambda}+A_1e^{R_2z^*}-F$$

$$e^{\frac{z^*}{1-\alpha}} \frac{b_6}{1-\alpha} = \frac{\underline{n}^\alpha}{r+\lambda-\mu-\frac{\sigma^2}{2}} e^{z^*} + A_1 R_2 e^{R_2 z^*}.$$

This characterizes entirely the solution. It is easy to verify that this yields the same results as those obtained in the main text using the theoretical results of Stokey (2008).

B.3 Derivation of the Stationary Cross-Sectional Distribution

To solve for f, first note that the general solution of the ODE 7 is

$$f(z,0) = D_0 e^{\beta_1 z} + D_1 e^{\beta_2 z},$$

where $\beta_1 < 0 < \beta_2$ are the two real roots of the characteristic equation:

$$\lambda = -\mu X + \frac{\sigma^2}{2} X^2.$$

This equation must be solved separately on each interval. Given that f is a density, the exponential terms which do not go to 0 must disappear. This yields the following simpler form:

$$\begin{aligned} f(z,0) &= C_1 e^{\beta_2 z}, \text{ for } z < z, \\ &= C_2 e^{\beta_1 z} + C_3 e^{\beta_2 z}, \text{ for } z^* > z > z, \end{aligned}$$

and

$$\begin{aligned} f(z,1) &= C_4 e^{\beta_2 z}, \text{ for } z < z^*, \\ &= C_5 e^{\beta_1 z}, \text{ for } z > z^*. \end{aligned}$$

The boundary conditions can then be expressed as a system of five linear equations in five unknowns. First, f is a p.d.f., i.e. its integral is one:

$$\frac{C_1}{\beta_2} e^{\beta_2 z} + \frac{C_2}{\beta_1}\left(e^{\beta_1 z^*} - e^{\beta_1 z}\right) + \frac{C_3}{\beta_2}\left(e^{\beta_2 z^*} - e^{\beta_2 z}\right) + \frac{C_4}{\beta_2} e^{\beta_2 z^*} - \frac{C_5}{\beta_1} e^{\beta_1 z^*} = 1.$$

Second, $f(.,0)$ is continuous at z :

$$C_1 e^{\beta_2 z} = C_2 e^{\beta_1 z} + C_3 e^{\beta_2 z},$$

Third, $f(.,0)$ is continuous at z^* :

$$C_2 e^{\beta_1 z^*} + C_3 e^{\beta_2 z^*} = 0,$$

Fourth, $f(.,1)$ is continuous at z^* :

$$C_5 e^{\beta_1 z^*} = C_4 e^{\beta_2 z^*}.$$

And finally the boundary condition at z^* :

$$-\frac{\sigma^2}{2}\left(C_2\beta_1 e^{\beta_1 z^*} + C_3\beta_2 e^{\beta_2 z^*}\right) = \lambda\left(\frac{C_4}{\beta_2}e^{\beta_2 z^*} - \frac{C_5}{\beta_1}e^{\beta_1 z^*}\right).$$

This system of equations can be solved analytically using Mathematica, yielding the results in the main text.