The Empirical Distribution of Firm Dynamics and Its Macro Implications*

Nir Jaimovich
University of California, San Diego and *NBER, CEPR*
nijaimovich@ucsd.edu

Stephen J. Terry
Boston University and NBER
stephent@bu.edu

Nicolas Vincent
HEC Montréal
nicolas.vincent@hec.ca

September 2022

Abstract

Heterogeneous firm models are central to the modern economics literature. In this context, we revisit a central feature of these models: the nature of idiosyncratic shocks faced by firms. Using firm-level data, we document that the commonly assumed parametric distributions of the shocks that firms face differ in important ways from observed dynamics. We embed these findings into a general equilibrium heterogeneous firm model and show they have a first-order effect on the response of the economy to aggregate disturbances.

Keywords: firm dynamics, non-parametric shocks, selection, policy

*Financial support from the Fondation HEC Montréal and the Social Sciences and Humanities Research Council of Canada (Vincent) is gratefully acknowledged. Financial support from the Swiss National Science Foundation (grant number: 100018 − 197799) is gratefully acknowledged (Jaimovich).

"The overall effect on aggregate employment seems ambiguous, depending on the stochastic structure of firm-level shocks. This being the case, evidence on the firm-level stochastic environment is necessary." Hopenhayn and Rogerson (1993)

1 Introduction

Heterogeneous firm models are at the center of modern macroeconomics and are nowadays routinely used both for positive and normative analysis. Since the seminal contribution of Hopenhayn (1992), this type of framework has increasingly been employed in macroeconomics to study, among others, the contributions of entry and exit to aggregate job creation and destruction (Hopenhayn and Rogerson, 1993); the cyclical implications of firm entry and exit (Bilbiie et al., 2012; Clementi and Palazzo, 2016; Lee and Mukoyama, 2018); the decline in business dynamism (Decker et al., 2016, 2020; Karahan et al., 2022); the role of firm heterogeneity in shaping the dynamics of aggregate investment (Khan and Thomas, 2008, 2013; Winberry, 2021); the propagation of financial frictions (Moll, 2014; Midrigan and Xu, 2014; Ottonello and Winberry, 2020); the role of uncertainty shocks (Bloom et al., 2018); and the drivers and consequences of the (mis)allocation of resources (Restuccia and Rogerson, 2008; Hsieh and Klenow, 2009; Bento and Restuccia, 2017; Kehrig and Vincent, 2020). This broad family of models has also been highly influential in the trade literature, drawing on the work of Melitz (2003).

In this paper, we revisit the standard assumptions regarding the key driving force in these models, namely the idiosyncratic shocks. Their importance was already recognized by Hopenhayn and Rogerson (1993), as evidenced by the quote at the top of this page. Specifically, we show that the common assumptions made in the literature to parameterize these shocks leads to firm-level distributions and dynamics that differ in important ways from what is observed in the data. Furthermore, we demonstrate that these differences have a first-order impact on allocations when analyzing the response of the model economy to aggregate perturbances. In what follows, we describe more broadly these findings.

Leveraging a large, representative firm-level dataset, we first extract rich non-parametric distributions and transition dynamics for revenue, and compare them to similar objects obtained under the standard autoregressive process widely used in the literature, a Gaussian AR(1). We document a number of facts that prove to be highly relevant for the modeling exercise. Broadly speaking, we note that while the ergodic distribution of firm revenue is not strikingly different between the two models (besides the well-known presence of fat tails in the empirical distribution), the *dynamics* are. For instance, a firm initially in the middle of the revenue distribution has a much higher probability of staying around the median *or* moving to the tails than implied under an autoregressive specification. Conversely, conditional on firm revenue being in the tails of the distribution, the data reveal a greater probability of returning towards the center than in an AR(1) model. These differences result in the distribution of revenue *growth* in the data being leptokurtic, unlike that generated from the common

AR(1) assumption.

What are the economic implications of these statistical findings? In any model, the expected continuation value is ultimately the object that shapes the firm's optimal decisions, such as whether to exit or how much to hire and invest. Hence, in the second part of our data-driven exercise, we generate model-free distributions of firm values (present discounted value of lifetime expected revenue) based on the transition dynamics extracted empirically. We show that these approximations are very different between the non-parametric and parametric models. In particular, we note that firm values are much more clustered (that is, the probability density function is much steeper) for the non-parametric version than under the AR(1) assumption. The reason is intuitive: the higher probability of moving away from the tails and remaining around the center of the revenue distribution creates a more pronounced compression of the distribution of firm values under the empirical, non-parametric specification.

In the last part of the paper, we incorporate these new non-parametric findings into a canonical general equilibrium heterogeneous firm dynamics model. The rich environment we consider features shocks to firm-level profitability, fixed costs of operation, endogenous exit, and sunk entry costs. For the non-parametric version of the model, our quantitative approach allows us to perfectly match a number of empirical objects: (i) the transition matrix of revenue for incumbents; (ii) the exit hazard; and (iii) the relative size distribution of entrants. For the parametric version, we follow the literature in calibrating the AR(1) shock process and other parameters to match a number of moments. Aside from the modeling of the idiosyncratic shock process, the two models are identical.

We then compare the impact in the two models to two types of policies: a fixed subsidy to each operating firm or a subsidy to entrants.[1] In both cases, the exit rate is much more responsive in the non-parametric version. This is a direct result of the fact that firm values are more clustered for low-revenue states, where exit is more likely to occur. Yet, we show that the response of aggregate output to the fixed cost subsidy is 45% *larger* under the AR(1) shock process, but about one-third *lower* in the case in which the subsidy is aimed at new entrants. That is, the impact of the higher exit rate sensitivity on macro aggregates depends on the policy implemented. As we show, the reason has to do with the role of selection.

Consider first the case of the subsidy to operating firms, which increases firms' values, leading to a decline in the likelihood of exit. As a result, selection worsens, as more low-productivity firms survive. This negative selection effect is more pronounced in the non-parametric model, since the density of firms around the exit thresholds is higher, which dampens the positive output response to the subsidy relative to the AR(1) version. On the other hand, the entry subsidy generates an *increase* in the steady-state rate of exit. As a result, selection improves, boosting output. Since the exit rate responds more in the non-parametric version, so does output through the selection effect.

We view this paper as closely linked to three main strands in the literature. First it naturally

[1] These experiments mirror those found in Hopenhayn (1992).

relates to the theory and empirics of firm dynamics (Dunne et al., 1989; Hopenhayn, 1992; Davis and Haltiwanger, 1992; Kehrig, 2015; Clementi and Palazzo, 2016; Karahan et al., 2022). This literature has investigated how taking into account firm heterogeneity can help rationalize some empirical stylized facts at both the micro and macro levels. Second, it relates to the literature contrasting empirical and "conventional" distributions (Midrigan, 2011; Carvalho and Grassi, 2019; Forneron, 2020; Guvenen et al., 2021; Sterk et al., 2021). We show that the common parametric assumptions used in heterogeneous firm models are poor approximations of reality.

Third, our work relates to the allocative implications of policy and shocks in the presence of firm heterogeneity (Hopenhayn and Rogerson, 1993; Guner et al., 2008; Restuccia and Rogerson, 2008; Hsieh and Klenow, 2009; Davies and Eckel, 2010; Gourio and Miao, 2010; Asker et al., 2014; Garicano et al., 2016; Catherine et al., 2018; Kehrig and Vincent, 2020; Ottonello and Winberry, 2020; Bils et al., 2021; Sraer and Thesmar, 2021). We show that our empirical findings alter in a quantitatively significant way the impact of different policies.

The rest of the paper is organized as follows. We start in Section 2 with the empirical analysis. In Sections 3 and 4, we describe our model economy and discuss the approach used to incorporate a non-parametric characterization of the idiosyncratic shock process. Section 5 contrasts the comparative statics in our non-parametric model economy vis-à-vis the commonly calibrated parametric economy. Section 6 presents supportive evidence in the cross-section of the model's key prediction regarding the positive relation between the clustering of firms' present value distributions and exit probabilities Finally, 7 concludes.

2 Empirical analysis of firm dynamics

In this section, we discuss the data used and our empirical approach. Drawing on a number of stylized facts, we show that the commonly used parametric assumption has very different implications about the firm-level distribution and dynamics of revenue than those generated from the data. These stylized facts will be used to calibrate the heterogenous firms model presented in Section 3.

2.1 Data sources

We briefly present here the key aspects of the data used for the analysis and refer the reader to Appendix A for a detailed discussion.

Our objective of capturing the rich firm-level heterogeneity requires us to have access to a dataset that is representative of the population of firms. For our purposes, we rely on Bureau van Dijk' ORBIS database, accessed through a NBER data initiative. Drawing mainly from government tax records, it contains a large number of firm-level economic and financial variables at yearly frequency for a number of European countries. While in principle the database includes both private and publicly listed companies, coverage and representativeness varies greatly across countries. For presentation purposes, our focus in this paper is on Spain between 2005 and 2014, but we also present results for Italy (2008-2014), Portugal (2006-2014), France (2005-2014), and Norway (2007-2014).[2] In earlier work, Kalemli-Ozcan et al. (2022) and Bajgar et al. (2020) have shown that for these countries, the information contained in ORBIS is consistent with the indicators derived from other firm-level datasets. A particularly relevant comparison is with COMPNET, which provides micro-aggregated moments based on firm-level datasets maintained by national statistical agencies or central banks. While for presentation purposes our focus is generally on Spain in the body of the paper, we provide evidence that our main findings are robust to the use of alternative countries.

2.2 Data preparation

The main variable of interest throughout our analysis is firm-level operating revenue.[3] We proceed in three steps. First, we estimate the revenue dynamics of incumbent firms from the data and compare them to those obtained from a standard, properly calibrated AR(1) assumption. Second, we provide evidence on the behavior of entrants and exiters. Finally, in the next section, we exploit these moments to generate a distribution of firm values, computed as the present discounted flow of revenue. The stylized facts we document, in addition of being of interest by themselves, will be crucial to ensure that our calibrated model perfectly matches firm-level revenue dynamics (see Section 4).

Demeaning We first demean log revenue at the industry/year level before extracting the empirical distribution and transitions. This ensures that our findings are not driven by systematic differences *across* sectors, but instead by dynamics *within industries*. In other words, when we later study firm-level revenue dynamics, we will implicitly be describing the behavior of firms located in a specific

[2]It is worth noting that while we in fact exploit data that goes beyond 2014, the latter years are only used to minimize the occurrence of "spurious exit" due to gaps in the data. For example, consider a Spanish firm for which data is available in the dataset for every year between 2005 and 2019, except between 2012 and 2016. In this case, the presence of information beyond 2016 would ensure that we do not consider that the firm exited between 2011 and 2012.

[3]The variable is variable *operating revenue turnover*. A natural alternative would have been to study the dynamics of cash flow (profits), as it is conceptually closer to the concept of flow return that is key to constructing the firm's value function, which will be central to our analysis. There are a number of reasons behind our decision. First, the operating revenue variable in ORBIS is more consistently populated than cash flow measures across countries. Second, using revenue will later allow us to transparently link the empirical object to the profitability shocks in the model. Third, we show in Appendix A.1 that revenue is a better predictor of future exit and hiring decisions than profits or earnings. This is consistent with the view in the literature that accounting profits are a poorly measured proxy for economic cash flows. Also note that ultimately, our model calibration in Section 4 will by definition ensure that the firm-level dynamics of *revenue* in the model are similar to those observed in the data. Finally, we also exploit additional information for various purposes, such as profit (PL), earnings ($EBIT$ and $EBTA$), employment ($EMPL$), cost of employees ($STAF$) or firm age. Later we provide additional information on our treatment of the data as well as the size of the sample.

portion of the distribution defined at the industry level. Specifically, let the level of revenue be Y, then demeaning is done by running the following regression:

$$\ln(Y_{it}) = y_{it} = \alpha + \gamma_{jt} + \epsilon_{it} \tag{1}$$

where y_{it} is the natural log of operating revenue of firm i in year t, and γ_{jt} corresponds to industry$\times$year fixed effects. The subsequent analysis is based on the residuals from this regression, which we will denote as $\hat{y_{it}}$.[4]

Recovering the transition matrices First, we divide the distribution of (residualized) revenue $\widehat{y}$ into a grid of size $(N_y \times 1)$ covering equal-weight intervals.[5] This allows us to directly extract the three non-parametric objects of interest from the data. Specifically, we compute:

1. The empirical transition distribution for incumbents

$$H(\hat{y}'|\hat{y}) \tag{2}$$

2. The distribution of revenue for entrants

$$H_E(\hat{y}) \tag{3}$$

3. The exit hazard for incumbents

$$\mathbb{P}(\text{Exit}|\hat{y}) \tag{4}$$

The non-parametric version of our quantitative model in Section 4 will exactly match these empirical revenue patterns by construction.[6]

Calibrating the AR(1) process Our objective is to compare the properties of the empirical revenue distribution to that generated by a parametric AR(1) specification $\hat{y}' = \rho\hat{y} + \sigma N(0,1)$ that is standard in the literature. To this aim, we calibrate two parameters of interest (the persistence, ρ, and standard deviation of innovations, σ) to match the autocorrelation and unconditional variance of $\hat{y}$ for incumbents. For Spain, we use $\rho = 0.9415$ and $\sigma = 0.1908$.

[4] We have verified that our conclusions are similar if we solely demean by industry, solely demean by year, or only take out the unconditional mean.

[5] We verified that the number of grid points N_y was high enough that our results were invariant to increasing or decreasing it in the vicinity of 101. As such we let $N_y = 101$.

[6] In order to embed H, H_E, and $P(\text{Exit}|\hat{y})$ into a canonical heterogeneous-firms model, each needs to be adjusted in order to satisfy some standard technical assumptions employed by that literature. These assumptions ensure monotonicity of firm value functions and internal consistency of firm exit policies. See Appendix A for more details and a comparison of the original and adjusted objects. The main takeaway is that the required adjustments are minor and in most cases barely discernible.

2.3 Firm-level revenue distribution and dynamics: data vs. AR(1)

We now turn to comparing the properties of the firm-level transition dynamics under the non-parametric and AR(1) specifications. Before doing so, we reproduce in Figure 1 the stationary distributions of the residualized logged revenue $\hat{y}$.

Figure 1: Ergodic distribution of firm-level revenue

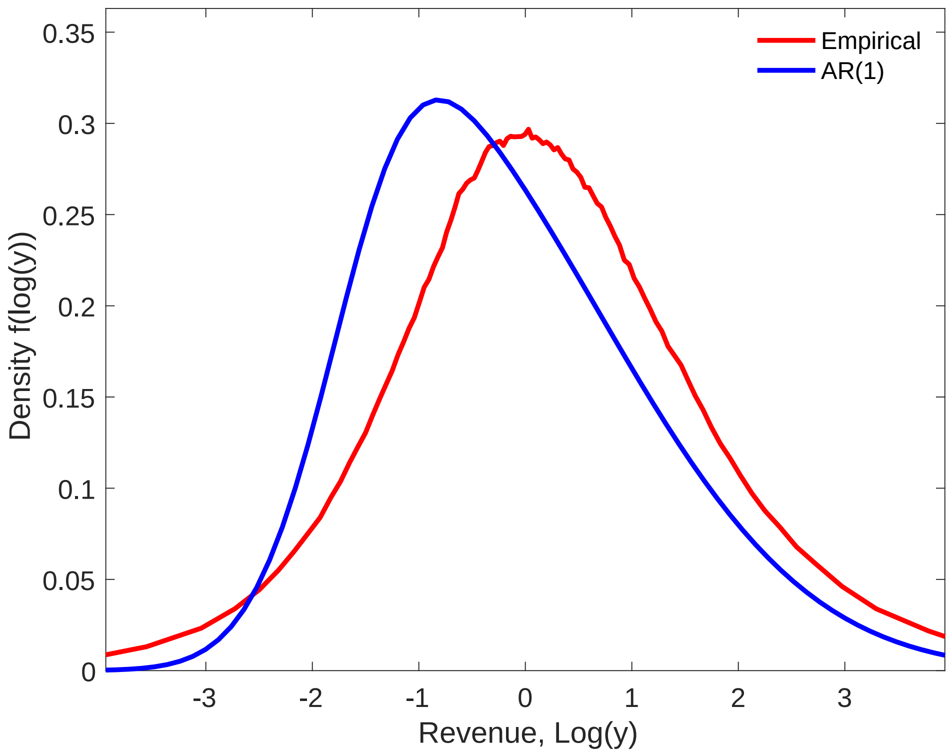

Notes: This figure depicts the ergodic distribution of $\log y$ in the empirical distribution and that implied by the $AR(1)$ parameterization.

While the empirical distribution displays somewhat fatter tails, the differences with its AR(1) counterpart are not particularly striking. Specifically, the standard deviation, skewness and kurtosis of $\log(\hat{y})$ are respectively 1.55, 0.025 and 4.196 for the empirical distribution compared to XX, 0 and 3 by definition under the autoregressive process. This suggests that the AR(1) specification is an acceptable approximation of the actual revenue distribution found in the data. We also show in Appendix A.3 that the empirical distribution of revenue features a Pareto tail, in line with evidence from the large literature on firm size distribution (e.g. Axtell (2001)). As we show next, differences are much more pronounced once we focus on transition dynamics.

2.3.1 Transition dynamics of incumbents

Figure 2 depicts the density of the transition probabilities of $H(\hat{y}'|\hat{y})$, conditional on $\hat{y}$ being at its median value. This figure shows that the empirical distribution is characterized by fat-tail leptokurtic dynamics; the differences between the empirical and AR(1) transitions are much more striking than for the ergodic distribution of log revenue.

Specifically, conditional on being initially at the median value of $\hat{y}$, the probability of remaining around the median is much higher than under the calibrated AR(1) specification. Yet, given the nature of the fat-tail dynamics, the empirical distribution also features a higher likelihood of moving to the tails from the median.

Figure 2: Transition away from the center of the revenue distribution

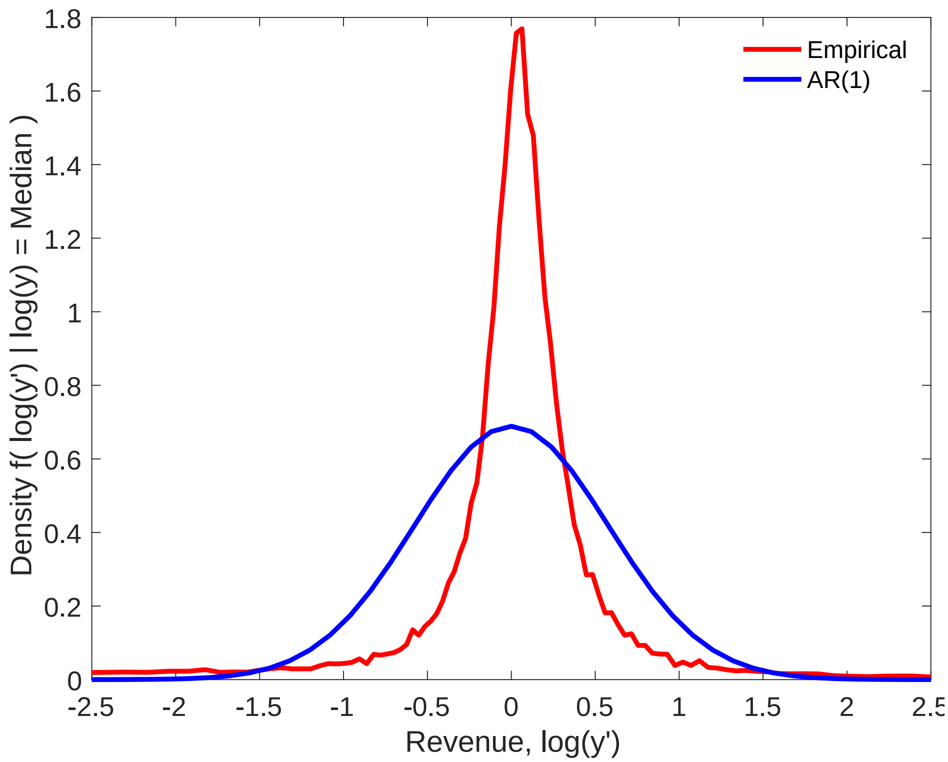

Notes: This figure depicts the density of $\log y'$ conditional on y being at its median value.

2.3.2 Transition from the tails

Next, we study the likelihood that a firm lands in the middle of the distribution given its initial position in the revenue distribution. Specifically, we plot in Figure 3 the probability that the firm transitions to the third quintile of the revenue distribution as a function of its current residualized log revenue, $\hat{y}$. The picture that emerges is clear: for firms currently around the median industry revenue (i.e. $\hat{y} = 0$), the probability of remaining in the third quintile is much higher in the empirical model.

Similarly, conditional on being in the tails of the distribution, the probability of moving towards the center of the distribution is significantly higher in the empirical model, particularly for adverse states.

Figure 3: Transition towards the center of the revenue distribution

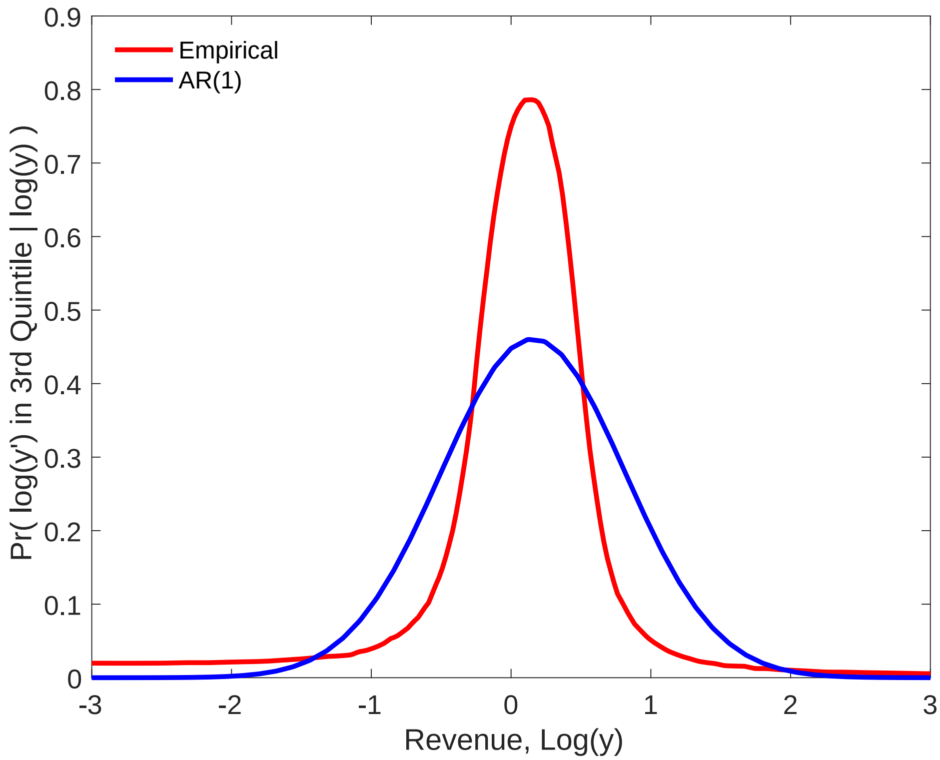

Notes: This figure depicts the probability that $\log y'$ is in the 3rd quintile as a function of y.

2.3.3 The distribution of growth rates

Finally, Figure 4 depicts the pooled distribution of the growth rates of revenue. It shows that the empirical dynamics discussed earlier generate a leptokurtic growth distribution: larger probabilities of very negative and positive growth rates, with higher density around zero growth than its AR(1) counterpart. More precisely, we find that the standard deviation, skewness and kurtosis of the revenue growth rates are respectively 0.656, -0.312 and 29.212 for the empirical distribution. In contrast, under the AR(1) calibration, the counterparts to these moments are XX, 0 and 3. Thus, while in levels the differences were not quantitatively meaningful, the distributions of growth rates are strikingly distinct. We show in Table A.2 of Appendix A.4 that this conclusion holds for a large array of alternative treatments of the data as well as other countries.

Figure 4: Revenue growth rates

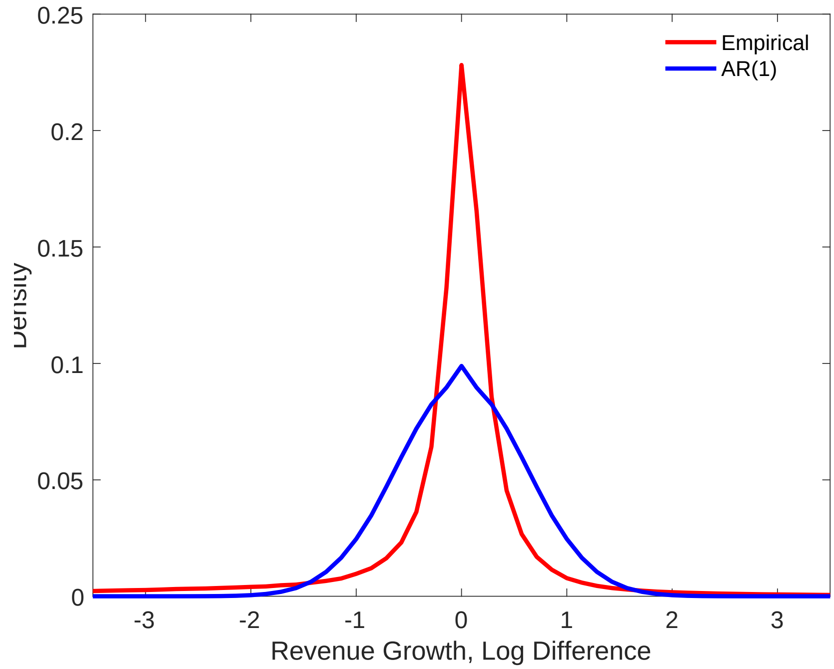

Notes: This figure depicts distribution of the growth rates of $log(y)$

2.3.4 A richer autoregressive specification

One potential concern is that the differences we document between the empirical and parametric cases are the result of features that are not embedded in the basic standard AR(1) model. In this section, we show that a richer specification that includes (1) firm fixed effects (permanent types) and (2) i.i.d. measurement error is not sufficient to allow this type of model to perform as well as the nonparametric counterpart.

Specifically, we consider an augmented AR(1) specification in which the evolution of log revenue is given by

$$\hat{y}_i{}' = \tilde{y}_i{}' + \tilde{\sigma} N(0,1) \quad \text{where} \quad \tilde{y}_i{}' = \alpha_i + \rho \tilde{y}_i + \hat{\sigma} N(0,1). \tag{5}$$

In the equation above, α_i denotes a firm fixed effect – which we assume follows a Pareto distribution – and $\tilde{\sigma}$ corresponds to the standard deviation of the measurement error. We estimate this augmented AR(1) model with an SMM exercise targeting the autocorrelation of revenue, the variances of revenue and revenue growth, the top 1% revenue share, and mean revenue.

The first column of Table 1 reports the average root mean squared errors (RMSEs) of the one-year ahead forecast for the three specifications: empirical, AR(1) and augmented AR(1). While the predictive performance of the augmented AR(1) is better (lower RMSE) than that of the basic AR(1) specification, it still clearly underperforms our nonparametric statistical model.[7]

[7]Note that by definition, the nonparametric specification predicts perfectly log revenue at a one-year horizon.

For the results found in the second column, we instead computed the one-year-ahead log predictive score (LPS) for each specification. The LPS allows us to account for the predictive fit across the whole density, and not only the mean. In this case, a higher LPS denotes better forecasting performance. Again, while the augmented model outperforms the standard autoregressive specification, it falls short of the nonparametric statistical model.

To summarize, our findings indicate that augmenting the AR(1) specification with off-the-shelf, standard features such as firm fixed effects and measurement error does not eliminate the superior performance of the nonparametric version. Next, we show that these statistical differences have first-order consequences on the key object of interest in heterogeneous firm models: the distribution of firm values.

Table 1: Relative Model Performance at One-year Horizon

Model	RMSE	LPS
Nonparametric	1	-3.25
AR(1)	1.033	-3.7
Augmented AR(1)	1.021	-3.6

Notes: This table reports the average Root-Mean-Squared-Error (RMSE) and Log Predictive Score (LPS) for one-year-ahead forecasts computed using the empirical non-parametric, AR(1) and augmented AR(1) specifications (see the text for details). The RMSE values are normalized to equal 1 for the nonparametric model.

2.4 Firm values

In any firm optimization problem, the firm makes decisions to maximize its continuation value. As we will show later, exit decisions will be central in understanding the results of our comparative statics in a heterogeneous firms model. Yet, before going to the model, we exploit the empirical findings from Section 2.3 to construct a distribution of firm values, defined as the present discounted value of the expected stream of lifetime revenue. This exercise is solely based on information extracted from the data and, hence, is not model-dependent.

Let $W(y)$ denote the value of the firm, defined as the expected lifetime stream of revenue. It can be characterized by the following Bellman equation

$$W(\hat{y}) = \hat{y} + \left(\frac{1 - \mathbb{P}(\text{Exit}|\hat{y})}{R}\right) \int W(\hat{y}')dH(\hat{y}'|\hat{y}). \tag{6}$$

where $\frac{1}{R}$ denotes the firms' discount rate. That is, the firm value equals current (normalized) revenue, $\hat{y}$, plus the firm's continuation value, which is itself pinned down by expected revenue conditional on not exiting.[8]

The empirical value function $W(\hat{y})$ at each grid point can be obtained iteratively by relying on two objects that we extracted earlier from the data: the transition distribution $H(\hat{y}'|\hat{y})$ and the exit hazard $\mathbb{P}(\text{Exit}|\hat{y})$. To construct the ergodic distribution of firm values, all we need in addition to $W(\hat{y})$ is the ergodic distribution of the state variable $\hat{y}$ which can be computed as:

$$H(\hat{y}') = \int H(\hat{y}'|\hat{y})\,(1 - \mathbb{P}(\text{Exit}|\hat{y}))\,dH(\hat{y}) + \mathbb{P}(\text{Exit})H_E(\hat{y}') \tag{7}$$

where $\mathbb{P}(\text{Exit}|\hat{y})$ and $H_E(\hat{y}')$ match exactly their data counterparts. Note that for the last term, we rely on the fact that in steady state, $\mathbb{P}(\text{Exit}) = \mathbb{P}(\text{Entry})$.

To begin, Figure 5 plots the ergodic distribution of the lifetime revenue value $W(\hat{y})$ in the empirical versus the one implied by AR(1) incumbent transitions. Note that in order to highlight the role played by differences in the dynamics of incumbents, we feed to the AR(1) model the empirical entry/exit distribution. Hence, any change in the resulting lifetime revenue distribution is due to the differences in the incumbent transitions.

Figure 5 highlights our key empirical result: the empirical distribution of firm values is more clustered than the one implied by a standard AR(1) model, particularly to the left of the median: for lower firm values, the slope and height of the density function is much higher and steeper in the empirical version. This is particularly striking when we recall that the non-parametric and AR(1) ergodic distributions of revenue were found earlier to be quite similar (see Figure 1). This apparent disconnect is a consequence of the strong differences between the *transition dynamics* implied by an AR(1) process and those found in the data.

[8]Note that we rely on revenue, not cash flow as in a standard value function. As discussed earlier, the revenue variable is generally more populated in the data and, as shown in Appendix A.1, a better predictor of firms' exit and hiring decisions. Ultimately, we will show in Section 3 that the broad conclusions are similar once we analyze instead the present discounted value of cash flows from our structural model.

Figure 5: Ergodic distribution of firm values (PDV of revenue)

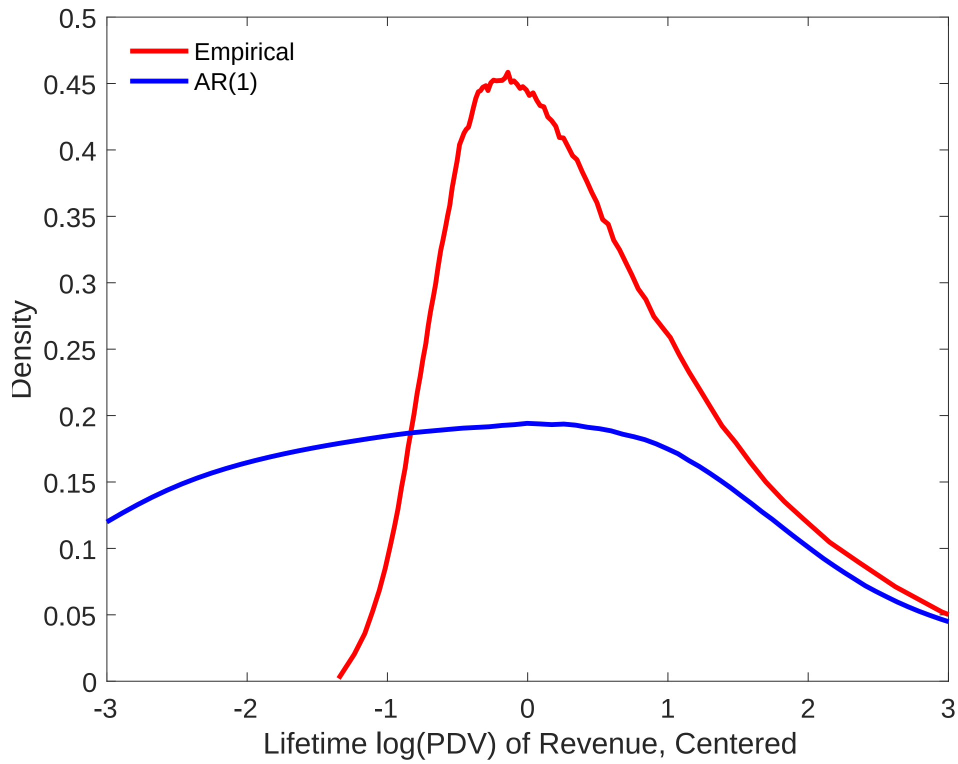

Notes: The figure depicts the density of the present discounted value (PDV) of firm revenue.

Finally, we present in Table 2 empirical that support the empirical validity of the firm values constructed above. Specifically, for each firm in our dataset, we regress its log market value as reported in the ORBIS dataset on current log revenue and our constructed lifetime revenue measure.

Table 2: Lifetime Revenue vs Market Value

	(1)	(2)	(3)	(4)
		log Market Value$_{it}$		
log Revenue$_{it}$	0.284***	0.141***	0.141***	-0.057
	(0.029)	(0.018)	(0.018)	(0.036)
log Lifetime Revenue$_{it}$				0.362***
				(0.076)
Fixed Effects	-	Ind.	Ind., Year	Ind., Year
Firm-Years	4273	4273	4273	4273

Note: The table reports OLS estimates of market value for Spanish firm i in year t on revenue and the expected lifetime revenue PDV measure. Ind. refers to four-digit industry codes. Unconditionally, corr(ln Rev., ln Mkt. Val.) = 0.24 and corr(ln Lifetime Rev., ln Mkt. Val.) = 0.27. Standard errors are clustered at the firm level. * = 10% level, ** = 5% level, *** = 1% level.

As is evident from Table 2, the contemporaneous log revenue (columns 1-3) is associated with the firm's market value. Yet, our constructed firm lifetime revenue variable is a much better predictor of the firm's log market value (column 4) and, once it is included in the regression, contemporaneous revenue ceases to be statistically significant. Overall we view these results as validating the empirical relevance of our firm value measure.

In sum, our results so far show that matching the distribution of revenue (or firm size) is not sufficient to ensure that the distribution of firm values, which is central to firm optimization, is close to its empirical counterpart.[9] As we illustrate in later sections, inadequately modeling this distribution of firm values has first-order consequences for the predictions of heterogeneous firms models.

[9] We note that this conclusion also applies to models that solely rely on permanent firm types in order to match the cross-sectional stationary distribution, such as in the trade literature: matching the size distribution of firms while disregarding the role of transition dynamics will *not* necessarily imply that the firm value distribution is adequately modeled.

3 Model

We begin with a stylized illustration of the main intuition built around the concept of distribution of firm values. We then proceed by describing our quantitative model.

3.1 A stylized example

To gain some intuition about the impact of the distribution of firms values on different comparative statics, we consider in Figure 6 a stylized example. Specifically, panel (a) depicts two distinct probability density functions of firm value: one labeled "Empirical" and the other "Parametric". The vertical line indicates the continuation value that makes a firm indifferent between exiting and continuing operating. The two distributions are constructed such that the fraction of exiters (the overall mass to the left of the vertical lines) is identical and equal to 10%. Note that the empirical distribution is more clustered around this exit region, in line with our empirical findings.

Figure 6: Stylized example of the impact of a subsidy

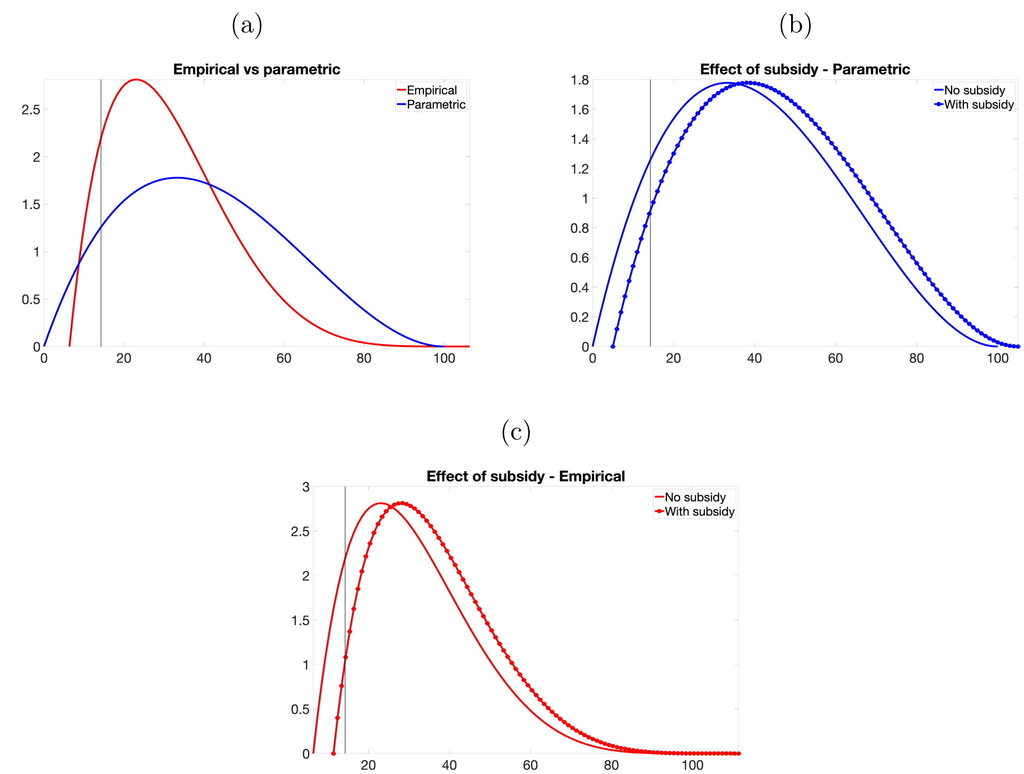

Next, consider the introduction of a government subsidy that is given to all operating firms. The subsidy simply shifts the distribution of firm values horizontally to the right. Such shifts are depicted in panel (b) for the more disperse (parametric) distribution, while panel (c) illustrates it for the bunched (empirical) one. As is clear visually, the mass of firms to the left of the vertical line, i.e. the exit rate, falls much further in the case of panel (c) than panel (b). As we will formally show later on, this sensitivity of the exit rate is directly linked to the level of the probability density function (pdf) in the area of the support of firm values where exit occurs.

Naturally, in our full general equilibrium model other active forces would vary across the two distributions in response to such a comparative change. Next, we describe and solve such framework.

3.2 A general equilibrium model with heterogeneous firms

Our model is in the spirit of Hopenhayn (1992), Hopenhayn and Rogerson (1993) and the literature that followed them. We first describe its various components before discussing calibration in the next section.

Types of firms The economy is populated by three types of firms: incumbents, entrants, and exiters.

The mass of operating firms in a given period is composed of incumbents and entrants and is denoted by M_O. Each operating firm produces a homogeneous numeraire output good and hires undifferentiated labor at a competitive equilibrium wage. Each period, it faces an idiosyncratic profitability shock z and must pay a fixed cost of operation. The shock z follows a first-order Markov chain with transition distribution $F(z'|z)$ and stationary distribution $F(z)$.

Entrants, whose mass is denoted by M_E, decide to enter by comparing the sunk entry cost to the expected value from producing. Upon entry, they draw a profitability level from a predetermined distribution and must produce for at least one period.

Finally, each period an operating firm decides whether to cease operations by comparing the fixed operating cost to the expected continuation value, which is itself a function of the current profitability shock and its expectation about the stochastic evolution of these shocks.

Dynamic problem of operating firms Next, we formally define the firms' dynamic problem and the household side of the model, before specifying the timing in this economy and its stationary equilibrium.

Specifically, consider an operating firm faced with an idiosyncratic profitability shock z; stochastic fixed costs iid distributed $\phi_F \sim G(\phi_F)$; an option to exit with value 0; an exogenous exit hazard $\delta \geq 0$; a decreasing-returns-to-scale production function with labor input n, which is given by zn^α; and the wage, W, which it takes as given. Then, its dynamic problem is given by

$$V(z) = \max_n \left[zn^\alpha - Wn + \mathbb{E}_{\phi_F} \max \left\{ 0, -\phi_F + \beta(1-\delta)\mathbb{E}\left(V(z')|z\right) \right\} \right]. \tag{8}$$

where next period is discounted at the discount rate β.

We note that this decreasing returns-to-scale production function specification is isomorphic to a monopolistic competition framework with love of variety. None of our results depend on which of the two formulations is chosen. This is also why while we could refer to z as "productivity" for the sake of simplifying the exposition and consistency with the literature, we prefer to use the broader term "profitability" for the z process.[10]

Entry A mass of potential entrants can pay a sunk cost $\phi_E > 0$ for a draw $z \sim F_E(z)$ to become incumbents. Free entry implies that with positive entry ($M_E > 0$)

$$\phi_E = \int V(z) dF_E(z). \tag{10}$$

where we note that firms are drawing from the entrants' distribution of profitability shocks, $F_E(z)$. Upon entry, firms produce for at least one period before deciding whether they wish to remain or exit.

Exit At the end of the period, after production has occurred, each firm draws a new operating cost, $\phi_F > 0$. The firm then optimally decides to exit if

$$\phi_F > \beta(1-\delta) \int V(z') dF(z'|z)$$

The exit condition above implies that for each value of z, there exists a threshold operating cost $\phi_F^*(z)$ defined as

$$\phi_F^*(z) = \beta(1-\delta) \int V(z') dF(z'|z) \tag{11}$$

such that exit occurs if $\phi_F > \phi_F^*(z)$. In addition, firms face an exogenous probability to exit δ. This feature allows us to match the fact that empirically, even very large firms face a positive probability of exit.

Households The economy is populated by a measure one of identical households. Households consume, pay taxes (if applicable) and supply labor inelastically in the total amount $\bar{N}$. In addition to labor income, they also receive dividends from operating firms. The representative household chooses its consumption to maximize log utility. Its dynamic problem is simply given by

$$S = \max_C \{\log(C) + \beta S'\} \tag{12}$$

with time discount rate $0 < \beta < 1$, subject to a standard budget constraint.

[10] In order to match the revenue dynamics we do *not* need to take a stand on whether the driving shocks are supply or demand shocks; the revenue function in such a model is given by

$$Revenue = z^{\nu-1} \daleth^{\nu} \text{Aggregates}, \tag{9}$$

where z and $\daleth$ respectively denote idiosyncratic productivity shocks and demand shocks.

3.3 Timing

To summarize, the timing of the model is as follows:

1. New entrant firms pay entry costs.
2. Incumbent firms and new entrants in period t receive an idiosyncratic productivity draw z. The entrants draw their productivity shocks from the distribution $F_E(z)$, while the incumbents draw them according to the transition distribution $F(z'|z)$.
3. The operating firms (incumbents and entrants) produce using z and labor n which is hired at the prevailing wage W.
4. Operating firms draw a fixed cost $\phi_F \sim G(\phi_F)$.
5. Operating firms form expectations of continuation values and choose either to exit at the end of period or to remain in operation for next period. Operating firms that choose to remain pay the fixed cost drawn.
6. A fraction δ of operating firms exogenously ceases to exist.
7. Surviving operating firms become incumbents for the next period.

3.4 Stationary equilibrium

A stationary equilibrium is a value function $V(z)$, threshold function $\phi_F^*(z)$, a distribution of operating firms $F_O(z)$, a mass of operating firms M_O, a mass of entrant firms M_E, aggregate output Y, aggregate labor N, aggregate fixed costs ϕ_F, aggregate sunk costs ϕ_E, aggregate fixed subsidy costs S^F, aggregate entrant subsidy costs S^E, and a wage W, such that the following conditions hold.

1. **Optimality of firms' decisions**: Taking as given the wage, W, and intertemporal price, p, together with the distribution of fixed costs $G(\phi_F)$ and the transition distribution for productivity z, operating firms optimize according to Equation 8.
2. **Optimality of exit decisions**: The fixed cost thresholds ϕ_F^* satisfy intertemporal optimization for the firms according to
$$\phi_F^*(z) = \beta(1-\delta)\int V(z')dF(z'|z)$$
3. **Stationarity of the distribution**: The operating distribution F_O replicates itself across periods according to
$$M_O F_O(z) = (1-\delta)M_O \int G(\phi_F^*(z_{-1}))F(z|z_{-1})dF_O(z_{-1}) + M_E F_E(z)$$

4. **Stability of operating firms**: The mass of operating firms is stable across periods, satisfying

$$M_O = (1-\delta)M_O \int G(\phi_F^*(z_{-1}))dF_O(z_{-1}) + M_E$$

5. **Free entry condition**: The free entry condition holds for the mass of entrants M_E, i.e.,

$$(\phi_E - s_E) \geq \int V(z)dF_E(z),$$

with equality whenever there is positive entry, $M_E > 0$.

6. **Household intertemporal Euler condition**: The representative household's intertemporal optimality implies,

$$R = \frac{1}{\beta}$$

7. **Definition of aggregate output**: Output Y is given by aggregation of output across the operating firm distribution

$$Y = M_O \int y(z, W)dF_O(z)$$

8. **Aggregate labor market clearing**: Labor N is given by aggregation of labor across the operating firm distribution

$$N = M_O \int n(z, W)dF_O(z)$$

and is equal to inelastically supplied labor supply $\bar{N}$.

9. **Definition of aggregate fixed cost**: The fixed cost aggregate ϕ_F reflects the mass of operating firms and their continuation decisions

$$\Phi_F = M_O \int \int_{\{\phi_F \leq \phi_F^*(z)\}} \phi_F dG(\phi_F)dF_O(z)$$

10. **Definition of aggregate sunk entry cost**: The sunk entry cost aggregate reflects the mass of entrant firms

$$\Phi_E = M_E \phi_E$$

11. **Aggregate resource constraint**: The resource constraint or goods market clearing conditions is satisfied

$$Y = C + \Phi_F + \Phi_E$$

4 Calibration

In this section, we describe our calibration strategy for both the empirical/non-parametric and parametric AR(1) models.

4.1 Non-parametric model

To guarantee consistency with the distributional dynamics found in the data in Section 2, there are three objects which we must match: (i) the revenue transition distribution, $H(\hat{y}'|\hat{y})$, (ii) the entrants' revenue distribution, $H_E(\hat{y})$, and (iii) the exit hazard $\mathbb{P}(\text{Exit}|\hat{y})$. We match these empirical patterns by manipulating the following model objects: (i) the profitability shock transition distribution $F(z'|z)$, (ii) the entrants' profitability distribution $F_E(z)$, (iii) the distribution of fixed cost shocks $G(\phi_F)$, and (iv) the exogenous exit rate δ.

The crucial step in mapping the observed revenue dynamics into statements about the driving profitability shock z it to note that in this specification of the model, there is a simple inversion from the profitability shock z to sales y based on the static labor optimality condition; incorporating the labor demand obtained from the static optimization problem

$$\max_{n} zn^{\alpha} - Wn$$

into the production function $y = zn^{\alpha}$, it follows that

$$\log z = (1 - \alpha) \log y + \text{Constant}. \tag{13}$$

Hence, all our findings regarding the dynamics of residualized log revenue $\hat{y}$ are also statements about the driving profitability shock in the model, $\log z$, up to a normalizing constant. Given the one-to-one mapping between z and $\hat{y}$, it follows that the exogenous transition distribution for incumbents ($F(z'|z)$) and the exogenous profitability distribution of entrants ($F_E(z)$) are direct analogues of the estimates obtained in Section 2, $H(\hat{y}'|\hat{y})$ and $H_E(\hat{y})$. We then match the empirical exit hazard $\mathbb{P}(\text{Exit}|\hat{y})$, which has an endogenous model counterpart $\mathbb{P}(\text{Exit}|z)$, by non-parametrically computing the fixed cost distribution $G(\phi_F)$ satisfying the identity $1 - (1 - \delta)G(\phi_F^*(z)) = \mathbb{P}(\text{Exit}|z)$ within the model, taking into account the firm continuation values $\phi_F^*(z)$ optimally implied by the full structure of the model in equilibrium. While computing $G(\phi_F)$, we assume that the largest firms in the data exit for only exogenous reasons, an assumption allowing us to recover $\delta = \mathbb{P}(\text{Exit}|\hat{y}_{N_y})$ directly from the data. We refer the reader to Appendix B for further details on our solution of the model and our associated recovery of the distribution $G(\phi_F)$.

By solving the non-parametric model in this manner, we recover a general equilibrium in which incumbent revenue transitions, the entrant revenue distribution, and the exit hazard within the model *exactly* match their empirical counterparts. Consequently, we also match all empirical moments of these distributions, e.g., the unconditional exit rate, the distribution of firm lifetime revenue values $W(\hat{y})$, the relative size of entrants, the auto-correlation and variance of revenue, etc.

Four parameters, α, β, ϕ_E, and $\bar{N}$, remain to be calibrated. We internally choose these parameters to match the mean labor share (equal to α), the mean real interest rate (pins down β), the employment rate (equal to $\bar{N}$) and mean employees per firm (pins down ϕ_E).[11] The column denoted "Empirical" in Table 3 reports these values.

4.2 Parametric AR(1) model

To calibrate the AR(1) model, we follow the standard approach in the literature. We assume that the shock z follows

$$\log z' = \rho \log z + \sigma \epsilon', \quad \epsilon' \sim N(0,1), \tag{14}$$

where ρ and σ are chosen to match the auto-correlation and unconditional variance of log revenue for incumbents. We assume the fixed operating cost ϕ_F is drawn from a uniform distribution

$$G(\phi_F) = U(0, \bar{\phi}_F), \tag{15}$$

Finally, we model the entry distribution as

$$\log z \sim N(\mu_E, \sigma^2). \tag{16}$$

In Appendix B, we plot the exit hazard and the entry distribution for the two models, parametric versus non-parametric.

We are left with three parameters to calibrate: $\bar{\phi}_F$ and μ_E, which are specific to this version of the model, as well as ϕ_E and $\bar{N}$ as above. They are chosen to jointly match the following three moments: the number of employees per firm; the average exit rate; and the mean difference in log revenue between entrants versus incumbents. The column denoted "AR(1)" in Table 3 reports these values.[12] In addition, the coefficients for the $\log z$ process are $\rho = 0.9415$ and $\sigma = 0.1908$.

[11] The two first parameters are commonly calibrated in the macroeconomics literature. With respect to, ϕ_E, we note that "the mean employees per firm" is informative about it because increases in this sunk cost lead to a higher firm value via the free entry condition. This rise in the firm value must be accompanied by a decline in the wage rate, which raises the firm-intensive margin of labor demand and thus affects the mean employees per firm. And setting the value of total labor supply $\bar{N}$ to equal the aggregate employment rate is essentially a normalization defining the units of labor.

[12] Three parameters (α, β and δ) are common to the two models and set externally, before solving the model.

Table 3: Values of the calibrated parameters under both models

Moments		Parameters	
Description	Value	Empirical	AR(1)
Labor share	2/3	$\alpha = 2/3$	$\alpha = 2/3$
Real interest rate	4%	$\beta = 1/1.04$	$\beta = 1/1.04$
Exogenous exit rate	3.8%	$\delta = 0.038$	$\delta = 0.038$
Average exit rate	6.9%	(by construction)	$\bar{\phi}_F = 2.3$
Employment rate	59.7%	$\bar{N}$	$\bar{N}$
Employees per firm	12.3	$\phi_E = 22.9$	$\phi_E = 5.18$
Relative profits of entrants	-26.7%	(by construction)	$\mu_E = -0.44$

4.3 Model-implied distribution of firm values

Before turning to the comparative statics, we plot in Figure 7 the distribution of firm values (that is, the continuation values $\int V(z)dF_E(z)$) under the two calibrated models. The difference is striking and the message is in line with that from the model-free, revenue-based lifetime value estimates that we obtained earlier (see Figure 7): at low firm values, where most exit occurs, the model-implied distribution is much more "bunched" (the probability density function is much higher and steeper) in the empirical than the AR(1) version. As we show next, the shape of the distribution of firm values plays a crucial role in explaining the differences between the two models' response to aggregate disturbances.[13]

[13]We note that in Figure 5 we plotted the lifetime firms revenues under the non-parametric process (the red line) against the one generated when only the incumbent transitions are assumed to follow an AR(1) process. In contrast, in the calibration of the AR(1) model, we calibrate the exit process and entry distributions in a way that is comparable with the approaches used in the existing literature that follows parametric assumptions.

Figure 7: Ergodic distribution of firm values in the benchmark models

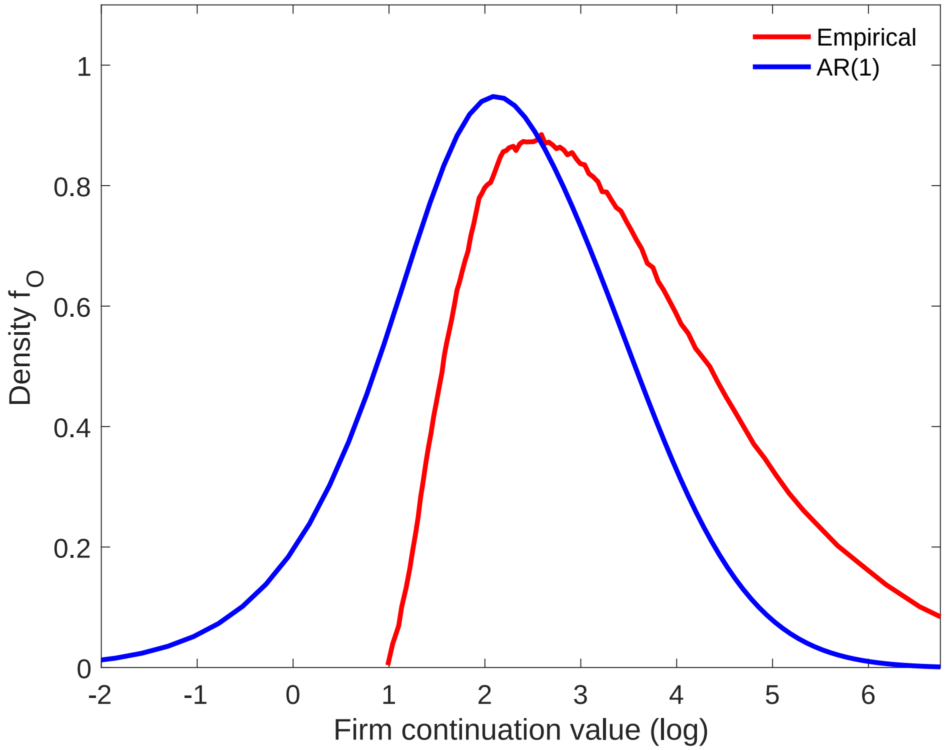

Notes: This figure depicts the ergodic distribution of firms PDV in the benchmark non-parametric (red line) and AR(1) models (blue line).

5 Comparative statics

Our goal in this section is to assess the quantitative relevance of our empirical findings for the predictions of heterogeneous firms macroeconomic models. We consider the impact on the steady state of two simple experiments: (i) a subsidy to all operating firms, and (ii) a subsidy only to new entrants. These experiments mirror the changes in the fixed operating cost and the sunk entry studied in Hopenhayn (1992). Under each scenario, we compute and analyze the response of various aggregates, taking into account general equilibrium effects.

Via these experiments we show that our empirical findings have, quantitatively, a first order impact on the response of aggregate to changes in the economic environment. In essence, and with an abuse of a definition, we show that our non-parametric model, changes the "elasticity" of the model; this elasticity depends on the interactions between the various forces at play and, in particular, the selection effect. Thus, the change in the elasticity has a significant impact on the transmission of an aggregate perturbance on the economy. Specifically, we show that (i) this elasticity change has a first order impact quantitatively, and moreover that (ii) whether it magnifies or attenuates the aggregate perturbation depends on the exact experiment we consider; hence this change in elasticity alters the relative responses across the two models and the two experiments we consider.

5.1 Subsidy to fixed operating costs

In our first experiment, we introduce a fixed subsidy that is given to all firms at each period. The subsidy, which ultimately lowers the fixed operating cost, is financed through lump sum taxes on households.

5.1.1 Basic mechanism

In both models, the mechanism behind the response of the economy is qualitatively similar. We start by providing a general overview of the forces at play.

First, the subsidy, s_F, naturally raises the value of every operating firm, $V(z, W, \phi_F - s_F)$. This leads to a decrease in the exit rate and a negative selection effect as lower-z firms now survive. Consider next the expression for the free-entry condition:

$$\phi_E = \int V(z, W, \phi_F - s_F) dF_E(z)$$

The entry sunk cost ϕ_E on the left-hand side is constant in this experiment. Therefore, in order to satisfy the free-entry condition, the wage W has to increase to counter the rise in firm value from the fixed cost subsidy s_F.[14]

The equation below, which characterizes the labor market clearing condition, has in turn implications for the mass of operating firms. This equation shows that the overall labor employed in the economy can be characterized as a product of the mass of operating firms, M_O and the average labor per firm, $\int n(z, W) dF_O(z)$,

$$N = M_O \int n(z, W) dF_O(z).$$

The higher wage depresses labor demand at each level of z, $n(z, W)$. In addition, the negative selection pushes down average labor per firm.[15] Together, these two forces lower the term under the integral. Since the labor level N is fixed in equilibrium given fixed labor supply, the mass of operating firms, M_O, must rise with the subsidy to compensate for the fall in the number of employees per firm.

In this economy with fixed labor supply, the wage, output and TFP must move proportionately.[16]

[14] Recall that firms enter based on an expected continuation value: only after entry do they learn their productivity level, produce and then can choose to exit. For this reason, there is no selection through entry.

[15] Formally, with one cutoff exit this would manifest itself as a reduction in the lower bound of the integral. Given that we have exit occurring at different levels of z, due to the randomness in the fixed cost of operation, this argument maps to lower levels of z surviving for a given level of the fixed cost of operation.

[16] To see why, note that the optimality condition of the firm dictates that $n = \alpha y/W$. Integrating over the mass of operating firms, we therefore obtain in the aggregate that $Y = \frac{WN}{\alpha}$, where N is equal to the fixed labor supply in equilibrium. Moreover, $TFP = \frac{Y}{N^\alpha}$.

Hence, it follows that TFP increases with the fixed cost subsidy, despite the negative selection effect through lower exit; this is because the economy is now composed of a larger mass of operating firms, but of smaller size; through decreasing returns to scale, the presence of smaller firms pushes up TFP.

5.1.2 Model comparison

Despite a similar underlying mechanism, the two models produce quantitatively very different responses. This can be readily observed from Figure 8, which depicts the response of various aggregate variables to different levels of the firm subsidy (on the x-axis, as a percentage of the fixed operating cost).[17] The response of output, for example, is about 45% larger in the parametric AR(1) model than under the non-parametric specification. As we show next, the key reason behind this discrepancy is the role of selection through the response of the exit rate, which is much more pronounced in the non-parametric case.

Figure 8: Impact of a fixed operating cost subsidy

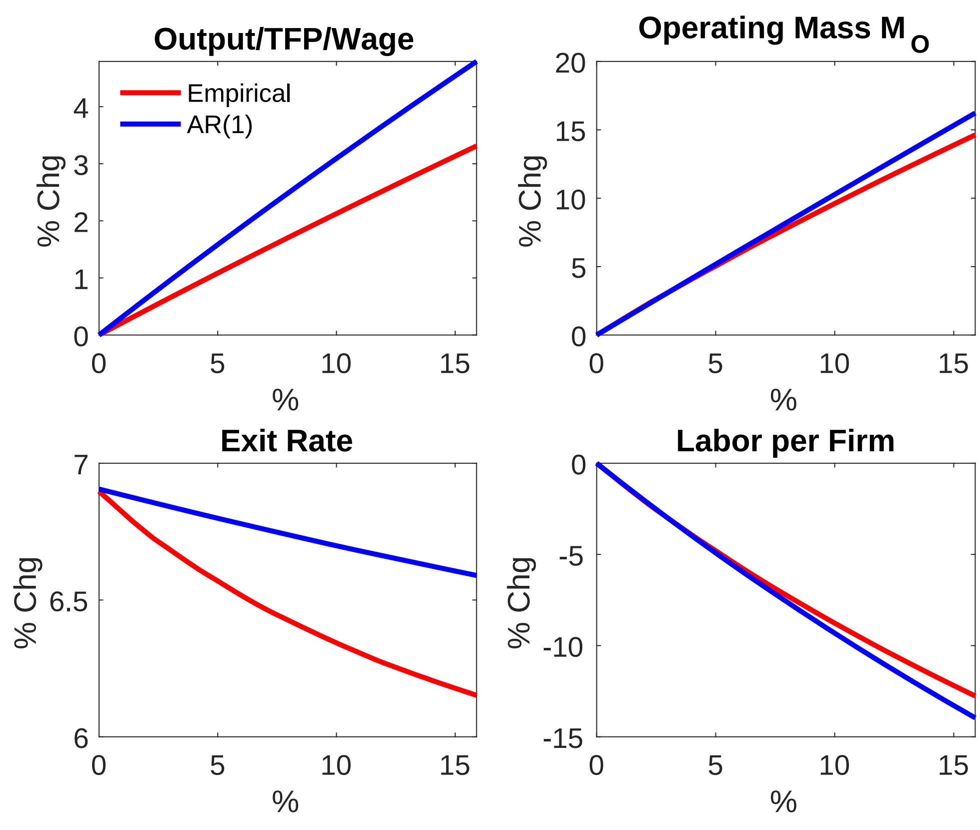

Note: This figure depicts the impact of a fixed operating cost subsidy in the non-parametric model (red line) and the AR(1) models (blue line).

To gain intuition, we turn to Figure 9. Both panels consider the partial equilibrium response of

[17]Note that the size of the fixed cost as a fraction of output is the same in the two economies.

the economy (by holding the wage fixed) following the fixed cost subsidy. Specifically, the panels plot three main objects following the subsidy introduction: the original distribution of firm value (blue), the distribution following a 15% fixed cost subsidy (green) and the exit hazard (red). In addition, the vertical bar depicts the average firm value implied by the corresponding value distribution. The top and bottom panels correspond to the non-parametric and parametric specification, respectively.

Figure 9: Value distributions and exit hazards - Fixed operating cost subsidy

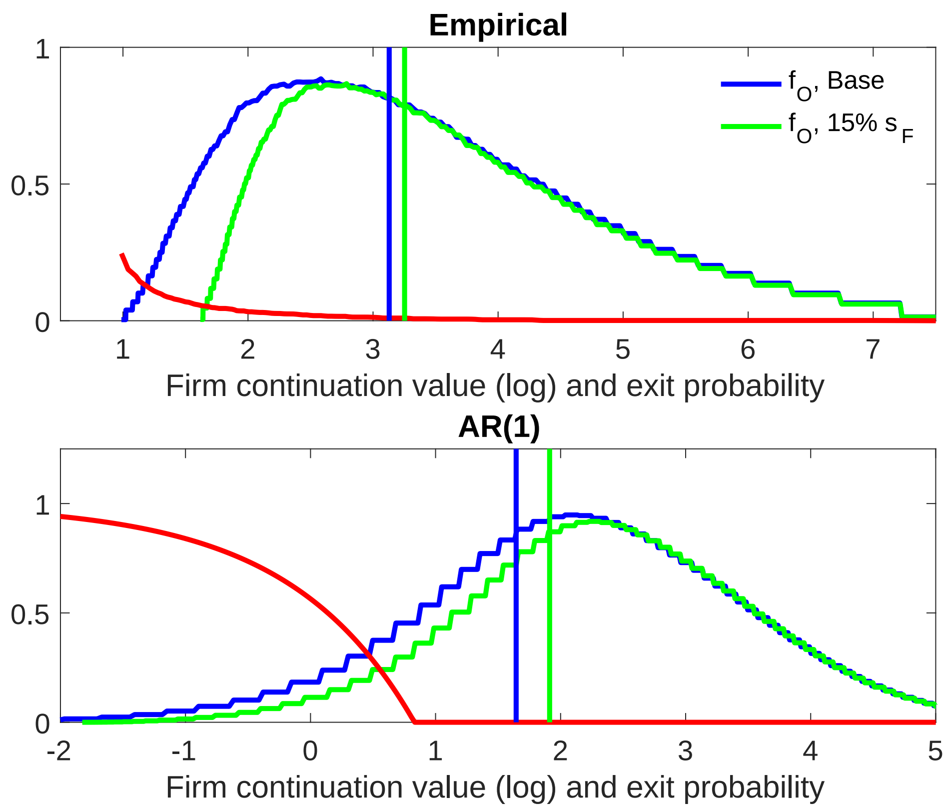

Note: For each model: original distribution of firm value (blue), distribution following a 15% fixed cost subsidy (green), exit hazard (red) and average firm value (blue and green vertical bars).

The subsidy leads to a rightward shift of the value distribution. As a result, the average firm value rises in both models, an increase represented by the horizontal shift from the blue to the green vertical bars. The increase, however, is much more pronounced in the parametric case, even though the subsidy is the same in both cases. The reason has to do with the very different impact of selection in the two models: as discussed above, in both models the subsidy leads to a lower rate of exit and the survival of more low-productivity firms. This negative selection effect, however, is more acute for the non-parametric model: through their interaction, the shapes of the exit hazard and value distribution imply a stronger decline of the exit rate in the empirically-relevant version and therefore a *larger negative selection effect* which, in turn, significantly reduces the increase in the average firm value. In contrast, the selection effect is much weaker in the AR(1) specification, allowing average firm value to rise by more thanks to the subsidy. This implies different wage responses between the two models: in order for the free-entry condition to hold, the wage increase has to be more pronounced in the parametric model since the fixed-wage average firm value rises by more. This is what can be observed in Figure 8.

As discussed earlier, both the rise in the wage and negative selection reduce the average labor per firm. Since the labor supply is fixed, the implication is that the reaction of the mass of operating firms is also comparable: in both cases, M_O increases sharply with the size of the subsidy, where re remind the reader that the labor market clearing condition is given by,

$$N = M_O \int n(z, W) dF_O(z).$$

But how does the response differ across models? On the one hand, the larger wage increase in the parametric model implies relatively lower demand for labor on average across operating firms. On the other hand, the negative selection effect, which pushes down average labor per firm, is less pronounced. It turns out that for this exercise, these two opposing forces work in a way that they almost perfectly cancel each other (bottom-right panel of Figure 8). Hence, the impact of the subsidy on the integral in the equation above is very similar across the two models, leading to similar increases in M_O (top-right panel of Figure 8).

We are now equipped with all the necessary elements to understand the forces behind the response of TFP and therefore output, which are proportional to each other. For each model, we can decompose the response of aggregate TFP into two margins. Formally, TFP in this model can be written as:

$$TFP = \underbrace{M_O^{1-\alpha}}_{\text{Mass of operating firms}} \underbrace{\left(\int z^{\frac{1}{1-\alpha}} dF_O(z) \right)^{1-\alpha}}_{\text{Selection}}. \tag{17}$$

In Figure 10, we plot the contribution of each of these two components to the change in TFP (red line). As discussed above, the response of the mass of operating firms to the subsidy is very similar across the two models, and this is confirmed in the figure (see the line depicting M_O). Selection, however, has a much more pronounced negative impact in the non-parametric model, due to the stronger fall in the exit rate. This stronger negative selection effect in the non-parametric model leads to a muted response in TFP and output vis-a-vis the parametric model.

Figure 10: Aggregate TFP Decomposition - Fixed operating cost subsidy

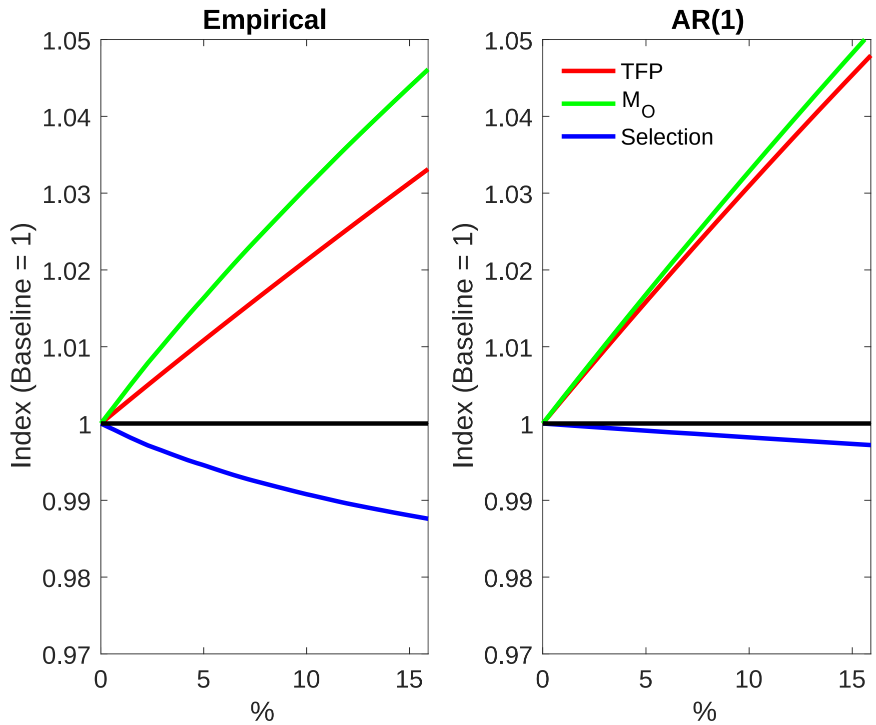

Notes: This figure depicts the contribution of various margins to the change of TFP in the two models under various values of the fixed operating cost subsidy.

5.2 Entry subsidy

Our next comparative experiment is one where a subsidy (financed by lump sum taxes) is given to entrants. As we discuss below, the relative aggregate response across the two models is flipped vis-a-vis the case of the fixed operating cost subsidy analyzed in the previous section.

5.2.1 Basic mechanism.

Consider the expression for the free-entry condition:

$$\phi_E - s_E = \int V(z, W, \phi_F) dF_E(z)$$

The subsidy s_E effectively lowers the cost of entry. Hence, for the free-entry condition to be satisfied, the wage W must increase to bring V down such that the right-hand side (the value of entry) is lower. Higher wage and lower V, in turn, makes exit more likely for low-z operating firms, generating a positive selection effect.

Next, let us turn to the labor market clearing condition:

$$N = M_O \int n(z, W) dF_O(z)$$

The higher wage depresses labor demand at each firm, $n(z, W)$. On the other hand, negative selection pushes up average labor per firm. Under our benchmark parameterization, the wage effect is stronger, lowering the term under the integral. Since the labor supply N is fixed, the implication is that the mass of operating firms, M_O, must rise with the subsidy to compensate for the fall in the number of employees per firm.

Overall, the economy is now composed of a larger mass of operating firms, accompanied by a fall in average firm size, and a positive selection. Recalling that TFP is given by

$$TFP = \underbrace{M_O^{1-\alpha}}_{\text{Mass of operating firms}} \underbrace{\left(\int z^{\frac{1}{1-\alpha}} dF_O(z)\right)^{1-\alpha}}_{\text{Selection}},$$

both the increase in the mass and the positive selection lead to an increase in TFP, and hence output.

5.2.2 Model comparison

Figure 11 plots the response of various aggregates to the entry subsidy under both models, expressed as a fraction of the sunk entry cost ϕ_E. While in the previous exercise the response of output under the parametric model was stronger, in this experiment the response is in fact close to 50% larger in the non-parametric specification. As we show next, selection plays again a central role.

Figure 11: Impact of a subsidy to entry

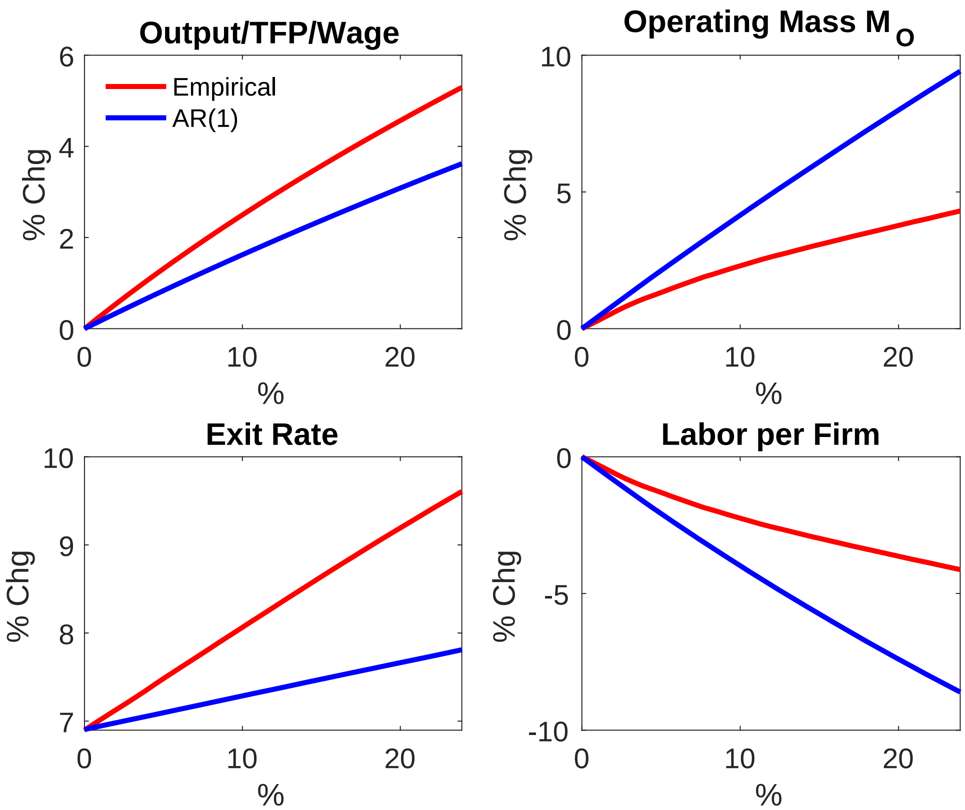

Note: This figure depicts the impact of an entry subsidy (as a portion of the sunk entry cost) in the non-parametric model (red line) and the AR(1) models (blue line).

Figure 12 plots for both models the original distribution of firm value (blue); the distribution following a 5% wage increase (green); the average firm value, before and after the subsidy; and the exit hazard (red). The wage increase lowers the value of the average operating firm which pushes, in both models, lower-z firms to exit. However, due to the shape of the underlying value distribution, the exit rate is more sensitive in the non-parametric model. This stronger positive selection effect undoes some of the wage effect, which explains why the average firm value falls by less. The implication is that for the average expected value of entry to match the drop in the net cost of entry, the wage has to rise by more under the empirically-relevant specification.

Figure 12: Value distributions and exit hazards - Entry subsidy

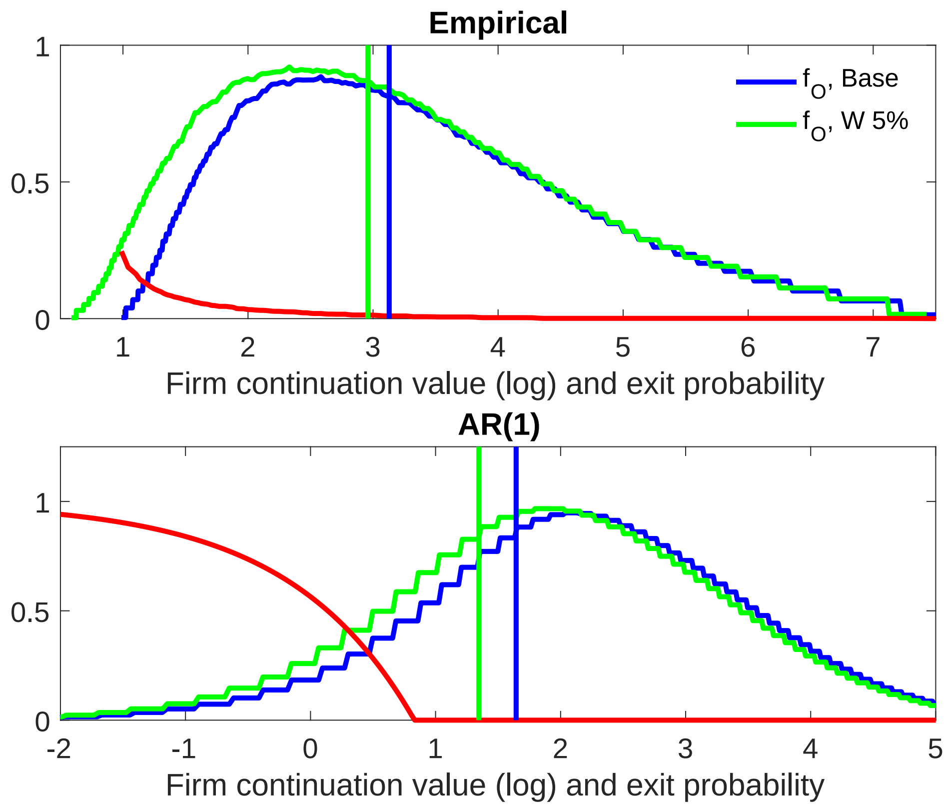

Note: For each model: original distribution of firm value (blue), distribution following a 5% increase in the wage (green), exit hazard (red) and average firm value (blue and green vertical bars).

To understand the relative change in the mass of operating firms between the two models, we turn our attention to the labor market clearing condition:

$$N = M_O \int n(z, W) dF_O(z)$$

The higher response of the wage in the empirical model depresses more drastically labor demand at each firm. The selection channel, however, has the opposite effect: with relatively more low-z firms exiting in the empirical model following the subsidy, average labor demand is higher in this specification, all else equal. The net effect is therefore ambiguous. In our parameterization, it turns out that the average labor demand per operating firm falls more with the subsidy in the non-parametric model (bottom-right panel of Figure 12). As a result, the mass of operating firms M_O increases by more, to ensure that labor demand equals labor supply (top-right panel).

Finally, we plot in Figure 13 the contribution of each of selection and the mass of operating firms to the change in TFP (red line). As discussed earlier, the shape of the distribution implies a stronger positive selection effect in the non-parametric specification. The operating mass, however, rises more starkly in the parametric version. On net, the selection channel is stronger, resulting in a larger impact of the subsidy on TFP and output in the empirical model.

Figure 13: Aggregate TFP decomposition - Entry subsidy

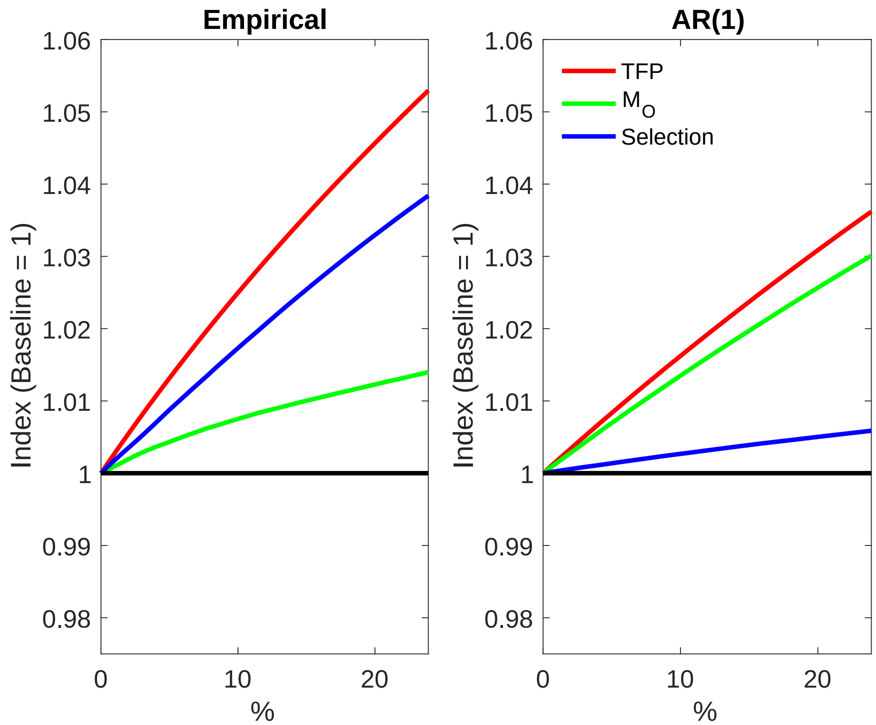

Notes: This figure depicts the contribution of various margins to the change of TFP in the two models under various values of the entry subsidy.

5.3 Robustness

Finally, we consider different robustness exercises for the fixed cost subsidy exercise. First, in terms of the modelling assumptions, we consider (i) other values of the labor share, α, and (ii) endogenous labor supply. Second, we repeat the analysis for additional countries (Italy, Portugal, France, and Norway, which are four other countries with excellent ORBIS coverage) and other treatments of the data, such as subsamples or different handling of outliers, eliminating MAs, etc. The results of these comparative statics, summarized in Table C.1 of Appendix C, show that our conclusions are robust to various alternatives.

6 The empirical relation between bunching and exit

The analysis so far has highlighted the relationship between the degree of clustering of the firm value distribution and the reaction of exit rates to variations in profitability. In this section we verify that this relation holds empirically by exploiting sectoral differences in the shape of the firm value distribution.

First, recall from above that we can obtain a distribution $H(W)$ of lifetime revenue W by relying the transition matrix of incumbents as well as the entry and exit profiles. Moreover, we can decompose exit as

$$\mathbb{P}(\text{Exit}) = \int \mathbb{P}(\text{Exit}|W)dH(W).$$

In this *purely statistical* model, the partial equilibrium, short-run exit rate predicted after a windfall increase ϵ in lifetime revenue W is given by

$$\int \mathbb{P}(\text{Exit}|W+\epsilon)dH(W).$$

The sensitivity of exit, which is linked to the degree of bunching of the firm value distribution, can then be feasibly computed as the statistic

$$\text{B} = -\frac{\partial \mathbb{P}(\text{Exit})}{\partial \epsilon}|_{\epsilon=0} = -\int \frac{\partial \mathbb{P}(\text{Exit}|W)}{\partial W}dH(W)$$

measuring coincidence of high density or bunching with steep exit hazards.

We compute our bunching statistic for the nonparametric lifetime revenue distributions of each 2-digit sector in our data. Table D.1 in Appendix D reports that this statistic varies widely across sectors. Construction, real estate, professional services and retail trade are characterized by much larger bunching statistics than health care, transportation or manufacturing.

6.1 Estimation

We then use industry data for year t for 4-digit industry j within 2-digit sector s to estimate versions of the following specification

$$\mathbb{P}(\text{Exit})_{jst} = \alpha + \beta\Delta\text{Revenue}_{jst} + \gamma\Delta\text{Revenue}_{jst} \times \text{B}_s + \delta\text{B}_s + \varepsilon_{jst}$$

where the model predicts $\gamma < 0$ if more bunching is linked to higher exit sensitivity. The maintained assumption is that the degree of bunching at the 4-digit level is relatively homogeneous within a given 2-digit sector and stable over our sample period.[18]

Results are presented in Table 4. The first column shows that high revenue growth at the 4-digit industry level is related to lower exit probabilities. The interaction term indicates that this negative relationship is stronger in sectors that feature a higher degree of bunching of the firm value distribution, consistent with the central mechanism of the model. The second and third columns show that this conclusion is robust to the inclusion of year fixed effects (column 2) as well as the addition of sectoral fixed effects (column 3). In both cases, the interaction term continues to be negative and statistically significant at the 10% level. In the last column, we replace the linear interaction assumption with

[18] Because output in our model is stationary while, naturally, it exhibits positive growth in the data, note that we are linking the exit rate to the *growth rate* of sectoral revenue, and not its level. This allows the empirical test to be consistent with the interpretation of the model.

a categorical approach: we define a high-bunching sector to be one that has a bunching statistic in the upper quartile of the bunching distribution. The interaction term is significant at the 5% level, emphasizing again that bunching affects the mapping between revenue growth and exit rate.

Indeed, as Table 4 shows across different specifications, industries featuring lifetime revenue distributions that are more bunched exhibit higher exit sensitivity to changes in revenue growth, consistent with our key model mechanism.

Table 4: Bunching and exit

	(1)	(2)	(3)	(4)
	Exit Rate$_{jt}$			
Δ Revenue$_{jt}$	-0.045***	-0.046***	-0.050***	-0.039***
	(0.008)	(0.012)	(0.011)	(0.011)
Δ Revenue$_{jt}$ $\times$ Bunching$_s$	-0.011*	-0.013*	-0.014*	
	(0.006)	(0.006)	(0.007)	
Bunching$_s$	0.428**	0.413**		
	(0.211)	(0.209)		
Δ Revenue$_{jt}$ $\times$ I(High Bunching$_s$)				-0.052**
				(0.022)
Fixed Effects	-	Year	Year, Sector	Year, Sector
Industry-Years	1584	1584	1584	1584
Years	2006-13	2006-13	2006-13	2006-13

Note: The table reports OLS estimates of Spanish four-digit industry-level exit rates on industry revenue growth and two-digit standardized sectoral bunching statistics. High bunching is an indicator for sectoral bunching above the 75% percentile. Standard errors are clustered at the industry level. * = 10% level, ** = 5% level, *** = 1% level.

7 Conclusion

In this paper, we show that the standard parametric assumption for firm-level shocks used in the macro heterogeneous firms literature is not an appropriate representation of the true data generating process. In particular, we find that non-parametrically solving a model consistent with the firm-level revenue dynamics we observe in the data has a large impact on the behavior of the model at the macro level. In particular, the standard parametric model implies a firm value distribution which is far too dispersed relative to the firm value distribution consistent with empirical firm dynamics. As a result, the empirical, non-parametric model generates substantially higher sensitivity of the exit rate to a set of standard policy experiments. The stronger extensive margin reaction in our non-parametric model also drives strong selection effects serving to amplify or dampen the response of aggregate output, depending upon the exact details of the underlying policy. As a result, we conclude that the standard parametric assumptions adopted in this model are far from innocuous but directly change the macro implications of firm-level mechanisms.

References

Asker, J., A. Collard-Wexler, and J. De Loecker (2014). Dynamic Inputs and Resource (Mis)Allocation. *Journal of Political Economy 122*(5), 1013–1063.

Axtell, R. L. (2001). Zipf Distribution of U. S. Firm Sizes. *Science CCXCIII*, 1818 –1820.

Bajgar, M., G. Berlingieri, S. Calligaris, C. Criscuolo, and J. Timmis (2020). Coverage and Representativeness of Orbis Data. Working paper.

Bento, P. and D. Restuccia (2017). Misallocation, Establishment Size, and Productivity. *American Economic Journal: Macroeconomics 9*(3), 267–303.

Bilbiie, F. O., F. Ghironi, and M. J. Melitz (2012). Endogenous Entry, Product Variety, and Business Cycles. *Journal of Political Economy 120*(2), 304–345.

Bils, M., P. J. Klenow, and C. Ruane (2021). Misallocation or Mismeasurement? *Journal of Monetary Economics 124*, S39–S56.

Bloom, N., M. Floetotto, N. Jaimovich, I. Saporta-Eksten, and S. J. Terry (2018). Really uncertain business cycles. *Econometrica 86*(3), 1031–1065.

Carvalho, V. M. and B. Grassi (2019). Large Firm Dynamics and the Business Cycle. *American Economic Review 109*(4), 1375–1425.

Catherine, S., T. Chaney, Z. Huang, D. Sraer, and D. Thesmar (2018). Quantifying Reduced-Form Evidence on Collateral Constraints. Forthcoming, *Journal of Finance.*

Clementi, G. L. and B. Palazzo (2016). Entry, Exit, Firm Dynamics, and Aggregate Fluctuations. *American Economic Journal: Macroeconomics 8*(3), 1–41.

Davies, R. B. and C. Eckel (2010). Tax Competition for Heterogeneous Firms with Endogenous Entry. *American Economic Journal: Economic Policy 2*(1), 77–102.

Davis, S. J. and J. Haltiwanger (1992). Gross Job Creation, Gross Job Destruction, and Employment Reallocation. *Quarterly Journal of Economics 107*(3), 819–863.

Decker, R. A., J. Haltiwanger, R. S. Jarmin, and J. Miranda (2016). Declining business dynamism: What we know and the way forward. *American Economic Review 106*(5), 203–07.

Decker, R. A., J. Haltiwanger, R. S. Jarmin, and J. Miranda (2020). Changing business dynamism and productivity: Shocks versus responsiveness. *American Economic Review 110*(12), 3952–90.

Dunne, T., M. J. Roberts, and L. Samuelson (1989). The Growth and Failure of US Manufacturing Plants. *Quarterly Journal of Economics 104*(4), 671–698.

Forneron, J.-J. (2020). A Sieve-SMM Estimator for Dynamic Models. Working paper.

Garicano, L., C. Lelarge, and J. Van Reenen (2016). Firm Size Distortions and the Productivity Distribution: Evidence from France. *American Economic Review 106*, 3439–3479.

Gourio, F. and J. Miao (2010). Firm Heterogeneity and the Long-Run Effects of Dividend Tax Reform. *American Economic Journal: Macroeconomics 2*(1), 131–68.

Guner, N., G. Ventura, and Y. Xu (2008). Macroeconomic Implications of Size-dependent Policies. *Review of Economic Dynamics 11*(4), 721–744.

Guvenen, F., F. Karahan, S. Ozkan, and J. Song (2021). What Do Data on Millions of US Workers Reveal About Lifecycle Earnings Dynamics? *Econometrica 89*(5), 2303–2339.

Hopenhayn, H. (1992). Entry, Exit, and Firm Dynamics in Long Run Equilibrium. *Econometrica 60*(5), 1127–1150.

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

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

Kalemli-Ozcan, S., B. Sorensen, C. Villegas-Sanchez, V. Volosovych, and S. Yesiltas (2022). How to Construct Nationally Representative Firm Level Data from the Orbis Global Database: New Facts and Aggregate Implications. Working paper.

Karahan, F., B. Pugsley, and A. Sahin (2022). Demographic Origins of the Startup Deficit. Working paper.

Kehrig, M. (2015). The Cyclical Nature of the Productivity Distribution. Working paper.

Kehrig, M. and N. Vincent (2020). Good Dispersion, Bad Dispersion. Working paper.

Khan, A. and J. K. Thomas (2008). Idiosyncratic Shocks and the Role of Nonconvexities in Plant and Aggregate Investment Dynamics. *Econometrica 76*(2), 395–436.

Khan, A. and J. K. Thomas (2013). Credit Shocks and Aggregate Fluctuations in an Economy with Production Heterogeneity. *Journal of Political Economy 121*(6), 1055–1107.

Lee, Y. and T. Mukoyama (2018). A Model of Entry, Exit, and Plant-level Dynamics Over the Business Cycle. *Journal of Economic Dynamics and Control 96*, 1–25.

Melitz, M. J. (2003). The Impact of Trade on Intra-industry Reallocations and Aggregate Industry Productivity. *Econometrica 71*(6), 1695–1725.

Midrigan, V. (2011). Menu Costs, Multiproduct Firms, and Aggregate Fluctuations. *Econometrica 79*(4), 1139–1180.

Midrigan, V. and D. Y. Xu (2014). Finance and Misallocation: Evidence from Plant-level Data. *American Economic Review 104*(2), 422–58.

Moll, B. (2014). Productivity Losses from Financial Frictions: Can Self-financing Undo Capital Misallocation? *American Economic Review 104*(10), 3186–3221.

Ottonello, P. and T. Winberry (2020). Financial Heterogeneity and the Investment Channel of Monetary Policy. *Econometrica 88*(6), 2473–2502.

Restuccia, D. and R. Rogerson (2008). Policy Distortions and Aggregate Productivity with Heterogeneous Establishments. *Review of Economic dynamics 11*(4), 707–720.

Sraer, D. and D. Thesmar (2021). A Sufficient Statistics Approach for Aggregating Firm-Level Experiments. Working paper.

Sterk, V., P. Sedláček, and B. Pugsley (2021). The Nature of Firm Growth. *American Economic Review 111*(2), 547–79.

Tauchen, G. (1986). Finite State Markov-chain Approximations to Univariate and Vector Autoregressions. *Economics Letters 20*(2), 177–181.

Winberry, T. (2021). Lumpy Investment, Business Cycles, and Stimulus Policy. *American Economic Review 111*(1), 364–396.

Appendix

A Data

A.1 Revenue vs profits

The variable of interest throughout the analysis is firm-level operating revenue (*operatingrevenue-turnover*). As mentioned in the text, the facts that (i) the revenue variable is often more populated and that (ii) it can be more directly linked to the concept of shocks in the model are two of the reasons we ended using it. Moreover, as can be seen in Table A.1, revenue is a better predictor of both exit and hiring decisions than alternatives based on accounting profits.

Table A.1: Revenue vs profits

	Regressor					
	ln *opre*	*opre*	Profit	EBIT	Profit/*opre*	EBIT/*opre*
Regressand	(1)	(2)	(3)	(4)	(5)	(6)
Exit	-0.021***	-9.43e-10***	-3.82e-10**	-1.08e-09**	-5.37e-07	-0.001***
	(0.001)	(7.92e-11)	(1.60e-10)	(5.30e-10)	(5.04e-07)	(0.002)
R^2	0.037	0.020	0.020	0.020	0.020	0.027
Empl. growth	0.021***	9.86e-10***	4.24e-10	1.35e-0.9***	2.25e-07***	0.0011***
	(0.001)	(7.39e-11)	(2.66e-10)	(5.02e-10)	(2.92e-07)	(0.0013)
R^2	0.022	0.017	0.017	0.017	0.017	0.018

Notes: This table reports the results of regressing a Spanish firm's exit event (top panel) or employment growth (bottom panel) at time t on revenue or various measures of accounting profit at time $t-1$. $year \times industry$ fixed effects are included throughout. "*opre*" is operating revenue; "Profit" is profit/loss before tax; and "EBIT" is earnings before interest and taxes. Standard errors are clustered at the industry level. *, **, and *** indicate significance at the 10%, 5%, and 1% level, respectively.

A.2 Outliers and adjustments to the transitions

In order to avoid that our results are driven by outliers, we trim the observations at the 0.1% and 99.9% thresholds. Also, we need to guard against the possibility that exit or entry might be driven by missing information in a specific year. First, if for a given year the firm has data for at least one of the most populated variables (e.g. employment and payroll) but not revenue, then it is dropped altogether. Second, we ensure that data "holes" do not generate spurious entry or exit by verifying that the firm is not ever present in the dataset before (after) the candidate entry (exit) year.[19]

In order to embed H, H_E, and $P(\text{Exit}|\hat{y})$ into a canonical heterogeneous-firms model, each needs to be adjusted in order to satisfy some standard technical assumptions employed by that literature. These assumptions ensure monotonicity of firm value functions and internal consistency of firm exit policies.

The raw transition distribution H is modified to ensure first-order stochastic dominance: for two states such that $y_2 \geq y_1$, we require that $H(y'|y_2) \leq H(y'|y_1)$ for all y'. Figure A.1 compares the raw (red) and adjusted (blue) transition matrices for residualized revenue. More specifically, it plots the density of the distribution $H(\hat{y'}|\hat{y})$ at different value of $\hat{y}$.

Figure A.1: Transition matrix of incumbents: original vs. adjusted

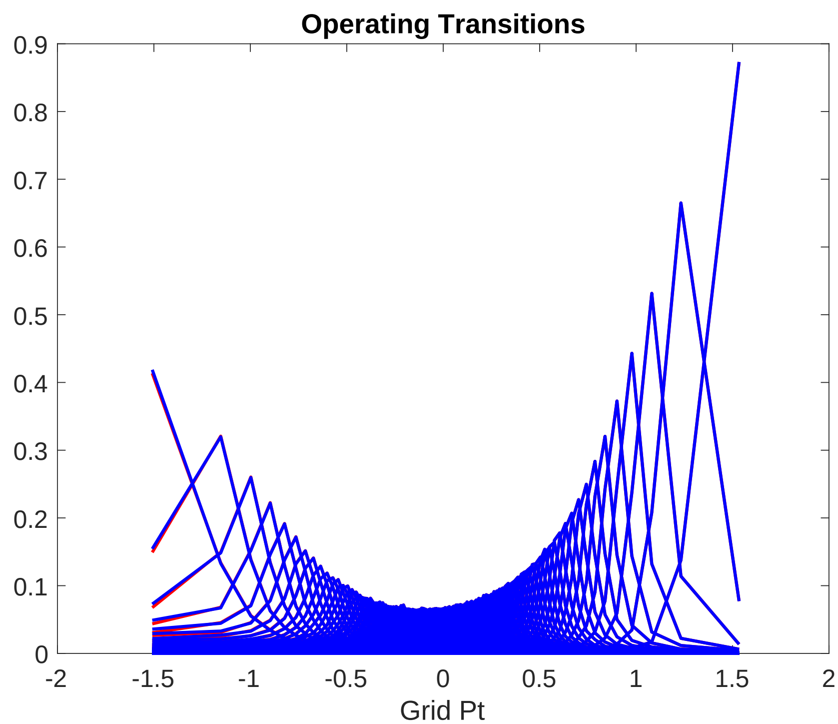

Notes: The figure compares the raw (red) and adjusted (blue) transition matrices for residualized revenue.

Second, we require that both entry and exit $\mathbb{P}(\text{Exit}|\hat{y})$ must be weakly decreasing in $\hat{y}$. Figure A.2 compares the raw and processed objects.

[19]For this reason, our sample period never goes beyond 2014. This allows us to verify that a firm does not show up again between 2015 and 2019, which is the end of the dataset in the vintage used.

Figure A.2: Entry and exit: original vs. adjusted

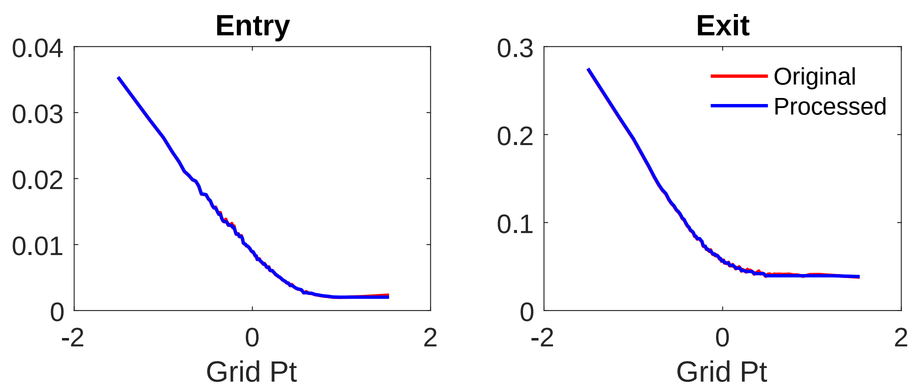

In both cases, the main takeaway is that the required adjustments are minor and, in some cases barely discernible. In other words, the raw data from ORBIS are already very close to satisfying the minimal requirements of our model.

A.3 Firm Size Distribution

The following figure shows the counter CDF of log revenue in the Spanish data. The tail of the distribution

Figure A.3: Counter CDF of Revenue

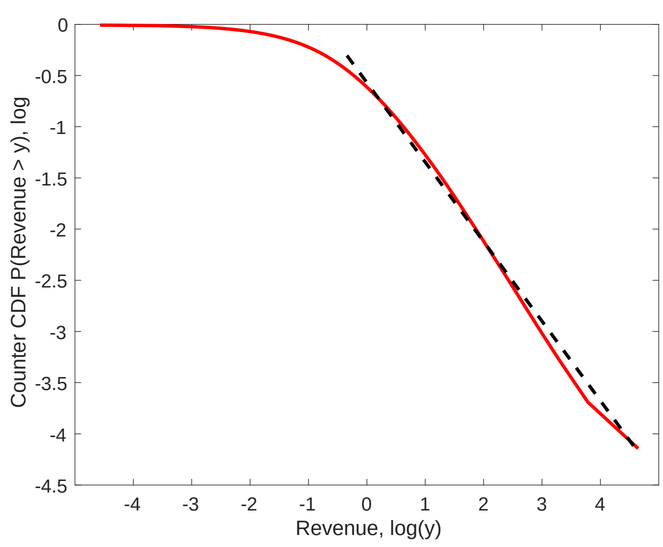

A.4 Empirical Moments

We report in Table A.2 the standard deviation, skewness and kurtosis of demeaned log revenue and the first-difference of demeaned log revenue under various alternative treatments of the data as well

Table A.2: Empirical moments under alternative data treatments

	Std dev	Skewness	Kurtosis	Std dev	Skewness	Kurtosis
Baseline	1.548	0.025	4.196	0.656	-0.312	29.212
Before 2009	1.515	0.007	4.182	0.692	-0.071	29.920
After 2009	1.583	0.025	4.224	0.652	-0.482	26.974
Mfg	1.561	0.056	3.862	0.480	-0.829	41.572
Non-Mfg	1.546	0.021	4.250	0.679	-0.285	27.767
Unconsolidated only	1.527	-0.031	4.125	0.655	-0.321	29.160
No M&A	1.545	0.028	4.214	0.655	-0.320	28.937
Year Effects Only	1.699	0.131	3.756	0.660	-0.415	28.765
No Trimming	1.416	-0.008	3.336	0.586	-0.201	21.646
1% Trimming	1.416	-0.008	3.336	0.586	-0.201	21.646
Italy	1.764	-0.660	6.070	0.956	-0.165	33.889
Portugal	1.514	-0.122	5.077	0.736	0.347	28.683
France	1.379	-0.152	6.483	0.604	0.637	72.380
Norway	1.705	-0.233	4.274	0.681	-0.153	26.698

Note: Basic moments of the distributions of demeaned log revenue and the first-difference of demeaned log revenue under alternative treatments of the data.

as other countries. The conclusions obtained for the benchmark Spanish dataset are unchanged.

B Solving and Calibrating the Model

In this appendix, we discuss in more detail our approach to solving and quantifying the model.

B.1 Solving the Empirical/Non-parametric Model

Note that the static optimality condition for the input n in equation (13), together with the residualized log revenue grid $\hat{y}_i$, implies a quantile-based grid for profitability shocks z_i, $i = 1, ..., N_z$, where $N_z = N_y$ and $\log z_i = (1-\alpha)\hat{y}_i$ for all i. Similarly, the empirical objects $H(\hat{y}'|\hat{y})$, $H_E(\hat{y})$, and $\mathbb{P}(Exit|\hat{y})$ imply an incumbent profitability transition $F(z'|z)$, an entry distribution $F_E(z)$, and an exit hazard $\mathbb{P}(\text{Exit}|z)$ on the profitability grid z_i.

We assume that exit occurs for the highest profitability firms in our sample for only exogenous reasons, i.e., that $\delta = \mathbb{P}(\text{Exit}|z_{N_z})$. In our baseline sample in Spain, the resulting exogenous exit rate is $\delta = 3.8\%$. The remaining parameters to be calibrated in our non-parametric model include only the labor share α, the household's rate of time preference β, the fixed labor supply $\bar{N}$, and the sunk entry cost ϕ_E. Given a parameterization of the model, i.e., a list of these parameters, we solve the model with an outer loop-inner loop approach as follows.

1. **Outer Loop on GE Objects** Guess values for the wage W and the entry mass M_E, and fix a GE tolerance $\epsilon^{GE} > 0$.

    (a) **Inner Loop on Firm Value Function** Initialize $k = 0$, guess a value function $V^{(k)}(z)$, and fix a value function error tolerance $\epsilon^V > 0$.

    i. Compute the implied continuation values ${\phi_F^*}^{(k)}(z)$ via equation (11) and using $V^{(k)}(z)$.

    ii. Infer the distribution $G^{(k)}(\phi_F)$ of fixed cost shocks ϕ_F consistent with ${\phi_F^*}^{(k)}(z)$, $V^{(k)}(z)$, and the empirical exit hazard by using the mapping

    $$G^{(k)}({\phi_F^*}^{(k)}(z)) = \frac{1 - \mathbb{P}(\text{Exit}|z)}{1-\delta}.$$

    iii. Compute an updated value function $V^{(k+1)}(z)$ via the Bellman equation

    $$V^{(k+1)}(z) = \left\{ \begin{array}{c} \max_n (zn^\alpha - Wz) \\ -\int_0^{{\phi_F^*}^{(k)}(z)} \phi_F dG(\phi_F) \end{array} + \beta(1-\delta) \int V^{(k)}(z') dF(z'|z) \right\}.$$

    iv. If the error in the Bellman equation $\max_z |V^{(k+1)}(z) - V^{(k+1)}(z)|$ is smaller than ϵ^V, then the firm value function $V(z) = V^{(k)}(z)$, continuation values $\phi_F^*(z) = {\phi_F^*}^{(k)}(z)$, and the fixed cost distribution $G(\phi_F) = G^{(k)}(\phi_F)$ are computed. Otherwise, set $k = k+1$ and return to step (1(a)i).

    (b) **Inner Loop on Firm Distribution** Initialize $k = 0$, guess an operating distribution $F_O^{(k)}(z)$ for firms, guess a mass $M_O^{(k)}$ of operating firms, and fix a tolerance $\epsilon^F > 0$ for distributional convergence.

i. Compute the implied mass of operating firms $M_O^{(k+1)}$ via

$$M_O^{(k+1)} = (1-\delta)M_O^{(k)} \int G(\phi_F^*(z))dF_O^{(k)}(z) + M_E.$$

ii. Compute the implied distribution of operating firms $F_O^{(k+1)}(z)$ via

$$F_O^{(k+1)}(z') = (1-\delta)\frac{M_O^{(k)}}{M_O^{(k+1)}} \int G(\phi_F^*(z))F(z'|z)dF_O^{(k)}(z) + \frac{M_E}{M_O^{(k+1)}}F_E(z').$$

iii. If the errors in the operating mass update $|M_O^{(k+1)} - M_O^{(k)}|$ and distributional update $\max_z |F_O^{(k+1)}(z) - F_O^{(k)}(z)|$ are both less than ϵ^F, then the operating mass $M_O = M_O^{(k)}$ and operating distribution $F_O(z) = F_O^{(k)}(z)$ are computed. Otherwise, set $k = k+1$ and return to step (1(b)i).

2. Compute the implied value to entry V_E via

$$V_E = \int V(z)dF_E(z).$$

3. Compute the implied labor demand N via

$$N = M_O \int n^*(z)dF_O(z),$$

where $n^*(z)$ is optimal static labor demand for an individual firm with profitability z.

4. If the error in the free entry condition $|V_E - \phi_E|$ and the error in the labor market clearing condition $|N - \bar{N}|$ are both less than the GE tolerance ϵ^{GE}, then the model is solved. Otherwise, update your guesses for the wage and entry mass and return to step (1).

When the algorithm above is complete, the non-parameteric version of our model is solved in a manner not only consistent with general equilibrium but also, by construction, with the observed revenue transitions, the entry distribution, and the exit hazard measured non-parametrically.

A few additional technical details are useful. We implement all of the calculations above continuously, linearly interpolating value functions, fixed cost distributions, operating distributions, and continuation values on the grid z_i. Where integration is required, we use Simpson quadrature with densities $f_O(z)$, $f_E(z)$, and $f(z'|z)$ consistent with linear interpolation of the CDFs $F_O(z)$, $F_E(z)$, and $F(z'|z)$ in a manner which preserves the empirical weight on equal-mass intervals containing the revenue quantiles $\hat{y}_i$. Because the free entry condition is separable from the entry mass M_E, we first employ bisection on the aggregate wage W to ensure that the free entry condition is satisfied, then we update M_E so that (3) is exactly satisfied. In our baseline, we employ $N_y = N_z = 101$ grid points or quantiles, and on a 2017 iMac Pro model solution takes around a minute or two in MATLAB without requiring aggressive parallelization.

B.2 Solving the AR(1)/Parametric Model

In our AR(1) or parametric model version, the parameters to be calibrated include the labor share α, the household's rate of time preference β, the fixed labor supply $\bar{N}$, the sunk entry cost ϕ_E, the upper bound $\bar{\phi}_F$ of the fixed cost distribution $G(\phi_F) = U(0, \bar{\phi}_F)$, the persistence of the lognormal AR(1) profitability process ρ, the conditional variance of the lognormal AR(1) profitability process σ^2, and the mean of the lognormal entry distribution μ_E. The exogenous exit hazard δ is carried over identically from our non-parametric model solution as described above. Given a parameterization of the model, i.e., a list of these parameters, we solve the model with an outer loop-inner loop approach as follows.

1. **Outer Loop on GE Objects** Guess values for the wage W and the entry mass M_E, and fix a GE tolerance $\epsilon^{GE} > 0$.

   (a) **Inner Loop on Firm Value Function** Initialize $k = 0$, guess a value function $V^{(k)}(z)$, and fix a value function error tolerance $\epsilon^V > 0$.

      i. Compute an updated value function $V^{(k+1)}(z)$ via the Bellman equation

$$V^{(k+1)}(z) = \left\{ \begin{array}{c} \max_n (zn^\alpha - Wz) \\ - \int_0^{\phi_F^{*(k)}(z)} \phi_F dG(\phi_F) \end{array} + \beta(1-\delta) \int V^{(k)}(z') dF(z'|z) \right\}.$$

      ii. If the error in the Bellman equation $\max_z |V^{(k+1)}(z) - V^{(k+1)}(z)|$ is smaller than ϵ^V, then the firm value function $V(z) = V^{(k)}(z)$ is computed. Otherwise, set $k = k + 1$ and return to step (1(a)i).

   (b) **Inner Loop on Firm Distribution** Initialize $k = 0$, guess an operating distribution $F_O^{(k)}(z)$ for firms, guess a mass $M_O^{(k)}$ of operating firms, and fix a tolerance $\epsilon^F > 0$ for distributional convergence.

      i. Compute the implied mass of operating firms $M_O^{(k+1)}$ via

$$M_O^{(k+1)} = (1-\delta) M_O^{(k)} \int G(\phi_F^*(z)) dF_O^{(k)}(z) + M_E.$$

      ii. Compute the implied distribution of operating firms $F_O^{(k+1)}(z)$ via

$$F_O^{(k+1)}(z') = (1-\delta) \frac{M_O^{(k)}}{M_O^{(k+1)}} \int G(\phi_F^*(z)) F(z'|z) dF_O^{(k)}(z) + \frac{M_E}{M_O^{(k+1)}} F_E(z').$$

      iii. If the errors in the operating mass update $|M_O^{(k+1)} - M_O^{(k)}|$ and distributional update $\max_z |F_O^{(k+1)}(z) - F_O^{(k)}(z)|$ are both less than ϵ^F, then the operating mass $M_O = M_O^{(k)}$ and operating distribution $F_O(z) = F_O^{(k)}(z)$ are computed. Otherwise, set $k = k + 1$ and return to step (1(b)i).

2. Compute the implied value to entry V_E via

$$V_E = \int V(z)dF_E(z).$$

3. Compute the implied labor demand N via

$$N = M_O \int n^*(z)dF_O(z),$$

where $n^*(z)$ is optimal static labor demand for an individual firm with profitability z.

4. If the error in the free entry condition $|V_E - \phi_E|$ and the error in the labor market clearing condition $|N - \bar{N}|$ are both less than the GE tolerance ϵ^{GE}, then the model is solved. Otherwise, update your guesses for the wage and entry mass and return to step (1).

Note that unlike in the empirical or non-parametric version of the model, the fixed cost distribution $G(\phi_F) = U(0, \bar{\phi}_F)$ is predetermined. Also note that the entry and transition distributions $F_E(z)$ and $F(z'|z)$ are parametric, following conventional lognormal processes converted to a uniform profitability grid as in Tauchen (1986). Just as in the non-parametric solution of the model, however, we continue to solve the model continuously, storing value functions via linear interpolation, computing integrals via Simpson quadrature, and evaluating entry, operating, and transition distributions using linear interpolation of the CDFs $F_E(z)$, $F_O(z)$, and $F(z'|z)$. In our baseline, we again employ $N_z = N_y = 101$ points for our interpolation procedures, and model solution takes around a minute or two on a 2017 iMac Pro in MATLAB without aggressive parallelization.

B.3 Calibrating the Model

There are two model parameters which we fix or calibrate externally before engaging in a moment-matching exercise. We set $\alpha = 2/3$ to generate a conventional labor share of 2/3, we set $\beta = 1/1.04$ to be consistent with a conventional 4% real interest rate and an annual solution of the model, and we set $\bar{N}$ to be equal to the aggregate employment rate (resulting in $\bar{N} = 0.5974$ in our baseline Spanish sample and comparable values for our other samples). As noted above, we also set the exogenous exit hazard δ based on the observed exit rate of the largest firms in our empirical sample (resulting in $\delta = 3.8\%$ for our baseline Spanish sample and comparable values for our other samples). Each of the versions of our model, non-parametric and parametric, is solved holding these externally calibrated parameters fixed.

Non-parametric Calibration With the externally calibrated parameters listed above fixed, only the sunk entry cost ϕ_E must be calibrated for the non-parametric model. We choose the value of ϕ_E to match the observed average number of employees per firm. The number of employees per firm declines in the wage W, which adjusts to satisfy the free entry condition as the parameter ϕ_E is shifted.

Parametric Calibration With the externally calibrated parameters above fixed, we must still fix the values of the lognormal AR(1) profitability process (ρ, σ^2), the mean of the lognormal entry distribution μ_E, the upper bound $\bar{\phi}_F$ of the fixed cost distribution $G(\phi_F) = U(0, \bar{\phi}_F)$, as well as the sunk entry cost ϕ_E. Following convention in the parametric firm dynamics literature, we first set ρ to the autocorrelation of the profitability process $\log z$ inferred from our observed revenue series $\hat{y}$, and we set σ^2 to match the observed variance of $\log z$.

Then, with ρ and σ^2 fixed, we choose the remaining three parameters $(\mu_E, \bar{\phi}_F, \phi_E)$ to jointly match three moments. As in the non-parametric model, we match (i) the observed average number of employees per firm. We also match (ii) the observed exit rate $\mathbb{P}(\text{Exit})$ which naturally moves with the fixed cost upper bound $\bar{\phi}_F$. Finally, we match (iii) the mean difference between log revenue for entering and operating firms, which naturally moves with the mean of the entry distribution μ_E. One might wonder why we did not target moments (ii) nor (iii) in our non-parametric model solution. But the non-parametric model matches both of these moments by construction, since both moments are implied by the combination of incumbent revenue transitions, exit hazards, and the entry distribution which are fully matched in the non-parametric model.

B.4 Comparison of entry and exit

We extract from the data a collection of incumbent revenue transitions $H(\hat{y}'|\hat{y})$, an entry distribution $H_E(\hat{y})$, and an exit hazard $\mathbb{P}(Exit|\hat{y})$ for a collection of residualized log revenue quantiles $\hat{y}_i$, $i = 1, ..., N_y$. The red lines in Figure B.1 plot the exit hazard (left panel) and the entry distribution (right panel), while the blue lines report comparable objects within our calibrated AR(1)/parametric version of the model discussed later.

Figure B.1: Exit and entry hazards

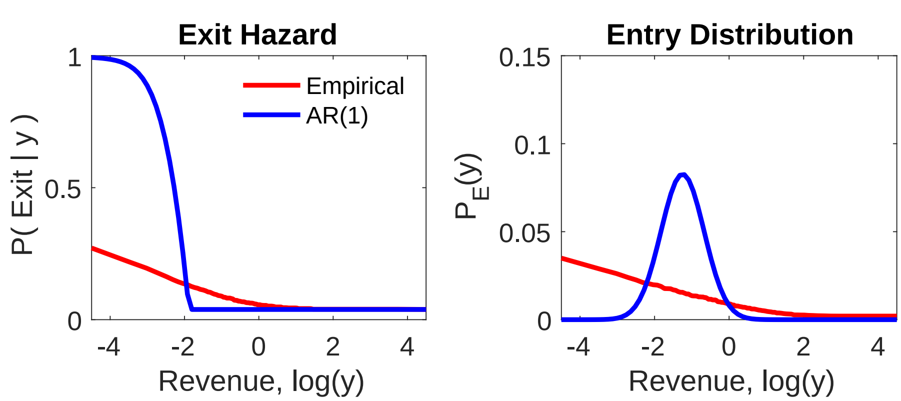

C Alternative datasets and approaches

Table C.1: Relative impacts of subsidizing 5% of fixed costs

	Exit Rate	Output
Baseline	3.0701	0.6845
Panel A: Alternative modelling assumptions		
Endogenous Labor Supply	3.0660	0.8125
Higher α	3.0728	0.6872
Lower α	3.0765	0.6859
Panel B: Alternative datasets		
Before 2009	2.5787	0.7095
After 2009	2.4338	0.6176
Manufacturing Only	3.6506	0.4211
Non-Manufacturing Only	2.7832	0.7108
Unconsolidated Accounts	3.1550	0.7109
Excluding M&A	3.0750	0.7161
Year Effects Only	2.8076	0.7015
No Trimming	3.3206	0.7883
Trimming at 1% and 99%	2.6755	0.7461
Italy	3.9723	0.7040
Portugal	5.8445	0.7905
France	1.8992	0.5814
Norway	2.8302	0.6560

Note: The table reports relative changes at the aggregate level from subsidizing 5% of mean fixed costs in the empirical model versus the parametric model. Panel A reports results under various alternative modelling assumptions, while Panel B considers calibrations based on different datasets. For each experiment, as listed in the first column, we calculate the change in the exit rate and output vis-a-vis the original values in the nonparametric and AR1 models. We then report the ratio of these two changes; the second column reports these values for the exit rate, while the third column reports these values for output.

C.1 Manufacturing vs non-manufacturing

Figure C.1: Manufacturing - Impact of a fixed operating cost subsidy

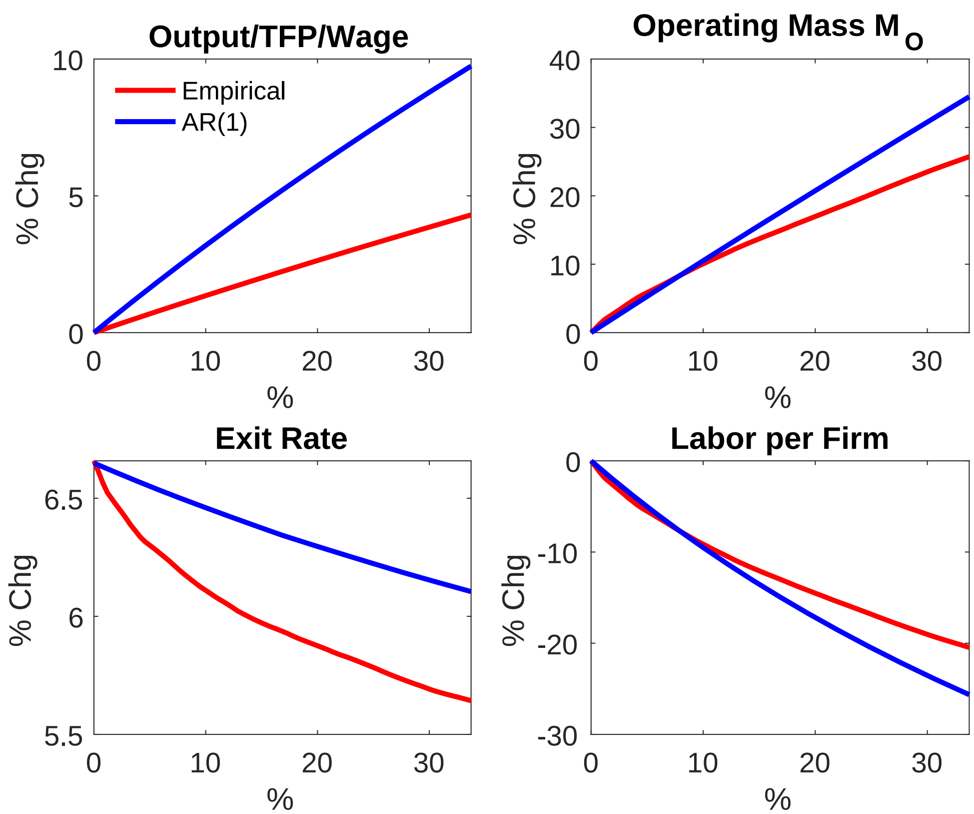

Note: This figure depicts the impact of a fixed operating cost subsidy in the non-parametric model (red line) and the AR(1) models (blue line).

Figure C.2: Non-manufacturing - Impact of a fixed operating cost subsidy

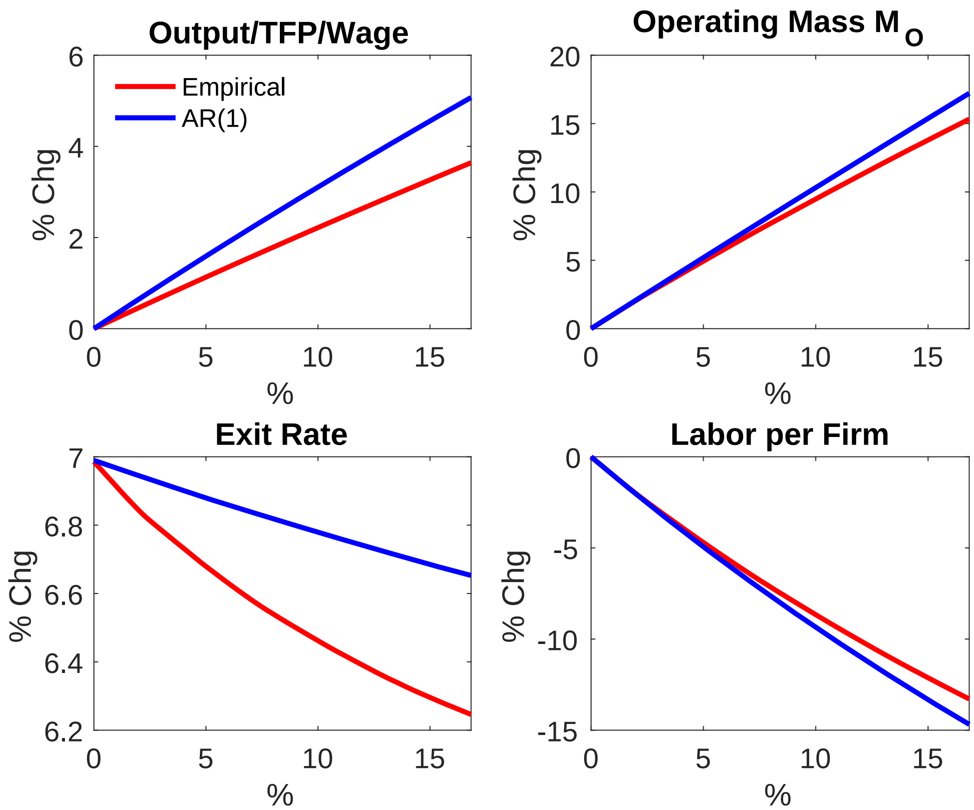

Note: This figure depicts the impact of a fixed operating cost subsidy in the non-parametric model (red line) and the AR(1) models (blue line).

C.2 Italy

Figure C.3: Italy - Impact of a fixed operating cost subsidy

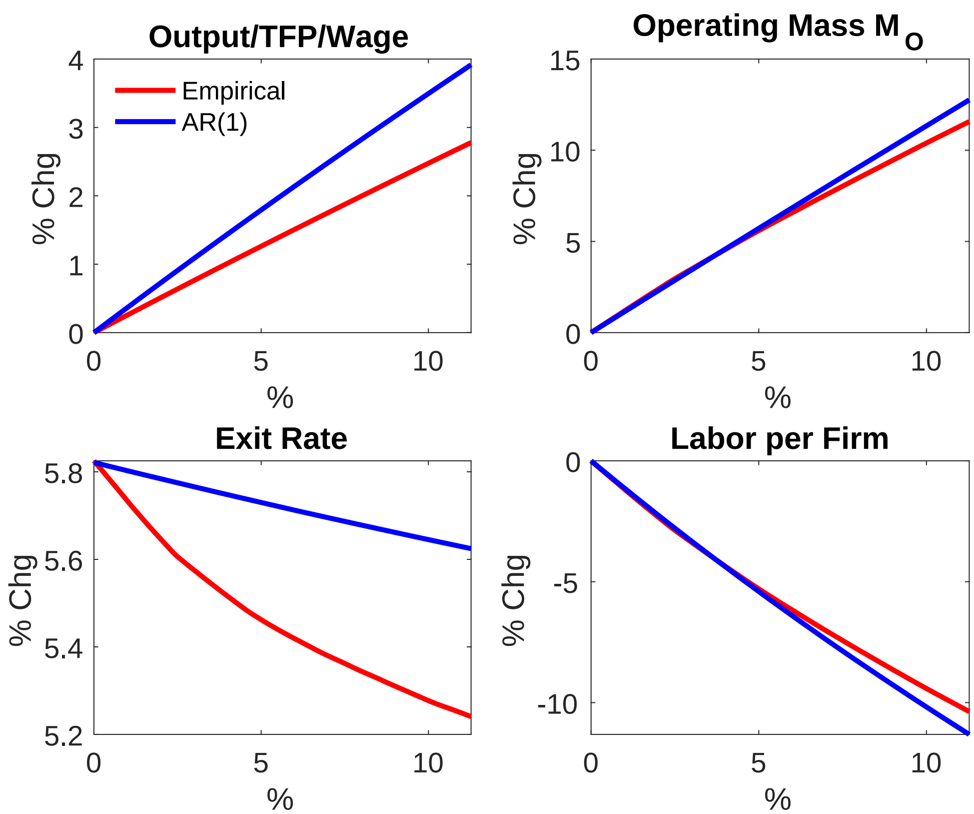

Note: This figure depicts the impact of a fixed operating cost subsidy in the non-parametric model (red line) and the AR(1) models (blue line).

Figure C.4: Italy - Impact of a subsidy to entry

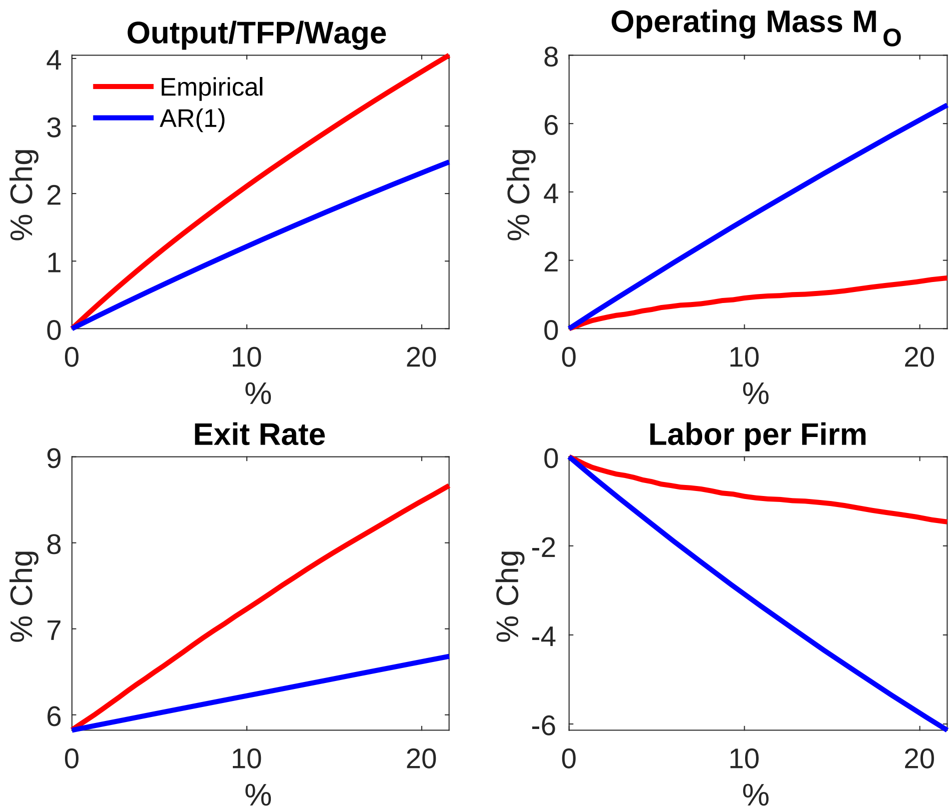

Note: This figure depicts the impact of an entry subsidy (as a portion of the sunk entry cost) in the non-parametric model (red line) and the AR(1) models (blue line).

C.3 Portugal

Figure C.5: Portugal - Impact of a fixed operating cost subsidy

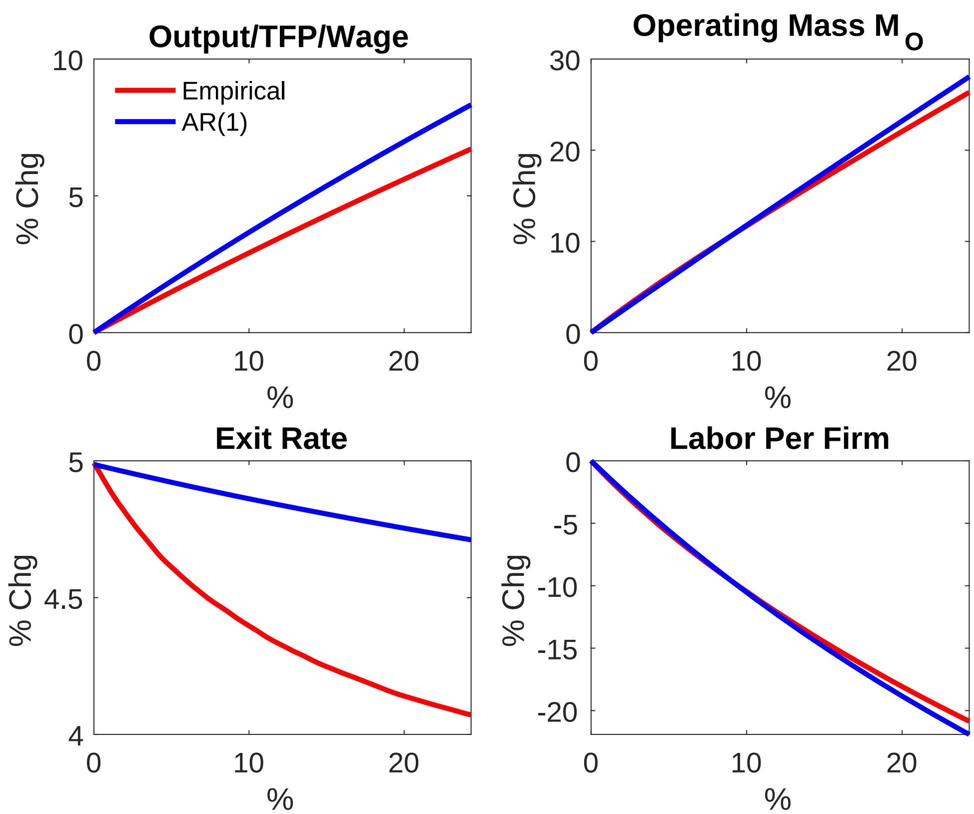

Note: This figure depicts the impact of a fixed operating cost subsidy in the non-parametric model (red line) and the AR(1) models (blue line).

Figure C.6: Portugal - Impact of a subsidy to entry

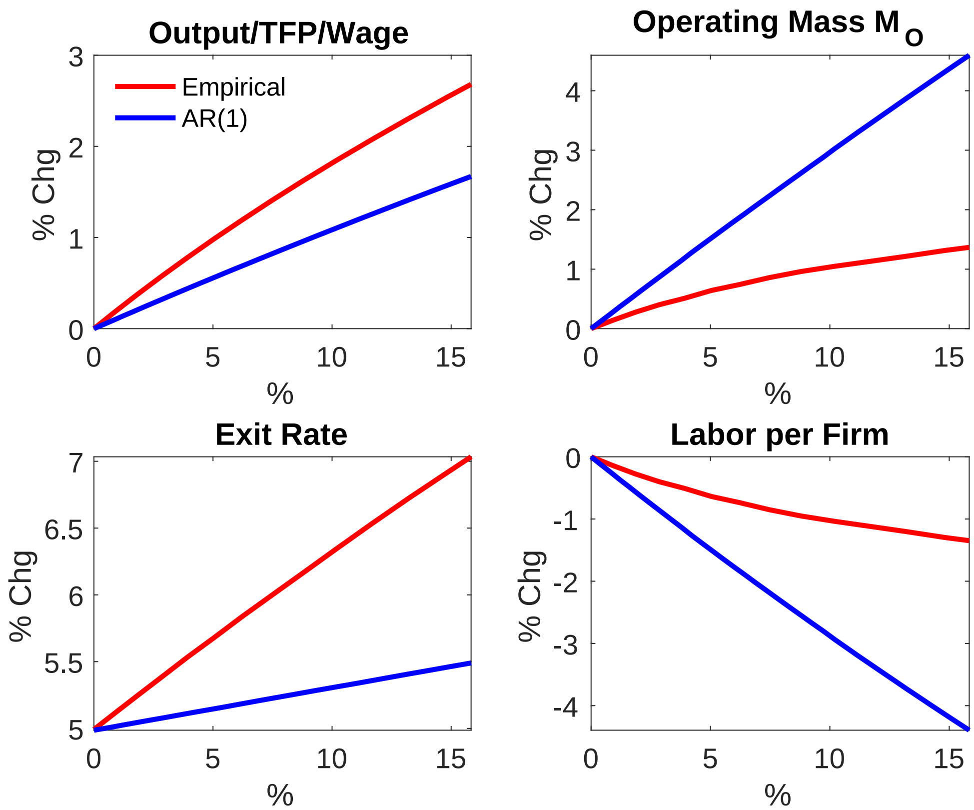

Note: This figure depicts the impact of an entry subsidy (as a portion of the sunk entry cost) in the non-parametric model (red line) and the AR(1) models (blue line).

D Bunching statistic

Table D.1: Bunching statistic at the sectoral level

Bunching statistic	NAICS Sector
0.091	Construction, 23
0.0543	Real Estate, 53
0.0492	Professional Technical Services, 54
0.0484	Retail Trade, 44
0.0453	Retail Trade, 45
0.0447	Information, 51
0.0401	Manufacturing, 33
0.0399	Wholesale Trade, 42
0.0399	Arts & Entertainment, 71
0.0392	Administrative Support Services, 56
0.0389	Accommodation and Food Services, 72
0.0376	Manufacturing, 32
0.0371	Educational Services, 61
0.0369	Other Services, 81
0.0351	Manufacturing, 31
0.0324	Transportation and Warehousing, 48
0.0288	Finance and Insurance, 52
0.0236	Health Care and Social Assistance, 62

Note: The bunching statistic at the sectoral level is computed as described in Section 6.