Firm Dynamics and the Granular Hypothesis*

Vasco Carvalho† Basile Grassi‡

October, 12 2013

[Link the the latest version]

Abstract

When large firms represent a disproportionate share of the economy, business cycles may be governed by idiosyncratic shocks to these large firms. We show – theoretically - in a standard firm dynamics setting (Hopenhayn, 1992) with a finite number of firms, each subject to persistent idiosyncratic productivity shocks, that this "granular hypothesis" (Gabaix 2011) leads to substantial aggregate fluctuations. A fat-tailed distribution of firm size arises because of large entrants and persistent shocks.

The model, calibrated to the US economy with a large number of firms, generates fluctuations of aggregate TFP (respectively output) of 1.1% (respectively 2.3%). The structure of the model allows us to study the micro and macro impact of a shock on the largest firm. Such a shock is contractionary at the aggregate level and expansionary at the idiosyncratic level. The conditional prediction of the model on firms' co-movement shows that the differential growth between large and small firms is pro-cyclical as it is in the data.

Keywords: Firm Dynamics; Granular Hypothesis; Firm Size Distribution; Aggregate Fluctuations

---

*The authors would like to acknowledge helpful comments from Axelle Arquié, Gian Luca Clementi, Antoine Coste, Florin Bilbiie, Xavier Gabaix, Jean Imbs, Boyan Jovanovic and Edouard Schaal. The authors thank seminar participants at NYU, PSE, UPF, CREST-INSEE, University of Essex and UCL. Part of this work was made during Basile Grassi's stay at the NYU Economics department and he is grateful for their hospitality. He also thanks the Fulbright commission for financial support.

†University of Cambridge, Centre de Recerca en Economia Internacional, Universitat Pompeu Fabra and Barcelona GSE. vcarvalho@crei.cat.

‡Paris School of Economics, Université Paris I Panthéon Sorbonne, and Centre de Recherche en Économie et STatistique. basile.grassi@gmail.com.

1 Introduction

When large firms represent a disproportionate share of the economy, aggregate fluctuations may be governed by idiosyncratic shocks to these large firms. This hypothesis – to which we refer as the "granular hypothesis" as proposed in Gabaix 2011 - opens the possibility of doing away with aggregate shocks, instead tracing back the origins of aggregate fluctuations to micro shocks hitting a small number of large firms. While this hypothesis has proven influential in the literature, it remains largely an accounting result - concerning the aggregation of firm level variability in static environments – whose relation to the extant theory of firm dynamics is left unspecified. Partly as a result, the granular hypothesis remains a possibility result whose quantitative relevance has not yet been established.

In this paper we cast the granular hypothesis in a standard firm dynamics setting. Building on the the setup of Hopenhayn 1992, we develop a quantitative theory of aggregate fluctuations arising only from idiosyncratic shocks to firm level productivity. This allows us to generalize the theoretical results in Gabaix 2011 to account for persistent micro-level shocks as well as endogenous firm entry and exit. Further, we provide a quantitative evaluation of this hypothesis and find that it yields aggregate fluctuations of the same order of magnitude as a standard representative-firm real business cycle model. A calibration of our model to the US economy with a large number of firms leads to sizable aggregate fluctuations: the standard deviation of aggregate TFP (respectively output) is 1.1% (respectively 2.3%).

Our model features a finite number of firms. Furthermore we do not rely on any "law of large number" arguments. Instead, we show how the dynamics of each firm affects the aggregate state of the economy. This can be summarized by a "reshuffling" shock on the firm size distribution - a shock that affect the whole distribution. The stochastic properties of this shock and the law of motion of the firm size distribution are governed by the idiosyncratic productivity process, the distribution of entrants and the endogenous exit decision.

We show theoretically that, as the number of firms increases, the aggregate volatility rate of convergence is smaller than what a simple central limit argument would predict. This result relies on the fat-tailedness of the firm size distribution as in Gabaix 2011. We generalize the latter by taking into account the persistence of the idiosyncratic shock process, firm entry and exit. In particular, we show that a low persistence of the idiosyncratic productivity process increases the rate of convergence of the aggregate volatility.

Our framework allows us to study the macro and micro dynamic effects of a negative productivity shock on the largest firm. Such a shock is contractionary at the aggregate level. By construction, the decrease in the productivity of the largest firm induces a smaller aggregate productivity, since the latter is shown to be an average of the idiosyncratic productivities. However, this shock induces the largest firm to reduce its size and to cut its labor demand. The latter pushes the wage down. Other firms benefit from this reduction of their cost. Their productivity remains the same while their cost decreases, so they expand. Overall, a negative shock on the largest firm is contractionary at the aggregate level but expansionary at the firm level. Our mechanism goes through the competition in the input market.

The upshot of this theory is that the business cycle, in our setting, is led by the large firms dynamics. In particular, if a small number of large firms face higher productivity than usual, this will lead to an aggregate expansion. To the econometrician, it would thus seem as if large firms are more cyclically sensitive, in that they co-move more with the aggregate echoing the recent findings in Moscarini and Postel-Vinay 2012. The latter find that the correlation between the differential net growth rate of large versus small firms is negatively correlated with the unemployment rate. We find a similar result on a simulated representative panel drawn from our calibrated model. Though our results are consistent with Moscarini and Postel-Vinay 2012 findings, causality is reversed : in our setting, it is not that large firms are more cyclically sensitive. Rather it is the business cycle itself that reflects large firms dynamics.

The paper relates to two distinct literatures: an emerging literature on the micro-origins of aggregate fluctuations and the more established firm dynamics literature. Gabaix 2011 describes the "granular hypothesis" and shows the possibility result that we extend to our framework. Other papers studying the micro-origins of aggregate fluctuations are Acemoglu et *al.* 2012, di Giovanni and Levchenko 2012, Carvalho 2010 and Carvalho and Gabaix 2013[1]. Relative to this literature, we contribute by grounding the granular hypothesis in a standard firm dynamics setting, extending the existing theoretical results to this setting and providing a first attempt at quantification.

[1]Some empirical evidence can be found in di Giovanni, Levchenko and Mejean 2012.

This paper is also related to the firm dynamics literature: Hopenhayn 1992, Campbell 1998, Veracierto 2002, Khan and Thomas 2003, 2008, Bachman and Bayer 2009. Some papers have studied aggregate fluctuations in an entry/exit framework as Lee and Mukoyama 2008, Clementi and Palazzo 2010 and Bilbiie et *al.* 2012. However they restrict their analysis to common, aggregate, shocks. Relative to this literature, we show that its standard workhorse model – once the assumption regarding a continuum of firms is dropped and the firm size distribution is fat tailed – already contains in it a theory of the business cycle. We show this both theoretically and quantitatively.

The paper is organized as follows. Section 2 derives the model. Section 3 shows our theoretical results on the rate of decay of aggregate fluctuations. Section 4 explains the numerical algorithm used to solve our framework. Section 5 explains the calibration, the origin of fat-tailed firm size distribution and exposes our first quantitative results. Section 6 displays the impulse response to a negative shock on the biggest firm. Section 7 looks at the cyclicality of large versus small firm in our model. Finally, section 8 concludes.

2 Model

We extend the Hopenhyan 1992 economy to allow for a finite (but large) number of firms. Consistent with this new feature we do not rely on any "law of large number" assumption, which implies that the firm size distribution becomes stochastic. Because of this crucial property, our model is able to generate aggregate fluctuations with only idiosyncratic shock of this crucial property.

Firms differ in their productivity level, which follows a discrete Markovian process. Incumbents have access to a decreasing return to scale technology using labor as the only input. They produce a unique good in a perfectly competitive market. They face an operating cost at each period, which generates endogenous exit. There is a large but finite number of potential entrants that differ in their productivity. To operate next period, potential entrants have to pay an entry cost. The economy is closed in a partial equilibrium fashion by specifying a labor supply that increases with the wage. One can think of this economy as one where households do not have access to savings and have a linear utility in consumption.

In this section, we describe the productivity process, the incumbents' problem and the entrants' problem. We then study the law of motion of the productivity distribution, the market clearing and aggregation. Finally, we define the stationary equilibrium.

2.1 Productivity Process

As stated above, the level of idiosyncratic productivity is discrete on a grid and follows a Markov chain with a transition matrix P. The productivity space is thus described by a n_s-uple $\Phi := \{\varphi_1, \ldots, \varphi_{n_s}\}$ such that $\varphi_1 < \ldots < \varphi_{n_s}$. A firm is in state (or productivity state) k when its idiosyncratic productivity is equal to φ_k. We denote $F(.|\varphi)$ the conditional distribution of the next period idiosyncratic productivity φ' given the current period idiosyncratic productivity φ.[2]

Although the productivity process is discrete, the number of states is large (123 in the baseline calibration) and is a discretization of an AR(1) process. So, to keep it simple, one can think of the idiosyncratic (log) productivity process as beeing:

$$\varphi_{t+1}^i = \rho \varphi_t^i + e_t^i, e_t^i \rightsquigarrow \mathcal{N}(0, \sigma_e)$$

To discretize this AR(1) process we follow the method described in Rouwenhorst 1995. We choose this method because the first-order autocorrelation, ρ, and the conditional variance, σ_e, are well defined and constant over the state space as shown in Kopecky and Suen 2010. For more details on this discretization method one can refer to appendix A.

2.2 Incumbents' problem

The only aggregate state variable of this model is the distribution of firms on the set Φ[3]. This distribution is represented by a vector μ_t which gives for each level of productivity the number of firms that have this productivity at a given period t. Given an aggregate state μ, and an idiosyncratic productivity level φ, the incumbent solves the following intra period problem:

$$\pi^*(\mu, \varphi) = \text{Max}\left\{\exp(\varphi) n^\alpha - wn - c_f\right\}$$

[2] Given a productivity level φ_s the distribution $F(.|\varphi_s)$ is given by the s^{th}-row vector of the matrix P.

[3] This will be shown at the end of this section.

where n is the labor input, w is the wage, which depends on the current aggregate state, and c_f is the operating cost that a firm should pay every period to operate. One can see that π^* is increasing in φ and decreasing in w for a given aggregate state μ. The output level is then $y(\mu,\varphi) = \exp(\varphi)^{\frac{1}{1-\alpha}} \left(\frac{\alpha}{w}\right)^{\frac{\alpha}{1-\alpha}}$. In what follows the size of a firm will refer to its output if not otherwise specified.

The incumbent timing is the following: she draws its idiosyncratic productivity φ at the beginning of the period, pays the operating cost c_f then hires labor, produces and decides to exit or not next period. It is worth emphasizing that once a firm decides to stay in operation it has to pay the operating cost next period, which starts by drawing a new idiosyncratic productivity. We denote the value of being an incumbent for a given aggregate state μ and idiosyncratic productivity level φ by $V(\mu,\varphi)$. Let us define the expected value of being an incumbent next period by

$$\mathcal{E}(\mu,\varphi) = \int_{\mu'\in\Lambda} \sum_{\varphi'\in\Phi} V(\mu',\varphi')F(d\varphi'|\varphi)\Gamma(d\mu'|\mu)$$

where $\Gamma(.|\mu)$ is the conditional distribution of μ', the tomorrow's aggregate state given today's aggregate state and where $F(.|\varphi)$ is the conditional distribution of φ', the tomorrow's idiosyncratic productivity given today's productivity.

For each aggregate state μ, since the instantaneous profit is increasing in the idiosyncratic productivity level and $F(.|\varphi)$ is decreasing in φ, there is a unique index $s^*(\mu)$ such that:

$$\mathcal{E}(\varphi_{s^*(\mu)},\mu) \geq 0 > \mathcal{E}(\varphi_{s^*(\mu)-1},\mu)$$

Thus for $\varphi \geq \varphi_{s^*(\mu)}$ the firm continues to operate next period and for $\varphi < \varphi_{s^*(\mu)}$ the firm exits[4]. When $\varphi = \varphi_{s^*(\mu)-1}$, we assume that the firm has a probability $\omega(\mu) = \frac{\mathcal{E}(\mu,\varphi_{s^*(\mu)})}{\mathcal{E}(\mu,\varphi_{s^*(\mu)})-\mathcal{E}(\mu,\varphi_{s^*(\mu)-1})}$ to stay in operation. In this case, the probability of staying in operation, $\omega(\mu)$, is greater when $\mathcal{E}(\varphi_{s^*(\mu)-1},\mu)$ is closer to zero. This assumption allows the entry rate to be continuous with respect to the aggregate state.

The Bellman equation associated with the incumbent's problem is then:

$$V(\mu,\varphi) = \pi^*(\mu,\varphi) + \beta \begin{cases} \mathcal{E}(\mu,\varphi) & \text{for } \varphi \geq \varphi_{s^*(\mu)} \\ \omega(\mu)\mathcal{E}(\mu,\varphi_{s^*(\mu)-1}) & \text{for } \varphi_{s^*(\mu)-1} = \varphi \\ 0 & \text{for } \varphi_{s^*(\mu)-1} > \varphi \end{cases}$$

[4]Given the state, the increasing instantaneous profit implies an increasing value function in φ and then an exit threshold. This result is shown in Hopenhyan 1992 and Clementi and Palazzo 2010.

After studying the incumbents' problem, we now turn to the problem of potential entrants.

2.3 Entrants' Problem

There is a constant and finite number of prospective entrants M. A share $M.G^q$ of them is of type φ_q, where φ_q lies within the idiosyncratic productivity level set Φ. It is a signal about their tomorrow's productivity. The number of entrants of type φ_q is deterministic.

If potential entrants decide to pay the entry cost c_e, then they produce next period with a productivity level drawn from $F(.|\varphi_q)$. The gross value of a successful entrant with a type φ_q given the aggregate state μ is thus:

$$\text{Max}\left\{0, \beta \int_{\mu' \in \Lambda} \sum_{\varphi' \in \Phi} V(\mu', \varphi') F(d\varphi'|\varphi_q) \Gamma(d\mu'|\mu)\right\} = \text{Max}\left\{0, \beta \mathcal{E}(\varphi_q, \mu)\right\}$$

A prospective entrant will enter if this value is greater or equal to the entry cost c_e. Since $\mathcal{E}(\varphi_q, \mu)$ is increasing in the signal φ_q, for any aggregate state μ there is a unique index $e^*(\mu)$ such that:

$$\mathcal{E}(\varphi_{e^*(\mu)}, \mu) \geq \frac{c_e}{\beta} > \mathcal{E}(\varphi_{e^*(\mu)-1}, \mu)$$

Thus for $\varphi_q < \varphi_{e^*(\mu)-1}$, the φ_q-potential entrant does not enter. Conversely, for $\varphi_q \geq \varphi_{e^*(\mu)}$, the q-potential entrant enters. As for the exit rule, when $\varphi_q = \varphi_{e^*(\mu)-1}$, we assume that the $\varphi_{e^*(\mu)-1}$-potential entrant has a positive probability, $\omega^e(\mu) := \frac{\mathcal{E}(\mu, \varphi_{e^*(\mu)}) - \frac{c_e}{\beta}}{\mathcal{E}(\mu, \varphi_{e^*(\mu)}) - \mathcal{E}(\mu, \varphi_{e^*(\mu)-1})}$, of staying in operation next period.

Given an aggregate state μ and a type φ_q, the value of an entrant, $V^e(\mu, \varphi_q)$, is then

$$V^e(\mu, \varphi) = \begin{cases} \beta \mathcal{E}(\mu, \varphi) - c_e & \text{for } \varphi \geq \varphi_{e^*(\mu)} \\ \omega^e(\mu)\Big(\beta \mathcal{E}(\mu, \varphi_{e^*(\mu)-1}) - c_e\Big) & \text{for } \varphi_{e^*(\mu)-1} = \varphi \\ 0 & \text{for } \varphi_{e^*(\mu)-1} > \varphi \end{cases}$$

To keep the mathematic as simple as possible, we will assume that the entry cost is normalized to zero: $c_e = 0$ which implies that $\varphi_{e^*(\mu)} = \varphi_{s^*(\mu)}$ and $\omega^e(\mu) = \omega(\mu)$.

2.4 Law of motion of the productivity distribution

In this section, we find the conditional distribution of the aggregate state tomorrow given the current aggregate state, i.e. what would be the next productivity distribution μ_{t+1} given the current one μ_t. It was denoted $\Gamma(.|\mu)$ in the previous section.

The distribution of firms μ_t across the discrete state space $\Phi = \{\varphi_1, \dots, \varphi_{n_s}\}$ is a $(n_s \times 1)$ vector equal to $(\mu_t^1, \dots, \mu_t^{n_s})$ such that μ_t^s is equal to the number of operating firms in state s at date t. The next period distribution is the sum of the evolution of incumbents and successful entrants.

In the following, we define two types of conditional distributions depending on firms' idiosyncrasy. The distribution of incumbent firms at date $t+1$ conditional on the fact that incumbents were in state s at date t is noted $f_{t+1}^{\cdot,s}$. This $(n_s \times 1)$ vector is such that for each state k in $\{1, \dots, n_s\}$:

$$\begin{aligned} f_{t+1}^{k,s} &= \text{ the } k^{th} \text{ element of } f_{t+1}^{\cdot,s} \\ &:= \text{ number of incumbents in state } k \text{ at } t+1 \text{ which were in state } \varphi_s \text{ at } t \end{aligned}$$

In the same way, let us define $g_{t+1}^{\cdot,s}$ the distribution of successful entrants at date $t+1$ given that they received the signal φ_s at date t. This vector is a $(n_s \times 1)$ vector such that for each state k in $\{1, \dots, n_s\}$:

$$\begin{aligned} g_{t+1}^{k,s} &= \text{ the } k^{th} \text{ element of } g_{t+1}^{k,s} \\ &:= \text{ number of entrants in state } k \text{ at } t+1 \text{ which received a signal } \varphi_s \text{ at } t \end{aligned}$$

The period $t+1$ firms distribution is the sum of all these conditional distributions and thus the vector μ_{t+1} satisfies:

$$\mu_{t+1} = \sum_{s=s^*(\mu_t)-1}^{n_s} f_{t+1}^{\cdot,s} + \sum_{s=s^*(\mu_t)-1}^{n_s} g_{t+1}^{\cdot,s} \tag{1}$$

It is important to emphasize the fact that $f_{t+1}^{\cdot,s}$ and $g_{t+1}^{\cdot,s}$ are multivariate random vectors which implies that μ_{t+1} also is a random vector.

At date $t+1$ for $s \geq s^*(\mu_t)$, $f_{t+1}^{\cdot,s}$ follows a multinomial distribution with two parameters: the integer μ_t^s and the $(n_s \times 1)$ vector $P'_{s,.}$ where $P_{s,.}$ is the s^{th} row vector of the matrix P (denoted $Multi(\mu_t^s, P'_{s,.})$). For $s < s^*(\mu) - 1$, $f_{t+1}^{\cdot,s}$ is equal to zero. For $s =$

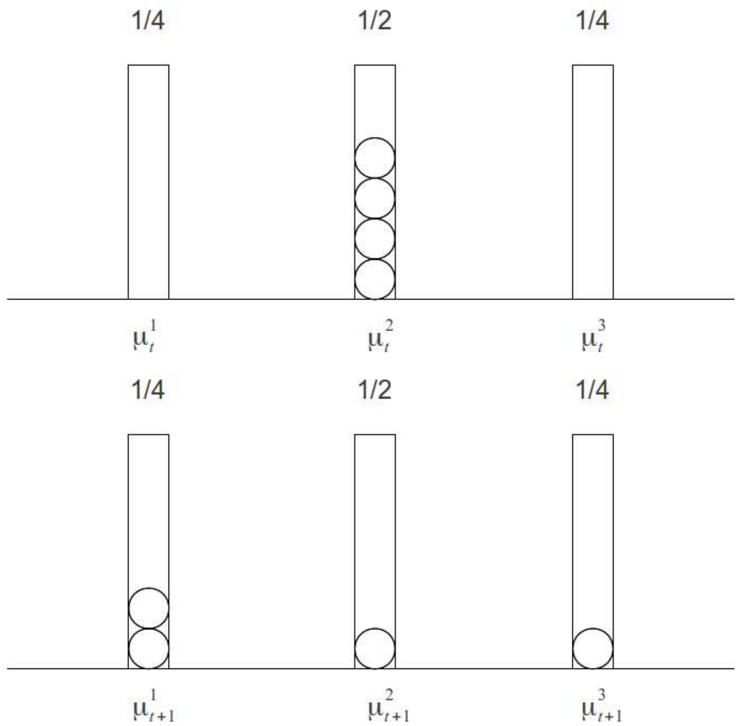

Figure 1: Why the vector $f_{t+1}^{.,s}$ follows a multinomial distribution.

$s^*(\mu)-1$, $f_{t+1}^{.,s^*(\mu)-1}$ follows a multinomial distribution with two parameters: the integer $\omega(\mu_t)\mu_t^{s^*(\mu)-1}$ and the $(n_s \times 1)$ vector $P'_{s^*(\mu)-1,.}$, i.e. $Multi\big(\omega(\mu_t)\mu_t^{s^*(\mu)-1}, P'_{s^*(\mu)-1,.}\big)$. Similarly, at date $t+1$ for $s > s^*(\mu_t)$, $g_{t+1}^{.,s}$ follows a multinomial distribution with two parameters: the integer MG^q and the $(n_s \times 1)$ vector $P'_{q,.}$, i.e. $Multi(MG^q, P'_{q,.})$. For $s = s^*(\mu_t)-1$, $g_{t+1}^{.,s^*(\mu_t)-1}$ follows $Multi\big(\omega(\mu_t)MG^{s^*(\mu_t)-1}, P'_{s^*(\mu_t)-1,.}\big)$.

To make clear the above statement, let us assume that there are only three levels of productivity ($n_s = 3$) and 4 firms. These firms are distributed according to the top panel of figure 1. Let us assume that the firms have a probability to go up (respectively down) on the productivity ladder of 1/2 and a probability to stay in the middle level of 1/4. If instead of 4 firms we had a continuum, the law of large number would hold and next period there would be exactly 1/4 of the firms at the first level, 1/2 at the middle level and 1/4 at the top level. This is not the case here, since the number of firms in each node is finite. For instance, a distribution of firms such as the one presented in the bottom panel of figure 1 is possible with a positive probability. In this particular case, the vector $(f_{t+1}^{1,2}, f_{t+1}^{2,2}, f_{t+1}^{3,2})'$ follows a multinomial distribution with a number of trials of 4 and an event probability vector $(1/2, 1/4, 1/2)'$.

The mean and variance-covariance matrix of a multinomial distribution $Multi(n, h)$ is respectively the $(n_s \times 1)$ vector nh and the $(n_s \times n_s)$ matrix $H = diag(h) - h'h$. So let us define $W_s = diag(P_{s,.}) - P_{s,.}P'_{s,.}$. The right hand side of equation 1 has a mean $m(\mu_t)$ and a variance-covariance matrix $\Sigma(\mu_t)$ with

$$m(\mu_t) = (P_t^*)'(\mu_t + MG)$$
$$\Sigma(\mu_t) = \sum_{s=s^*(\mu_t)}^{n_s} (MG^s + \mu_t^s).W_s + \omega(\mu_t)W_{s^*(\mu_t)-1}\left(\mu_t^{s^*(\mu_t)-1} + MG^{s^*(\mu_t)-1}\right)$$

where P_t^* is the transition matrix P with the first $(s^*(\mu_t) - 2)$ rows replaced by zeros and the $(s^*(\mu_t) - 1)$-th row multiplied by $\omega(\mu_t)$.

Equation 1 can be rewritten in a simple way as the sum of its mean and a zero-mean shock:

$$\mu_{t+1} = m(\mu_t) + \epsilon_{t+1} \tag{2}$$

where ϵ_{t+1} follows a multivariate distribution with mean zero and a covariance-variance matrix $\Sigma(\mu_t)$.

After taking into account all the dynamics of each incumbent firms and potential entrants, the law of motion of the aggregate state -the firm distribution- is quite simple: each period the distribution of firms is hit by a stochastic reshuffling shock and the persistence of this shock is governed by the transition matrix P_t^*.

Equation 2 shows the crucial difference of our model with a pure discretized Hopenhayn economy. In our setting, we make no use of any "law of large number" argument whereas it is largely used in this literature. Such argument would be equivalent to imposing the variance-covariance matrix to be zero, and so, the aggregate state μ_t would be non-stochastic, the model would not generate aggregate fluctuations. The "reshuffling" shock, ϵ_{t+1}, would be absent. Instead, in our framework, all the aggregate uncertainty comes from this new term, which is the main contribution of our model.

The importance of this "reshuffling" term relies on the deep parameters of the model, namely the first-order autocorrelation ρ, the conditional variance σ_e^2, the potential entrant distribution G and their number M. The first three combined determine the tail of the firm distribution that will govern the volatility of the aggregate variables as shown in following sections. As the number of potential entrants M goes to infinity (and thus so does the number of firms), the aggregate uncertainty becomes null but the rate of decay of this volatility is determined by this tail.

2.5 Market Clearing and Aggregation

If Y_t is the aggregate output, i.e. the sum of all individual firms' output, then $Y_t = A_t(L_t^d)^\alpha$ where L_t^d is the aggregate labor demand and A_t the aggregate total factor productivity, which is equal to:

$$A_t = \left(\sum_{i=1}^{N_t} \exp(\varphi_t^i)^{\frac{1}{1-\alpha}} \right)^{1-\alpha}$$

where φ_t^i is the productivity level at date t of the i^{th} firm among the N_t operating firms at date t. This can be rewritten by aggregating all firms that have the same productivity level:

$$A_t = \left(\sum_{s=1}^{n_s} \mu_t^s \exp(\varphi_s)^{\frac{1}{1-\alpha}} \right)^{1-\alpha} = (B'.\mu_t)^{1-\alpha}$$

where B is the $(n_s \times 1)$ vector of parameters $(\exp(\varphi_1)^{\frac{1}{1-\alpha}}, \ldots, \exp(\varphi_{n_s})^{\frac{1}{1-\alpha}})$ and where . is the matrix product.

The labor demand is $L^d(w_t) = \left(\frac{\alpha A_t}{w_t}\right)^{\frac{1}{1-\alpha}}$, the model behaves as a one factor model with aggregate TFP A_t. By Walras' law, only the labor market needs to be cleared. In a partial equilibrium fashion, we assume that the supply of labor at a given wage w is $L^s(w) = \underline{\mathrm{L}} w^\gamma$ with $\gamma > 0$. The market clearing condition then writes that labor supply equals labor demand, i.e. $L^s(w_t) = L_t^d$. Solving for the wage given the date t productivity distribution μ_t yields:

$$w_t = \left(\alpha^{\frac{1}{1-\alpha}} \frac{B'.\mu_t}{\underline{\mathrm{L}}} \right)^{\frac{1-\alpha}{\gamma(1-\alpha)+1}}$$

From this expression, one can see that the wage is fully pinned down by the distribution μ_t. Also, the distribution of productivity at t+1 depends only on the current distribution μ_t. It follows that the aggregate state at date t is μ_t.

2.6 Stationary Equilibrium

In this section we define a deterministic stationary equilibrium, which is similar to a deterministic steady state equilibrium. We define a stationary equilibrium as an equilibrium without aggregate uncertainty and thus where all variables are constant, that is to say with a deterministic aggregate state μ. The only source of uncertainty of μ is due to the fact that the $f^{\cdot,s}$ and $g^{\cdot,s}$ are random vectors. In a stationary equilibrium, we will assume that these variables are equal to their means $\mu^s P'_{s,\cdot}$ and $MG^q P'_{q,\cdot}$ respectively. This equilibrium is as if instead of considering a finite number of firms, we considered a continuum of firms. In the latter case, all the $f^{\cdot,s}$ and $g^{\cdot,q}$ are not stochastic and are equal to their mean.

Let us define the matrix P^* as the matrix P where the first $(s^* - 2)$ rows are replaced by zeros, and the $(s^* - 1)$-th row is multiplied by ω. The law of motion of μ implies that $\mu = P^{*'}.\mu + MP^{*'}G$. Solving for this matrix equation yields:

$$\mu = M(I - P^{*'})^{-1} P^{*'} G \tag{3}$$

We assume that this stationary distribution is fat-tailed, like is the case in the data as shown by Gabaix 2011. In the rest of the paper, we calibrate this distribution to be fat-tailed. From this, all other variables follow.

3 Aggregate Fluctuations

In this section we show analytically that the rate of decay of aggregate volatility depends on the tail of the stationary firm size distribution, and, on the persistence of the idiosyncratic productivity process. The dependence of aggregate volatility on the latter is one of the main contributions of our paper since it was absent from Gabaix 2011 analysis. A low persistence increases the aggregate volatility's rate of decay.

Results similar to Gabaix 2011 applies because μ_t is not too far from the stationary distribution μ^* as shown in Lemma 1. Thus μ_t can be considered as fat-tailed. The aggregate variable that pins down all other aggregate variables is

$$T_t = \sum_{i=1}^{N_t} \exp(\frac{1}{1-\alpha}\varphi_t^i)$$

The rate of decay at which T_t converges to its mean will be lower than $\sqrt{N_t}$ as shown in theorem 1.

We first need a lemma that shows that the tail of the firm size distribution at each date t is similar to the tail of the stationary distribution μ.

Lemma 1 *Assuming that the economy starts at date $t = 0$ with the stationary firm size distribution μ. Then at each period t the tail of the firm size distribution μ_t converges in probability to the tail of the stationary distribution μ as M, the number of potential entrants, goes to infinity.*

Proof See appendix B. $\square$

The theorem 1 describes the rate of decay of the aggregate volatility in our setting. It generalizes the result in Gabaix 2011 to the case of entry, exit and persistence of shocks.

Theorem 1 *Let ξ be the tail parameter of the firms size distribution and $\zeta' = \zeta(1-\alpha)$ the tail parameter of the entrants' size distribution. Assuming that $\xi/\rho < 2$ and $\zeta'/\rho < 2$ then*

$$\sigma(\frac{\Delta Y_t}{Y_t}) = \left(1 - \frac{\alpha}{\gamma(1-\alpha)+1}\right)\sigma(\frac{\Delta T_t}{T_t})$$

$$\sigma(\frac{\Delta T_t}{T_t}) \sim \frac{\sigma}{N_t^{1-\rho/\xi}} \frac{\left(\frac{N_t^l}{N_t}\right)^{\rho/\xi} u^{1/2}}{\bar{I}_t} \qquad \text{if } \zeta' > \xi$$

$$\sigma(\frac{\Delta T_t}{T_t}) \sim \frac{\sigma}{N_t^{1-\rho/\zeta'}} \frac{\left(\frac{N_t^E}{N_t}\right)^{\rho/\zeta'} w^{1/2}}{\bar{I}_t} \qquad \text{if } \zeta' < \xi$$

where $\bar{I}_t$ is a time dependent constant proportional to the incumbent average size at t, u and w are random variables with finite variance, σ is the standard deviation of $\exp(e_i^t/(1-\alpha))$ and N_t, N_t^l, N_t^E are respectively the number of incumbents, successful incumbents and successful entrants in period t.

Proof See appendix B. □

In this theorem, the parameters ρ and σ are well defined by the Markovian productivity process. We use the Rouwenhorst method (Rouwenhorst 1995) to define the latter process as the discretization of an AR(1). The process, as shown in Kopechy and Suen 2010, is characterized by a first order autocorrelation ρ and a conditional variance σ_e constant over the state space Φ. Thus, the parameters ρ and σ in theorem 1 are well-defined objects characterizing the productivity process.

This theorem states that the persistence of the productivity process affects the rate of decay of the aggregate volatility. Low persistence implies that large firms are getting back to the average size faster than in a high persistence case. So in the low persistence case, a given large firm would be smaller on average next period than in a high persistence case. Its impact on the aggregate is then smaller, its contribution to the aggregate volatility is smaller than in the high persistence case. This effect is absent in Gabaix 2011 since only transient shocks are considered.

4 Numerical Solution Algorithm

This section describes the algorithm used to solve numerically the full model.

The state variable of this model is only the distribution of productivity $\mu \in \mathbb{R}_+^{n_s}$. Since n_s is large, following the evolution of the distribution μ across time is not computationally feasible. To solve this model we use an algorithm similar to Krusell and Smith 1998, where we follow the evolution of the factor that matters for all the aggregate variables, namely T_t defined as:

$$T_t = \sum_{s=1}^{n_s} \mu_t^s \exp(\varphi_s)^{\frac{1}{1-\alpha}} = B'.\mu_t$$

where B is the $(n_s \times 1)$ vector $(\exp(\varphi_1)^{\frac{1}{1-\alpha}}, \ldots, \exp(\varphi_{n_s})^{\frac{1}{1-\alpha}})$.

The true evolution of T_t is:

$$T_{t+1} = B'.m(\mu_t) + B'.\epsilon_{t+1}$$

or

$$T_{t+1} = B'.m(\mu_t) + \sqrt{B'.\Sigma(\mu_t).B}\varepsilon_{t+1}$$

where ε_{t+1} is drawn from a standard univariate normal distribution. At the first order, the process followed by $\log(T_t)$ is:

$$\log(T_{t+1}) = \log(B'.m(\mu_t)) + \frac{\sqrt{B'.\Sigma(\mu_t).B}}{B'.m(\mu_t)}.\varepsilon_{t+1}$$

Assuming the following approximations:

$$\begin{aligned} \log(B'.m(\mu_t)) &= \alpha_0 + \alpha_1 \log(T_t) + w_t \\ \frac{\sqrt{B'.\Sigma(\mu_t).B}}{B'.m(\mu_t)} &= \beta_0 + \beta_1 \log(T_t) + v_t \end{aligned}$$

leads to the following approximate law of motion of T_t:

$$\log(T_{t+1}) = \alpha_0 + \alpha_1 \log(T_t) + \beta_0.\epsilon_{t+1} + \beta_1.\epsilon_{t+1} \log(T_t) + u_t \tag{4}$$

where u_t,v_t and w_t are error terms.

This approximate law of motion is used to compute expectations as in Krusell and Smith 1998. The coefficients are updated using estimation of this equation for a simulated series. We iterate until convergence. The algorithm is formally described bellow:

1. Guess some parameters $\alpha_0^0, \alpha_1^0, \alpha_2^0, \beta_0^0$ and β_1^0.
2. Solve jointly for the value function of an individual firm and the exit rule for all (φ, T) using the approximation law of motion 4 to compute the expectation.
3. Simulate a series of $\{T_t, \epsilon_t\}_{t=0...T}$ as follows:
   (a) Given a μ_0, compute $s^*(T_0)$ and $\omega(T_0)$ from the solution of step 2
   (b) Draw a ϵ_1 vector from its true distribution and use it to compute μ_1, and T_1, and using the law of motion of the productivity distribution described in equation (2).
   (c) Iterate from step 3a
4. Using the above simulated series, estimate the approximating rule 4.
5. Iterate from step 2 until convergence.

5 Simulations

The model is solved using the algorithm described in section 4. This section first describes the calibration procedure. Then, it assesses the quality of the numerical solution. A discussion of the source of the fat-tailedness of the firm size distribution follows. Finally, we compute some business cycle statistics.

5.1 Calibration

The (log of) productivity of incumbent firms follows a discretization of an AR(1) process as in Rouwenhorst 1995. This method is briefly described in appendix A. This discretization is convenient because the first order autocorrelation and the conditional variance are constant over the state space. On top of that, this method has been shown to be more accurate with highly persistent processes (Kopechy and Suen 2010). This is key to make sure that ρ and σ are well defined in theorem 1. This is not the case with other discretization methods such as the one proposed by Tauchen 1986.

To calibrate the model to the US economy, we first set the value of some deep parameters. The span of control parameter α is set at 0.8. This value is chosen to be on the lower end of recent estimates, such as Basu and Fernald 1997 and Lee 2005. The discount factor β is set at 0.95 so that the implied annual gross interest rate is 4%, a value in line with the business cycle literature. The entry cost is normalized to be zero to keep the mathematics as simple as possible. Finally, the potential entrant distribution is such that the distribution potential entrant on $\{\exp(\varphi_i)\}_{i\in[|1,n_s|]}$ is Pareto with tail parameter ζ. The economy scales with the number of potential entrants, M, and the parameter of the labor supply function, $\underline{\mathrm{L}}$. By increasing M and keeping the ratio $M/\underline{\mathrm{L}}$ unchanged the number of firms in equilibrium is increased as one can see in equation 3. So we choose the number of potential entrants such that the number of firms in equilibrium is $5e+03$. The labor supply elasticity parameter, γ, is chosen to be 2.

We jointly calibrate the operating cost c_f, the first order autocorrelation ρ and variance σ_e of the productivity process, and, the tail parameter ζ of the potential entrant distribution to match four moments: the entry rate, the entrant relative size, the maximum size of firms and the tail of the firm size distribution. Table 1 summarizes the result of this calibration procedure and the reference for each target. The first two targets are taken from Lee and Mukoyama 2008 where they use manufacturing firm level data

Statistic	Model	Data	References
Entry Rate	0.049	0.062	Lee and Mukoyama 2008
Entrants' relative size	0.60	0.60	Lee and Mukoyama 2008
Maximum size	2.1×10^6	2.1×10^6	Size of Wal-mart in 2010
Tail index of Firm size dist.	1.03	1.03	Gabaix 2011

Table 1: Targets for the calibration of parameters $\{c_f, \rho, \sigma_e, \zeta\}$

to compute these moments. The maximum size of a firm is the one of Wal-mart in 2010, which was the biggest firm in terms of employees. The tail index of the firm size distribution is estimated by Gabaix 2011. The corresponding moments in the model are computed for the stationary deterministic equilibrium. All the calibrated parameters are described in table 2.

The model is able to match most moments, except the entry rate, which is too low in our baseline calibration. There are not enough entrants in our baseline model although the difference is quite small. The first order autocorrelation and the conditional standard deviation of the productivity process are in line with previous estimates as in Lee and Mukoyama 2008 or in Castro et al 2013.

5.2 Numerical Solution

This section describes the numerical solution of the model by using the algorithm presented in section 4 for the calibration in table 2. At the last stage of the algorithm, the approximate law of motion for the state T_t is estimated to be:

$$\log(T_{t+1}) = 1.8649 + 0.8383 \log(T_t) + 1.4849\epsilon_{t+1} + u_t \tag{5}$$

The R^2 of the last step of the algorithm for this approximate law of motion is 0.99. However, it has been shown in Den Haan 2010 that this is not enough to assert the quality of the approximate law of motion. Figure 2 reproduces both the simulated path of $\log T_t$ and the path of the approximate law of motion (5). Despite some minor differences, one can see that the two paths coincide.

Parameters	Value	Description
ρ	0.9637	Autocorrelation of firms' level shocks
σ_e	0.1053	Std of idio shocks
n_s	123	Number of productivity levels
$[\varphi_1, .., \varphi_{n_s}]$	$[-4.35, .., 4.35]$	Productivity ladder
γ	2	Labor Elasticity
α	0.8	Production function
c_f	0.0033	Operating cost
c_e	0	Entry cost
β	0.95	Discount rate
M	1.45×10^{12}	Number of potential entrants
$\underline{\mathrm{L}}$	5.33	Parameters of the labor supply function
G	Pareto(ζ)	Entrant's distr. of $\exp(\varphi)$
ζ	4.2064	Parameter of the distr. G
$\zeta(1-\alpha)$	0.8413	Tail parameter of the entrants size distr.

Table 2: Baseline calibration

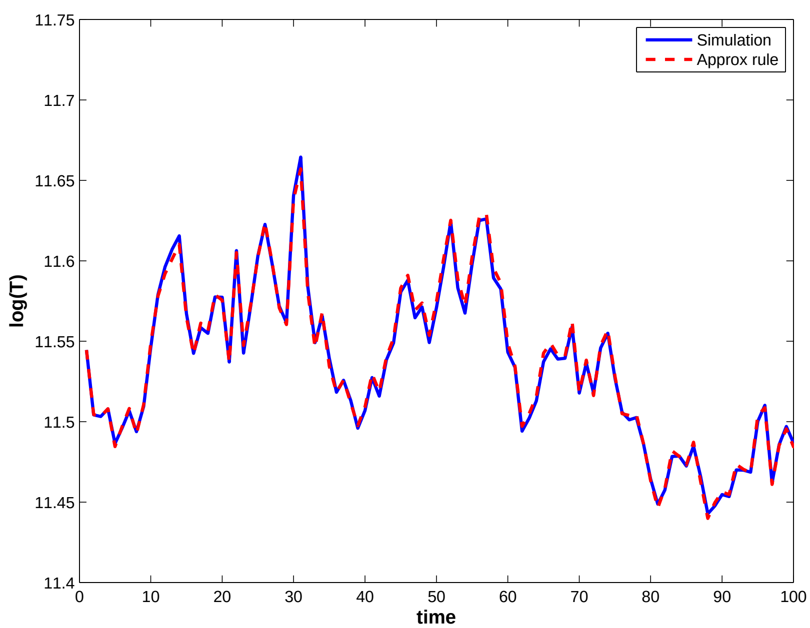

Figure 2: Simulated paths of the true and approximate evolution of T_t

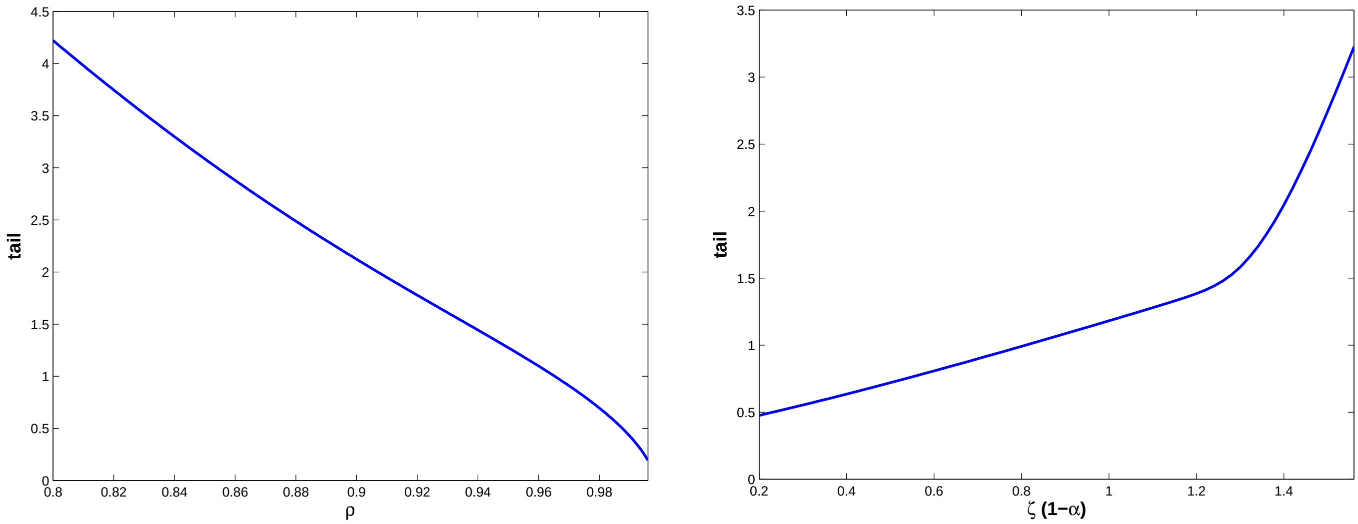

Figure 3: The tail of the stationary firm size distribution as a function of ρ (left) and $\zeta(1-\alpha)$, the tail of the potential entrant distribution G (right)

5.3 The tail of the firm size distribution

The calibration strategy is such that the model matches the fat tail of firm size distribution documented in Gabaix 2011. However the reason why the model is able to do so is different from the usual explanation provided in the firm dynamics literature. In this section, we show that the fat tail comes from the high persistence and the fat tail of the potential entrants distribution rather than Gibrat's law or imitation.

Power law distribution can come from proportional random growth also called Gibrat's law, as described in Gabaix 2009. The literature initiated by Luttmer (2007, 2010, 2012) build models that can generate fat-tailed distribution along a balanced growth path based on technology diffusion and spillover of technology by entrants. The flow of ideas literature (Lucas 2009, Lucas and Moll 2011 and Alvarez et al. 2008) also experiments growth and fat-tailed distribution along a balanced growth path.

In our framework, the fat-tailed distribution comes from a high persistence and large entrants. The persistence of the productivity process has an effect on the tail because with high persistence, firms benefit longer from a good shock and remain larger for a longer time than in a low persistence case. At the stationary equilibrium, there are more large firms and thus the tail is fatter. This can be seen in the left panel of figure

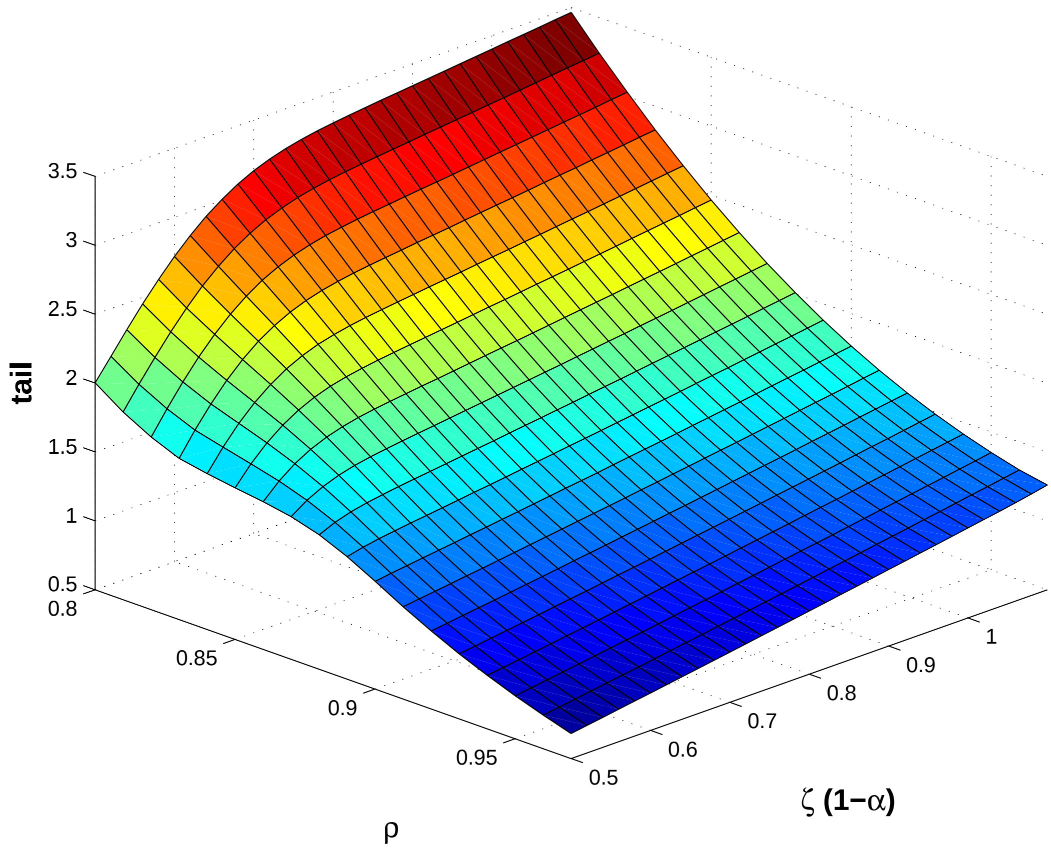

Figure 4: The tail of the stationary firm size distribution as a function of both ρ and $\zeta(1-\alpha)$, the tail of the potential entrant distribution G

3: as the persistence ρ increases, the tail parameter decreases which means that the tail becomes fatter.

The fatter the distribution (i.e the smaller the tail index) of potential entrants, the fatter the stationary firm size distribution as one can see on the right panel of figure 3. This is pretty intuitive, since bigger entrants also implies bigger firms and thus a fatter tail of the firm size distribution.

A high persistence and a fat-tailed distribution are both necessary to get fat-tailed firm size distribution at the stationary equilibrium. This can be seen in figure 4 which plots the tail of the stationary firm size distribution as a function of both the persistence of the productivity process and the tail of the potential entrant distribution. The tail of the firm size distribution gets close to one only for high persistence and fat tail of the potential entrants distribution.

	Model			Data		
	$\sigma(x)$	$\frac{\sigma(x)}{\sigma(y)}$	$\rho(x,y)$	$\sigma(x)$	$\frac{\sigma(x)}{\sigma(y)}$	$\rho(x,y)$
Output	2.3	1.0	1.0	2.2	1.0	-
Hours	1.5	0.67	1.0	1.8	0.83	0.85
Agg. Productivity	1.1	0.46	1.0	na	na	na

Table 3: Business Cycle Statistics

NOTE: These statistics are computed for the baseline calibration (cf. Table 2) for an economy simulated during 20,000 periods.

5.4 Business Cycle Statistics

After solving the model using the algorithm described in section 4, we compute the business cycle statistics. We simulate time series for output, hours and aggregate TFP using the law of motion (2) of the productivity distribution. These statistics are presented in table 3.

First, the standard deviation of output is 2.3%, which is in line with a real business cycle framework. To assess the performance of the model in producing aggregate fluctuations without any aggregate shock, a better statistic is the volatility of aggregate TFP. For the baseline calibration, the standard deviation of aggregate productivity is 1.1%, which is non-negligible. We are able to generate the same aggregate fluctuations as in the standard representative firm real business cycle (RBC) model with only idiosyncratic shocks and no aggregate shocks. The table 3 shows that a model with only idiosyncratic shocks can account for a large share of aggregate fluctuations.

The firm level dynamics of the baseline model is different from the one of a standard RBC model. In the latter all firms, or the representative firm, face similar productivity and have the same dynamics. In the baseline model, it is the aggregation of idiosyncratic shocks that gives rise to the aggregate fluctuations and thus each firm has its own dynamic.

The standard deviation of hours is 67% of the one of output whereas the same number is 83% in the data. Our model fails to match this number as it is the case in standard RBC model. It is the case because our model, without any aggregate shocks, evolves as a one-factor model. One could choose a higher value of the labor supply elasticity to match the relative volatility of hours but we choose to stay closer as possible to the RBC literature.

6 Shock on the biggest firm

In this section, we will describe the micro and macro consequences of a shock on the biggest firm. First, we will describe the methodology and then the impulse response to such shock.

6.1 Methodology

We assume that the biggest firm suffers a one standard deviation negative shock on its productivity. The standard deviation is σ_e and its calibrated value is 0.1053. The initial "reshuffling" shock on the firm size distribution ϵ_0 is only a vector of zero where at the highest level we subtract one and add one to the level of productivity corresponding to a one standard deviation on the left. Figure 5 shows on the left panel the initial distribution (blue) and the shocked distribution (dashed red). The difference between the two distributions can hardly be seen on the graph. The right panel of figure 5 gives a closer view at the right end of the difference between the shock and the initial distribution. In this figure, one can see that a mass of firms has been moved from the highest level to the left by a one standard deviation. This figure displays the initial shock ϵ_0.

From the structure of the model, computing the impulse response is straightforward. As one can see in equation 2, the transition between date t and date $t+1$ firm size distribution is a linear operator. So after computing the initial "reshuffling" shock ϵ_0, we do not need to simulate a large number of paths and to take the average. Instead, we assume ϵ_t to be zero for $t \geq 1$ and thanks to the linearity of transition described by equation 2 the result is exactly the same.

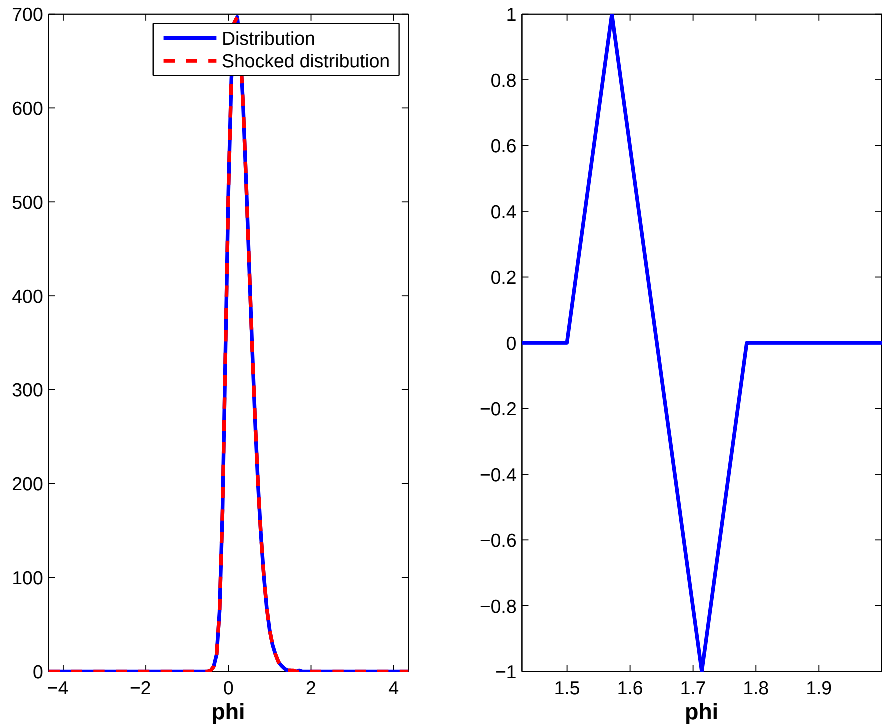

Figure 5: The shock and initial distribution (left panel) and the initial shock (right panel)

6.2 Impulse Responses

The top left panel of figure 6 shows the responses of aggregate productivity and aggregate output to this shock. Since our model behaves as a one-factor model, the dynamics of output follows closely the one of aggregate productivity. After this negative shock on the biggest firm, aggregate productivity decreases by 0.5% which is non negligible. This decrease in aggregate productivity has a non-negligible effect on the aggregate output. This decrease in output in turn decreases the aggregate labor demand and thus pushes the equilibrium wage down as can be seen in the top right panel of figure 6.

The intuition is as follows. Since the biggest firm becomes less productive, its optimal size becomes smaller. Since there is no frictions in adjustment of labor, the biggest firm shrinks by cutting its labor demand. The aggregate productivity, which is only an average of the productivity of each firm, decreases because one firm, the biggest, becomes less productive. Due to the decrease in the wage, the second biggest firm increases its size in response to the negative shock on the biggest firm. Since the wage that this second biggest firm is now facing is smaller and since its productivity remains the same, its optimal size is now bigger. The output of the second biggest firms increases as one can see on the bottom left panel of figure 6. Other firms benefit from the shrink of the biggest firm through competition in the labor market. This latter effect reduces the impact on aggregate output, since others firms produce more. However, this effect is not strong enough to mitigate the decrease in aggregate output because labor is reallocated towards less productive firms. The overall effect of a negative shock on the biggest firm is contractionary at the aggregate level.

In this setting firms that are subject to exit are the smallest firms. Because of the decrease in their cost, the wage, small firms can suffer bigger negative shocks while keeping their profitability positive. There are fewer exiters, so on impact the number of exiters decreases. In the meantime, entry becomes more profitable as the wage is smaller. The number of entrants becomes larger. However, there are more small entrants. Some potential entrants start producing whereas it would have not been the case if the wage had not decreased. The bottom right panel of figure 6 describes the dynamics of the number of incumbents, exiters and entrants. As the wage returns to its stationary value, the profitability of these entrants and the small firms that have survived, thanks to this decrease in the wage, shrinks. The number of small firms that

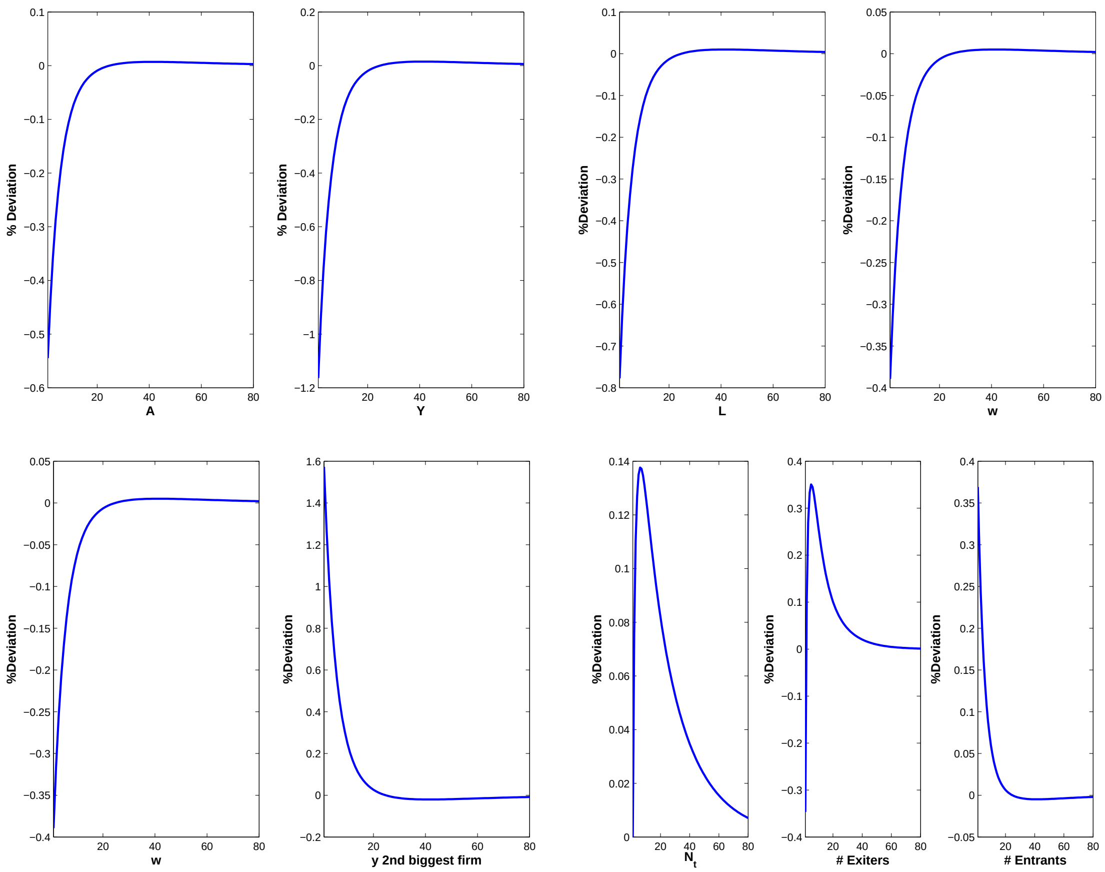

Figure 6: Impulse response to a one standard deviation negative productivity shock on the biggest firm

are subject to exit is thus larger. Overall the number of exiters exibits a hump-shaped response to this negative shock on the biggest firm.

The number of exiters hump-shaped pattern implies an hump-shaped response of the number of incumbents. Even if the maximum number of incumbents is somehow small - it becomes only bigger of about 0.14% of the its stationary value - it implies a small hump-shaped responses of aggregate variables.

Overall, we find that such a negative event on the biggest firm productivity is contractionary at the macro level and expansionary at the micro level with a small hump-shaped pattern.

7 Co-movement

In this section, we study the relative co-movement of large and small firms and their link with the cycle. We analyse the correlation between the differential growth rate of large versus small firms $\Delta\widehat{g}_{t-1,t}$ with aggregate conditions.

We follow the method of Moscarini and Postel-Vinay 2012 to define large and small firms. In their paper, small firms are defined as firms with less than 50 employees and large firms with more than 1000 employees. The differential net growth rate between large and small firms is defined as the difference between the employment growth rate of large firms minus the employment growth rate of small firms. Moscarini and Postel-Vinay 2012 finds that the correlation between this differential growth rate and the deviation of unemployment from an HP trend is negative and equal to -0.52. During "good" times, when unemployment is low, the differential growth rate is higher i.e. big firms create more employment than small firms.

To compute the corresponding number in our baseline model, we extract a representative sample of firms and follow a balanced panel. From this balanced panel we compute the same statistics as in Moscarini and Postel-Vinay 2012. The closest variable to unemployment in our setting is the hours L_t. We compute then the correlation between our differential growth rate and the hours. The correlation is positive and equal to 0.24. The correlation is then of the same sign as in the data. During "good" times, when the number of hours is high, large firms create more employment than small firms.

Moscarini and Postel-Vinay 2012 US BDS data	Baseline model Simulated data
$\Delta\widehat{g}_{t-1,t}$ Correlation with $\widehat{u}_t$	$\Delta\widehat{g}_{t-1,t}$ Correlation with L_t
-0.52 (0.003)	0.24 (0.03)

Table 4: Correlation of the differential growth rate between large and small firms and the cycle indicator (p-value in parenthesis)

Moscarini and Postel-Vinay 2012 argues that this correlation supports the existence of labor adjustment cost. In their approach the causality is that aggregate conditions determine the faster growth of big firms relatively to small firms because of these labor adjustment costs. In our model, the causality is reversed: the fast growth of large firms determines the aggregate condition, and so these large firms are naturally more cyclicality sensitive than small firms.

8 Conclusion

We build a quantitative firm dynamics model in which we cast "the granular hypothesis". It features a finite number of firms and we do not rely on any "law of large" number assumption. Our model is able to generate aggregate fluctuations from idiosyncratic shocks only. We show analytically that aggregate fluctuations do not die out as the number of firms increases. A calibrated version of our framework to the US economy implies sizable aggregate volatility. We also look at the macro and micro effect of a negative shock on the biggest firm. Finally, we show that our model can reproduce the relative cyclicality of large firms compared to small firms'.

In our framework, a firm does not internalize its own effect on the aggregate wage. In others models of firms dynamics, the assumption that the "law of large" numbers holds justifies thinking about firms as price takers. However, we are showing in a standard firm dynamics framework that large firms actually have an effect on the aggregate and thus the price taker assumption should be taken carefully. Our work is a first step toward this objective.

References

[1] Daron Acemoglu, Vasco M. Carvalho, Asuman Ozdaglar, and Alireza Tahbaz-Salehi. The network origins of aggregate fluctuations. *Econometrica*, 80(5):1977–2016, 09 2012.

[2] Fernando E. Alvarez, Francisco J. Buera, and Jr. Robert E. Lucas. Models of idea flows. NBER Working Papers 14135, National Bureau of Economic Research, Inc, June 2008.

[3] Ruediger Bachmann and Christian Bayer. Firm-specific productivity risk over the business cycle: Facts and aggregate implications. CESifo Working Paper Series 2844, CESifo Group Munich, 2009.

[4] Susanto Basu and John G Fernald. Returns to scale in u.s. production: Estimates and implications. *Journal of Political Economy*, 105(2):249–83, April 1997.

[5] Florin O. Bilbiie, Fabio Ghironi, and Marc J. Melitz. Endogenous entry, product variety, and business cycles. *Journal of Political Economy*, 120(2):304 – 345, 2012.

[6] Jeffrey Campbell. Entry, exit, embodied technology, and business cycles. *Review of Economic Dynamics*, 1(2):371–408, April 1998.

[7] Vasco M Carvalho. Aggregate fluctuations and the network structure of intersectoral trade. mimeo, 2010.

[8] Vasco M. Carvalho and Xavier Gabaix. The great diversification and its undoing. *American Economic Review*, Forthcoming, 2013.

[9] Rui Castro, Gian Luca Clementi, and Yoonsoo Lee. Cross-sectoral variation in firm-level idiosyncratic risk. *Journal of Industrial Economics*, Forthcoming, 2013.

[10] Gian Luca Clementi and Dino Palazzo. Entry, exit, firm dynamics, and aggregate fluctuations. Working Paper Series 27 10, The Rimini Centre for Economic Analysis, January 2010.

[11] Wouter J. Den Haan. Assessing the accuracy of the aggregate law of motion in models with heterogeneous agents. *Journal of Economic Dynamics and Control*, 34(1):79–99, January 2010.

[12] Julian di Giovanni and Andrei A. Levchenko. Country size, international trade, and aggregate fluctuations in granular economies. *Journal of Political Economy*, 120(6):1083 – 1132, 2012.

[13] Julian di Giovanni, Andrei A. Levchenko, and Isabelle Mejean. Firms, destinations, and aggregate fluctuations. CEPR Discussion Papers 9168, C.E.P.R. Discussion Papers, October 2012.

[14] Xavier Gabaix. Power laws in economics and finance. *Annual Review of Economics*, 1(1):255–294, 05 2009.

[15] Xavier Gabaix. The granular origins of aggregate fluctuations. *Econometrica*, 79(3):733–772, 05 2011.

[16] Hugo A Hopenhayn. Entry, exit, and firm dynamics in long run equilibrium. *Econometrica*, 60(5):1127–50, September 1992.

[17] Aubhik Khan and Julia K. Thomas. Nonconvex factor adjustments in equilibrium business cycle models: do nonlinearities matter? *Journal of Monetary Economics*, 50(2):331–360, March 2003.

[18] Aubhik Khan and Julia K. Thomas. Idiosyncratic shocks and the role of nonconvexities in plant and aggregate investment dynamics. *Econometrica*, 76(2):395–436, 03 2008.

[19] Karen Kopecky and Richard Suen. Finite state markov-chain approximations to highly persistent processes. *Review of Economic Dynamics*, 13(3):701–714, July 2010.

[20] Per Krusell and Anthony A. Smith. Income and wealth heterogeneity in the macroeconomy. *Journal of Political Economy*, 106(5):867–896, October 1998.

[21] Yoonsoo Lee. The importance of reallocations in cyclical productivity and returns to scale: evidence from plant-level data. Working Paper 0509, Federal Reserve Bank of Cleveland, 2005.

[22] Yoonsoo Lee and Toshihiko Mukoyama. Entry, exit and plant-level dynamics over the business cycle. Working Paper 0718, Federal Reserve Bank of Cleveland, 2008.

[23] Robert E. Lucas. Ideas and growth. *Economica*, 76(301):1–19, 02 2009.

[24] Robert E. Lucas and Benjamin Moll. Knowledge growth and the allocation of time. NBER Working Papers 17495, National Bureau of Economic Research, Inc, October 2011.

[25] Erzo G. J. Luttmer. Selection, growth, and the size distribution of firms. *The Quarterly Journal of Economics*, 122(3):1103–1144, 08 2007.

[26] Erzo G.J. Luttmer. Models of growth and firm heterogeneity. *Annual Review of Economics*, 2(1):547–576, 09 2010.

[27] Erzo G.J. Luttmer. Technology diffusion and growth. *Journal of Economic Theory*, 147(2):602–622, 2012.

[28] Giuseppe Moscarini and Fabien Postel-Vinay. The contribution of large and small employers to job creation in times of high and low unemployment. *American Economic Review*, 102(6):2509–39, October 2012.

[29] K.G. Rouwenhorst. Asset pricing implications of equilibrium business cycle models. In T.F. Cooley, editor, *Frontiers of Business Cycle Research*, pages 294–330. Princeton University Press, 1995.

[30] George Tauchen. Finite state markov-chain approximations to univariate and vector autoregressions. *Economics Letters*, 20(2):177–181, 1986.

[31] Marcelo L. Veracierto. Plant-level irreversible investment and equilibrium business cycles. *American Economic Review*, 92(1):181–197, March 2002.

A Discretization of AR(1) by the Rouwenhorst method

Consider the AR(1) process $\widetilde{\varphi}_{t+1} = \rho\widetilde{\varphi}_t + e_t$ where $e_t^i \rightsquigarrow \mathcal{N}(0, \sigma_e)$. We approximate this process by a discrete process $\{\varphi_t\}$ over the evenly distributed state space $\{\varphi_1, \dots, \varphi_{n_s}\}$. For $p, q \in (0, 1)$ and each n_s, we can defined recursively the matrix:

$$
\begin{aligned}
P_2 &= \begin{pmatrix} p & 1-p \\ 1-q & q \end{pmatrix} \text{ if } n_s = 2 \\
P_{n_s} &= p \begin{pmatrix} P_{ns-1} & 0 \\ 0' & 0 \end{pmatrix} + (1-p) \begin{pmatrix} 0 & P_{ns-1} \\ 0 & 0 \end{pmatrix} + (1-q) \begin{pmatrix} 0 & 0 \\ P_{ns-1} & 0 \end{pmatrix} \\
&\quad + q \begin{pmatrix} 0 & 0 \\ 0 & P_{ns-1} \end{pmatrix} \text{ if } n_s > 2
\end{aligned}
$$

and normalizing all but the first and last rows in order that they sum to one. These matrix is a transition matrix for a discrete Markovian process. Rouwenhorst 1995 and Kopecky and Suen 2010 show that for $p = q = \frac{1+\rho}{2}$ and $\varphi_{n_s} = \sqrt{n_s - 1}\frac{\sigma_e}{\sqrt{1-\rho^2}}$ we have:

$$
\begin{aligned}
\mathbb{E}(\varphi_{t+1}|\varphi_t = \varphi_s) &= \rho\varphi_s \\
\mathbb{V}ar(\varphi_{t+1}|\varphi_t = \varphi_s) &= \sigma_e^2
\end{aligned}
$$

Thus the Markovian process associated with the transition matrix P_{n_s} have a first order autocorrelation and a conditional variance equal to ρ and σ_e^2 respectively. Kopecky and Suen 2010 also shows that this discretization of an AR(1) is preferable for highly persistent process.

B Proof

B.1 Proof of the lemma 1

The law of motion of the firm size distribution as stated by equation 2 is $\mu_{t+1} = m(\mu_t) + \epsilon_t$ or

$$\mu_{t+1} = (P_t^*)'(\mu_t + MG) + \epsilon_t$$

The stationary distribution is defined as the solution of the following equation:

$$\mu = (P^*)'\mu + M(P^*)'G$$

Let us assume that $s_t^* = s^*$ and $\omega_t = \omega$ and thus that $P_t^* = P^*$ at each date. This assumption does not affect the results[5] and is just made for sake of simplicity of the proof.

Subtracting the two previous equations yields:

$$\mu_{t+1} - \mu = (P^*)'(\mu_t - \mu) + \epsilon_{t+1}$$

Recall that at $t = 0$ the $\mu_0 = \mu$, let us iterate recursively on the above equation:

$$\begin{aligned}
\mu_{t+1} - \mu &= (P^*)'(\mu_t - \mu) + \epsilon_{t+1} \\
&= (P^{*'})^2(\mu_{t-1} - \mu) + P^{*'}\epsilon_t + \epsilon_{t+1} \\
&= \dots \\
&= (P^{*'})^{t+1}(\mu_0 - \mu) + \sum_{j=0}^{t}(P^{*'})^j\epsilon_{t+1-j} \\
\mu_{t+1} - \mu &= \sum_{j=0}^{t}(P^{*'})^j\epsilon_{t+1-j} \text{ since } \mu_0 = \mu \qquad (6)
\end{aligned}$$

We need to define the quantity that we interested in, the tail. The tail of the distribution is just the behavior of the counter cumulative mass function. By showing that two counter cumulative mass functions are similar, we show formally that the associated distribution have the same tail.

Let us define $\hat{\mu}_t = \frac{\mu_t}{M}$. The economy scale with M (assuming that the labor supply also increases[6]). When M goes to infinity so does μ_t^s for all s and t, but $\hat{\mu}_t^s$ does not. The counter cumulative mass function that we are interested in is thus:

$$\forall k \in [|1, n_s|] \qquad \left|\sum_{i=k+1}^{n_s}\hat{\mu}_{t+1}^i - \sum_{i=k+1}^{n_s}\hat{\mu}^i\right| = \left|\sum_{i=k+1}^{n_s}\frac{\mu_{t+1}^i}{M} - \sum_{i=k+1}^{n_s}\frac{\mu^i}{M}\right|$$

The left term in the left hand side equation is the counter cumulative mass function of $\hat{\mu}_{t+1}$, the right term is the same function for $\hat{\mu}$.

[5]This assumption only affects the left-end of the firm size distribution and thus not the right tail.
[6]The wage is kept constant by increasing $\underline{\text{L}}$ as M increases: the ratio $M/\underline{\text{L}}$ is constant.

$$\forall k \in [|1, n_s|] \qquad \left|\sum_{i=k+1}^{n_s} \hat{\mu}^i_{t+1} - \sum_{i=k+1}^{n_s} \hat{\mu}^i\right| = \left|\sum_{i=k+1}^{n_s} \frac{\mu^i_{t+1}}{M} - \sum_{i=k+1}^{n_s} \frac{\mu^i}{M}\right|$$

$$= \left|\sum_{i=k+1}^{n_s} \sum_{j=0}^{t} \Big[\frac{(P^{*'})^j \epsilon_{t+1-j}}{M}\Big]_i\right|$$

$$= \Bigg|\sum_{i=k+1}^{n_s} \sum_{j=0}^{t} \sum_{s=s^*}^{n_s} \Big[\frac{(P^{*'})^j f^{\cdot,s}_{t+1-j}}{M} + \frac{(P^{*'})^j g^{\cdot,s}_{t+1-j}}{M} - \mathbb{E}\left[\frac{(P^{*'})^j f^{\cdot,s}_{t+1-j}}{M} + \frac{(P^{*'})^j g^{\cdot,s}_{t+1-j}}{M}\right]\Big]_i\Bigg|$$

Soit $\varepsilon > 0$,

$$\mathbb{P}\left(\left|\sum_{i=k+1}^{n_s} \hat{\mu}^i_{t+1} - \sum_{i=k+1}^{n_s} \hat{\mu}^i\right| \geq \varepsilon\right)$$

$$= \mathbb{P}\left(\left|\sum_{i=k+1}^{n_s} \sum_{j=0}^{t} \sum_{s=s^*}^{n_s} \Big[\frac{(P^{*'})^j f^{\cdot,s}_{t+1-j}}{M} + \frac{(P^{*'})^j g^{\cdot,s}_{t+1-j}}{M} - \mathbb{E}\left[\frac{(P^{*'})^j f^{\cdot,s}_{t+1-j}}{M} + \frac{(P^{*'})^j g^{\cdot,s}_{t+1-j}}{M}\right]\Big]_i\right| \geq \varepsilon\right)$$

$$\leq \mathbb{P}\left(\sum_{j=0}^{t} \sum_{s=s^*}^{n_s} \left|\sum_{i=k+1}^{n_s} \Big[\frac{(P^{*'})^j f^{\cdot,s}_{t+1-j}}{M} + \frac{(P^{*'})^j g^{\cdot,s}_{t+1-j}}{M} - \mathbb{E}\left[\frac{(P^{*'})^j f^{\cdot,s}_{t+1-j}}{M} + \frac{(P^{*'})^j g^{\cdot,s}_{t+1-j}}{M}\right]\Big]_i\right| \geq \varepsilon\right)$$

$$\leq \sum_{j=0}^{t} \sum_{s=s^*}^{n_s} \mathbb{P}\left(\left|\sum_{i=k+1}^{n_s} \Big[\frac{(P^{*'})^j f^{\cdot,s}_{t+1-j}}{M} + \frac{(P^{*'})^j g^{\cdot,s}_{t+1-j}}{M} - \mathbb{E}\left[\frac{(P^{*'})^j f^{\cdot,s}_{t+1-j}}{M} + \frac{(P^{*'})^j g^{\cdot,s}_{t+1-j}}{M}\right]\Big]_i\right| \geq \varepsilon\right)$$

$$\leq \sum_{j=0}^{t} \sum_{s=s^*}^{n_s} \frac{1}{\varepsilon^2} \mathbb{V}ar\left[\sum_{i=k+1}^{n_s} \Big[\frac{(P^{*'})^j f^{\cdot,s}_{t+1-j}}{M} + \frac{(P^{*'})^j g^{\cdot,s}_{t+1-j}}{M}\Big]_i\right] \text{ by the Bienaymé-Chebyshev inequality}$$

$$\leq \sum_{j=0}^{t} \sum_{s=s^*}^{n_s} \frac{1}{\varepsilon^2} \mathbb{V}ar\left[\sum_{i=k+1}^{n_s} \Big[\frac{(P^{*'})^j f^{\cdot,s}_{t+1-j}}{M}\Big]_i\right] + \mathbb{V}ar\left[\sum_{i=k+1}^{n_s} \Big[\frac{(P^{*'})^j g^{\cdot,s}_{t+1-j}}{M}\Big]_i\right] \qquad (7)$$

since the $f^{i,s}_t$ and $g^{i,s}_t$ are independent

Recall that $f^{i,s}_{t+1-j} \rightsquigarrow Multi\Big(\mu^s_{t-j}, P_{s,.}\Big)$ and $g^{i,s}_{t+1-j} \rightsquigarrow Multi\Big(MG^s, P_{s,.}\Big)$.

It follows that:

$$\mathbb{V}ar\left[\sum_{i=k+1}^{n_s}[\frac{(P^{*'})^j f^{\cdot,s}_{t+1-j}}{M}]_i\right] = \mathbb{V}ar\left[\frac{(e_k)'(P^{*'})^j f^{\cdot,s}_{t+1-j}}{M}\right]$$
$$= \mu^s_{t+1-j}\frac{(e_k)'(P^{*'})^j W_s (P^*)^j (e_k)}{M^2}$$
$$= \hat{\mu}^s_{t+1-j}\frac{(e_k)'(P^{*'})^j W_s (P^*)^j (e_k)}{M} \longrightarrow 0$$

this last term goes to zero when M goes to infinity. Here, e_k is a (n_s) vector with the first k elements equal to zero.

Since all the sums in equation 7 are finite, we can conclude that

$$\forall k \in [|1, n_s|], \forall \varepsilon > 0, \mathbb{P}\left(\left|\sum_{i=k+1}^{n_s}\hat{\mu}^i_{t+1} - \sum_{i=k+1}^{n_s}\hat{\mu}^i\right| \geq \varepsilon\right) \longrightarrow 0 \text{ as } M \to \infty$$

This just means that the counter cumulative mass function, the tail, of μ_{t+1} converge in probability toward the tail of μ as M goes to infinity. □

B.2 Proof of the theorem 1

Let us first compute the aggregate output Y_t as a function of only T_t:

$$Y_t = \sum_{i=1}^{N_t} y^i_t = \sum_{i=1}^{N_t}\exp(\varphi^i_t)^{\frac{1}{1-\alpha}}\left(\frac{\alpha}{w_t}\right)^{\frac{\alpha}{1-\alpha}} = \left(\frac{\alpha}{w_t}\right)^{\frac{\alpha}{1-\alpha}} T_t$$

Recall that:

$$w_t = \left(\alpha^{\frac{1}{1-\alpha}} T_t\right)^{\frac{1-\alpha}{\gamma(1-\alpha)+1}}$$

Substituting this expression of the wage in the latter equation and taking the growth rate yields:

$$Y_t = \alpha^{\frac{\alpha\gamma}{\gamma(1-\alpha)+1}}\left(T_t\right)^{1-\frac{\alpha}{\gamma(1-\alpha)+1}}$$
$$\frac{\Delta Y_t}{Y_t} = \left(1 - \frac{\alpha}{\gamma(1-\alpha)+1}\right)\frac{\Delta T_t}{T_t}$$

This shows the growth rate of the output is proportional to the growth rate of T_t.

Let us evaluate the standard deviation of the growth rate of T_t. To do so, let us have a close look at this growth rate. Note that

$$T_{t+1} = \sum_{l \text{ successful incumbent at } t} \exp(\frac{\varphi_{t+1}^l}{1-\alpha}) + \sum_{e \text{ successful entrant at } t} \exp(\frac{\varphi_{t+1}^e}{1-\alpha})$$

and

$$T_t = \sum_{l \text{ successful incumbent at } t} \exp(\frac{\varphi_t^l}{1-\alpha}) + \sum_{x \text{ exiters at } t} \exp(\frac{\varphi_t^x}{1-\alpha})$$

Let us define $Z_t^l = \exp(\frac{\varphi_t^l}{1-\alpha})$ for l-successful incumbent at t, $E_{t+1}^e = \exp(\frac{\varphi_{t+1}^e}{1-\alpha})$ for e successful entrant at t and $X_t^x = \exp(\frac{\varphi_t^x}{1-\alpha})$ for x exiters at t.

The growth rate of T_t is:

$$\frac{\Delta T_t}{T_t} = \frac{1}{T_t}\left(\sum_l \Delta Z_t^l + \sum_e E_{t+1}^e - \sum_x X_t^x\right)$$

The variance of the Z_t^l follows from the productivity process. For a firm i, we have $\varphi_{t+1}^i = \rho\varphi_t^i + e_{t+1}^i$, where e_t^i is random variable with mean zero and variance σ_e^2. This is true since the first order autocorrelation and condition volatility is well defined and constant across the idiosyncratic state space. By taking the exponential and subtracting, this yields:

$$\exp(\frac{\varphi_{t+1}^i}{1-\alpha}) - \exp(\frac{\varphi_t^i}{1-\alpha}) = \exp(\frac{\varphi_t^i}{1-\alpha})^\rho \exp(\frac{e_t^i}{1-\alpha}) - \exp(\frac{\varphi_t^i}{1-\alpha})$$

Furthermore the e_t^i are i.i.d. across firms and time. Let us define $\sigma = \sqrt{\mathbb{V}ar\left(\exp(\frac{e_t^i}{1-\alpha})\right)}$. It follows that:

$$\mathbb{V}ar_t\left(\Delta Z_t^l\right) = (Z_t^l)^{2\rho}\sigma^2$$

and, by the same argument,

$$\mathbb{V}ar_t\left(E_{t+1}^e\right) = (E_t^e)^{2\rho}\sigma^2$$

where $E_t^e = \exp(\frac{\varphi_t^e}{1-\alpha})$ with φ_t^e the signal at t of the successful entrant e.

This leads to

$$\mathbb{V}ar_t\frac{\Delta T_t}{T_t} = \frac{\sigma^2}{(T_t)^2}\left(\sum_l (Z_t^l)^{2\rho} + \sum_e (E_t^e)^{2\rho}\right)$$

since the variance conditional on date t of X_t^x is equal to zero.

Denoting N_t^l, N_t^E and N_t^X the number of successful incumbents, successful entrants and exiters at date t respectively. According to the law of large number, we have:

$$(N^l)^{-1}\sum_l Z_t^l \to \mathbb{E}Z_t^l := \bar{Z}_t$$
$$(N^X)^{-1}\sum_l X_t^x \to \mathbb{E}X_t^x := \bar{X}_t$$

It is straightforward that

$$N^{-1}T_t \sim \frac{N_t^l}{N_t}\bar{Z}_t + \frac{N_t^X}{N_t}\bar{X}_t := \bar{I}_t$$

the average of incumbent size at date t (both successful and exiters).

The distribution of the random variable Z_t^l has a power law tail with parameters ξ, the tail parameter of firm size distribution (since only small firms exit). The distribution of the random variable Z_t^l has a power law tail with parameters $\zeta' = \zeta(1-\alpha)$ the tail parameter of entrant size distribution (since only big entrants are successful).

Since $\xi/\rho < 2$ and $\zeta'/\rho < 2$ and using the Lévy theorem of the appendix of Gabaix 2011, we have

$$(N_t^l)^{-2\rho/\xi}\sum_l (Z_t^l)^{2\rho} \to^d u$$
$$(N_t^X)^{-2\rho/\zeta'}\sum_x (X_t^x)^{2\rho} \to^d w$$

where u and w are standard Lévy distribution with parameters $\xi/2\rho$ and $\zeta'/2\rho$ respectively.

Computing the two above results yields

$$\mathbb{V}ar_t\frac{\Delta T_t}{T_t} \sim N^{-2}\frac{\sum_l (Z_t^l)^{2\rho} + \sum_x (X_t^x)^{2\rho}}{(\bar{I}_t)^2}$$

Note the numerator of the right hand side is equivalent to

$$N^{-2+2\rho/\xi}\left(\left(\frac{N_t^l}{N_t}\right)^{2\rho/\xi} u\right) \text{ if } \zeta' > \xi$$

since $N^{2\rho(1/\zeta-1/\xi)} \to 0$ in this case.

Similarly:

$$N^{-2+2\rho/\zeta'}\left(\left(\frac{N_t^E}{N_t}\right)^{2\rho/\zeta'}u\right) \text{ if } \zeta' < \xi$$

This gives the results:

$$\sigma(\frac{\Delta T_t}{T_t}) \sim \frac{\sigma}{N_t^{1-\rho/\xi}}\frac{\left(\frac{N_t^l}{N_t}\right)^{\rho/\xi}u^{1/2}}{\bar{I}_t} \qquad \text{if } \zeta' > \xi$$

$$\sigma(\frac{\Delta T_t}{T_t}) \sim \frac{\sigma}{N_t^{1-\rho/\zeta'}}\frac{\left(\frac{N_t^E}{N_t}\right)^{\rho/\zeta'}w^{1/2}}{\bar{I}_t} \qquad \text{if } \zeta' < \xi$$

□