The distribution of firm growth, and business cycles

Carlos Daniel Santos[*,†,‡]

December 2022

Abstract

Large firms can magnify aggregate fluctuations from microeconomic shocks. However, it is not only the size *level* that exhibits heavy tails. Using Portuguese census data we show that *growth* rates have: (i) heavy tails; (ii) volatility decreasing with size; (iii) kurtosis increasing with size. These shocks have aggregate implications. Following evidence on earnings' dynamics, we introduce consumer heterogeneity into a CES-demand model to explain our findings qualitatively and quantitatively. Larger firms diversify their sales risk by selling to more customers. However, customer concentration limits diversification and increases tail risks. This explains the stronger responsiveness of small firms and the heavier tails for larger firms.

Keywords: Business cycles, Firm dynamics, Granularity

JEL Codes: C23, E32, L16, L22

[*]Financial support by FCT (Fundação para a Ciência e a Tecnologia), Portugal, is gratefully acknowledged. This article is part of projects PTDC/EGE-ECO/28814/2017. I would like to thank comments from Dave Donaldson, Luis Costa, Helmut Luetkepohl, Paulo Rodrigues, and participants at the Conference on Heavy Tails and Robust Estimation and the Conference on New Trends and Developments in Econometrics. All remaining errors are my own.

[†]REM Research in Economics and Mathematics, UECE Research Unit on Complexity and Economics, Rua Miguel Lupi 20, 1249-078 Lisboa, Portugal.

[‡]ISEG - Lisbon School of Economics & Management, Universidade de Lisboa, Rua do Quelhas 6, 1200-781 Lisboa, Portugal.

1 Introduction

"What causes aggregate fluctuations?" (Carvalho and Gabaix, 2013). Recent research has demonstrated that with the observed granularity (very large firms), microeconomic shocks may generate aggregate fluctuations (Gabaix, 2016). Models that ignore such amplification may deliver incorrect policy multipliers and misrepresent the mechanism for shock transmission (Ascari, Fagiolo, and Roventini, 2015).

To explain business cycle fluctuations, the literature has either used firm size distributions with heavy tails (Gabaix, 2011) or the asymmetry of input-output relations (network structure) of the economy (Acemoglu, Carvalho, et al., 2012) to show that "idiosyncratic shocks do not die out in the aggregate". However, the existence of either very large firms or networks is insufficient to obtain the tail risk for macroeconomic variables. Furthermore, "aggregating normally distributed shocks always results in normally distributed shocks" (Acemoglu, Ozdaglar, and Tahbaz-Salehi, 2017). It is necessary to combine microeconomic tail risks with granularity (large firms/sectors) to obtain a slower convergence rate for the law of large numbers. A slow convergence implies that aggregating different tail risks at the micro level does not (quickly) disappear at the macro level.[1]

Our contribution is two-fold: a set of empirical regularities and a theoretical microfoundation. We first use census data for Portugal to show that firm growth rates (i) exhibit heavy tails, (ii) the volatility decreases with firm size, and (iii) the kurtosis increases with firm size. The smaller volatility of large firms limits the potential for the transmission of shocks (Yeh, 2021). However, the larger kurtosis counteracts this and further amplifies the transmission of microeconomic shocks. Consequently, besides the heavy tails in firm size distribution (levels) already documented in the literature, we show that growth rates exhibit heavy tails that vary with firm size.[2]

The shocks have aggregate implications. For example, the kurtosis of firm-level growth rates is strongly correlated with GDP growth (correlation of 0.78). When put together, the results challenge the notion of risk diversification. Kramarz, Martin, and Mejean (2020) find that the concentration of sales in a few large customers limits

[1] Technically, the conditions for Cramér's theorem are violated, resulting in a much slower converge of aggregate behavior.

[2] The results are robust to the exclusion of expansionary/recessive years, large/small firms, and the definition of size.

the potential for diversification, even for large firms. Because of granularity, large positive and negative shocks may not cancel out. Since sales are concentrated in a few top customers, increasing the number of firms leads to a smaller gain from diversification. Our result suggest that besides volatility, one must also consider the behavior of the tails. The larger magnitude of the shocks (kurtosis) counteracts the potential attenuation of the reduction in volatility for the larger firms. As such, larger firms do not have their risks substantially lowered, further contributing to aggregate fluctuations.[3]

Our second contribution is to develop an heterogeneous consumer model that explains the heavy tails in the distribution of firm growth rates and the reduction in volatility/increase in kurtosis for larger firms.[4] Recent evidence shows that the distribution of workers' income volatility also exhibits large variances and kurtosis (Guneven et al., 2021). We use this to motivate the introduction of customer level shocks. Furthermore, in models with a representative consumer, any shock to individuals will generate proportional effects to firms, even if there is firm heterogeneity. We must thus have heterogeneous consumers to generate asymmetric shocks. This departure from the representative agent is important as it also addresses the Barro-King criticism (Barro and King, 1984). Consumers' asymmetric responses during the business cycle can generate positive co-movement in employment, consumption and production (Beaudry and Portier, 2014).

In a nutshell, microeconomic discreteness on the demand side coupled with a skewed income distribution can explain our empirical findings. The volatility of shocks decreases with firms that sell to more customers (Fact 2). It is not immediately obvious why the kurtosis increases with firm size (Fact 3). Customer concentration emerging from the income distribution is an important feature to explain this tail behavior. As companies sell to more customers, volatility decreases. However, if sales are concentrated in a few firms, the behavior of the tail becomes more prevalent, thus leading to an increase in kurtosis.

The demand-side formulation is a departure from the standard real business cycle

[3]When the source of volatility is due to input-output linkages, we expect a larger volatility in manufacturing (more intermediate input intensive) and smaller in services (less intermediate input intensive) as in Moro (2012). We find no such effect. Volatility is stable across different levels of intermediate input intensity.

[4]Technically, heavy tails emerge from a mixture of normally distributed shocks with varying variances.

literature. In this case, shocks to consumer preferences and choices can generate and amplify fluctuations. While this formulation is consistent with recent empirical findings (di Giovanni, Levchenko, and Mejean (2014)), this is still a departure from the theoretical literature that has focused on the granularity of supply side shocks as an explanation for business cycle fluctuations, following a long tradition on RBC models (e.g. Carvalho and Grassi, 2019). In the standard supply side model, TFP shocks scale across all of the firm's production. Sales are proportional to TFP and, in this setting, a larger firm will not have a smaller volatility or a larger kurtosis of sales growth. However, there is a mismatch between theory and empirics (Mankiw, 1989). *Measured* productivity (or TFPR - total factor productivity in revenues) is not a clean measure as the theory postulates (Foster, Haltiwanger, and Syverson, 2016). *Measured* TFP shocks are a mixture of demand and supply side shocks. We thus reconcile theory and empirics by endogenizing firm sales volatility from consumer heterogeneity and preference dynamics - a demand-driven real business cycle model.

2 Empirical Results

Our data consists of yearly census data for non-financial firms in 2004-2019 (Statistics Portugal, 2019b). The dataset is collected by the Statistics Portugal and can be accessed by any accredited researcher (Statistics Portugal, 2019b). It contains over 5 million observations for about 300 to 380 thousand firms per year. This is a particular advantage of our dataset given the difficulty (precision) in estimating higher order moments in small samples. For each firm-year we obtain revenues, labor costs, capital, and employment.[5]

2.1 Sales growth distribution

Fact 1. *The distribution of sales growth exhibits heavy tails.*

To describe the tail behavior, let ζ be the tail index of a power law variable with cumulative distribution $1 - F(x) \sim Cx^{-\zeta}$. This can be estimated by OLS in log-log regression $\ln(1 - F(x)) = c - \zeta \ln(x)$ as in Gabaix and Ibragimov (2011).

As reported in Figure 1, while the tail is well approximated with a linear function, it also contains a slight curvature. The tail is estimated at about 3.3 using a threshold

[5]Further details about the data can be found in the replication file and appendix.

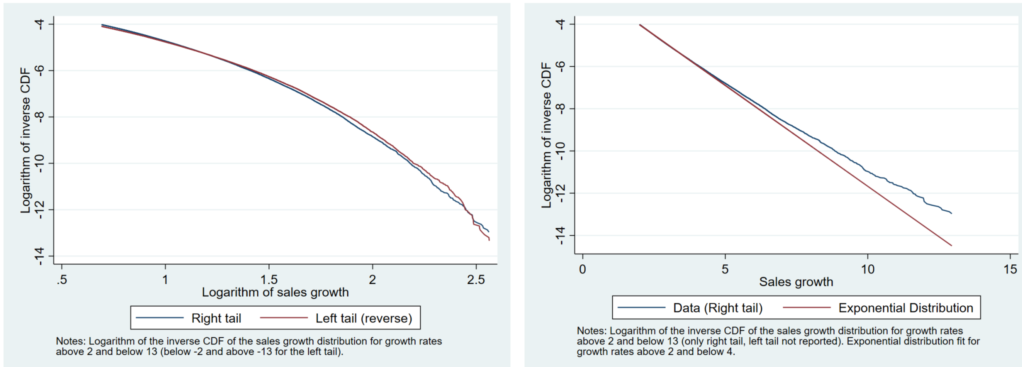

Figure 1: Log-log and log-level density-growth plot and tail exponent estimates.

for sales growth set at -2/2 for the left/right tail. If we use a threshold of -1/1 for the left/right tail, we obtain a smaller coefficient, confirming the curvature.[6] This suggests that the empirical tails decay faster than predicted by the power law distribution: ζ is not constant.[7]

Besides the log-log plot (power-law), we also report a log-level (exponential) plot in Figure 1.. The logarithm of the inverse CDF of an exponential distribution is linear in levels. While the distribution is close to exponential, it presents slower decay. This is consistent with heavy tails (the definition of heavy tails is that of a distribution that decay slower than an exponential distribution). Furthermore, the kurtosis is larger than 9, which is the kurtosis of the exponential distribution.

An alternative method to evaluate the "fatness" of a distribution is the probability of a shock more than three standard deviations away from the mean (tail risk). In the data this correspond to a probability of 2.6%, that compares to the probability in the normal distribution of 0.27%.[8] Large shocks (three standard deviations above the mean) are much (10 times) more frequent than what we expect from the normal distribution.

[6] Table A.3 reports the more detailed results.

[7] We also note that the shock distribution is similar across years with simultaneous long positive and negative tails both in expansions as well as recessions (Figure A.2) with small overall shifts in the cumulative density function (Figure A.3).

[8] Table A.4.

2.2 Size-volatility

To analyze the size relation, we discretize size into 5% intervals thus creating 20 bins {5%, 10%, .., 90%, 95%, 100%}. We use four definitions of size: sales, employment, capital stock, and labor productivity. [9]

Fact 2. *Sales growth volatility decreases with firm size.*

To compute volatility, we estimate the following dummy variable regression

$$\left[dy_{it} - \sum_s \mu_{dy(s)} 1(i \in s) \right]^2 = c + \sum_s \alpha_s 1(i \in s) + \epsilon_{it}$$

where s is a size bin, $1(i \in s)$ is a dummy equal to 1 if firm i is in bin s, $\mu_{dy(s)}$ is the average sales growth of firms in bin s, ϵ is a residual, and c,$\{\alpha_s\}_{s=1}^{S}$are parameters. The regression estimates the variance for each bin. As reported in Figure 2, the standard deviation of sales growth is not constant across firms of different size. The pattern of heteroskedasticity is markedly size dependent, with small firms more volatile than large firms.[10]

Sales growth volatility decreases from over 1.5 for the smallest 10% of firms to about 0.35 for the largest 10% firms. This is robust to the definition of size used (labor costs or capital). The fall in volatility is slightly smaller if we use the capital stock as a measure of size. This may signal the stock of capital is not aligned with firm size, thus reflecting potential distortions in capital allocation.[11] Finally, there is a slight u-shaped relation between volatility and labor productivity. That is, volatility decreases with labor productivity for firms in the bottom half of the productivity distribution but increases with labor productivity for firms in the top half of the productivity distribution. In particular, the largest firms have a sales volatility of 0.35 while the most productive firms have a sales volatility almost three times larger. This suggests that (labor) productivity is not proportional to size.

Gabaix (2011) posits the following log-linear relation

[9] To avoid mean-reversion, for variable x, we show the results are robust to the use of symmetric growth rates Haltiwanger, Jarmin, and Miranda (2013): $2(x_t - x_{t-1})/(x_t + x_{t-1})$, and the average size: $\bar{x} = (x_t + x_{t-1})/2$ (Figure A.6).

[10] Table A.7 contains regression results.

[11] This size effect is not driven by sectoral differences (Figure A.7).

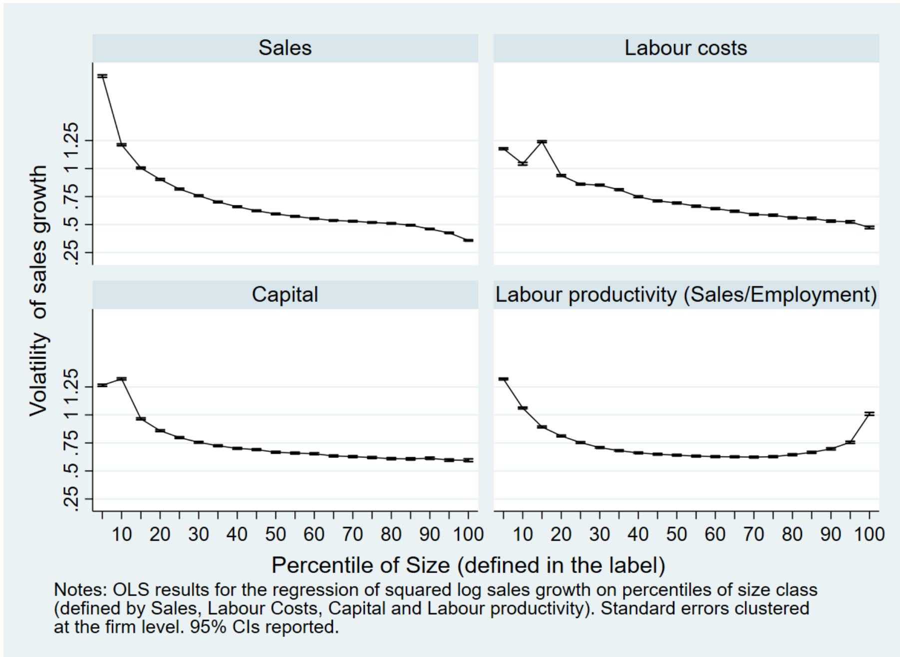

Figure 2: Volatility of sales growth, by percentile of firm size

$$\ln std.dev(dy) = \beta_\sigma - \alpha_\sigma \ln Y.$$

We calculate the standard deviation by aggregating across each of the 20 bins of firm size. Estimating the slope of this line separately for smaller (sales below 150,000€) and larger (sales above 150,000€) firms, we find the coefficient (α_σ) varies between -0.337 for smaller and -0.113 for larger firms.[12] This rejection of the "Gibrat's law for variances" is consistent with the evidence from previous studies (Dunne, M. Roberts, and Samuelson, 1989, Stanley et al., 1996,). Given the samples used with different firm sizes, it may also explain why the literature has found coefficients ranging from -0.5 to -0.15. For example, Calvino et al. (2018) finds an elasticity of -0.18, while Sutton (2007) finds an elasticity of -0.58. The closest result to ours is reported in Yeh (2021), who explains that the "negative log-linear relationship between firm-level volatility and its size is robust to a variety of specifications and the estimation results indicate that some traditional conjectures can be ruled out for explaining the size-variance relationship. This includes narratives on output, input or product diversification, and firm learning".

[12]Table A.5 and Figure A.5 in the Appendix.

2.3 Size-kurtosis

We now evaluate the tail behavior for firms of different size using the kurtosis. The kurtosis measures the fourth power of deviations normalized by the fourth power of the standard deviation. The intuition is that shocks below one standard deviation have very small weight when raised to the fourth power and shocks above one standard deviation get a larger weight. While there is some debate if the kurtosis measures the tails or the peakness, our results suggest strong correlations with the tails. The correlations with the 1st and 99th percentiles of the growth rate per year are 0.81 and 0.57, respectively.

Fact 3. *The kurtosis of firm growth rates increases with firm size.*

To compute the kurtosis, we estimate the following dummy variable regression

$$\left[\frac{\left(dy_{it}-\sum_s \mu_{dy(s)}1(i\in s)\right)^2}{\sum_s \sigma^2_{dy(s)}1(i\in s)}\right]^2 = c + \sum_s \alpha_s 1(i\in s) + \epsilon_{it}$$

where s is a size bin (percentile), $1(i \in s)$ is a dummy equal to 1 if firm i is in bin s, $\mu_{dy(s)}$ and $\sigma^2_{dy(s)}$ are the mean and variance of sales growth of firms in bin s, ϵ is a residual, and c,$\{\alpha_s\}_{s=1}^{S}$are parameters. This regression estimates the kurtosis for each bin. As reported in Figure 3, the distribution of sales growth is leptokurtic (kurtosis above 3) across all firm sizes. The pattern is markedly size dependent, with the kurtosis increasing with firm size.[13]

The relation of kurtosis with firm size follows an inverse relation from the one observed for the standard deviation: larger firms have larger kurtosis. Again, this is true for any of the measures of firm size: sales, capital, labor costs, and labor productivity. However, for the firms in the top 10% of the size distribution this excess kurtosis peaks very strongly, particularly if we use labor costs as our measure of firm size. The top 10% incorporates the very large firms and they are much larger on average than the 9th decile. This suggests that large firms are more likely to face more extreme shocks. Notice these shocks are not larger than shocks to smaller firms. What the kurtosis measures is that shocks are more extreme when compared

[13]Table A.9 contains regression results.

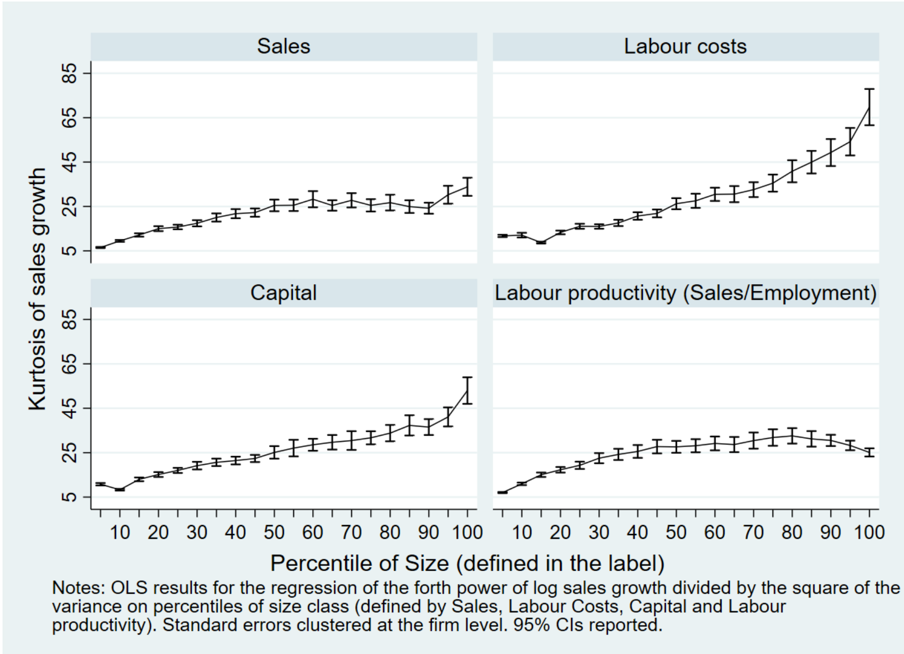

Figure 3: Kurtosis of sales growth, by percentile of firm size

to the "usual" shocks, as measured by the standard deviation (which is smaller for larger firms).

We also investigate the existence of a log-linear relation between the kurtosis and size[14]

$$\ln kurtosis(dy) = \beta_k + \alpha_k \ln Y.$$

Again, we calculate the kurtosis by aggregating across each of the 20 bins of firm size. Estimating the slope separately for smaller (sales below 150,000€) and larger (sales above 150,000€) firms, we find coefficient (α_k) varying between 0.419 for smaller and 0.072 for larger firms. The relation for the larger firms is more noisy (only significant at 5%), probably because of the lower precision for the estimated kurtosis.[15]

[14] To the best of our knowledge this has not been tested before.

[15] Table A.8 and Figure A.9.

3 Aggregate Implications

3.1 Volatility

As shown in Gabaix (2011, p.738) , the standard deviation of GDP growth can be written as

$$\sigma_{GDP} = \left(\sum_i \left(\frac{Y_{i,t-1}}{Y_{t-1}} \right)^2 \sigma_{i,t}^2 \right)^{1/2},$$

where Y_{it} is revenues (or value added), $Y_t = \sum_i Y_{it}$ is the aggregate value, and $\sigma_{i,t}^2$ is the firm-level volatility of value added growth. If all firms have the same volatility, $\sigma_i = \sigma$ for all i and

$$\sigma_{GDP} = \sigma \left(\sum_i \left(\frac{Y_{i,t-1}}{Y_{t-1}} \right)^2 \right)^{0.5}.$$

where $\left(\sum_i \left(\frac{Y_{i,t-1}}{Y_{t-1}} \right)^2 \right)^{0.5}$ is the Herfindal. The estimated firm-level volatility is

$$\hat{\sigma} = \left[\frac{1}{N} \sum_{i,t} \left(dva_{it} - \frac{1}{N} \sum_{i,t} dva_{it} \right)^2 \right]^{0.5} = 0.847.$$

However, the formula is incorrect when firm-level volatility is not constant. Given that volatility decreases with firm size, larger firms should carry more weight. This correction reduces the overall estimate to[16]

$$\hat{\sigma^w} = \left(\sum_{i,t} \frac{Y_{it}}{Y_t} \left(dva_{it} - \frac{1}{N} \sum_{i,t} dva_{it} \right)^2 \right) = 0.610$$

To obtain the aggregate formula we must also calculate the sales Herfindal[17]

[16] An alternative correction is to calculate the volatility for each bin. In this case, $\sigma_{GDP} = \left(\sum_s \sum_{i \in s} \left(\frac{Y_{i,t-1}}{Y_{t-1}} \right)^2 \sigma_s^2 \right)^{1/2}$. However, this requires calculating one volatility terms for each bin. Instead, the weighting correction requires only one volatility term.

[17] The value added Herfindal is similar, estimated at 0.038.

$$\hat{H} = \left(\sum_i \left(\frac{Y_{i,t-1}}{Y_{t-1}} \right)^2 \right)^{0.5} = 0.041.$$

We thus obtain an estimated aggregate volatility in the baseline case of

$$\hat{\sigma}_{GDP} = \hat{\sigma}\hat{H} = 0.847 * 0.041 = 0.035$$

while with the size correction we obtain

$$\hat{\sigma}_{GDP} = \hat{\sigma^W}\hat{H} = 0.610 * 0.041 = 0.025.$$

Both of these compare with the actual volatility of GDP estimated at 0.024.[18] We obtain a biased estimate if we don't use the correction. This illustrates the importance of obtaining representative samples. Using only the largest firms would result in an incorrect estimate. Given that previous studies have used datasets with varying degrees of firm size (mostly large firms), this difference may explain why they found smaller firm-level volatilities (ranging from 12% to 50%).

3.2 Kurtosis

We now show how the tail associates with macroeconomic behavior. We start by calculating the kurtosis across all firms for each year and compare it to GDP growth in Figure 4. GDP growth and kurtosis vary year-on-year declining from 2007 to 2012 and recovering from 2012 to 2019. In particular, there are two relevant declines: in 2009 (subprime) and 2011-2012 (sovereign debt). Overall, there is a correlation of 0.78 between the GDP growth rate and the kurtosis, which reveals a positive association between the existence of extreme events and GDP growth. As a comparison, the correlation of GDP growth with volatility is -0.44, 0.88 with the 1st percentile, and 0.72 with the 99th percentile of sales growth, thus confirming the relevance of the tails for the aggregate.

[18] The macro volatility in Portugal has been stable in the last 15/30 years.

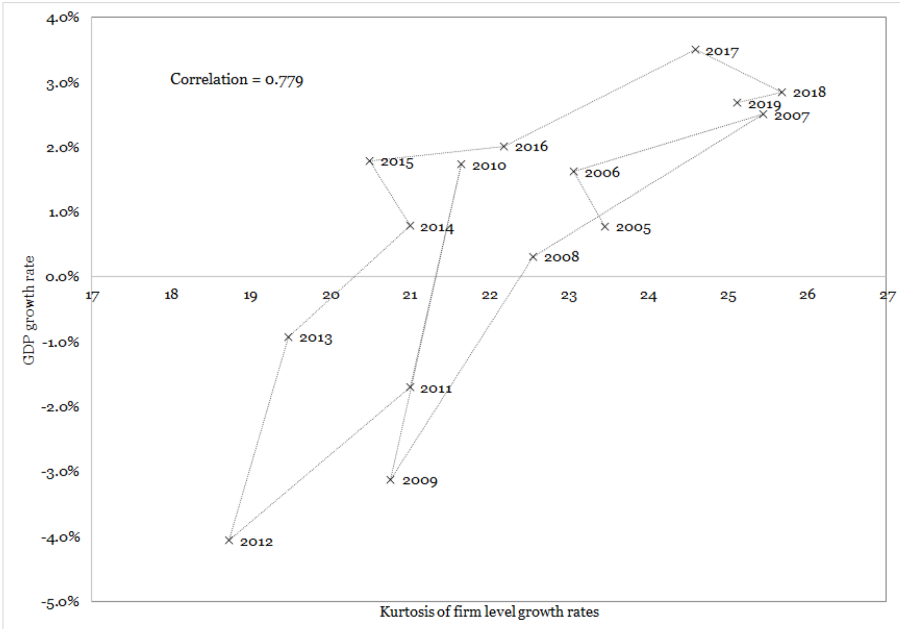

Figure 4: Kurtosis of firm size growth vs. GDP growth.

4 Model

What can explain the results? Recent evidence shows that "the distribution of earnings changes exhibits substantial deviations from log-normality, such as negative skewness and very high kurtosis" (Guneven et al., 2021). This motivates our introduction of consumer level shocks. We first provide the intuition for the role of customers.

Regarding volatility, the number of customers in models with customer markets is an implicit measure of firm size. For firms with a smaller customer base, the loss of one customer represents a large fraction of total sales. Larger firms spread risks by diversifying sales to a broader set of customers. Smaller firms are thus exposed to a larger risk.

For kurtosis, it is less straightforward to understand what explains the tail behavior. The distribution of shocks can be well approximated by a mixture probability consisting of a standard normal distribution, where the variance is drawn from a gamma distribution.[19] While this generates the heavy-tails, it tells nothing about the size-

[19] The approximation is better than with other heavy tailed distributions (Laplace/Cauchy). Section A.4.

kurtosis relation. To obtain a kurtosis that increases with size, requires a skewed customer distribution (large customers). In such case, even though the volatility is spread across many customers, the shocks to the largest customers continue to play a much larger role. Larger firms will thus have a smaller volatility and a larger kurtosis.

We derive our model from consumer preferences, aggregate first at the firm-level and next at the macroeconomic level. However, the baseline representative agent model with standard "CES preferences" fails to account for the decreasing volatility. This is because the consumer is representative and buys from all firms. So consumer shocks are proportional to all firms. We use a slight variation of the model to account for heterogeneity in a simple form: each customer has idiosyncratic income (y) and instead of having all customers consuming all products, we introduce consideration sets where each consumer only consumes N_i products.[20] We can think about this restriction emerging from location (travel) or information (search) frictions as in marketing and industrial organization models (e.g. Van Nierop et al., 2010). The main feature is the inexistence of a representative consumer. This exposes firms to idiosyncratic risks. For example, during a financial collapse a hotel located in the financial district will lose more customers than a beach resort. Its customers are not the same. While we understand that consideration sets emerge endogenously, to make progress we will take them as given (exogenous). The introduction of this simple (yet realistic) form of consumer heterogeneity is sufficient to generate the observed patterns.

There are $i = 1, ..., M$ consumers with CES preferences. Each only consumes a subset J_i (with size N_i) of all N products. There are j firms, each selling to M_j consumers. There are more consumers than products: $N << M$. In equilibrium the number of products demanded and produced are equal: $\sum_{i=1}^{M} N_i = \sum_{j=1}^{N} M_j$.

[20]The theory of consideration sets originates from the study of problem solving methods in Newell and Simon (1972), and has been used in Marketing (J. Roberts and Lattin, 1991). Consumers cannot cognitively evaluate all the products and must restrict to a subset.

4.1 Consumer problem

4.1.1 Dynamic problem

We use a standard consumption-savings model with the introduction of preference shifts $\{\xi_{it}\}_{t=1}^{\infty}$ and assume perfect foresight. To prevent non-stationary behavior, we restrict taste shocks to have zero expected growth: $E(\ln \frac{\xi_{it}}{\xi_{i,t-1}}) = 0$. The utility in every period is $\xi_{it} u(q_{it}, 1-l_{it})$, where $c_{it} = p_{it} q_{it}$ is the CES basket, p_{it} is a price index, q_{it} a consumption index (specified below), and the individual has one unit of time either allocated to work (l) or leisure ($1-l$). The consumer first solves the dynamic problem of consumption allocation over time by maximizing the value function

$$V(w_{it}) = \max_{q_{it}, l_{it}} \xi_{it} u(q_{it}, 1-l_{it}) + \beta V(w_{i,t+1})$$

$$s.t : w_{i,t+1} = y_{i,t} - p_{i,t} q_{i,t} + (1+r_t) w_{i,t}, \lim_{T\to\infty} (1+r_T)^{-T} w_{iT} \geq 0, y_{it} = \alpha_{it} \Pi_t + \omega_t l_{it}$$

where β is the discount factor, w is wealth that must be non-negative as $T \to \infty$, and the budget constraint must be satisfied. Income comes from employment ($\omega_t l_{it}$) or share of profits ($\alpha_{it} \Pi_t$). The share of profits is exogenous and known to all agents. The solution is the Euler equation[21]

$$\frac{\xi_{it} u_q(q_{i,t}, 1-l_{i,t})}{\xi_{i,t+1} u_q(q_{i,t+1}, 1-l_{i,t+1})} = \frac{p_{it}}{p_{i,t+1}} (1+r_{t+1}) \beta. \tag{1}$$

determining a non-linear first-order Markov process for consumption. If we let $u(q, 1-l) = \ln(q) + \ln \nu(1-l)$, this simplifies

$$c_{i,t+1} = (1+r_{t+1}) \beta \frac{\xi_{i,t+1}}{\xi_{it}} c_{it} \tag{2}$$

and $c_{it} = \frac{\xi_{it}}{\xi_{i0}} c_{i0}$. Let $\epsilon_{i,t+1} = \ln \frac{\xi_{i,t+1}}{\xi_{it}}$. Taking logs, we obtain a log-linear random walk

[21] The first order condition is $\xi u_q(q, 1-l) = \beta p V'(w')$ where $V'(w) = \xi u_q(q, 1-l) Q'(w) - \beta p V'(w') Q'(w) + (1+r) \beta V'(w') = (1+r) \beta V'(w')$. However, from the optimum $V'(w) = (1+r) \xi u_q(q, 1-l)/p$. The final solution is $\xi u_q(q, 1-l) = \beta(1+r') \xi' u_q(q', 1-l') p/(p')$.

$$\ln c_{i,t+1} = \ln(1+r_{t+1})\beta + \ln c_{it} + \epsilon_{i,t+1}.$$

Finally, the no-ponzi condition must be satisfied $\frac{c_{i0}}{\xi_{i0}} \sum_{t=0}^{\infty} \xi_{it} \leq \sum_{t=0}^{\infty} y_{it} + w_{i0}$, determining the starting level, c_{i0}. The optimal labor supply satisfies

$$\frac{\nu(1-l)}{\nu'(1-l)} = \frac{c_{it}}{\omega_t}.$$

4.1.2 Static problem

Let q_{ijt} denote the demanded quantity of product j, $c_{it} = p_{it}q_{it}$ the optimal consumption basket defined above, and p_{jt} the price. Choosing how to allocate the budget to consumption[22]

$$\max_{\{q_{ijt}\}_{j\in J_i}} \left(\sum_{j \in J_i} q_{ijt}^{(\eta-1)/\eta} \right)^{\eta/(\eta-1)} \quad s.t. \sum_{j\in J_i} p_{jt} q_{ijt} \leqslant c_{it} = p_{it} q_{it}$$

Demand is the standard CES function

$$q_{ijt} = \left(\frac{p_{jt}}{p_{it}} \right)^{-\eta} \frac{c_{it}}{p_{it}} \,\text{if}\, j \in J_i,\, q_{ijt} = 0 \,\text{otherwise}$$

where the price index faced by consumer i is $p_{it} = \left(\sum_{j\in J_i} p_{jt}^{1-\eta} \right)^{1/(1-\eta)}$ and the consumption bundle $q_{it} = \left(\sum_{j\in J_i} q_{ijt}^{(\eta-1)/\eta} \right)^{\eta/(\eta-1)}$. All consumers face the same set of prices but they consume different sets of products and thus face consumer specific price indices.

4.2 Firm problem

Each firm sells to M_j consumers and $q_{ijt} = 0$ if $j \notin J_i$. Using the consumer demand derived above, firm-level sales are

$$q_{jt}(p_{jt}) = \sum_{i=1}^{M_j} q_{ijt}(c_{it}, p_{jt}).$$

[22] From monotonicity, this is equivalent to maximizing $\xi_{it} u(q_{it}, 1 - l_{it})$.

This representation is useful to calculate firm-level demand elasticity and set the optimal price even in the presence of heterogeneous consumers without having to keep track of each consumer. On the supply side, the sole input in production is labor, supplied by consumers at wage rate ω_t, with constant returns to scale: $q_{jt} = a_{jt} l_{jt}$. The marginal cost of production is thus $mc_{jt} = \frac{\omega_t}{a_{jt}}$. Firms then set prices

$$\max_{p_{jt}} p_{jt} q_{jt} - mc_{jt} q_{jt},$$

$$foc : q_{jt} + p_{jt} \frac{\partial q_{jt}}{\partial p_{jt}} - mc_{jt} = 0$$

We follow the usual assumption in monopolistic competition models and assume that firms ignore the effect of own price on the price index.[23] Formally, $\frac{\partial p_{it}}{\partial p_{jt}} = 0$. The elasticity of demand is thus

$$\frac{\partial q_{jt}}{\partial p_{jt}} = \sum_{i=1}^{M} \frac{\partial q_{ijt}(c_{it}, p_{jt}, p_{it})}{\partial p_{jt}} = -\eta \frac{q_{jt}}{p_{jt}},$$

and we obtain the optimal pricing rule

$$p_{jt} = \frac{\eta}{\eta - 1} mc_{jt}.$$

Marginal costs are determined by the level of productivity and input prices. Firms' revenues equal $p_{jt} q_{jt} = (\frac{\eta}{\eta-1}) \omega_t l_{jt}$.

4.3 Equilibrium

We now define the equilibrium for the labor, financial, and product markets. We can solve the model at the aggregate level and abstract from specifying how wealth, wages or capital returns are allocated to each consumer. Firms' revenues are equal to wages plus profits and equal to the income flow to individuals. Aggregating across all consumers

[23] In our model, each consumer's price index depends on the set of goods consumed. Given that price indices are consumer specific, not assuming this would require the firms to keep track of each consumer's consideration set.

$$\sum_{i=1}^{M} y_{it} = \sum_{i=1}^{M} (\frac{\eta}{\eta - 1})\omega_t l_{it} = \sum_{i=1}^{M} \sum_{j \in J_i} p_{jt} q_{ijt}.$$

This determines the equilibrium labor market clearing wage[24]

$$\omega_t = \frac{\eta - 1}{\eta} \frac{\sum_{j=1}^{N} p_{jt} q_{jt}}{\sum_{i=1}^{M} l_{it}}.$$

In this economy there is no investment and we assume perfect capital markets. Each consumer can lend and borrow freely at the interest rate, r_t. Market clearing implies that the supply of funds must equal the demand for funds: $\sum_{i=1}^{M} w_{it} = 0$ at all periods, t. From the budget constraint, consumption equals income, $\sum_i y_{i,t} - c_{i,t} = 0$. Finally, this implies that aggregate supply equals aggregate demand and the product market clears[25]

$$\sum_{j=1}^{N} p_{jt} q_{jt} = \sum_{i=1}^{M} c_{it}.$$

4.4 Discussion

Before presenting the computational results, we discuss two concerns. The first is how can there be consumption shocks at the aggregate level without changing productivity (or inputs). The second is that the model predicts a similar volatility for consumption and income, while this is at odds with empirical evidence.

In a closed economy, agents can only consume what they produce. So how can consumption increase without changing productivity? To answer this, notice that heterogeneity implies that aggregate consumption and aggregate labor supply can either be positively or negatively correlated (or uncorrelated). This is because with heterogeneity, the agents that increase consumption are not the same that increase

[24] By leaving the parametric form of $(\nu(1-l))$ unspecified, we sidestep the debate on the labor supply elasticity (e.g. King and Rebelo, 1999 and Chetty et al., 2011). Labor market clearing implies that, at the given wage, supply equals demand: $\sum_{i=1}^{M} l_{it} = \sum_{j=1}^{N} l_{jt}$. The model is compatible with: (i) fixed labor supply, in which case trivially $\sum_{i=1}^{M} l_{it} = \bar{l}$ or (ii) fully flexible labour supply (exogenous wages), $\sum_{i=1}^{M} l_{it} = \frac{\eta-1}{\eta} \frac{\sum_{j=1}^{N} p_{jt} q_{jt}}{\omega_t}$.

[25] Aggregate income equals aggregate revenues $\sum_{i=1}^{M} y_{it} = \sum_{i=1}^{M} \sum_{j \in J_i} p_{jt} q_{ijt}$, while from perfect capital markets $\sum_{i=1}^{M} y_{it} = \sum_{i=1}^{M} c_{it}$ and the result follows.

labor supply. Because of granularity, these are not canceled by aggregation. For example, individuals with large wealth may increase consumption while individuals with low wealth increase labor supply. This breaks the representative agent model's prediction of Barro and King (1984) that aggregate consumption and labor supply must be negatively correlated as a response to a demand shock.

Besides the theoretical implication, there is also an empirical implication that results in mismatch between the model and the data. To understand this, it is important to distinguish again measured TFP (TFPr) from technical TFP (TFPq). While technical TFPq may stay fixed, measured TFPr will be endogenously determined. Empirically, productivity shocks are still at the center of the fluctuations as in standard Real Business Cycle models but, due to mismeasurement, these shocks have a demand origin. Researchers that estimate the shocks will misinterpret them as supply side shocks. Translating to our model, there are two polar cases to consider depending on labor supply. If labor supply is inelastic (e.g. long run), TFPq stays constant while wage and price increases generate an increase in measured TFPr. If labor supply is infinitely elastic (e.g. short run, particularly at the intensive margin), wages are given, TFPq stays constant, while employment increases. For anything in between, there is a combination of the two. Given the empirical evidence of small firms' price responses to demand shocks (Bonomo et al., 2020; DellaVigna and Gentzkow, 2019; Santos, Costa, and Brito, 2022), this is consistent with a model of elastic labor supply. Particularly, if we consider the intensive margin, where worker effort can vary (to some extent) so that the same number of workers can actually translate into more effective labor units resulting in a measured TFPr increase. This is one reason why intermediate inputs are a better proxy to measure marginal costs (and markups) than either labor or capital.

The second concern is the prediction that the volatility of consumption and income are the same. We can reconcile this with the empirical evidence by considering the extension to two types of goods: consumption (e.g. food) and investment (e.g. car). In this case shocks to preferences for investment goods are more volatile than shocks to preferences for consumption goods. I.e., purchases of investment goods (e.g. cars) are more volatile than purchases of consumption goods (e.g. food).

4.5 Computational results

From equation 1, take a log-linear approximation to the Euler equation of consumer i

$$\ln c_{it} = \rho_0 + \rho_1 \ln c_{i,t-1} + \epsilon_{it}.$$

The process is the same for all consumers and let r_t be constant. Assume the shock ($\epsilon_{it} = \ln \frac{\xi_{i,t}}{\xi_{i,t-1}}$) is drawn from a variance-gamma distribution - a mixture of a zero-mean normal with variance drawn from a gamma distribution ($Gamma(k, \theta)$). This gives log-normal customer sales with heavy-tailed shocks. The distribution of consumption becomes skewed and customer sales become concentrated. Together with consideration sets, this guarantees that when we sum across customers, the volatility and kurtosis of growth rates will vary with firm size. This matches the firm-level distributions observed in our data and will be fundamental to generate an increasing kurtosis. The process is stationary if $|\rho_1| < 1$. We calibrate our parameters to

M	N	$\frac{\eta}{\eta-1}mc$	ρ_0	ρ_1	k	θ	
123.4 million	350,000	1	1.3	0.8	3.2	0.25	

The model is simulated for $T = 150$ periods with a burn-in of 50 periods. To obtain similar aggregation effects, we set the number of firms (N) to be approximately equal to the Portuguese economy: 350,000. We then split the firms into 10 bins of equal size and for each we set the number of customers equal to $M_j \in [4, 9, 15, 22, 32, 45, 70, 110, 220, 3000]$. The number of customers is not obtained from data but is instead calibrated to obtain a similar average sales per decile observed in the data. As a simplification, we will assume that each consumer only buys one product so that $N_i = 1$ (imagine that consumers have multiple selves). This gives a total of $M = 123.4$ million individual consumption decisions (e.g. 12.3 selves for a Portuguese population of 10 million). We assume that all firms face the same marginal cost and normalize $\frac{\eta}{\eta-1}mc = 1$. This allows us to normalize all prices equal to 1 and simplify the notation so that $q_{ijt} = \frac{c_{it}}{N_i}$. Finally, the dynamic parameters $(\rho_0, \rho_1, k, \theta)$ are calibrated to replicate the standard deviation and kurtosis across the 10 deciles of firm size.

Qualitatively, Table 1 shows that the model replicates the patterns for standard deviation and kurtosis from empirical facts 1 to 3. Quantitatively, we obtain good approximations for the volatility (except for the very small firms in the first two deciles) and we slightly underestimate the kurtosis at almost all size levels. At the aggregate level, the model generates an aggregate volatility of sales of 3.7% that compares with an aggregate volatility of sales in the data of 5.3%.

While the reduction in volatility for larger firms is straigthforward to understand as a simple diversification argument, the increase in kurtosis in this model depends on the level of customer concentration. We show this by limiting the size of each customer and imposing an upper bound to the consumption of each individual. We use 3 different values: $c_{it} < \{10^6, 10^7, 10^8\}$. The results in Table 1 show how the results change with different bounds on customer size. The kurtosis decreases as we reduce the size of the largest customers. Furthermore, the volatility is also reduced, confirming that risk diversification is limited by the existence of large customers. If we impose a sufficiently small limit to the size of the largest customer, the kurtosis no longer increases with firm size as reflected in column (v). This demonstrates the importance of large customers to match the empirical results regarding the volatility-kurtosis-size relation.

5 Concluding remarks

In this article, we have shown that both the volatility and kurtosis of the growth rate of sales varies with firm size. This has implications for the aggregate responsiveness to the business cycle. We have microfounded the observed firm-level shocks with demand side heterogeneity. Obtaining the theoretical microfoundations requires two conditions: (i) each customer buys from a small set of firms, and (ii) the shocks to individual customers are normally distributed with varying variances. This guarantees that the model is able to replicate the empirical results qualitatively and quantitatively.

The model has several economic implications. For example, the magnification channel of the consumption side means that what may start as a small supply side shock (e.g. oil prices) is amplified through the effect to the agents consumption decisions (e.g. postponing the purchase of a new automobile). A credit crunch is also potentially amplified. In the model, credit is allocated from customers with negative demand shocks to customers with positive demand shocks. As credit is restricted, customers with a positive demand shock (e.g. that need to buy a car) cannot bor-

Table 1: Model vs. Data

	Revenues					Std. Dev.					Kurtosis				
Deciles	(i)	(ii)	(iii)	(iv)	(v)	(i)	(ii)	(iii)	(iv)	(v)	(i)	(ii)	(iii)	(iv)	(v)
1	11,734	12,222	12,064	11,511	10,451	1.55	0.71	0.71	0.71	0.71	8.0	7.7	7.8	7.7	7.1
2	28,448	28,264	27,333	25,714	23,493	0.96	0.63	0.63	0.63	0.62	13.4	9.5	9.4	9.1	7.8
3	46,862	47,286	44,725	42,858	39,054	0.79	0.59	0.59	0.59	0.58	16.5	10.7	10.4	10.0	8.0
4	69,973	72,399	66,142	62,809	57,227	0.68	0.56	0.56	0.56	0.55	21.0	11.9	11.5	10.7	8.0
5	100,384	107,992	96,736	90,981	83,421	0.61	0.54	0.54	0.53	0.52	24.3	12.8	12.4	11.2	7.9
6	144,414	146,262	136,060	127,890	117,010	0.56	0.52	0.52	0.51	0.49	26.7	13.7	13.5	11.7	7.6
7	216,249	223,478	210,310	199,510	182,060	0.53	0.49	0.49	0.49	0.45	26.9	15.2	14.7	12.3	7.1
8	355,375	363,090	333,050	312,950	286,340	0.52	0.47	0.47	0.46	0.41	26.2	17.5	16.3	12.5	6.3
9	723,377	734,845	664,470	625,830	572,660	0.48	0.44	0.44	0.42	0.35	24.6	20.9	17.9	12.3	5.1
10	9,482,151	9,936,518	9,064,800	8,535,800	7,880,500	0.39	0.36	0.34	0.26	0.13	33.6	35.1	17.8	6.0	3.1

Notes: Column (i) reports the data. Column (ii) is the unconstrained model while columns (iii) to (v) impose a constraint to the size of the the largest customer of 10^8, 10^7, and 10^6, respectively.

row from the customers with a negative demand shock (e.g. that want to sell a house). This leads to an overall depression in aggregate demand. Changes to agents' expectations can have similar amplification effects.

We have been able to replicate the main microeconomic empirical features in the data, despite the fact that we have abstracted from any type of cost side shocks, labor or capital market frictions, and input-output linkages. These are important mechanisms of the Portuguese economy that may further amplify the transmission of shocks. For example, shocks may be correlated at the customer level due to either aggregate events or networks (input-output linkages). If we allow for input-output linkages, the (unweighted) volatility of GDP using Domar weights (sales divided by GDP) becomes 0.065, almost twice the estimated volatility of 0.035 using revenue weights. This suggests an amplification of almost two times originated from input-output linkages.

We highlight some topics for future research. First, preference/consumption shocks are the drivers of fluctuations in our model. There is evidence of income shocks that are consistent with the formulation we attribute to preference/consumption shocks, namely, the large kurtosis (Guneven et al., 2021). However, consumption dynamics may be different from income dynamics. Hence, it is important to understand if shocks to consumption also exhibit similar heavy tails. This requires consumption data. Second, our simple model of firm dynamics with customer heterogeneity allows a link of business cycles with economic growth through productivity dynamics (e.g. Luttmer, 2007). We believe that models of granularity can help in the development of new theories of endogenous business cycles.

References

Acemoglu, D., V. Carvalho, et al. (2012). "The Network Origins of Aggregate Fluctuations". *Econometrica* 80.5, pp. 1977–2016.

Acemoglu, D., A. Ozdaglar, and A. Tahbaz-Salehi (2017). "Microeconomic Origins of Macroeconomic Tail Risks". *American Economic Review* 107.1, pp. 54–108.

Ascari, G., G. Fagiolo, and A. Roventini (2015). "Fat-Tail Distributions and Business-Cycle Models". *Macroeconomic Dynamics* 19.2, pp. 465–476.

Barro, R. and R. King (1984). "Time-separable preferences and intertemporal-substitution models of business cycles". *Quarterly Journal of Economics* 99.4, pp. 817–839.

Beaudry, P. and F. Portier (2014). "Understanding Noninflationary Demand-Driven Business Cycles". *NBER Macroeconomics Annual* 28.1, pp. 69–130.

Bonomo, M. et al. (2020). "Multi-Product Pricing: Theory and Evidence from Large Retailers in Israel". *Bank of Canada Staff Working Papers 20-12* 2020-12.

Calvino, F. et al. (2018). "Growth volatility and size: a firm-level study". *OECD Working Paper* 26.

Carvalho, V. and X. Gabaix (2013). "The Great Diversification and Its Undoing". *American Economic Review* 103.5, pp. 1697–1727.

Carvalho, V. and B. Grassi (2019). "Large Firm Dynamics and the Business Cycle". *American Economic Review* 109.4, pp. 1375–1425.

Chari, V., P. Kehoe, and E. McGrattan (2007). "Business Cycle Accounting". *Econometrica* 75(3), pp. 781–836.

Chetty, R. et al. (2011). "Are Micro and Macro Labor Supply Elasticities Consistent? A Review of Evidence on the Intensive and Extensive Margins?" *American Economic Review* 101.3, pp. 471–75.

De Loecker, J. (2011). "Recovering Markups from Production Data". *International Journal of Industrial Organization* 29, pp. 350–355.

De Loecker, J. et al. (2016). "Prices, Markups, and Trade Reform". *Econometrica* 84(2), pp. 445–510.

DellaVigna, S. and M. Gentzkow (2019). "Uniform Pricing in US Retail Chains". *Quarterly Journal of Economics* 134, pp. 2011–2084.

di Giovanni, J., A. Levchenko, and I. Mejean (2014). "Firms, Destinations, and Aggregate Fluctuations". *Econometrica* 82.4, pp. 1303–1340.

Dunne, T., M. Roberts, and L. Samuelson (1989). "The growth and failure of us manufacturing plants". *The Quarterly Journal of Economics* 104.4, pp. 671–698.

Foster, L., J. Haltiwanger, and C. Syverson (2016). "The Slow Growth of New Plants: Learning about Demand?" *Economica* 83(329), pp. 91–129.

Gabaix, X. (2011). "The Granular Origins of Aggregate Fluctuations". *Econometrica* 79.3, pp. 733–772.

– (2016). "Power laws in economics: An introduction". *Journal of Economic Perspectives* 30.1, pp. 185–206.

Gabaix, X. and R. Ibragimov (2011). "Rank 1/2: A Simple Way to Improve the OLS Estimation of Tail Exponents". *Journal of Business and Economic Statistics* 29.1, pp. 24–39.

Guneven, F. et al. (2021). "What Do Data on Millions of U.S. Workers Reveal about Life-Cycle Earnings Dynamics?" *Econometrica* 89.5, pp. 2303–2339.

Haltiwanger, J., R. Jarmin, and J. Miranda (2013). "Who Creates Jobs? Small versus Larges versus Young". *The Review of Economics and Statistics* 95.2, pp. 347–61.

Karabarbounis, L. (2014). "The Labor Wedge: MRS vs. MPN." *Review of Economic Dynamics* 17.2, pp. 206–223.

King, R. and S. Rebelo (1999). "Resuscitating real business cycles". In: *Handbook of Macroeconomics, vol. 1.* Ed. by J.B Taylor and M. Woodford. Elsevier, pp. 927–1007.

Kramarz, F., J. Martin, and I. Mejean (2020). "Volatility in the small and in the large: The lack of diversification in international trade". *Journal of International Economics* 122.

Luttmer, E. (2007). "Selection, Growth, and the Size Distribution of Firms". *The Quarterly Journal of Economics* 122.3, pp. 1103–1144.

Mankiw, N.G. (1989). "Real Business Cycles: A New Keynesian Perspective". *Journal of Economic Perspecives* 3.3, pp. 79–90.

Moro, A. (2012). "The structural transformation between manufacturing and services and the decline in the US GDP volatility". *Review of Economic Dynamics* 15.3, pp. 402–415.

Newell, A. and H.A. Simon (1972). *Human Problem Solving.* Englewood Cliffs, NJ: Prentice Hall.

Roberts, J.H. and J.M. Lattin (1991). "Consideration: Review of Research and Prospects for Future Insights". *Journal of Marketing Research* 34.August, pp. 406–410.

Santos, C.D. (2020). "Identifying Demand Shocks from Production Data". *Journal of Business and Economic Statistics* 38(1), pp. 93–106.

Santos, C.D., L. Costa, and P. Brito (2022). "Demand, Supply, and Markup Fluctuations". *The Economic Journal* 132.644, pp. 1620–1645.

Stanley, M. et al. (1996). "Scaling Behavior in the Growth of Companies". *Nature* 379, pp. 804–806.

Statistics Portugal (2019a). *Contas Nacionais e Regionais*. Database. URL: `https://www.ine.pt/xportal/xmain?xpid=INE&xpgid=ine_base_dados(last%20accessed%2017%20March%2021)`.

– (2019b). *Sistema de Contas Integradas da Empresa*. Database.

Sutton, J. (2007). "Market Share Dynamics and the Persistence of Leadership Debate". *American Economic Review* 97.1, pp. 222–241.

Van Nierop, E. et al. (2010). "Retrieving Unobserved Consideration Sets from Household Panel Data". *Journal of Marketing Research* 47.1, pp. 63–74.

Yeh, C. (2021). "Revisiting the origins of business cycles with the size-variance relationship". *unpublished*.

A Online Appendix

A.1 Data Description

Our dataset consists of yearly census level data for all non-financial firms over the period 2004-2019, covering the two main crisis - sub-prime and euro sovereign debt. It contains balance sheet, P&L and employment information. The dataset is collected by the national statistics office (INE) and can be accessed by any accredited researcher (Statistics Portugal, 2019b). We then obtain for each individual firm the following variables: Sales, intermediate inputs, labor costs, value added (sales minus intermediate inputs), employment (number of workers), and capital stock. Capital stock is the only variable that has to be constructed. We use data on investment and fixed assets to reconstruct the capital stock using the perpetual inventory formula.[26] Finally, we obtain aggregate data for the real GDP from Statistics Portugal (2019a).

Table A.1: Descriptive statistics

Variable	Obs	Mean	Std. Dev.	Min	Max
Sales	4,290,567	1,121,813	24,300,000	1	9,630,000,000
Intermediate inputs	4,290,567	871,240	22,000,000	0	9,220,000,000
Capital Stock	4,284,586	724,535	31,000,000	1	10,700,000,000
Employment	4,290,567	9	95	1	26,857
Labor costs	4,290,567	162,358	2,107,706	0	694,000,000
Sales growth rate (ln)	4,290,567	0.02	0.78	-17.70	14.11

Notes: All variables in euros except for employment (number of workers). Sales growth rate is calculated as the time-difference of log sales.

Table A.2: Aggregate statistics

Year	Real GDP - growth rate	Sales - average growth rate	Number of firms
2004	-	-	291,924
2005	0.78%	2.01%	300,783
2006	1.62%	1.82%	304,373
2007	2.51%	5.38%	308,218
2008	0.32%	0.37%	311,075
2009	-3.12%	-5.85%	307,780
2010	1.74%	0.69%	319,079
2011	-1.70%	-9.13%	318,351
2012	-4.06%	-12.17%	310,430
2013	-0.92%	-1.98%	309,873
2014	0.79%	4.05%	316,575
2015	1.79%	5.41%	323,897
2016	2.02%	5.34%	331,716
2017	3.51%	9.42%	343,008
2018	2.85%	6.97%	357,713
2019	2.68%	7.17%	379,626

Notes: Sales growth rates calculated as log differences.

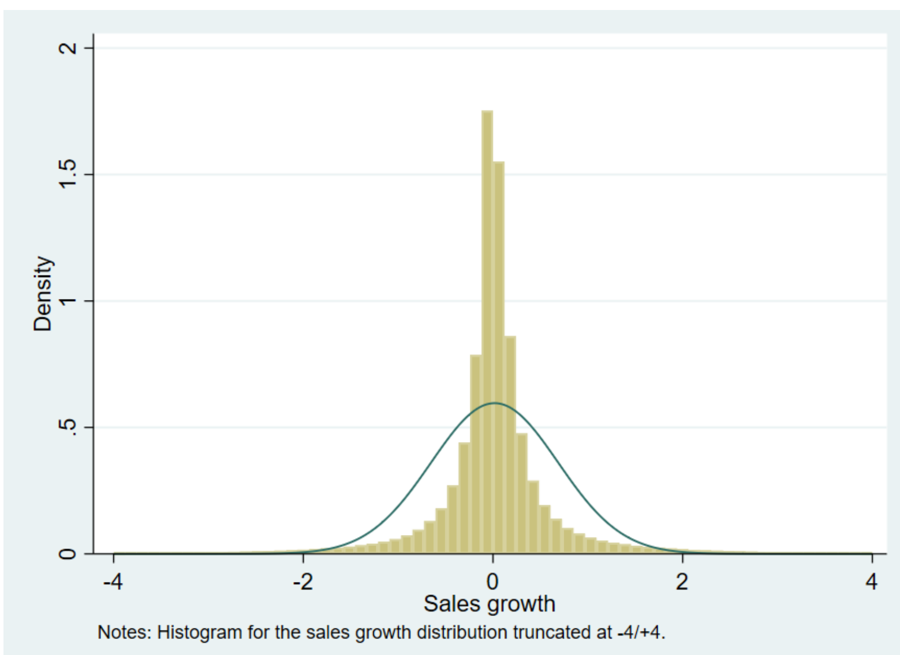

Figure A.1: Distribution of sales growth and normal density.

Table A.3: Tail exponent estimates

	(i) Rank	(ii) Rank-corrected	(iii) Rank
Right tail exponent estimate	-3.279	-3.280	-2.273
s.d.	0.0022	0.0022	0.0012
Left tail exponent estimate	-3.132	-3.128	-2.205
	0.0023	0.0023	0.0012

Notes: Tail exponent OLS. Columns (i) and (iii) use rank regression with sales growth threshold of 2 and 1 (-2 and -1 for the left tail), respectively. Column (ii) uses the rank regression with the Gabaix and Ibragimov (2011) correction.

Table A.4: Tail risk

	Data	Normal distribution	Ratio
2004	2.595%	0.268%	9.68

Notes: Probability of a shock three standard deviations away from the mean.

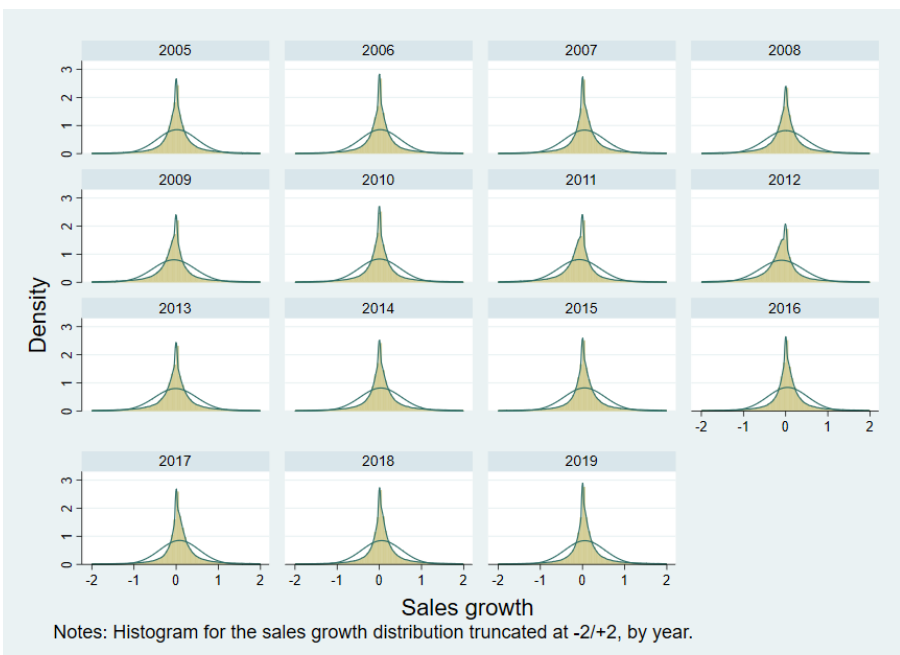

Figure A.2: Distribution of sales growth and normal density, by year.

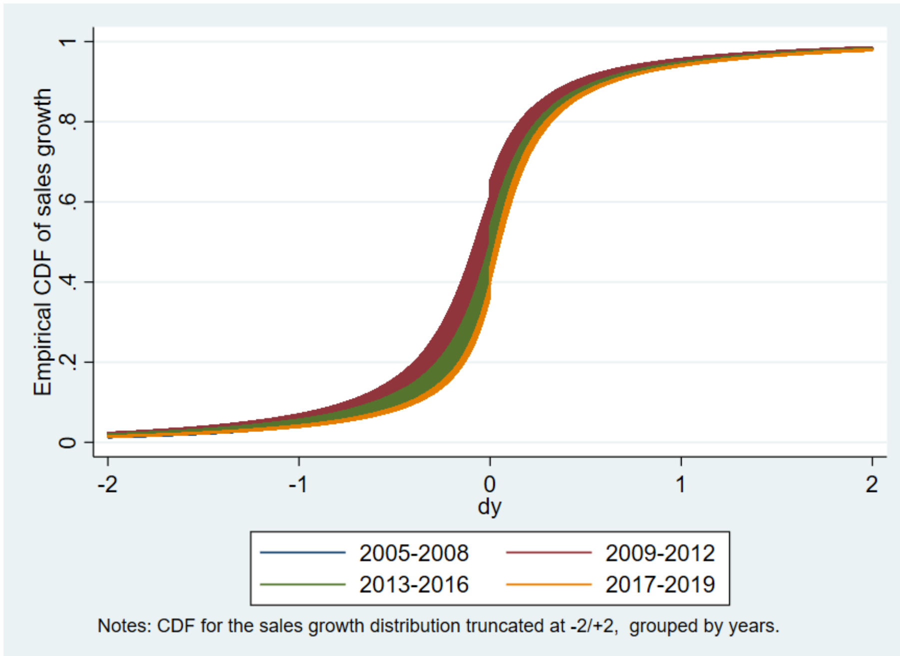

Figure A.3: CDF of sales growth, grouped by years.

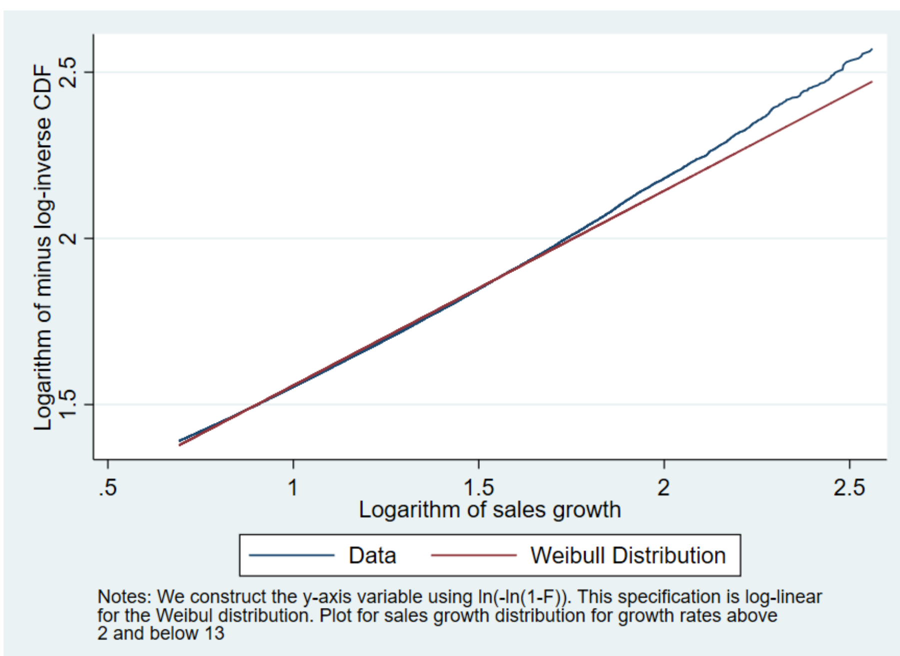

Figure A.4: Weibull plot

	Sales Below 150,000€	Sales Above 150,000€
Slope	-0.3378***	-0.1129***
s.d	(0.0052)	(0.0042)
Intercept	3.3715***	0.7645***
s.d.	0.0549	(0.0551)
R-squared	0.9979	0.9906
N	11	9

Table A.5: OLS results for the relation of between standard deviation and sales

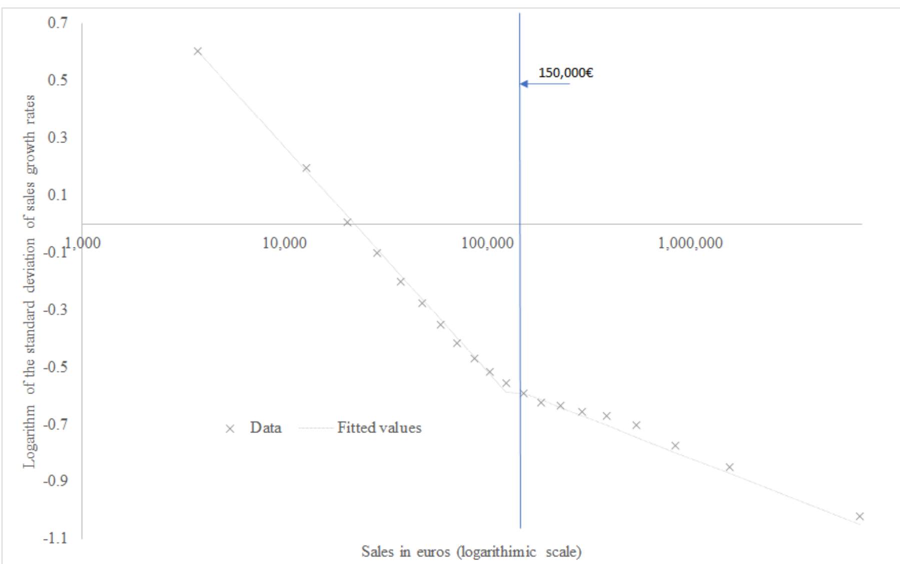

Notes: Fitted values (OLS) with a slope of α=0.337 (1/2.96) for firms with less than 150,000€ in sales and an α=0.113 (1/8.85) for firms with more than 150,000€ in sales.

Figure A.5: Volatility of sales growth vs. firm size (in logarithms).

A.2 Heavy tails of sales growth distribution

A.3 Size-volatility relation

A.3.1 Volatility of sales growth using log difference growth rates

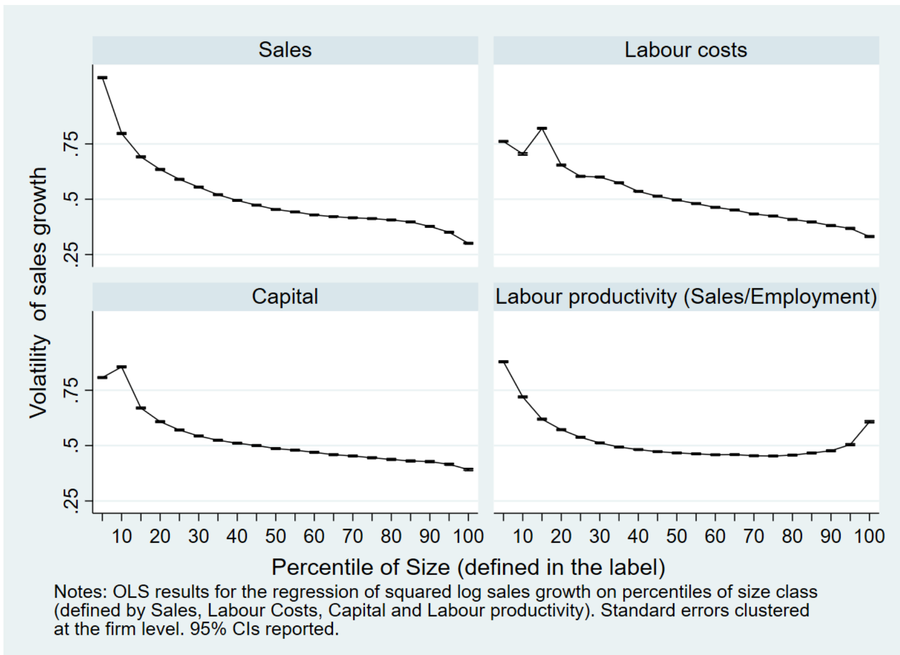

Figure A.6: Volatility of sales growth (log differences), by percentile of firm size

A.3.2 Volatility of sales growth by 3 digit CAE code

The standard deviation of sales growth can be computed for individual sectors. Overall, Figure A.7 shows that there is no pattern. That is, services are not more/less volatile than industry. We isolate two particular industries which contain many firms and that exhibit a large volatility: 411 - Property development and 681 - Commerce of real estate. Both are related to the real estate sector and exhibit a sales growth volatility above 0.9. Removing these two industries does not alter the main results as reported in Figure A.8.

[26]The formula for the capital stock is $K_t = (1-\delta)K_{t-1} + I_t$, where we use an estimated depreciation rate $\delta = 0.058$, K_1 is the first observation of the capital stock of firm i equals fixed assets and I is the net investment (investment minus disinvestment).

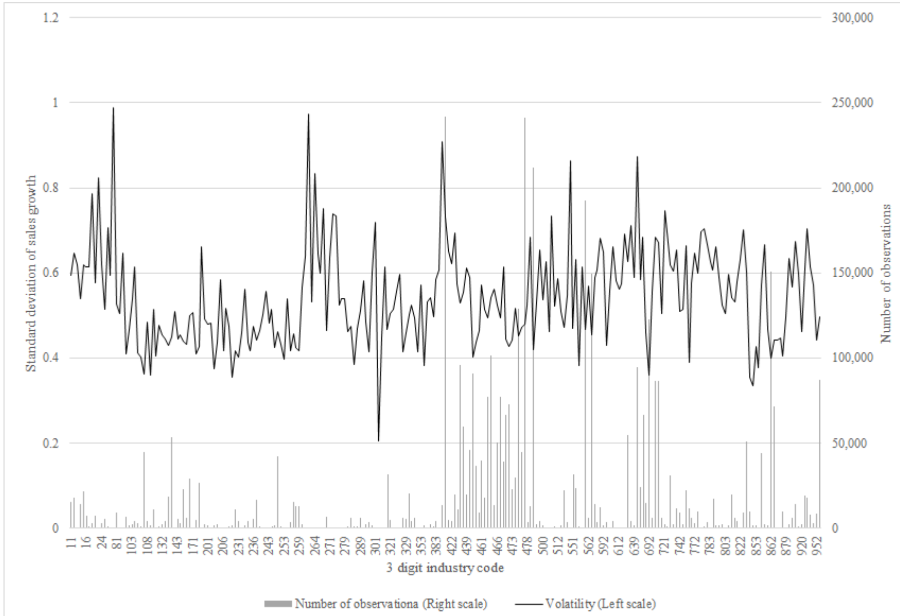

Figure A.7: Volatility of sales growth - by 3 digit sector.

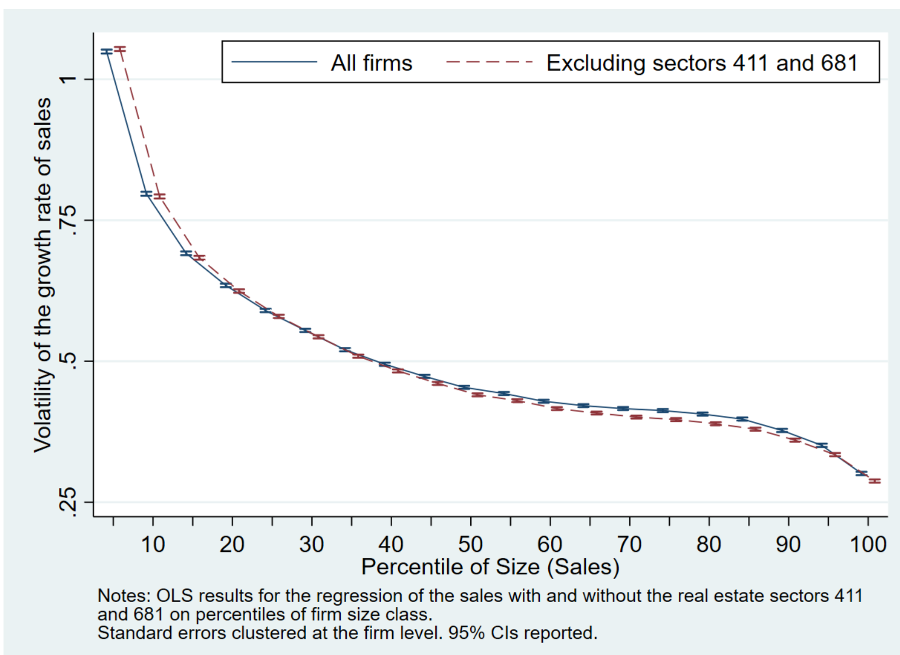

Figure A.8: Volatility of sales growth - by size (all firms and excluding real estate sectors 411 and 681).

Dependent variable: $\lvert dy - \mu_{dy} \rvert$				
Percentiles of:	Sales	Labor Costs	Capital	Labor Productivity
5th percentile	1.209***	0.685***	0.770***	0.859***
	(308.18)	(290.03)	(306.24)	(314.90)
10th percentile	0.731***	0.602***	0.847***	0.621***
	(258.82)	(181.78)	(298.00)	(267.01)
15th percentile	0.584***	0.767***	0.566***	0.491***
	(256.35)	(312.53)	(263.47)	(253.24)
20th percentile	0.516***	0.538***	0.488***	0.437***
	(261.08)	(263.37)	(253.39)	(251.99)
25th percentile	0.465***	0.479***	0.442***	0.401***
	(264.26)	(260.01)	(247.48)	(253.12)
30th percentile	0.425***	0.476***	0.410***	0.376***
	(262.51)	(266.49)	(243.22)	(255.01)
35th percentile	0.390***	0.446***	0.391***	0.358***
	(260.89)	(260.78)	(243.90)	(255.26)
40th percentile	0.363***	0.404***	0.376***	0.348***
	(258.41)	(255.59)	(242.82)	(255.16)
45th percentile	0.342***	0.380***	0.367***	0.341***
	(257.38)	(253.16)	(238.44)	(257.12)
50th percentile	0.324***	0.365***	0.353***	0.336***
	(254.25)	(247.96)	(236.44)	(254.72)
55th percentile	0.315***	0.348***	0.346***	0.333***
	(255.28)	(243.44)	(231.52)	(255.19)
60th percentile	0.302***	0.332***	0.337***	0.328***
	(252.42)	(238.73)	(224.81)	(252.63)
65th percentile	0.293***	0.320***	0.327***	0.329***
	(251.22)	(235.81)	(222.33)	(252.32)
70th percentile	0.289***	0.304***	0.321***	0.324***
	(244.87)	(234.33)	(215.54)	(245.34)
75th percentile	0.285***	0.295***	0.314***	0.322***
	(243.27)	(224.34)	(208.67)	(238.08)
80th percentile	0.279***	0.281***	0.306***	0.325***
	(236.24)	(214.40)	(199.11)	(228.54)
85th percentile	0.270***	0.271***	0.300***	0.333***
	(230.68)	(202.98)	(189.29)	(220.32)
90th percentile	0.253***	0.256***	0.297***	0.341***
	(222.63)	(194.50)	(178.69)	(204.84)
95th percentile	0.231***	0.246***	0.285***	0.367***
	(207.94)	(182.64)	(170.30)	(192.12)
100th percentile	0.193***	0.216***	0.259***	0.492***
	(191.74)	(168.75)	(142.23)	(171.64)
Observations	4290591	4290591	4278746	4290591

Notes: OLS results for the regression of absolute deviation of log sales growth on percentiles of size class (defined by sales, labor costs, capital and labor productivity). Standard errors clustered at the firm-level. 95% CIs reported.

Table A.7: Volatility of sales growth - by size

	Sales Below 150,000€	Sales Above 150,000€
Slope	0.4193***	0.0724**
s.d	(0.0163)	(0.0289)
Intercept	-1.6345***	2.3631***
s.d.	(0.1729)	(0.3832)
R-squared	0.9865	0.4721
N	11	9

Table A.8: OLS results for the relation of between kurtosis and sales

A.4 Size-kurtosis relation

A.5 Distributional approximation

The financial literature has long studied the existence of (semi-)heavy tails for the distribution of returns on financial assets. These distributions are different from strong heavy tailed distributions, such as the Cauchy distribution, because they have well defined moments. For example, the heavy-tailed Cauchy distribution may not even have a first moment. We have shown before in Figure 1 that the tail of the growth rate distribution slightly departs from a power law distribution. Instead, it seems a more moderate distribution with tails just slightly larger than the exponential (and smaller than a Weibull). In Figure A.10 we compare it with a two parameter variance mixing distribution. In this case, we mix a standard normal distribution ($N(0, 1)$) with a gamma distribution ($\Gamma(0.8, 0.65)$ for the variance. This is also known as a variance-gamma distribution, a special case of the generalized hyperbolic distribution. The quality of the specification is quite impressive given that this is a two parameter distribution as reflected by the quantile plot in Figure A.10 containing 99.5% of the data and confirmed in Figure A.11 reporting the quantile plot for 99.999% of the data. This is a convenient specification where growth rate shocks are still drawn from a normal distribution while the variance is nonconstant.[27]

[27] As a comparison we also report a Cauchy distribution in Figure A.12 and a Laplace distribution in Figure A.13. The Cauchy distribution has two parameters the location x_0 and scale γ. We compute the parameters of the Cauchy distribution from the median $x_0 = median$ and the interquartile range $\gamma = IQR/2$. A Laplace distribution also shares some features with the data, such as the heavy tails, with the benefit of having well defined first and second moments. We compare the data with a Laplace distribution with the same mean and variance as our data. The Cauchy distribution reproduces very closely all the percentiles of the growth rate distribution (see the qq plot), while the Laplace distribution is not a good approximation. In particular, the Laplace distribution tends

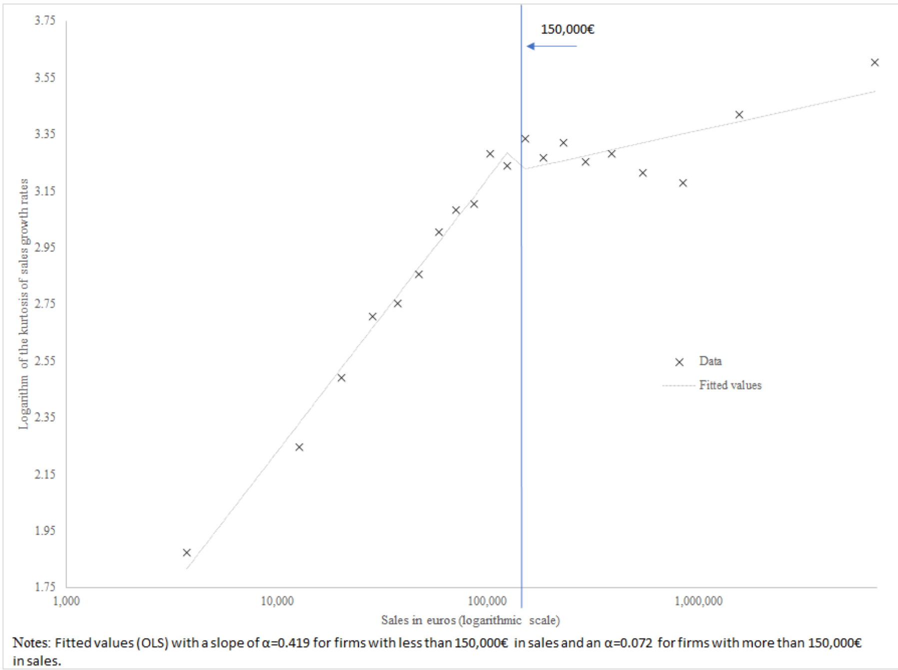

Figure A.9: Kurtosis of sales growth vs. firm size (in logarithms).

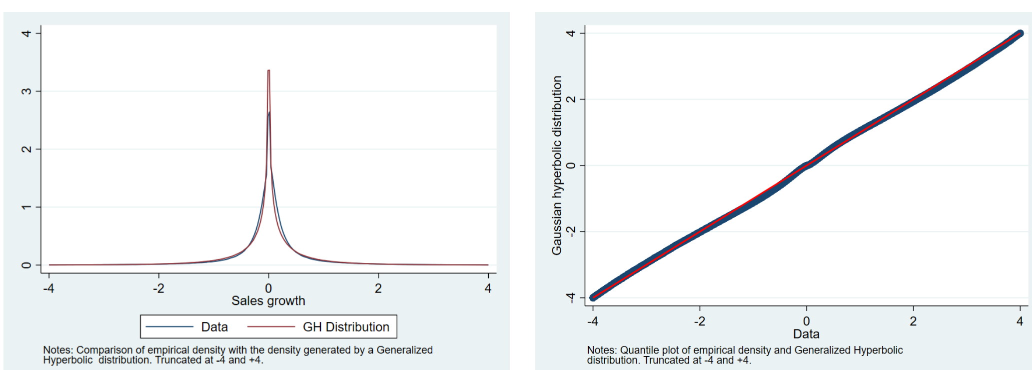

Figure A.10: Distribution of sales growth *vs.* generalized hyperbolic density and quantile-quantile plot (truncated at 0.25% for each tail).

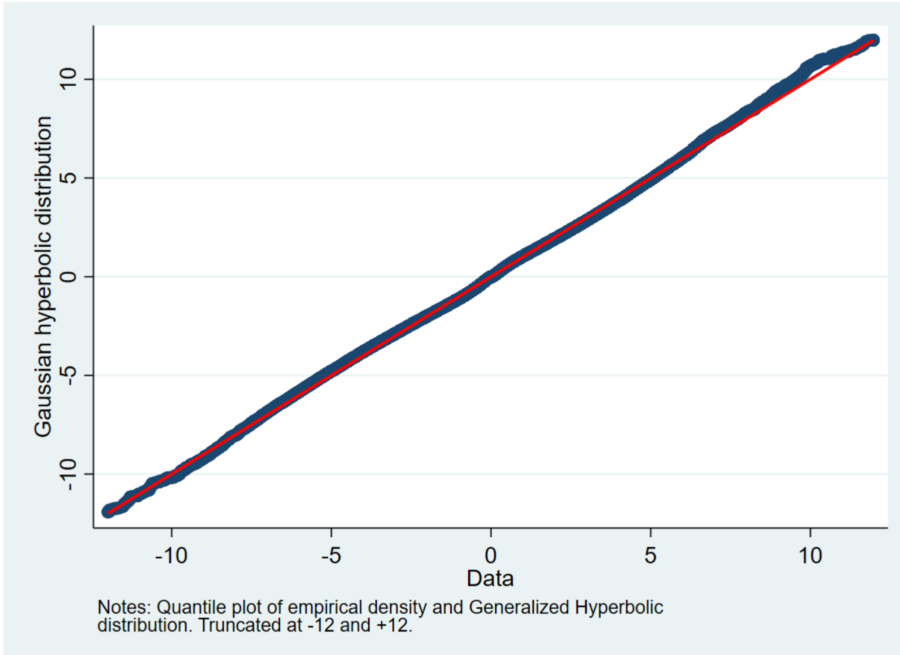

Figure A.11: Quantile-quantile plot for data vs. generalized hyperbolic distribution (truncated at 0.0005% for each tail).

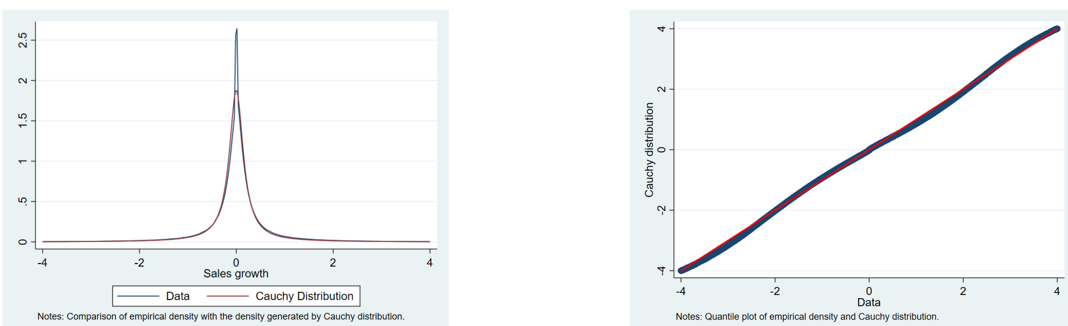

Figure A.12: Distribution of sales growth and Cauchy distribution and quantile-quantile plot.

---

to have tails that are much finer than in the data. While the Cauchy distribution approximates the data well, it performs poorly at the extreme tails (0.1% and 99.9%). Another disadvantage of the Cauchy distribution is that it may produce no defined mean or variance.

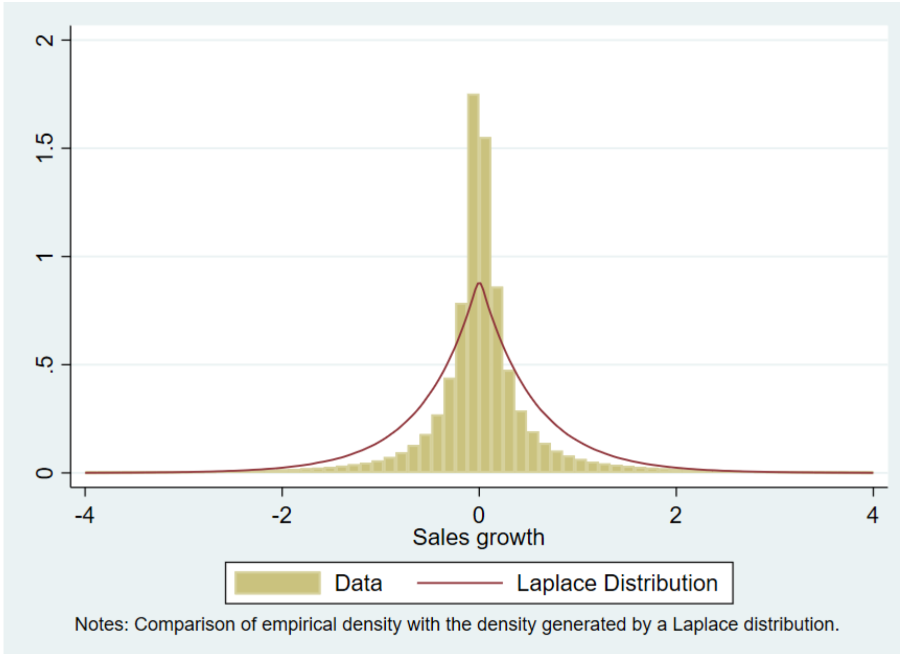

Figure A.13: Distribution of sales growth and Laplace density.

A.6 Calculation of the growth rates of demand and markup

In the case of a Cobb-Douglas production function, markups can be recovered up to a constant by calculating the inverse of the input share. This is because at the optimum, the input share must equate the output elasticity. Deviations from this equality reveal a product-level wedge (Chari, Kehoe, and McGrattan, 2007). Using intermediate inputs as the flexible input, we can thus write the markup

$$\mu = \theta^M \frac{PQ}{P_M M}$$

where P is the price, Q the quantities, P_M the price of intermediate inputs, M the quantity of intermediate inputs used, and θ^M is the input elasticity.

Santos (2020) shows how the demand shock can be recovered as a function of the flexible and pre-determined inputs. The logic is that the variation in the ratio signal unexpected shocks to the demand. As such, one can write the demand shock as a function of the input expenditures and the marginal rate of technical substitution.

$$\nu = \frac{\theta^L}{\theta^M} \frac{P_M M}{wL}$$

where w is the wage, L the employment level, and β the input elasticity of labor. Finally, TFP in revenues can be calculated in the Cobb-Douglas case using

$$a = \ln(PQ) - \alpha_k \ln K - \alpha_l \ln L - \alpha_m \ln P_M M$$

where the input shares α_j are estimated as the share of revenues for every 2 digit industry $\alpha_l = \frac{\sum_{i \in 2dig} wL_i}{\sum_{i \in 2dig} PQ_i}, \alpha_m = \frac{\sum_{i \in 2dig} P_M M_i}{\sum_{i \in 2dig} PQ_i}$ and imposing constant returns to scale $\alpha_k = 1 - \alpha_l - \alpha_m$. Finally, TFPq is calculated as the residual from the regression of TFPr on the markups and demand shocks.

A.7 Some robustness checks

A.8 The role of adjustment costs

In this section we investigate the role of adjustment costs in explaining the evidence. One reason why smaller firms have more volatile sales could be that smaller firms face less frictions in adjusting production. This is true independently of the origin of the shock. That is, it would be true for the adjustment to a supply or a demand shock. Let us now consider the volatility patterns for the productive inputs (capital, labor costs, intermediates) in this case. As a baseline reference, in the frictionless case under constant markups and Cobb-Douglas production technology, optimality conditions from the cost minimization imply that volatility of costs (e.g. expenditure in materials, total labor costs, or user cost of capital) should be proportional to the revenues. (De Loecker, 2011; De Loecker et al., 2016).

Figure A.14 shows that growth rate of labor costs, intermediates, and capital all exhibit a similar pattern to overall sales (decreasing volatility with firm size). In particular, the standard deviation of the growth rate is almost the same at about 0.3 to 0.35 for the largest 5% of firms, as expected in the frictionless baseline model. However, for the smaller firms, while the volatility of sales and intermediates is similar, capital is much less volatile and labor costs are much more volatile. This suggests that small firms depart from their optimal technical efficiency as a response to shocks, moving the capital-labor ratio (wL/rK) away from its optimal value.

	Dependent variable: $[(dy - \mu_{dy})/\sigma_{dy}]^4$			
Percentiles of:	Sales	Labor Costs	Capital	Labor Productivity
5th percentile	6.514***	11.74***	10.79***	7.031***
	(48.71)	(44.26)	(39.42)	(53.54)
10th percentile	9.494***	12.09***	8.300***	10.96***
	(44.36)	(22.56)	(42.77)	(37.39)
15th percentile	12.17***	8.718***	12.97***	15.05***
	(32.04)	(40.84)	(30.06)	(29.58)
20th percentile	15.01***	13.32***	15.15***	17.30***
	(26.28)	(31.41)	(26.99)	(26.77)
25th percentile	15.76***	16.08***	17.00***	19.30***
	(29.96)	(28.15)	(29.19)	(23.17)
30th percentile	17.43***	16.00***	19.15***	22.51***
	(24.51)	(30.89)	(21.92)	(19.23)
35th percentile	20.04***	17.60***	20.64***	24.22***
	(21.37)	(24.33)	(24.20)	(18.86)
40th percentile	21.77***	20.72***	21.45***	25.56***
	(20.85)	(24.00)	(24.12)	(17.35)
45th percentile	22.25***	21.87***	22.38***	27.75***
	(23.43)	(24.30)	(27.23)	(17.76)
50th percentile	26.63***	26.26***	25.12***	27.63***
	(15.14)	(20.55)	(17.45)	(20.17)
55th percentile	25.56***	27.56***	27.06***	29.14***
	(19.28)	(16.90)	(14.21)	(16.02)
60th percentile	28.27***	30.48***	28.56***	29.18***
	(15.22)	(20.08)	(20.79)	(18.32)
65th percentile	26.41***	30.57***	29.69***	28.66***
	(17.30)	(16.45)	(17.71)	(16.45)
70th percentile	27.78***	33.64***	30.46***	30.44***
	(16.91)	(16.81)	(14.17)	(16.42)
75th percentile	25.51***	40.32***	31.70***	31.82***
	(18.05)	(12.96)	(21.08)	(16.78)
80th percentile	26.74***	42.41***	33.80***	32.57***
	(14.81)	(14.31)	(17.99)	(18.36)
85th percentile	24.95***	51.10***	39.28***	31.22***
	(17.12)	(12.63)	(13.27)	(17.49)
90th percentile	24.23***	54.75***	41.23***	30.53***
	(19.17)	(12.95)	(12.22)	(23.61)
95th percentile	30.28***	65.17***	43.42***	28.26***
	(14.68)	(11.33)	(15.84)	(25.73)
100th percentile	36.44***	89.05***	62.33***	25.12***
	(12.95)	(7.87)	(11.12)	(26.76)
Observations	4290591	4290591	4278746	4290591

Notes: OLS results for the regression of the forth power of log sales growth divided by the square of the variance on percentiles of size class (defined by sales, labor costs, capital and labor productivity). Standard errors clustered at the firm-level. 95% CIs reported.

Table A.9: Kurtosis of sales growth - by size

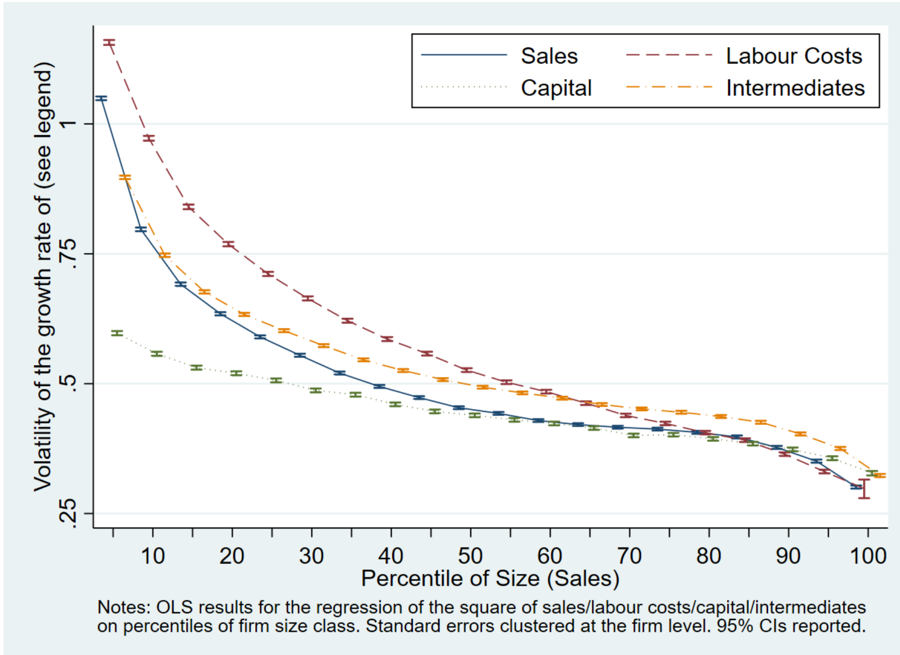

Figure A.14: Volatility of growth rate of sales, labor costs, capital, and intermediates - by percentile of sales

This is consistent with the macro-evidence (Chari, Kehoe, and McGrattan, 2007, Karabarbounis, 2014). On the one hand, capital is less volatile than sales. On the other hand, the fluctuations in production are adjusted by a large volatility of intermediate inputs and labor costs. This difference is particularly large for the smallest firms while for larger firms all the inputs and costs exhibit a similar volatility.

Given all the literature on adjustment costs and labor hoarding it is surprising to see labor costs adjusting so flexibly. We can investigate the individual components of labor adjustment using a subsample covering the period 2004-2009, for which we also observe hours worked. Figure A.15 shows that while employment/number of workers (extensive margin) is less volatile, hours worked (intensive margin) are more volatile than labor costs (labor hoarding). That is, firms seem to adjust to shocks using the intensive, rather than the extensive margin.

Summarizing, the evidence is therefore consistent with firms facing adjustment costs to capital (and the extensive margin of labor) which results in a smaller volatility for these inputs. To compensate this, intermediate inputs and hours worked have larger volatility. These differences are particularly noticeable for the smaller firms,

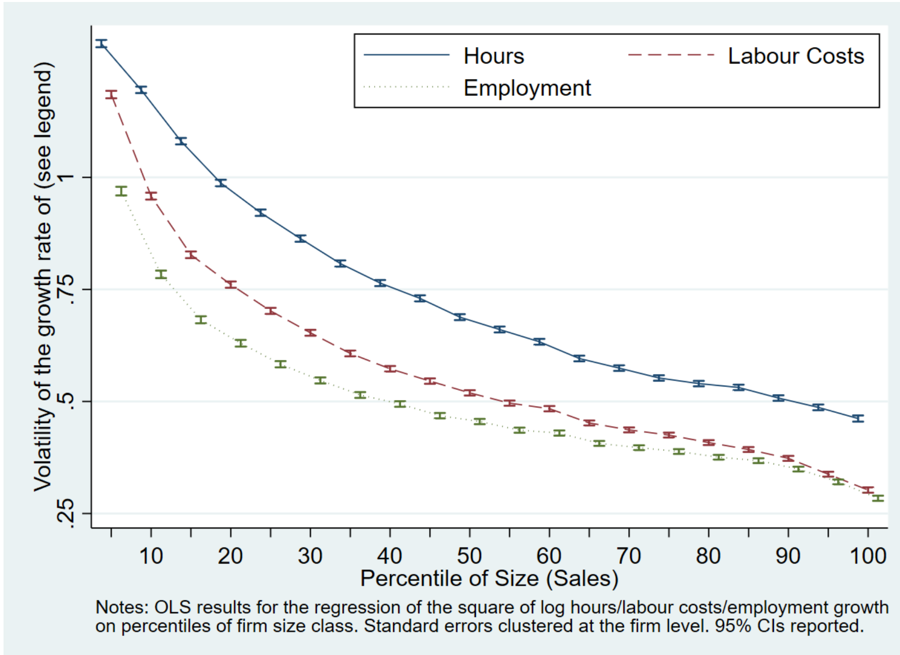

Figure A.15: Volatility of growth rate of labor costs, employment and hours - by percentile of sales

while large firms seem to behave as predicted by the frictionless model and adjust their costs more in line with their revenues. How the volatility of the different inputs vary with firm size is consistent with the baseline predictions above if we allow for adjustment costs (wedges) to capital and labor (extensive margin/hiring and firing costs). However, we also observe the same pattern of decreasing volatility for intermediates, that are unlikely to face adjustment costs. Hence, it seems that adjustment costs are insufficient to explain why volatility decreases with firm size. Furthermore, adjustment costs cannot explain the observed increase in kurtosis with firm size.