ISIGrowth

INNOVATION-FUELLED, SUSTAINABLE, INCLUSIVE GROWTH

Working Paper

Aggregate fluctuations and the distribution of firm growth rates

Giulio Bottazzi
Institute of Economics, Scuola Superiore Sant'Anna
Le Li
Institute of Economics, Scuola Superiore Sant'Anna and IBIMET-CNR
Angelo Secchi
Paris School of Economics, University of Paris 1

26/2017 **October**

This project has received funding from the European Union Horizon 2020 Research and Innovation action under grant agreement No 649186

INSTITUTE
OF ECONOMICS

Scuola Superiore
Sant'Anna

LEM | Laboratory of Economics and Management

Institute of Economics
Scuola Superiore Sant'Anna

Piazza Martiri della Libertà, 33 - 56127 Pisa, Italy
ph. +39 050 88.33.43
institute.economics@sssup.it

LEM

WORKING PAPER SERIES

Aggregate fluctuations and the distribution of firm growth rates

Giulio Bottazzi °
Le Li §°
Angelo Secchi ¶

° Institute of Economics, Scuola Superiore Sant'Anna, Pisa, Italy
§ IBIMET-CNR, Florence, Italy
¶ Paris School of Economics, University of Paris 1, France

2017/24 September 2017
ISSN(ONLINE) 2284-0400

Aggregate fluctuations and the distribution of firm growth rates*

Giulio Bottazzi†, Le Li†‡, and Angelo Secchi°

†Scuola Superiore Sant'Anna
°PSE, Universitè Paris 1 Panthèon-Sorbonne
‡IBIMET - CNR

First Draft - October 19, 2015
Present Version - September 25, 2017

Abstract

We propose an aggregate growth index that explicitly accounts for non-normality in the micro-economic distribution of firm growth rates and for the presence of a negative scaling relation between their volatility and the size of the firm. Using Compustat data on US publicly traded company, we show that the new index tracks aggregate fluctuations better than the sample average, confirming that the statistical properties characterizing the micro-economic dynamics of firms are relevant for the dynamics of the aggregate. To better characterize the origins of aggregate fluctuations, we decompose the index in two parts, describing respectively the modal (typical) value of growth rates and the tilt (asymmetry) of their distribution. Regression analysis shows that models based on this decomposition, despite their simplicity, possess a remarkable explanatory and predictive power with respect to the aggregate growth.

Keywords: Firm growth rates asymmetry and volatility; Aggregate economic fluctuations and business cycles; Aggregation of non-normal variables.

JEL Classification: C13, D22, E3, L25

*Angelo Secchi acknowledges financial support from the European Union Horizon 2020 research and innovation program under grant agreement No. 649186 (ISIGrowth).

1 Introduction

The distribution of output, employment, productivity, profitability and, in general, of all measures of firms performance is characterized by a high and persistent level of heterogeneity.[1] The same heterogeneity is present also in the distribution of the corresponding rates of change. In particular, the distribution of firms growth rate persistently displays tails fatter than those of a normal distribution (Stanley et al., 1996; Bottazzi and Secchi, 2003b, 2006a) and its dispersion is related with firms size (Hymer and Pashigian, 1962) through a negative scaling relation with an exponent approximately equal to -0.2 (Stanley et al., 1996; Amaral et al., 2001; Bottazzi and Secchi, 2006b; Criscuolo et al., 2016). These properties, which appear robust across countries and sectors, suggest that idiosyncratic shocks at firm level cannot be considered merely as disturbances or noise around a common trend but, rather, represent factors directly shaping the observed patterns of industrial evolution. Their widespread presence rises interesting questions about the link between micro behaviors and aggregate dynamics supporting the intuition put forward in Haltiwanger (1997) that changes in macro aggregates can be better understood by looking at the evolution of the cross sectional distribution of activity and of their rates of change. Inspired by these considerations, Higson et al. (2002) show that the variance and skewness of the growth rate distribution in terms of sales display a countercyclical behavior while kurtosis seems, on the contrary, procyclical. This link between micro properties and aggregate dynamics was confirmed to be quite robust in later studies finding the dispersion of the rates of change of productivity, employment, prices and business forecasts to be countercyclical while the dispersion of investments rates to be procyclical (see Bachmann and Bayer, 2014, and the references therein). However, with the only exception in Holly et al. (2013),[2] all these studies focus only on central moments of the micro-economic distributions and do not take explicitly into account neither the heteroskedastic nature of firm growth rates nor the fat tails of their distribution.

In this work we attempt to overcome this limitation and we show that exploiting the richer statistical structure of the firm growth rates distribution developed in the recent years in the field of industrial dynamics (Amaral et al., 2001; Bottazzi and Secchi, 2006b,a) improve our understanding of the economic dynamics observed in the aggregate. With this aim we develop a theoretical micro-founded index, that we call H^2, able to synthetically account for both the non-normality and the scaling of volatility of the distribution of firms growth rate. Using Compustat data on publicly traded firms operating in US from 1960 to 2013, we show that the index H^2, while remaining simple to compute, tracks the observed aggregate growth better than the first central moment of the growth distribution, that is the sample average. With respect to the existing literature this

[1] The effect of this heterogeneity on macroeconomic fluctuations has been the subject of several recent papers: Gabaix (2011) suggests that a significant part of aggregate fluctuations is explained by idiosyncratic shocks hitting few large firms, Carvalho and Gabaix (2013) investigate the possibility that the microeconomic structure explains the swings in macroeconomic volatility and explore the role of idiosyncratic shocks due to input-output linkages across the economy.

[2] Indeed, they characterize parametrically the firm growth rates distribution showing that both its shape and its scale co-move with the business cycle and contribute to the observed volatility of aggregate growth.

first result highlights the importance of properly characterizing the rich micro-economic statistical structure of firms growth and the complex economic phenomena embedded in it in order to better understand the aggregate dynamics.

Not surprisingly the agreement between our micro-economic index H^2 and aggregate growth is far from being perfect pointing at the existence of important economic phenomena not captured by the fat tails and the scaling of variance characterizing the growth rates distribution. In searching for a further improvement of the performance of our index H^2 and consistently with our empirically-driven approach, we revert to data and consider the following simple facts. In 1973, at the end of a three years robust expansion of the US economy, the average growth rate of the US public companies was around 15%.[3] In the same year the modal, or typical, growth rate was much milder, around 6%, but only 27% of all publicly traded firms were performing worse than that. Two years later, in 1975, at the end of the 73-75 recession, the average growth was -6.6%, a huge contraction, but the typical one was a more modest -1.4%. However, in that year, almost 60% of US publicly traded companies were performing worse than the modal value. A similar picture can be observed at the end of the Great Recession. These facts suggest that the large aggregate fluctuations often observed are not exclusively due to a simple common shift in the performances of individual firms, but they tend to be linked to the change of status of large groups of firms passing from being relative over-performers to be relative under-performers, and *vice versa*. These facts show that one often observes the existence of a wedge between the mean and the mode of the firm growth rate distribution, implying that the observed typical growth rate tends to be different from the average one. Despite its simplicity, this observation managed to escape to most of the previous investigations.

Inspired by these remarks and in line with our distributional approach, we decompose the micro-economic index H^2 into two parts: one part representing the modal growth rate and a residual part, that we call "distributional tilt". The former captures, by definition, the most probable growth rate one can observe in an economy at a given point in time. In this sense, it represents the growth rate of the typical firm. The latter identifies the share of firms performing better or worse than the typical one and it represents a measure of the asymmetry of the distribution. Beyond its simplicity, this decomposition possesses three distinctive features. Firstly, it allows to separate, along the business cycle, the change in the typical company behavior from the effects of moving firms above and below the modal threshold. Secondly, being based on the mode, our statistics is robust to the presence of extreme growth rates, which are likely to emerge in presence of fat tails. This choice allows us to avoid any trimming of data and to keep the likely important information embedded in extreme values.[4] Thirdly, the distributional tilt captures a form of distributional asymmetry different from the one captured by the more widely adopted skewness (i.e. third central moment),

[3]Here we focus on publicly traded firms since they are those covered by the data source we use in the present paper.

[4]Trimming or winsoring the data is common in this literature (cfr among others Higson et al., 2002; Gabaix, 2011; Holly et al., 2013). Any procedure of trimming/winsoring extreme observations, in any case disputable, becomes highly problematic in presence of fat tails.

with a likely different informational content.

We explore the validity of the mode-tilt decomposition of the index H^2 with regression analysis. Using the Compustat database, we show that the mode of the firm growth rate distribution tracks the growth rate of the aggregate output more closely than the average firm growth rate. Both the mode and the distributional tilt display a rather strong procyclical nature, manifesting a quite satisfactory explanatory power of the aggregate growth. In particular, the latter possesses a predictive power significantly better than that of the average growth rate. Finally and more interestingly, we show that these explanatory and predictive powers remain significant if we replace the Compustat aggregate growth rate with more general measures of macroeconomic growth, such as the growth rate of the Real Gross Domestic Product. In summary, the single micro-index, by blending together the typical growth rate and the distributional tilt, confounds important but diverse aspects of the relation between micro heterogeneity and aggregate dynamics. Conversely, the proposed decomposing improves our capability of tracking aggregate fluctuations and open new possibilities to better understand the micro-macro linkages.

The remainder of this paper is organized as follows. Section 2 describes the data and defines the main variables. Section 3 introduces the micro-index capturing the statistical properties of the firms growth rate distribution and discuss its decomposition into the typical growth and the distributional tilt. Section 4 contains the regression analysis while Section 5 concludes.

2 Data

The firm level analysis in this paper is based on US publicly traded companies as collected in the Compustat North America database and covers the period 1960-2014.[5] Firm size is measured in terms of Net Sales[6] expressed in millions of US dollars and deflated using the GDP Implicit Price Deflator index (base year is 2009), as reported in FRED (Federal Reserve Economic Data). We denote with $S_{i,t}$ the size of firm i at time t, with $G_{i,t} = S_{i,t+1}/S_{i,t} - 1$ the net growth rate and with $g_{i,t} = \log(S_{i,t+1}/S_{i,t})$ the corresponding logarithmic growth rate.

Economic activity at the macro level is defined using the real Gross Domestic Product (GDP) and the real Final Sales of Domestic Products (FSDP) expressed in billions of chained 2009 US dollars, as reported in FRED. We will denote with G_t^{GDP} and G_t^{FSDP} their respective net growth rate at date t. Notice that micro data from Compustat are organized in fiscal years while aggregate data are provided in calendar year.[7] In order to compare them we express everything in terms of the fiscal year (see Appendix A.1 for details).

Differently from previous works, we did not perform any trimming or winsorizing of the firm

[5]Standard & Poor's Compustat North America is a database of financial, statistical, and market information covering publicly traded companies in the United States and Canada. Canadian firms are excluded from our study.

[6]Net sales represents gross sales (the amount of actual billings to customers for regular sales completed during the period) reduced by cash discounts, trade discounts, and returned sales and allowances for which credit is given to customers. The result is the amount of money received from the normal operations of the business.

[7]This simple fact seems to went unobserved in the literature and might be responsible for a spurious dependence in lagged variables. For example in Higson et al. (2002), Holly et al. (2013) and Gabaix (2011).

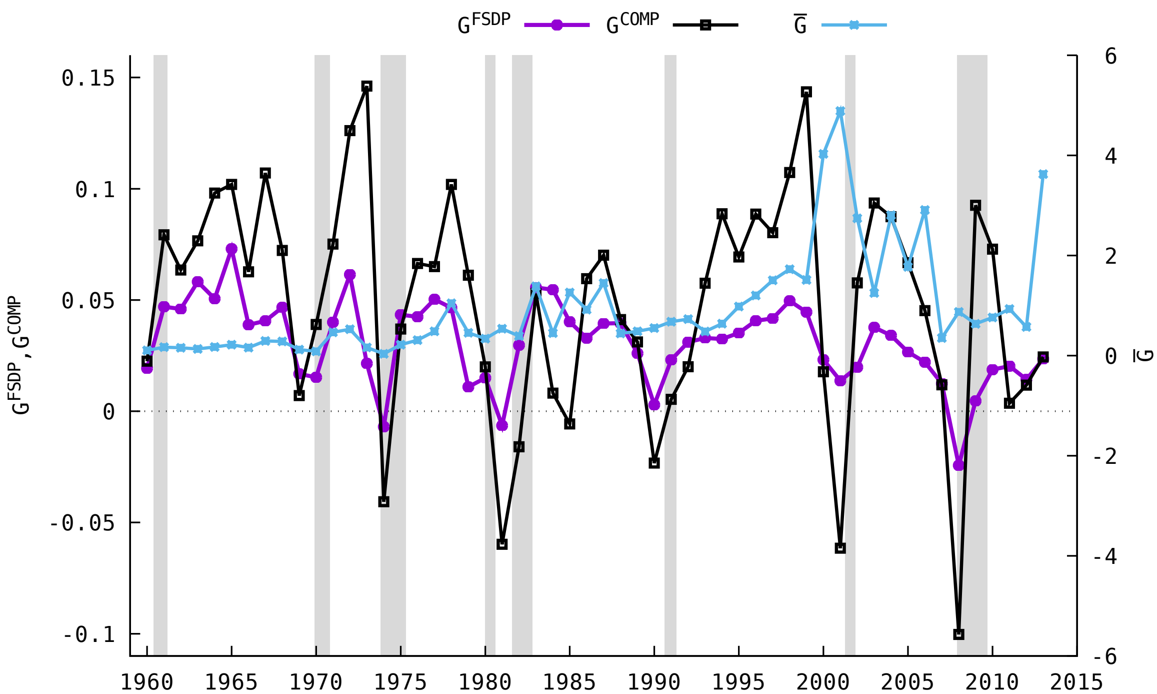

Figure 1: Time evolution over the period 1960-2013 of the net growth rate of the real FSDP G_t^{FSDP}(dark-violet line with filled circles), of the Compustat aggregate G_t^{COMP}computed according to (2) (black line with empty squares) and of the average net growth rate (cyan line with asterisks) of Compustat companies $\bar{G}_t$ computed according to (1). The reference scale of the latter is on the right y-axis. Shaded areas represent recessions according to the NBER business cycle dates.

growth rates distribution.[8] These procedures are generally adopted to avoid mixing of organic growth and external growth via e.g. merger and acquisition. However, we checked that the extreme growth rates present in our database represent perfectly legitimate events of the normal life of a business firm (See Appendix A.2 for a deeper discussion of this point). Hence we prefer not to exclude them from the analysis.

3 Aggregate growth rate and firms dynamics

Obviously, if all the companies of the US economy grew at the same annual net growth rate G_t, then the aggregate total sales, as measured for example by the real FSDP, would growth at the same rate. Does this trivial equivalence extend to the average micro-level growth rate when a population of heterogeneous firms is considered? In other words, is the average net growth rate of companies a good approximation of the aggregate growth rate observed for the whole economy? To answer this question, Figure 1 reports the time evolution of G_t^{FSDP}, the net rate of change of the real FSDP,

[8]Higson et al. (2002) and Holly et al. (2013) trim growth rates at $(-25\%, 25\%)$ and $(-50\%, 50\%)$ respectively while Gabaix (2011) winsorize them at 20%.

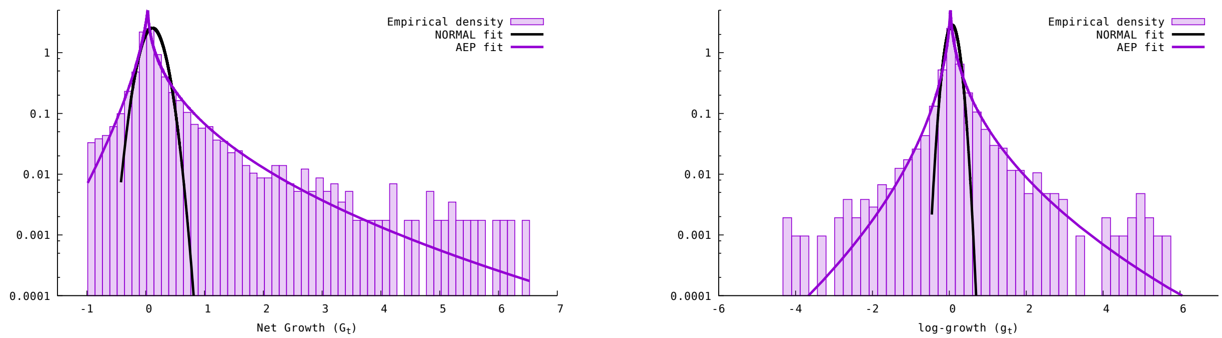

Figure 2: Total Sales net growth rates (G_t) distribution in 2013 together with a Gaussian and Asymmetric Exponential Power (AEP) fit (left panel). Total Sales log-growth rates (g_t) distribution in 2013 together with a Gaussian and Asymmetric Exponential Power (AEP) fit (right panel).

together with the average net firm growth rate of our sample defined as

$$\bar{G}_t = \frac{1}{N} \sum_i \frac{S_{i,t+1} - S_{i,t}}{S_{i,t}} \ . \tag{1}$$

As can be seen, the difference between the average growth rate of US publicly traded business companies (reference scale on the right y-axis) and the macro growth rate (reference scale on the left y-axis) is huge, often spanning two orders of magnitude. And this appears true across highly diverse historical periods. One possible explanation for the observed difference might be the limited coverage of the Compustat database, which only includes the relatively few companies that are publicly traded. One might suspect, indeed, that when averaging over a larger group of US firms, a stronger agreement between $\bar{G}_t$ and G_t^{FSDP} will emerge. Due to the lack of data, we cannot increase the number of firms we consider, but we can do a similar test by reducing the scope of the aggregate variable. To this end we define the Compustat net growth rate as

$$G_t^{\text{COMP}} = \frac{\sum_i S_{i,t+1}}{\sum_i S_{i,t}} - 1 \ , \tag{2}$$

which is basically equivalent to the G_t^{FSDP} but it is built considering only the publicly traded companies included in Compustat.[9] Its time evolution is reported in Figure 1 (reference scale on the left y-axis). Even if G_t^{COMP} fluctuates significantly more than the aggregate growth rate G_t^{FSDP}, the two quantities have the same order of magnitude and they are highly correlated, with a Spearmann rank statistics of 0.58. However, $\bar{G}$ does not seem to track G_t^{COMP} any better than G_t^{FSDP}. The average growth rate $\bar{G}_t$ constitutes a poor and uninformative approximation not only of the macro-economic growth rate, but also of the growth rate of the Compustat aggregate. Thus, we can conclude that the difference between the micro-economic average and the aggregate measure

[9] Our definition of $G_{i,t}$ requires to observe the same firm in two consecutive years. For consistency in building G_t^{COMP} we consider only those firms that are present in both t+1 and t. Since this might be associated with an attrition bias Appendix A.3 provides evidence that this bias is not very large.

persists even when the whole universe of firms contributing to the aggregate measure is used in computing the average.

Why do we observe such a poor agreement between micro and aggregate growth rates? The high and persistent heterogeneity observed in firm growth rates (Stanley et al., 1996; Bottazzi and Secchi, 2003a) surely plays a role in this mismatch. The left panel of Figure 2 displays the empirical density of the net growth rates for Compustat firms in 2013. Notice its extremely skewed shape. This skweness implies that the net growth rate at firm level can be extremely diverse and this diversity is in fact responsible of the high volatility of $\bar{G}_t$ observed in Figure 1. The average values are in fact driven by a few extreme observations and they are in general a poor and unreliable approximation of the net growth rate of the typical firm. In presence of such extreme growth events, the log growth rate $g_{i,t}$ (see the right panel of Figure 2) seems better suited than the net growth rate $G_{i,t}$ to represent the growth dynamics of firms. Indeed, the density of log growth rates presents an apparent smoother and more symmetric behavior.[10] This difference in the shape of the two densities is not peculiar of 2013 but it is common across all the years of our database. However, the statistical issue posed by the extremely skewed nature of the net growth rates distribution is not the only phenomenon responsible for the poor agreement between micro and aggregate growth rates. As discussed in the next section, a more fundamental role is played by the multiplicative nature of the firm growth process.

Heteroskedasticity and fat tails

As a large amount of empirical studies has made clear, the best synthetic description of the dynamics of firms is the so called Gibrat's Law,[11] which postulates that a firm's growth dynamics can be characterized as a geometric Brownian motion $S_{i,t+1} = \epsilon_{i,t} S_{i,t}$ where $\epsilon_{i,t}$ is a random variable shocking a firm's initial size $S_{i,t}$ in a multiplicative way.[12]

In order to exploit the multiplicative nature of the firm growth process and the observed relative higher stability of the log growth rate distribution we rewrite the Compustat aggregate net growth rate G_t^{COMP} in terms of firm log growth rates $g_{i,t}$ as

$$G_t^{\text{COMP}} = \frac{\sum_i S_{i,t} e^{g_{i,t}}}{\sum_i S_{i,t}} - 1 \,. \tag{3}$$

Applying the expectation operator and using the definition of the cumulant generating function

[10]Note that fitting an Asymmetric Exponential Power distribution (cfr. Bottazzi and Secchi, 2011) on $g_{i,t}$ suggests that the empirical distribution is neither perfectly symmetric nor Gaussian in the tails. We will discuss and exploit these two features in the next Section.

[11]Sutton (1997) is a complete even if rather old review of the literature on the Gibrat's legacy. See Lotti et al. (2003) for an update. See also Fu et al. (2005).

[12]An alternative, additive, model would be $S_{i,t+1} = S_{i,t} + \epsilon_{i,t}$. This model would imply that the average and the standard deviation of firm growth rates decrease linearly with the size of the firm, a prediction which is strongly violated by data. Notice also that Gibrat's original idea was that a firm's growth rate is independent from its size. This is only partly true, as we will discuss below.

one obtains

$$E[G_t^{\mathrm{COMP}}] = \frac{\sum_i S_{i,t} E[e^{g_{i,t}}]}{\sum_i S_{i,t}} - 1 = \frac{\sum_i S_{i,t} e^{\sum_{n=1}^{\infty} C_n[g_{i,t}]/n!}}{\sum_i S_{i,t}} - 1 \ , \tag{4}$$

where $C_n[g_{n,t}]$ represents the n-th central cumulant of the distribution of $g_{i,t}$.

Even if the support of distribution of firm log growth rate $g_{i,t}$ is smaller and more stable than the support of net growth rate $G_{i,t}$, the distribution itself still possesses a significant level of variance and we have to account for it. If we assume that the growth shocks $g_{i,t}$ are independently and normally distributed, we can truncate the cumulant expansion at the second order[13] and define H^1 as

$$H_t^1 \equiv E[G_t^{\mathrm{COMP}}] = e^{C_1[g_{i,t}] + \frac{1}{2} C_2[g_{i,t}]} - 1 = e^{\mu_t + \sigma_t^2/2} - 1 \ , \tag{5}$$

where μ_t and σ_t are respectively the mean and standard deviation of log growth rates of Compustat companies at date t. The expression in (5) takes into account the contribution of the variance of a Gaussian random variable to the expected value of its exponential. The time profile of $G_t^{\mathrm{COMP}}$and of H_t^1 are both reported in Figure 3. The variable are similar in magnitude, even if seemingly diverging. Their Spearman rank correlation is 0.32, suggesting a moderate correlation. A stronger agreement between the two quantities would have been observed, at least on average, with possible small discrepancies mainly due to sampling errors in the estimate of μ_t and σ_t, if the assumptions of normality and independence of the firm growth shocks were valid.[14] In fact, as we will discuss below, both these assumptions are violated.

Firstly, as it is apparent from Figure 2 (right panel), the empirical density of $g_{i,t}$ presents tails substantially fatter than those of a Gaussian distribution also within the Compustat database.[15] This is in line with Stanley et al. (1996) and Bottazzi and Secchi (2006a).

Secondly, in (5) we did not take into account the dependence of the volatility of a firm's growth rates on its size, a relation robustly observed in the literature. This dependence is clearly illustrated in Figure 4, were a binned plot reports the standard deviation of growth rates in 2013 as a function of firm (log) size S_i in the same year.[16] It is clear that the former declines with latter, confirming that the growth rates of small firms are more volatile than those of large companies. This negative relation displays an approximate exponential decay with an exponent of about $-0.23(0.01)$, a value very similar to that found in previous investigations (Stanley et al., 1996; Amaral et al., 2001; Bottazzi and Secchi, 2006b; Criscuolo et al., 2016). The bottom panel in the same figure reports the estimate of the scaling exponent in all the years under investigation. The exponent is characterized by a remarkable stability which confirms that the scaling property is not a peculiar feature of any

[13] In this case, indeed, $C_n[g_{i,t}] = 0$ for n>2.

[14] The weight $S_{i,t}$ being very skewed, the firms in the sample contribute in different ways to the determination of the sample average and the error does not generally decrease with $\sqrt{N}$. This is basically the central argument of the "granularity" literature (Gabaix, 2011).

[15] Results of the Maximum Likelihood estimation of the Asymmetric Exponential Power family on the growth rates distribution strongly confirm this statement and are available upon request.

[16] We rank firms according to their size in a specific year, then we split them in equipopulated bins, compute the standard deviation of growth rates of firms in each bin for that year, and plot this standard deviation against the average log-size of the bin. This procedure can be repeated for each year separately. Notice that the bin each firm belongs to can change in different years.

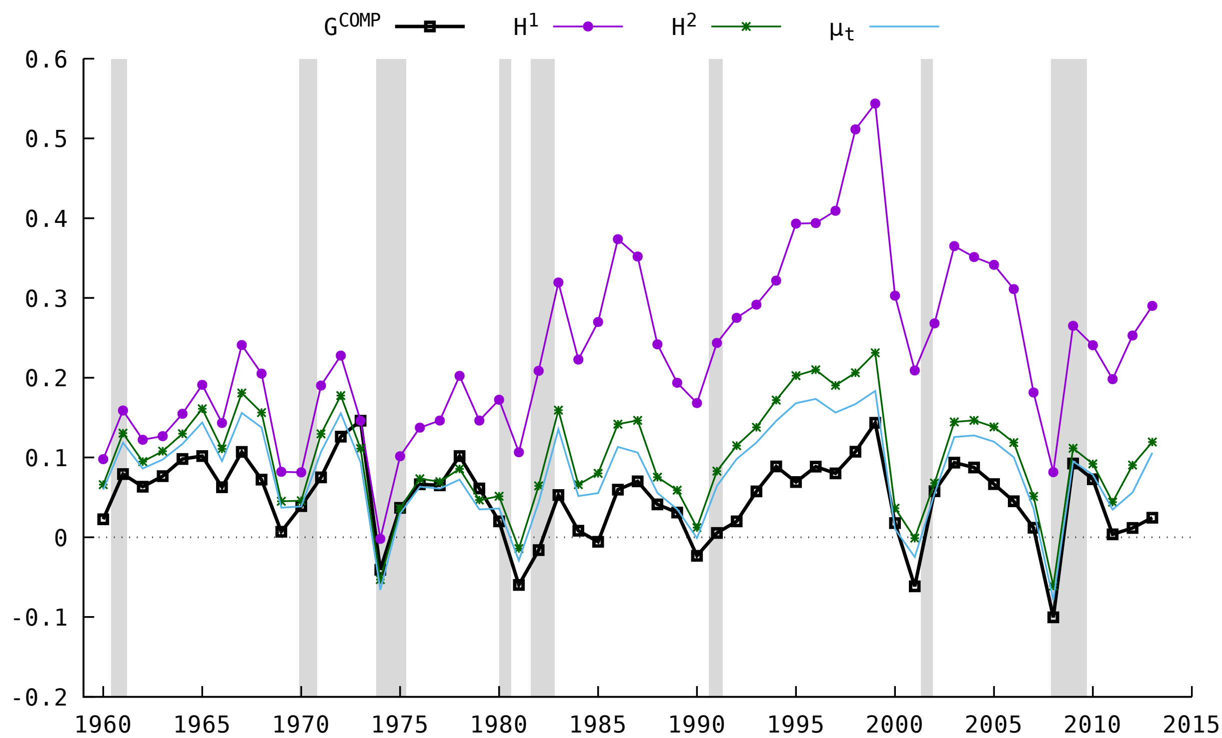

Figure 3: Time evolution over the period 1960-2013 of the net growth rate of Compustat aggregate G_t^{COMP}(black line with empty squares) and of the two approximations, H^1 and H^2, computed according to (5) (dark-violet line with filled circles) and (6) (dark-green line with asterisks) respectively. The average log growth rate of Compustat μ_t firms is also reported (cyan line with no point). Shaded areas represent recessions according to the NBER business cycle dates.

particular year but it is rather persistent in time.

This evidence implies that in order to improve our approximation of $E[G_t^{\mathrm{COMP}}]$, we should include both the observed non-normality of the log growth rate distribution and the heteroskedastic relation between the size of the firm and the volatility of its growth rates. We again start with the relation in (4) but we follow a different approach in deriving an approximation for $E[G_t^{\mathrm{COMP}}]$. First, given the fat-tailed nature of the growth rate distribution, we are forced to retain all the central cumulants since, in general, it will be $C_n[g_{i,t}] \neq 0$ for any n. Second, in order to capture the heteroskedastic effect, we assume that the second cumulant, the variance, displays an exponential relation with size with a characteristic exponent β while all the others cumulants remain independent from size. Formally, $C_n[g_{i,t}] = C_{n,t}$ for $n = 1$ and $n > 2$, while $C_2[g_{i,t}] = (\bar{S}/S_{i,t})^{2\beta_t}\bar{C}_{2,t}$, where $\bar{S}$ and $\bar{C}_{2,t}$ represent a reference firm size and the variance of growth rates of firms of that size.[17] With these

[17]The choice of a specific reference size is irrelevant for the argument. However, for statistical reliability, it is better to chose a moderate value. A too large value would imply a small sample, as the number of firms of larger size are fewer. Conversely, a too small value would be sensitive to the lower fringe of the size distribution, which is rather turbulent due to the continuous exit of incumbent and entry of new firms.

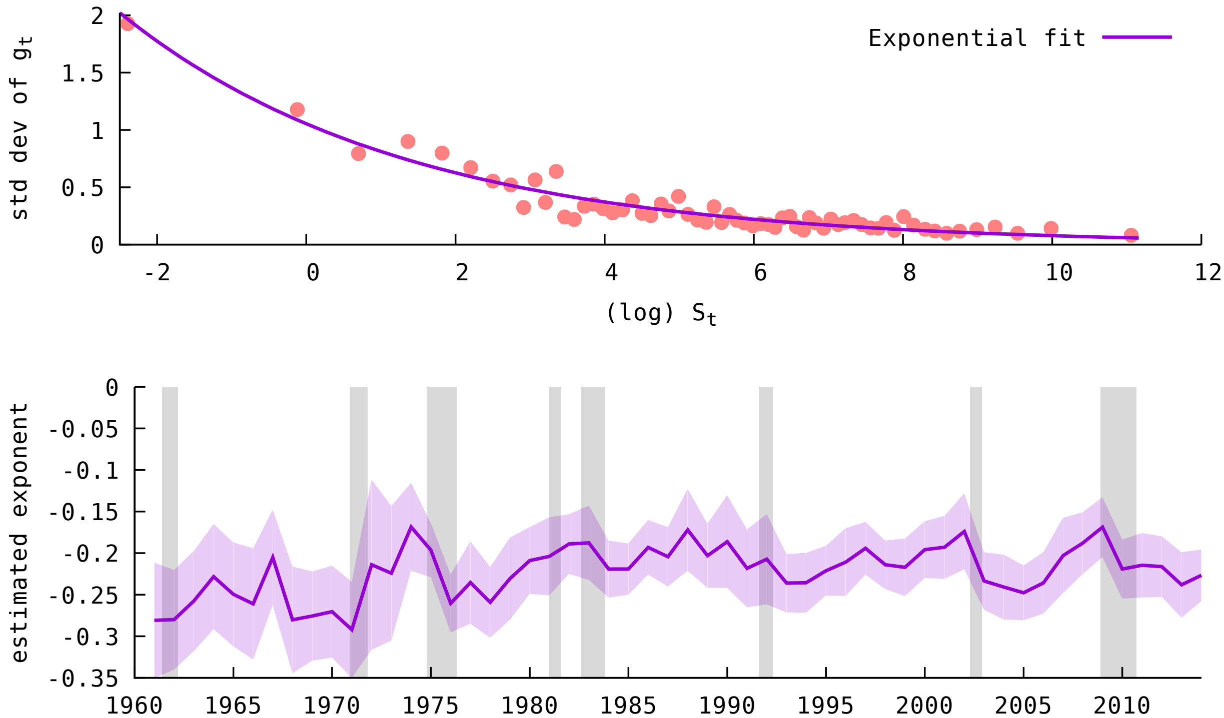

Figure 4: Standard deviation of growth rates σ_t as a function of initial (log) net sales $S_{i,t}$ together with an exponential fit (**top panel**). Estimated exponent is -0.23(0.01). Time evolution of the estimated exponent together with a 99% confidence band (**bottom panel**)

assumptions, the expression in (4) can be rewritten as

$$E[G_t^{\mathrm{COMP}}] = \frac{\sum_i S_{i,t} e^{(C_2[g_{i,t}] - \bar{C}_{2,t})/2}}{\sum_i S_{i,t}} \; e^{C_{1,t} + \bar{C}_{2,t}/2 + \sum_{n=3}^{\infty} C_{n,t}/n!} - 1 \; ,$$

where we have factorized the size-dependent variance term. Reorganizing the terms in the last expression we obtain a new approximation of the aggregate growth rate

$$E[G_t^{\mathrm{COMP}}] = \Theta_t E[e^{g(\bar{S})}] - 1 \equiv H_t^2 \; , \tag{6}$$

where $g(\bar{S})$ is log growth rate of a firm of size equal to the reference level $\bar{S}$ and

$$\Theta_t = \frac{\sum_i S_{i,t} e^{((\bar{S}/S_{i,t})^{2\beta_t} - 1)\bar{C}_{2,t}/2}}{\sum_i S_{i,t}} \tag{7}$$

is a correction term that takes into consideration the scaling of the growth rates variance with size. In order to estimate H_t^2 we follow a three-step procedure. We begin by estimating the scaling relation between the standard deviation of growth rates and firm size, thus obtaining an estimate of β_t for each year in the database (cfr. the bottom panel in Figure 4). Then, we split the sample

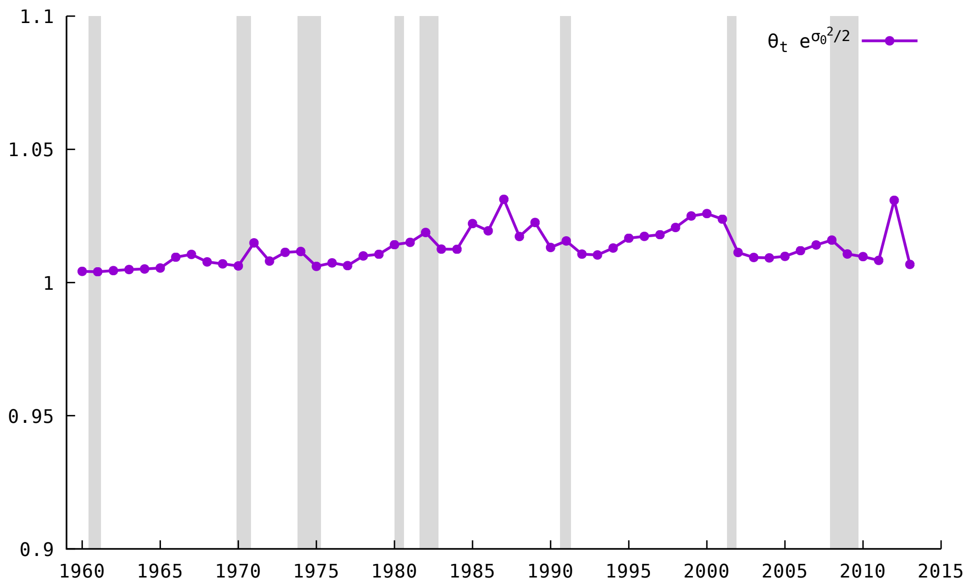

Figure 5: Time evolution over the period 1960-2013 of the factor $\Theta_t e^{\sigma_t^2(\bar{S})/2}$ appearing in (6) computed with $\bar{S} = 1$, that is assuming Total Sales equal to 1 million dollar in real terms.[20] Shaded areas represent recessions according to the NBER business cycle dates.

of firms in equally populated size classes and we compute the expected log growth rate of firms belonging to the size class including $\bar{S}$.[18] In doing this we make the assumption that the distribution of growth rates for firms in the size class $\bar{S}$ does not display a large variance, so that we can safely assume that $E[e^{g(\bar{S})}] = e^{\mu_t + \sigma_t^2(\bar{S})/2}$.[19] Finally, we compute the correction factor Θ using the entire firm size distribution, summing across all firms, each weighed with its observed size $S_{i,t}$. Notice that in the expression for Θ the smallest firms, which are weighted less for their reduced size, have in fact an enhanced effect due to the scaling of the standard deviation. If the latter were not present, that is if $\beta_1 = 0$, then it would be $\Theta = 1$ and one would get back to the previous approximation $H_t^2 \sim H_t^1$.

The performance of H_t^2 in tracking the observed aggregate growth rate G_t^{COMP}can be judged once again from Figure 3. Three comments are in order. First, H_t^2 is substantially better than H_t^1 in its capability of tracking G_t^{COMP}, with an almost doubled Spearmann correlation of about 0.77. Second, the improvement associated with H_t^2 becomes more important starting from the 70s when a well known compositional change of the Compustat database, due to the listing of younger and smaller firms, began. This observation lends support to our approach: it is precisely when the firms in the sample become potentially more diverse that explicitly taking into account the distributional properties of their sizes, their growth rates, and the relationship between the two, becomes more

[18]The number of classes should be large enough for the firms in each class to have reasonably similar sizes and small enough to provide a reasonable sample size for the computation of the average growth rate.

[19]Assuming normality for the distribution of log-growth rates in a single size class is very different, and much less demanding, than assuming normality for the distribution of log-growth rates in the entire sample, as we did for the derivation of H_t^1 in (5).

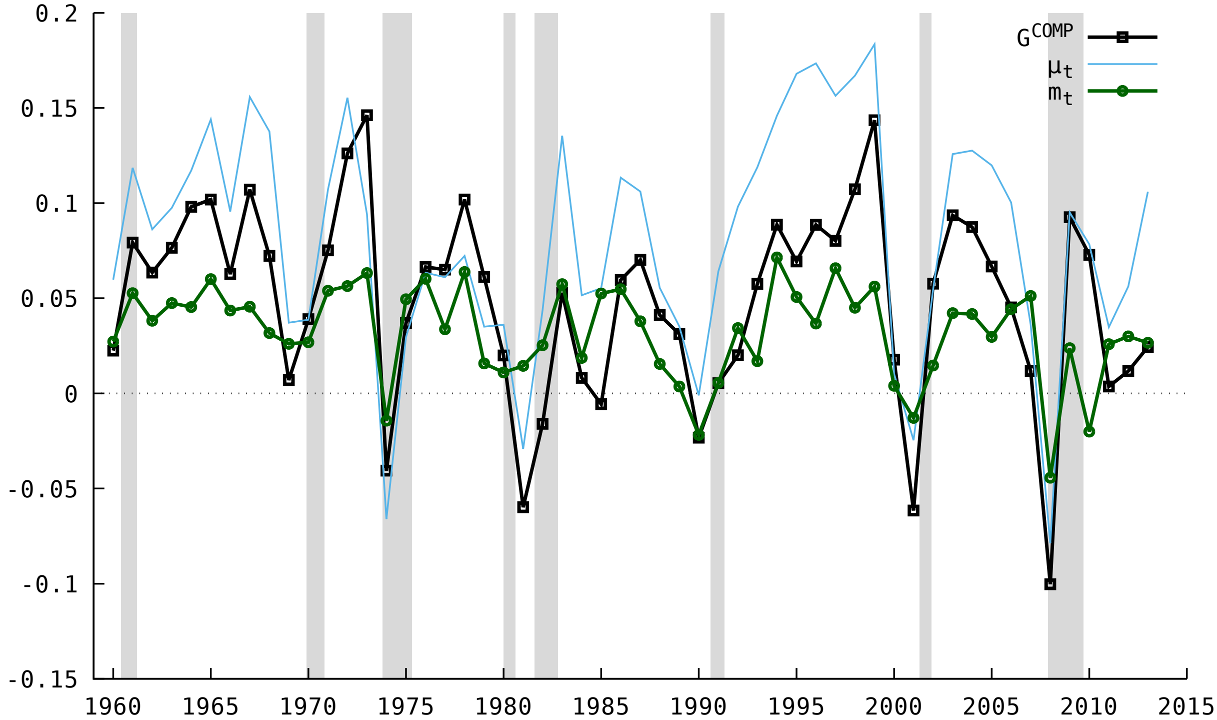

Figure 6: Time evolution over the period 1960-2013 of the aggregate Compustat growth rate G_t^{COMP}(black line with empty squares) and of the average μ_t (cyan line with no point) and of the modal m_t (dark-green line with empty circles) firms growth rate defined as the (log) difference of their Total Sales. Shaded areas represent recessions according to the NBER business cycle dates.

important. Third, the estimated value $\Theta_t\, e^{\sigma_t^2(\bar{S})/2}$ in (6) turns out to be almost 1 in every year of our data set, as shown in Figure 5. As a consequence, the expression for H_t^2 can be simplified to read

$$H_t^2 \simeq e^{\mu_t} - 1 \sim \mu_t \ .$$

This means that when one uses COMPUSTAT data the information contained in H_t^2 in terms of the cumulants of the underlying growth rates distribution are to a large extent captured by the simple mean log growth rate. This result is unexpected and not obvious and it derives from the interplay between the values of the scaling coefficient β and of the variance of the reference size class $\sigma^2(\bar{S})$ in the different years.

Asymmetry

While H_t^2 tracks better than more naive alternatives the aggregate behavior of G_t^{COMP}, it has been obtained under the rather restrictive assumption that only the second cumulant of the firms growth rate distribution depends on firm size while all the others do not. To refine H^2 and further improve

[20]The plot does not change significantly if other reference sizes are adopted.

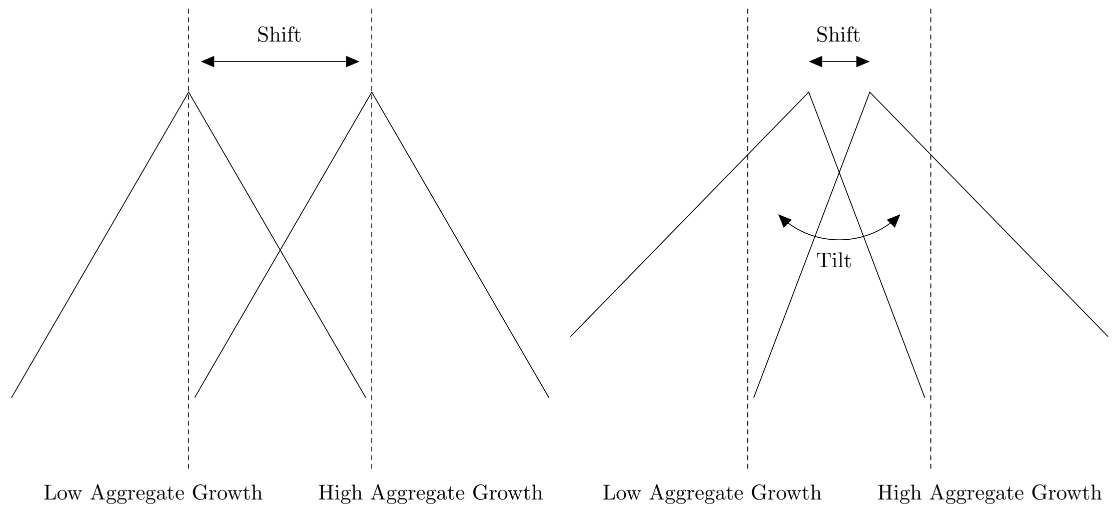

Figure 7: Firm level growth rate distribution for high and low levels of observed aggregate growth rate (dashed lines). Left: the difference is a shift of the mean. Right: in addition to a shift, also the asymmetry (tilt) of the distribution changes.

its ability to track G_t^{COMP}one would need to estimate at least a few of the infinite higher order cumulants $C_n[g_t]$ together with their possible relation with size. This turns out to be an unworkable strategy, the main reason being the relative small size of our sample of firms. Indeed, as higher cumulants are considered, the sample size required to obtain proper estimates of their values in each size class increases and, consequently, the number of size classes available for estimating the scaling coefficient reduces. This makes obtaining a reliable fit of the higher-order scaling relations impervious.

At the same time, however, it is apparent that the fluctuation of the mean and the scaling of the variance do not capture entirely the temporal evolution of the cross-sectional growth rates distribution. Consider Figure 2 (right panel), which reports the probability density of the log growth rates of COMPUSTAT companies in 2013. Contrary to what one might conclude form a superficial visual inspection, in that year the empirical density is asymmetric and its average $\mu_t = 0.106$ overestimates the typical modal growth rate, m_t, which is equal to 0.027. This is not a specific feature of the year 2013: mean and mode tend to be significantly different over the whole period under analysis. To show this, Figure 6 reports the time evolution of both quantities together with G_t^{COMP}. The mode appears much less volatile than the mean and it tends to stay on the opposite side of the aggregate growth rate G_t^{COMP}. Thus, the growth rates distribution is characterized by a changing but persistent asymmetry, which somehow seems to tack the aggregate fluctuations. Indeed the diverse dynamics of mode and mean reinforce the idea that the distributional properties of firm growth rates cannot be simply captured by μ_t and that the aggregate growth rate is linked to the dynamics of individual companies in ways more complex of those that one single central tendency measure might capture.

These considerations suggest that we may improve in tracking $G_t^{\mathrm{COMP}}$by supplementing the index H_t^2 with some measure of asymmetry. Given the strong similarity between H_t^2 and μ_t, we can decompose the index in the sum of two component

$$H_t^2 \sim m_t + p_t \ , \tag{8}$$

where m_t represents the mode of the distribution and $p_t = \mu_t - m_t$ a residual term that we identify as the "distributional tilt". Technically, the mode-tilt decomposition allows us to identify and separate, inside the average growth rate observed in one specific year as captured by H_t^2, the typical, modal, value of the log-growth rate from the movement of the probability mass between the two regions below and above the mode. The distributional tilt represents a measure of the observed asymmetry of the distribution which is alternative with respect to the more widely adopted skewness.

In the next section we show that by considering the modal value and the tilt as separate and complementary observations, we can build regression models with remarkable explanatory and predictive power with respect to the aggregate growth G_t^{COMP}. However, before moving to study the performance of the decomposition in (8), we want to conclude this section by providing a simple economic interpretation of what the wedge between the mean and the mode of the firms growth rate distribution means in terms of the underlying firms dynamics and their response to macroeconomic and idiosyncratic shocks.

Figure 7 reports a stylized representation of the firm growth rates distribution characterized by the peculiar tent shape appearing in Figure 2 and robustly observed in the literature. The peak of these tents represent the mode of the distribution, that is the most common, or typical, growth rate observed in the sample considered. Clearly, if the distribution of the growth rates is symmetric, the peak represents also the mean. In both the left and the right panel of Figure 7 we depict two of such notional distributions associated with two different growth regimes: a low growth and a high growth ones. In the scenario represented in the left panel we assume that, while influenced by idiosyncratic factors, all firms react homogeneously to the macroeconomic shocks hitting the economy. For the law of large numbers, this assumption would result in a simple shift of the distribution in the two regimes with the mean and the mode of the distribution moving together and ultimately sharing the same relation with the aggregate growth dynamics.

The scene appears different if one allows not only for idiosyncratic individual shocks, but also for possible heterogeneous responses of individual companies to the aggregate shocks. In this case, since some group of firms might over-react while other firms might under-react, together with the shift we are likely to observe a change in the distribution of probability mass around the modal value, as firms with diverse characteristics move across the modal threshold (in both directions), breaking the symmetry and separating the average and the modal growth rate. This scenario is represented in the right panel of Figure 7 where, in moving from a low growth to a high growth regime, the firm growth rates distribution is modified both by a shift and by a "tilt" of its shape.

This second scenario is more flexible and it provides a more general framework, as the shift and tilt movements of the micro distribution are allowed to take place over different time scales, to have different degree of persistence and, ultimately, to exert different effect on, or differently react to, the time evolution of the aggregate growth rate. This scenario is also more suitable to accommodate the observed large differences in the behavior of individual companies, that are continuously hit by idiosyncratic shocks and that due to differences in their internal structure or in the market environment in which they operate, are plausibly differently affected by the economic opportunities or downturns. All this is in tune with the empirical evidence built in the last few decades about the sectoral specificity of firm growth dynamics (cfr. for example the discussion in Haltiwanger, 1997), as the mass of probability moving around the modal value of the aggregate distribution might well represent groups of firms belonging to the same or similar sectors.[21] In the next section we compare the ability of the mean and the skewness, on one side, and the mode and the tilt, on the other, to track G_t^{COMP}within a regression framework.

4 Regression analysis

So far we have been content of assessing the goodness of our index in (6) and the further decomposition in (8) simply through correlation measures and visual inspections. In this Section we want to go beyond those simple analysis and try to asses, on a more quantitative basis, how good our micro variables are in tracking the cyclical behavior of the aggregate economic activity. We will perform a series of regression analysis to measure both the explanatory and the predictive power or the former with respect to the latter. The aggregate quantity we consider here, that is our dependent variable, will be G_t^{COMP}. This is a an informative exercise since, in this case, we know that the sample of firms we consider contains, by definition, all the firms contributing to G_t^{COMP}. As the regression analysis will made clear, however, this information is not trivial to extract and the choice of the statistics used to capture the properties of the the micro level distribution is likely to affect the quality of the results.

Explanatory power

To investigate the correlation between the aggregate net growth rate $G_t^{\mathrm{COMP}}$and our statistics based on the micro log-growth rate distribution of Compustat firms we consider the following specification

$$G_t^{\mathrm{COMP}} = \alpha + \sum_{\tau=0}^{T} \beta^m_{t-\tau} m_{t-\tau} + \sum_{\tau=0}^{T} \beta^p_{t-\tau} p_{t-\tau} + \epsilon_t \quad , \tag{9}$$

where m_t and p_t represent the mode and the distributional tilt $(\mu_t - m_t)$ defined in the previous Section.[22] In what follows the mode m_t is estimated using the Half Sample Method (HSM) devel-

[21]The relevance of the sectoral decomposition of the aggregate growth rate inside this more flexible framework is an inviting subject for further research which, due to space constraints, we decide not to pursue here.

[22]One could use H_t^2 instead of μ_t but, due to the similarity of the two quantities, we would not observe any significant difference. Since the computation of the former is more complicated, in what follows we will use its simpler approximation.

Table 1: EXPLANATION AND PREDICTION - Compustat AGGREGATE

							Decomposition of R^2					
	(1)	(2)	(3)	(4)	(5)	(6)	(1)a	(2)a	(3)a	(4)a	(5)a	(6)a
m_t	1.008***		0.893***				48%		39%			
	(0.201)		(0.187)									
m_{t-1}			0.157		0.290				4%		15%	
			(0.151)		(0.316)							
m_{t-1}			0.297*		0.367				2%		6%	
			(0.156)		(0.231)							
p_t	0.602***		0.586***				52%		43%			
	(0.099)		(0.113)									
p_{t-1}			0.067		0.575***				8%		49%	
			(0.125)		(0.191)							
p_{t-2}			-0.254**		-0.559***				5%		30%	
			(0.110)		(0.170)							
μ_t		0.761***		0.714***				96%		78%		
		(0.065)		(0.073)								
μ_{t-1}				0.043		0.373**				9%		53%
				(0.095)		(0.156)						
μ_{t-2}				-0.100		-0.290**				2%		16%
				(0.068)		(0.110)						
γ_t		-0.004		-0.004				4%		3%		
		(0.003)		(0.003)								
γ_{t-1}				0.005		0.009*				7%		30%
				(0.004)		(0.005)						
γ_{t-2}				0.000		0.000				0%		1%
				(0.003)		(0.004)						
$\bar{R}^2$	0.720	0.711	0.739	0.708	0.253	0.210						
Obs.	54	54	52	52	52	52						

m_t, p_t, μ_t, γ_t represent the mode, the tilt as defined in equation (8), the mean and the skewness of the firms growth rate distribution over the time span 1960-2013. The dependent variable is G_t^{COMP}, the net growth rate of the Compustat aggregate. Robust standard error in parenthesis. As usual ***, **, * denotes coefficients statistically significant at the 1%, 5% and 10% respectively.

oped in Bickel and Frühwirth (2006) and briefly described in Appendix B.1. This specification is then compared with

$$G_t^{\mathrm{COMP}} = \alpha + \sum_{\tau=0}^{T} \beta_{t-\tau}^{m} \mu_{t-\tau} + \sum_{\tau=0}^{T} \beta_{t-\tau}^{p} \gamma_{t-\tau} + \epsilon_t \quad , \tag{10}$$

where μ_t is the mean and γ_t the skewness of the firm log-growth rate distribution. Both specifications contain a measure of central tendency, the mode in the first case and the mean in the second, and a measure of distributional asymmetry, the tilt and the skewness respectively. In both models T stands for the number of lags allowed for in the model and ϵ is an error term. The results of the two regressions are reported in Table 1 and discussed below.

Let us start by estimating via OLS the simple benchmark model obtained setting $T = 0$. Column (1) shows that both m_t and p_t display a robust procyclical behavior. They both have a highly significant power in explaining G_t^{COMP}: the overall goodness of fit of this simple model is good with an adjusted R-squared $\bar{\mathrm{R}}^2$ of about 72%. This explanatory power is almost evenly distributed between the m_t and the p_t: as reported in column (1-2)a they both account for about half of the explained variance of G_t^{COMP}.[23] This as to be confronted with the same benchmark case obtaining estimating (10). In this case the result is reported in Column (2). The overall goodness of fit is slightly lower than with the previous model and, more importantly, the explanatory power of this second model is entirely due to μ_t, while no significant correlation emerges between $G_t^{\mathrm{COMP}}$and γ_t. The skewness and the tilt, while in principle capturing similar effects, perform differently in practice. This suggests that the way in which these two measures capture the observed asymmetry is in fact different.[24]

Next we turn our attention to the more general case by estimating (9) setting $T = 2$. Adding two lags to the benchmark model improves its overall explanatory power (Column 3): both m_t and p_t contribute to explain the dependent, and lagged variables turn out to be relevant. The comparison with the model using the average and the skewness (Column 4) confirms the interest of our decomposition, as in this case only the average is significant. We verified that adding further lagged values of the mean and the skewness does not improve at all the quality of the model: none of the extra regressors emerge as statistically significant.

From this analysis we discover that, as expected, the modal firm growth rate is procyclical. But from the positive contemporaneous correlation coefficient, we also observe that during an economic boom not only the typical firm grows more, but one also observes a larger "mass" of firms which perform better than the typical one. Conversely, during an economic downturn, the typical firm tends to grow at a lower pace, eventually negative, and at the same time more firms perform worse

[23]This decomposition of the R^2 is obtained using the "lgm" metric described in Chevan and Sutherland (1991). To perform the decomposition we use the "Relaimpo" R package (Grömping, 2006). Note that this decomposition applies to R^2 and not to the adjusted R^2; for our purpose of within model comparisons this discrepancy is irrelevant.

[24]This is a consideration which might well have implications going beyond the exercise presented here. For instance Imbs (2007), in studying the relation at the sectoral level among output growth, its volatility and investments, note that the skewness of output growth does not show any significant correlation with output growth. On the base of this evidence he concludes that investments are not as lumpy as usually expected. Our result suggest that when sample of heterogeneous firms are considered, the way in which one measures the asymmetry could be relevant. An analogous exercise conducted using the tilt instead of the skewness might lead to different conclusions.

than it. The double relation between $G_t^{\text{COMP}}$and the distribution of micro growth rates cannot be devised by the simple analysis of the mean value and it is not revealed if one adopt the skewness as a measure of asymmetry. The presence of a strong relation with the lagged tilt hints to a possible predictive power of this statistics and lead us to the analysis of the next section.

Predictive power

In this section we move from an explanatory to a forecasting exercise and investigate the power of m and p in predicting the future values of G_t^{COMP}, again comparing the results with what can be obtained using μ_t and γ_t. To do that, we remove from (9) and (10) the values of the regressors contemporaneous to the dependent variable, that is $\tau = 0$, and we estimate the remaining lagged variables setting again $T = 2$. The results of an OLS estimate are reported in columns (5) and (6) of Table 1.

First, considering its simplicity, the goodness of fit of the model with mode and tilt is rather remarkable, with an $\bar{R}^2$ of about 25%. Almost 80% of this predictive power is due to the role of the distributional tilt, whose two lag values emerge as highly statistically significant. Hence, the movement of the probability mass of firm growth rates around the typical, modal, value in a year, p_t, represents a new and apparently useful predictor of the observed aggregate fluctuations. The performance of the model with mean and skeweness is lower, with an $\bar{R}^2$ of 21%. As expected from the previous analysis, the skewness index has a relatively minor predictive power.

Robustness checks

To check the robustness of our regression results and hence the reliability of our interpretations in this section below we perform 3 sets of tests whose results are reported in Table 2.

A first concern is related with our choice of not removing extreme growth events from the database in connection with the use of an estimation technique, the OLS, known to be sensitive to them. To deal with the fact that extreme growth rates might represent either legitimate events of the life of a business firm or peculiar ones, like mergers and acquisitions, that one would like to clean out we perform two complementary exercises. First we estimate our benchmark model with a technique more robust to the presence of extreme observations, namely the Least Absolute Deviation (LAD) regression. Second we estimate the benchmark by trimming our individual firm growth rates at 25%. Using the LAD approach, instead of the OLS, does not seem to have any impact on our results, the estimated coefficients are unaffected.[25]. Trimming the database seems to have a bigger quantitative impact on estimates without, however, changing qualitatively our story. This was expected since the trimming procedure “artificially” changes the shape of the cross-sectional distribution of individual growth rates directly impacting m_t and p_t. So the shape of this distribution ends up to depend on the trimming threshold; a dependence that support once again our skepticism in sample cleaning of this sort.

[25] Note that the $\bar{R}^2$ for the LAD regression is not directly comparable with the adjusted $\bar{R}^2$

Table 2: EXPLANATION AND PREDICTION - ROBUSTNESS CHECKS

	LAD Regression	25% Trimming	Parametric Mode and Tilt				Including macro controls	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
m_t	1.042***	1.201***	1.313***	1.266***			0.907***	
	(0.280)	(0.060)	(0.173)	(0.184)			(0.161)	
m_{t-1}					0.509*	0.447		0.126
					(0.295)	(0.317)		(0.368)
m_{t-2}					-0.057	-0.011		0.471*
					(0.278)	(0.294)		(0.261)
p_t	0.602***	1.072***	0.535***				0.624***	
	(0.128)	(0.104)	(0.072)				(0.105)	
p_{t-1}					0.427**			0.464**
					(0.189)			(0.212)
p_{t-2}					-0.313*			-0.422**
					(0.160)			(0.207)
$(a_r - a_l)_t$				0.794***				
				(0.121)				
$(a_r - a_l)_{t-1}$						0.665**		
						(0.300)		
$(a_r - a_l)_{t-2}$						-0.491*		
						(0.270)		
$\bar{R}^2$	0.458	0.862	0.775	0.749	0.156	0.144	0.787	0.395
Obs	54	54	54	54	52	52	54	52

m_t, p_t represent the mode (estimated with the HSM method in column (1) and (2) and AEP from column (3) to (6)) and the tilt as defined in equation (8) of the FGRD over the time span 1960-2013. $(a_r - a_l)_t$ represents one alternative way to measure tilt using scaling parameters of AEP. The dependent variable is G_t^{COMP}, the growth rate of the Compustat aggregate, except for column (2) where it is the aggregate growth computed using all firms with growth rates between $[-0.25, 0.25]$. Robust standard error (bootstrap standard errors for LAD regression) in parenthesis. $\bar{R}^2$ represents Pseudo-R^2 for LAD regression and adjusted-R^2 for others. As usual ***, **, * denotes coefficients statistically significant at the 1%, 5% and 10% respectively.

A second concern is related to our non-parametric estimate of the mode of the distribution, as said based on the Half Sample Method (HSM) discussed in (Bickel and Frühwirth, 2006). To check the robustness of our results in this respect we estimate the mode and the distributional tilt using a parametric method based on the Asymmetric Power Exponential distribution and whose details are reported in Appendix B.2. We then estimate again equation (9) with OLS reporting the results in Columns (3)-(6) of Table 1. Our main results, namely the existence of a good deal of explanatory and predictive power of the typical growth rates and of a proper index of asymmetry with the latter playing a dominant role, emerge as robust.

A final concern regards the lack in our regression models of some of the most common explanatory factors and predictors of aggregate growth. To check if the omission of these variables affects our results we follow Gabaix (2011) and we estimate our benchmark models adding controls for oil and monetary shocks, a short term interest rate and a short term spread.[26] Not surprisingly adding these factors improves the explanatory and predictive power of our regression models. However their inclusion does not kill the statistical significance of the mode and of the distributional tilt and their overall contribution to the explained variance remains substantial: more than 80% in column (7) and almost 40% in column (8)).

Predictive power with macroeconomic aggregates

So far we have provided evidence that describing the structure of the firm growth rate distribution using its mode m_t and the corresponding distributional tilt p_t yields a higher explanatory and predictive power of the net growth rate of the Compustat aggregate than when the mean and the skewness are used. In this section we conclude our investigation with a harder test: we check if the explanatory and predictive power of m_t and p_t is confirmed when we replace the aggregate net growth rate computed with Compustat firms (G_t^{COMP}) with a macroeconomic measure describing the behavior of the whole economy, namely the net growth rate of the real Gross Domestic Product (GDP) and the net growth rate of the real Final Sales of Domestic Products (FSDP).[27] The spirit of this exercise is to verify to what extent our simple way to characterize the growth rate distribution of an important but selected sample of firms, i.e. those that are publicly traded, is effective in explaining and predicting the dynamics of the whole economy.

With this aim we estimate equations (9) and (10) replacing the dependent variable with the net growth rate of the real GDP and of the real FSDP respectively. Results are reported in Table 3. We start by commenting results for the real GDP. Column 1 shows that contemporary values of m_t and p_t both have a statistically significant explanatory power for the net growth of the real GDP, even if as expected the overall goodness of fit ($\bar{R}^2$ is 58%) of this model is lower than what we got in Table 1 with the same regressors and the aggregate Compustat growth G_t^{COMP}on the left hand

[26]Details on these proxies are provided in Appendix A.4.

[27]In principle FSDP differs from GDP because the former takes into account the change in private inventories. In practice the growth rate of the two are highly correlated justifying the similarity of the results. With this respect the correlation of the net growths of GDP and FSDP with the aggregate Compustat growth G_t is not as high, being 0.68 and 0.66 respectively.

Table 3: EXPLANATION AND PREDICTION - MACROECONOMIC AGGREGATES

Dep. variable	RGDP net growth rate				RFSDP net growth rate				Decomposition of R^2							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(1)a	(2)a	(3)a	(4)a	(5)a	(6)a	(7)a	(8)a
m_t	0.486***				0.425***				63%				65%			
	(0.081)				(0.068)											
m_{t-1}		0.064		0.026		0.158		0.112		7%		2%		27%		7%
		(0.123)		(0.097)		(0.109)		(0.084)								
m_{t-2}		0.186*		0.197**		0.134		0.134		10%		6%		7%		4%
		(0.099)		(0.089)		(0.090)		(0.080)								
p_t	0.171***				0.137***				37%				35%			
	(0.043)				(0.036)											
p_{t-1}		0.181**		0.241***		0.150*		0.191***		31%		12%		33%		13%
		(0.087)		(0.075)		(0.077)		(0.063)								
p_{t-2}		-0.248***		-0.199**		-0.184***		-0.123*		52%		8%		33%		4%
		(0.071)		(0.075)		(0.056)		(0.065)								
μ_{t-1}			0.105*				0.123**				41%				62%	
			(0.057)				(0.051)									
μ_{t-2}			-0.116*				-0.091*				36%				19%	
			(0.066)				(0.054)									
γ_{t-1}			0.002				0.002				22%				18%	
			(0.002)				(0.002)									
γ_{t-2}			0.000				0.000				1%				1%	
			(0.003)				(0.002)									
Macro controls	No	No	No	Yes	No	No	No	Yes								
$\bar{R}^2$	0.583	0.172	0.068	0.575	0.562	0.184	0.110	0.574								
Obs.	54	52	52	52	54	52	52	52								

m_t, p_t, μ_t, γ_t represent the mode, the tilt as defined in equation (8), the mean and the skewness of the firms growth rate distribution over the time span 1960-2013. The dependent variable is G_t^{GDP}, the net growth rate of Real Gross Domestic Product (RGDP), in columns 1 to 4 and G_t^{FSDP}, the net growth rate of Real Final Sales of Domestic Products (RFSDP), in columns 5 to 8. Robust standard error in parenthesis. ***, *, * denotes coefficients statistically significant at the 1%, 5% and 10% respectively.

side ($\bar{R}^2$ was 72%). Around two thirds of this 58% is associated with the mode and the remaining part to the distributional tilt. Once again the combined explanatory power of m_t and p_t is higher than that of mean and skewness with this latter not significant in practically any of our regressions.

Next columns (2) and (3) focus on their predictive power, and consequently they include only lagged values of the regressors. Three considerations are relevant. Firstly, lagged values of m_t and p_t appear to have a predictive power with respect to real GDP net growth; this power is stronger for the distributional tilt both in terms of statistical significance and in term of their contribution to the $\bar{R}^2$ which amounts to around 83%. Secondly, lag values of mode and tilt display a stronger predictive power than the mean and the skewness: $\bar{R}^2$ of the former is 10 percentage points higher. Thirdly, this predictive power is not killed by the introduction in the regression of a set of usual GDP predictors. Column 4 indeed show that m_{t-2}, p_{t-1} and p_{t-2} are all highly statistical significant and they count for around one fourth of the explained variance. Results for the growth of RFSDP are very similar and all the comments above remain valid.

5 Conclusions

In this paper we explore the relation between the collective growth dynamics of firms and that observed in the aggregate. Indeed this relation is found to be not trivial: even with a rather homogeneous sample of firms, those publicly traded in the US, the average firm growth rate is only weakly correlated and much more volatile than the aggregate one. We show that the two major factors driving the wedge between micro and macro dynamics are the heteroskedastic nature of firm growth rates, whose variance is significantly smaller for larger firms, and the fat-tailed nature of their distribution. Consistently we build a synthetic index, H^2, that embeds these two distributional properties. Using H^2 index, the tracking of aggregate fluctuations is improved by almost 100%, with a correlation between the index and the aggregate growth rate reaching 0.77. However, the agreement is still far from being perfect. This is due to the fact that H^2 doe not take into account the apparent asymmetry of the distribution of firm growth rates. A fact that has been remarkably ignored by the previous literature. Thus, inspired by the available empirical evidence, we propose a parsimonious way to account for the observed asymmetry by decomposing the mean into the mode and the distributional tilt defined as the difference between the mean and the mode.

Then, we show that a simple regression model for the Compustat aggregate growth rate which includes as regressors the mode and the distributional tilt of the firms growth rate distribution, possesses a quite satisfactory explanatory and predictive power. In particular, both the explanatory and the predictive powers are higher than those obtained using the mean and the skewness of the same distribution. Remarkably this improved performance is retained when we replace as the dependent variable in our regressions the aggregate growth rate of the Compustat sample with the aggregate growth rate of the whole economy, as measured for example in terms of GDP. This result is consistent with the existence of differences in the way firms react to economic booms and downturns. An increase of the aggregate growth rate is more likely associated with a "tilt effect" induced by a group of firms that outperform the typical, modal, firm rather than by a “shift effect” due to a change in the average growth rate of the whole population. We interpret this evidence as suggesting that the economic mechanisms driving the shift and the tilt effects are diverse, and that they are likely to manifest themselves to some extent independently and possibly on different time scales. The use of the average growth rate to describe the common behavior of the sample of firms contributing to the definition of the economic aggregate, a widespread approach, mistakenly mixes together two different aspects of the distributional dynamics: the movement of the typical growth rate and that of the probability mass around it. This mixing might hide what, on the contrary, can emerge as important separate factors influencing the growth dynamics in the aggregate.

Despite the simplicity and the limits of our approach, we believe that our analysis provides new supports to the idea that, in order to improve our understanding of the macroeconomic dynamics, it is crucial to take into consideration the structure of the persistent heterogeneity observed at the firm level, as captured by the distributional properties that represent the natural description of this heterogeneity. One has to go beyond the simplistic view that the complex phenomenon of economic growth can be described by taking simple averages and, rather, build upon the specific knowledge provided by an increasing range of studies on firm dynamics, at the same time exploiting the amount of micro-economic data nowadays available.

References

Amaral, L., P. Gopikrishnan, V. Plerou, and H. Stanley (2001): “A model for the growth dynamics of economic organizations,” *Physica A*, 299, 127–136.

Bachmann, R. and C. Bayer (2014): “Investment Dispersion and the Business Cycle,” *American Economic Review*, 104, 1392–1416.

Bickel, D. R. and R. Frühwirth (2006): “On a fast, robust estimator of the mode: Comparisons to other robust estimators with applications,” *Computational Statistics & Data Analysis*, 50, 3500–3530.

Bottazzi, G. (2004): “Subbotools user's manual,” *LEM Papers Series*, 14.

Bottazzi, G. and A. Secchi (2003a): “Properties and Sectoral Specificities in the Dynamics of U.S. Manufacturing Companies,” *Review of Industrial Organization*, 23, 217–232.

——— (2003b): “Why are distributions of firm growth rates tent-shaped?” *Economics Letters*, 80, 415 – 420.

——— (2006a): “Explaining the Distribution of Firms Growth Rates,” *The RAND Journal of Economics*, 37, 235–256.

——— (2006b): “Gibrat's Law and diversification,” *Industrial and Corporate Change*, 15, 847–875.

——— (2011): "A new class of asymmetric exponential power densities with applications to economics and finance," *Industrial and corporate change*, 20, 991–1030.

Carvalho, V. and X. Gabaix (2013): "The Great Diversification and Its Undoing," *American Economic Review*, 103, 1697–1727.

Chevan, A. and M. Sutherland (1991): "Hierarchical Partitioning," *The American Statistician*, 45, 90–96.

Criscuolo, C., F. Calvino, C. Menon, and A. Secchi (2016): "Volatility scaling: a cross-countries investigation," *mimeo*.

di Giovanni, J., A. A. Levchenko, and I. Mejean (2014): "Firms, Destinations, and Aggregate Fluctuations," *Econometrica*, 82, 1303–1340.

Fu, D., F. Pammolli, S. V. Buldyrev, M. Riccaboni, K. Matia, K. Yamasaki, and H. E. Stanley (2005): "The growth of business firms: Theoretical framework and empirical evidence," *Proceedings of the National Academy of Sciences of the United States of America*, 102, 18801–18806.

Gabaix, X. (2011): "The granular origins of aggregate fluctuations," *Econometrica*, 79, 733–772.

Grömping, U. (2006): "Relative Importance for Linear Regression in R: The Package relaimpo," *Journal of Statistical Software*, 17, 1–27.

Haltiwanger, J. C. (1997): "Measuring and analyzing aggregate fluctuations: the importance of building from microeconomic evidence," *Federal Reserve Bank of St. Louis Review*.

Hamilton, J. D. (2003): "What is an oil shock?" *Journal of econometrics*, 113, 363–398.

Higson, C., S. Holly, and P. Kattuman (2002): "The cross-sectional dynamics of the US business cycle: 1950–1999," *Journal of Economic Dynamics and Control*, 26, 1539–1555.

Holly, S., I. Petrella, and E. Santoro (2013): "Aggregate fluctuations and the cross-sectional dynamics of firm growth," *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 176, 459–479.

Hymer, S. and P. Pashigian (1962): "Firm Size and Rate of Growth," *Journal of Political Economy*, 70, 556.

Imbs, J. (2007): "Growth and volatility," *Journal of Monetary Economics*, 54, 1848–1862.

Lotti, F., E. Santarelli, and M. Vivarelli (2003): "Does Gibrat's Law hold among young, small firms?" *Journal of Evolutionary Economics*, 13, 213–235.

Romer, C. D. and D. H. Romer (2004): "A New Measure of Monetary Shocks: Derivation and Implications," *The American Economic Review*, 94, 1055–1084.

Stanley, M. H., L. A. Amaral, S. V. Buldyrev, S. Havlin, H. Leschhorn, P. Maass, M. A. Salinger, and H. E. Stanley (1996): "Scaling behaviour in the growth of companies," *Nature*, 379, 804–806.

Sutton, J. (1997): "Gibrat's legacy," *Journal of economic literature*, 40–59.

Appendixes

A Data

In this Appendix we provide details on the construction of the database used in this paper.

A.1 Synchronization

In the Compustat data base Net Sales are defined based on the fiscal year. Conversely, FRED series are reported in terms of calendar year. In this paper we decide to express all the variables according to fiscal years. So we apply to the series for the real Gross Domestic Product (GDP), the real Final Sales of Domestic Product and the GDP Implicit deflator the following filter

$$\mathrm{X}_t = \begin{cases} \mathrm{X}_{4,t-1} + \mathrm{X}_{1,t} + \mathrm{X}_{2,t} + \mathrm{X}_{3,t} & \text{if } t \geq 1976 \\ \mathrm{X}_{3,t-1} + \mathrm{X}_{4,t-1} + \mathrm{X}_{1,t} + \mathrm{X}_{2,t} & \text{if } t < 1976 \end{cases},$$

where $\mathrm{X}_{q,t}$ represents the value of the variable X in the quarter $q \in (1, \ldots, 4)$ of year t. A fiscal year in the US goes from July 1st to June 30th until 1975 and from October the 1st to September the 30th from 1976 onward.[28]

A.2 Extreme growth events

A firm's extreme (positive or negative) growth episode can represent either a "normal" event, like a demand shock or the granting of an important patent, or a "special" event, as those associated with mergers, acquisitions or other operation suddenly changing the structure of the firm. The former type of event is usually associated with the internal or organic growth of the firm, while the latter is considered somehow outside the explanatory domain of the regular dynamics of a business firm. To obtain some hints on the nature and frequency of these events in the Compustat database we perform an investigation using firms' annual reports (Form 10-K) obtained from the US Securities and Exchange Commission. We proceed as follows: we focus on a sample of recent years for which Forms 10-K are available and we select those firms reporting the highest and the lowest growth rates in that year. The selected firms are reported in Table 4 and here below we synthetically report a description of the reasons behind the observed extreme growth rates.

[28] See the Congressional Budget and Impoundment Control Act in 1974.

Table 4: EXTREME GROWTH - CASE STUDY

	Maximum Growth Rate				Minimum Growth Rate			
Year	Name	$S_{i,t}$	$S_{i,t+1}$	$g_{i,t}$	Name	$S_{i,t}$	$S_{i,t+1}$	$g_{i,t}$
2013	Quest Solution, Inc.(1)	0.0040	37.3100	9.1122	Neah P.S. Inc.(2)	0.1490	0.0020	-4.3531
2012	MediaShift, Inc.(3)	0.0130	6.9530	6.2661	Revett, Inc.(4)	59.2110	0.0730	-6.7160
2009	Dendreon Corp.(5)	0.1010	48.0570	6.1563	CollabRx, Inc.(6)	12.3980	0.0160	-6.6599
2006	Insite Vision Inc.(7)	0.0020	23.7610	9.3670	Odimo Inc.(8)	18.9840	0.0140	-7.2370

1. $S_{i,t+1}$ and $S_{i,t}$ represent nominal sale (millions of US$) while $g_{i,t}$ is the logarithmic difference growth rate for firm i at time t.

1. **Quest Solution, Inc.** (formerly Amerigo Energy, Inc.) The increase in revenue in 2014 is driven by the acquisitions of Quest Solution and Bar Code Solutions Inc. that was completed during this year.

2. **Neah Power System, Inc.** In 2013 this firm has signed a large multi-year development contract. Nothing similar happens in 2014.

3. **MediaShift, Inc.** This company is engaged in digital advertising technology services. In 2013, its subsidiary Travora Networks (owned at 100%) signed an asset purchase agreement with Travora Media to acquire Travora's digital advertising network business.

4. **Revett Mining Company, Inc.** In December 2012 this firm suspended operations in a silver and copper mine located in north-west Montana due to unstable ground conditions in large portions of the mine.

5. **Dendreon Corporation**. This firm is a typical biotechnology company focused on the discovery, development and commercialization of new drugs. On April 29, 2010, the U.S. Food and Drug Administration licensed *PROVENGE* (their first autologous cellular immunotherapy) and commercial sale began in May 2010.

6. **CollabRx, Inc.** This firm, now known as Rennova Health Inc., offers diagnostics and software solutions to health care providers. Starting in the third quarter of the fiscal year 2009 it experienced a sharp decline in revenues resulting from the collapse of the semiconductor capital equipment market and the global financial crisis.

7. **Insite Vision Inc.**. This firm develops ophthalmic products and it had total revenues of 23.8 million of $ in 2007. 22.1 million of these revenues represented the amortization of the license fee and payments for *AzaSite* (one sustained-delivery-technology azithromycin) received in February and April 2007, respectively.

8. **Odimo Inc.**. This firm is an online retailer of diamonds, jewelry, watches and other luxury goods. They ceased operations as an online retailer starting from December 31, 2006. Since then it does not record any sales other than commissions based on a percentage of gross

sales made to visitors to their homepage who were, then, redirected to websites owned and operated by others.

Without pretending too much about the representativeness and exhaustiveness of this list, it is apparent that there are different reasons behind these extreme growth episodes. They include, but are not limited to, the acquisition of new firms and purchase of new assets (case 1 and 3), accounting issues (case 2), true exogenous shocks (case 4 and 6), very specific cash flows for patents and licenses (case 5 and 7), and inactivity (case 8). All these event appear as genuine, albeit extreme, economic occurrences that might shape the history of a business firm and we see no reason to exclude them from the analysis, also considering that the minority of them can actually be considered due to "external" reasons.

A.3 Attrition bias

Including also exiting and entering firms, an alternative definition of G_t^{COMP} could be

$$AG_t^{\mathrm{COMP}} = \frac{\sum_i (S_{i,t+1}) + \sum_k (S_{k,t+1})}{\sum_i (S_{i,t}) + \sum_j (S_{j,t})} - 1 \; , \tag{11}$$

where $S_{i,t+1}$ and $S_{i,t}$ denote the size of firms with valid sales in both $t+1$ and t while $S_{j,t}$ and $S_{k,t+1}$ denote the size of firms observed only in t and only in $t+1$ respectively. Figure 8 reports the time evolution of the two measures. They do not diverge substantially.

A.4 Macroeconomic variables

Following Gabaix (2011) we consider 4 different macroeconomic variables as controls in our regressions: oil shocks, monetary shocks, interest rates and the term spread.

As far as the oil shock is concerned, we extend to 2013 the series built in Hamilton (2003). We start by the Monthly Producer Price Index-Crude petroleum (**WPU0561**, end of period) available at FRED then the quarterly oil shock is defined as the (log) amount by which the current oil price exceeds the maximum value over the past 4 quarters (quarterly oil shock is set to zero if it is negative). A yearly oil shock is then defined as the sum of 4 quarterly oil shocks within that year.

The monthly monetary shock comes from David Romer's web page and ranges from 1969 to 1996 (series **RESID**, Romer and Romer (2004)). Then following Gabaix (2011) the yearly shock is built as the sum of the 12 monthly shocks in that year. For the years not covered by the data, the value of the shock is assigned to be 0, the mean of the RESID series.[29]

The yearly interest rate is constructed by averaging, over one year, monthly observations of the 3-month nominal T-bill secondary market rate (TB3MS). Finally, the term spread is defined as the 5-year treasury constant maturity rate (GS5) minus the 3-month treasury bill secondary market rate (TB3MS). Further details are available upon request.

[29] Gabaix (2011) argues that this assignment does not bias the regression coefficient under simple conditions, for instance if the data are i.i.d. However it does lower the R^2 by the fraction of missed variance; fortunately, most large monetary shocks (e.g., of the 1970's and 1980's) are in the data set.

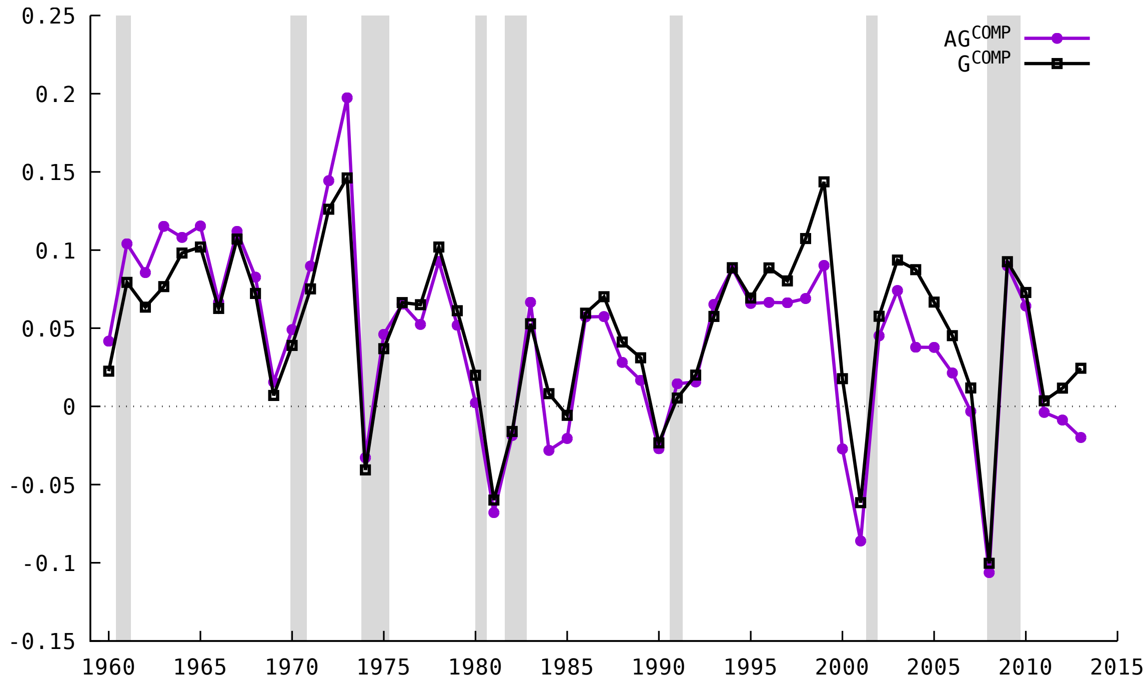

Figure 8: G_t^{COMP} and AG_t^{COMP} defined according to equation (2) and (11) respectively. Shaded areas represent recessions according to the NBER business cycle dates.

B Mode estimators

In this section, we briefly describe the two methods used to estimate m_t. The non-parametric and parametric estimation methods are based on Half-sample method (HSM) and the Asymmetric Exponential Power distribution (AEP) respectively.

B.1 Half-sample method (HSM)

Bickel and Frühwirth (2006) present the half-sample estimator of the mode for the distribution of a continuous random variable. [30]. Let $(x_i)_{i=1}^n$ be an ordered vector of n random numbers drawn from an unimodal distribution with mode M. Assume the sample size is an integer power of two ($n = 2^m$).[31] Then:

1. find the smallest interval that contains $n/2$ points from the sample. In other words, obtain the integer j_1 for which $x_{j_1+n/2-1} - x_{j_1}$ reaches a minimum, with $1 \leq j_1 \leq n/2 + 1$;

2. using only the data in that interval, find the smallest interval that contains $n/4$ points, i.e., obtain the integer j_2 for which $x_{j_2+n/4-1} - x_{j_2}$ reaches a minimum, with $j_1 \leq j_2 \leq n/4 + 1$;

[30] This procedure is implemented in the R package 'modeest' `https://cran.r-project.org/web/packages/modeest/modeest.pdf`.

[31] They also generalize this algorithm to allow n to be any positive integer. See the details in the appendix of their paper.

3. iterate this procedure until obtaining an interval with only two points, x_{j_m} and x_{j_m+1};

4. the estimated mode is the mean of these two values, $\widehat{M} = (x_{j_m} + x_{j_m+1})/2$.

The proponents show that this estimator performs relatively well under a wide range of conditions and, in particular, when the distribution is asymmetric with a large number of extreme observations, which is the case in our empirical investigation. The HSM estimator possesses another interesting property. Differently from methods based on density estimation, it does not require the arbitrary selection of a "bandwidth" , which is often problematic. Indeed if the bandwith is too large, then the mode cannot be precisely located, leading to a high bias. If, on the contrary, the bandwith is too small, then it will be likely that the interval with the highest empirical frequency does not contain the mode, leading to a high variance of the estimator. HSM avoids these issues by beginning with a large interval and progressively reducing its width.

B.2 The Asymmetric Exponential Power distribution (AEP)

The AEP distribution introduce by Bottazzi and Secchi (2011) can be used to obtain an alternative parametric estimate of the mode of the distribution. The AEP is a five-parameters family of distributions with probability density

$$f_{AEP}(x) = \frac{1}{C} e^{-\left(\frac{1}{b_l}\left|\frac{x-m}{a_l}\right|^{b_l}\theta(m-x)+\frac{1}{b_r}\left|\frac{x-m}{a_r}\right|^{b_r}\theta(x-m)\right)}$$

with $C = a_l A_0(b_l) + a_r A_0(b_r)$, $A_k(x) = x^{\frac{k+1}{x}-1}\Gamma(\frac{k+1}{x})$ and where m is the mode, b_l and b_r the shape parameter of the lower and upper tails respectively and a_l and a_r two width parameters associated with the probability mass below and above m. The mean of the AEP density is

$$\mu_{AEP} = m + \frac{1}{C}\left(a_r^2 A_1(b_r) - a_l^2 A_1(b_l)\right) \ .$$

Thus according to definition (8) the parametric distributional tilt is given by

$$\text{tilt}_{AEP} = \frac{1}{C}\left(a_r^2 A_1(b_r) - a_l^2 A_1(b_l)\right) \ .$$

The AEP parameters are estimated using Maximum Likelihood implemented in the package "Subbotools"[32].

[32] The latest version of *Subbotools* can be found at `http://cafim.sssup.it/~giulio/software/subbotools/`. And Bottazzi (2004) provides the detailed manual.