

Aggregate fluctuations and the distribution of firm growth rates*

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

[†]Institute of Economics and Department EMbeDS, Scuola Superiore Sant'Anna

[‡]Chuo University

[°]PSE, Université Paris 1 Panthéon-Sorbonne

First Draft - October 19, 2015

Present Version - March 14, 2019

Abstract

We propose an aggregate growth index that explicitly accounts for fat tails in the firm size distribution and for the negative scaling relation between the size of the firm and the volatility of its growth rates. Using Compustat data on US publicly traded company, we show that the new index tracks aggregate fluctuations much better than simpler measures of central tendency of the dynamics of firms, like the growth rates 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 log 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). Le Li is a JSPS International Research Fellow. The present version of the paper benefited from the comments of two anonymous referees. We are grateful for their insightful suggestions.

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: the dispersion of the rates of change of productivity, employment, prices and business forecasts was found 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. 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) we try to improve our understanding of the economic dynamics observed in the aggregate. With this aim, in Section 3, we develop a theoretical central tendency 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 scaling adjusted index H^2 , while remaining simple to compute, tracks the observed aggregate growth much better than the sample average of firm growth rates. In particular, our analysis shows that the presence of

¹The 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.

²Indeed, 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.

volatility scaling depresses the aggregate growth rate and pushes it toward the average of firm log-growth rates. With respect to the existing literature, this results 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 scaling adjusted index H^2 and aggregate growth is far from being perfect. In fact, we observed that there exists an important economic phenomena not captured by fat tails and volatility scaling. To see it, 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*. This is the reason why one often observes 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.

Starting from these considerations, we decompose the scaling adjusted 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 the loss of 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), with a likely different informational content.

In Section 4 we explore the validity of the mode-tilt decomposition of the index H^2 with re-

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

⁴Trimming 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.

gression 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 decomposition of the scaling adjusted central tendency index into the typical growth rate and the distributional tilt, improves our capability of tracking aggregate fluctuations and open new possibilities to better understand the micro-macro linkages.

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 Sales⁶ 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 growth rates distribution.⁸ These procedures are generally adopted to avoid mixing of organic growth and external growth generated by, for instance, mergers and acquisitions. 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.

⁵Standard & 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.

⁶Net 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.

⁷This 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).

⁸Higson 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%.

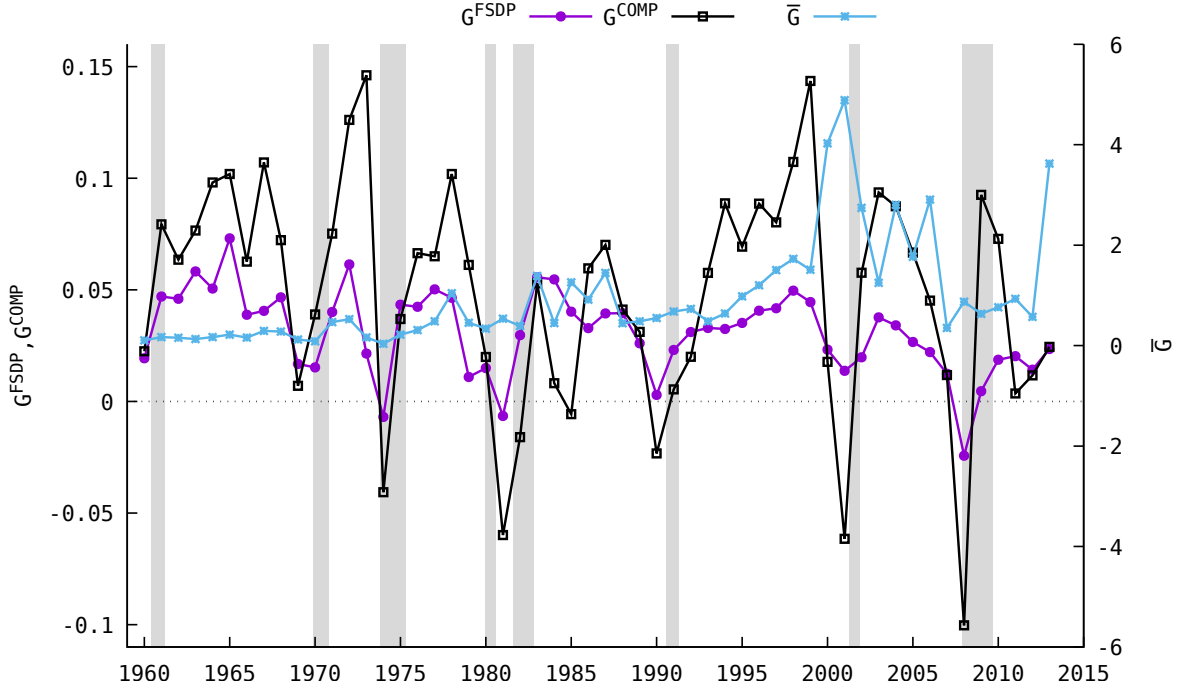


Figure 1: Time evolution over the period 1960-2013 of the net growth rate of the real FSDP G_t^{FSDP} (filled circles), of the Compustat aggregate G_t^{COMP} computed according to (2) (empty squares) and of the average net growth rate (asterisks) of Compustat companies \bar{G}_t computed according to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

3 Aggregate growth rate and firms dynamics

Obviously, if all the companies of the 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, together with the average net firm growth rate of our sample defined as

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

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 American companies

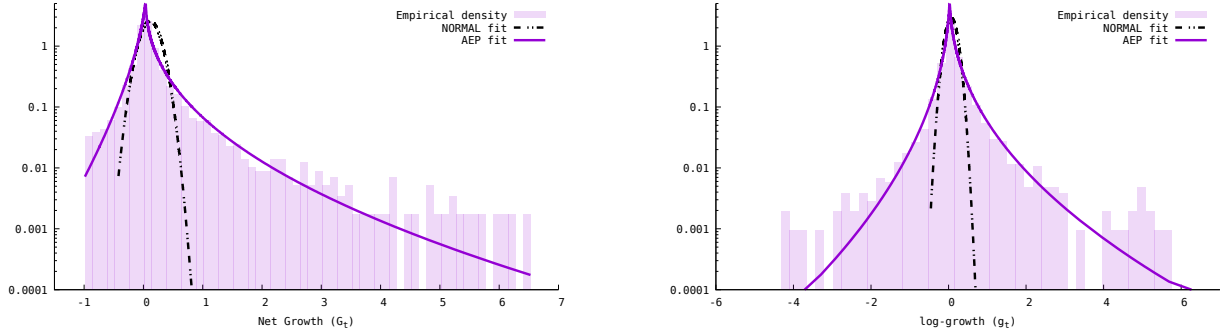


Figure 2: Total Sales net growth rates (G_t) distribution in 2013 together with a Gaussian and Asymmetric Exponential Power (AEP) fit (left panel). Total Sales log-growth rates (g_t) distribution in 2013 together with a Gaussian and Asymmetric Exponential Power (AEP) fit (right panel).

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

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

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 are positively correlated (the Spearman rank correlation about 0.58). In fact, \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 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 skewness implies that the net firm growth rates 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

⁹Our 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.

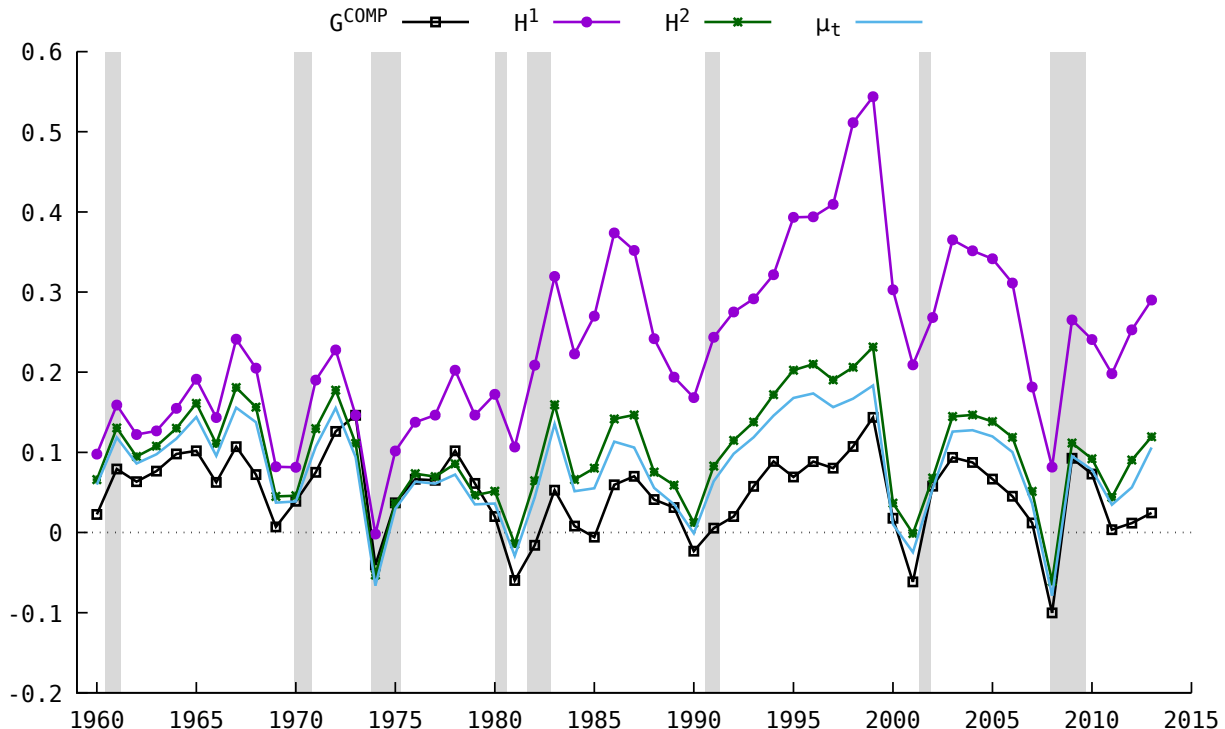


Figure 3: Time evolution over the period 1960-2013 of the net growth rate of Compustat aggregate G_t^{COMP} (black line with empty squares) and of the two approximations, H^1 and H^2 , computed according to (5) (dark-violet line with filled circles) and (7) (dark-green line with asterisks) respectively. 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.

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 a persistently smoother and more symmetric behavior.¹⁰ Indeed, the 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, firms grow following a multiplicative process $S_{i,t+1} = \epsilon_{i,t} S_{i,t}$ where $\epsilon_{i,t}$ is a random variable shocking a firm's initial size $S_{i,t}$.¹¹ In order to exploit the multiplicative nature of the firm growth process and the observed smoother shape of the log growth rate distribution one can rewrite the Compustat aggregate net growth rate G_t^{COMP} in terms of firm log growth rates $g_{i,t}$ as

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

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

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

where $C_n[g_{i,t}]$ represents the n -th cumulant of the distribution of $g_{i,t}$. Even if more symmetric than that of the net growth rate $G_{i,t}$, the distribution of the firm log growth rate $g_{i,t}$ possesses a significant level of variance that should be accounted for. To this aim, if we assume that the dynamics of firms follows the so called Gibrat's Law,¹² which postulates that growth shocks $g_{i,t}$ are independent of size and identically distributed, and if we further assume that the distribution of these growth shocks is Gaussian, we can truncate equation (4) at the second order¹³ and obtain an approximation H^1 as

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

where μ_t and σ_t represent 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^{COMP} and of H_t^1 are both reported in Figure 3. The variables are similar in magnitude, even if seemingly diverging. Their Spearman rank correlation is 0.32, suggesting a moderate correlation. If the assumptions of normality and independence of the firm growth shocks were valid, 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 .¹⁴ In fact, both these assumptions do

¹⁰Note 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.

¹¹An alternative, additive, model would be $S_{i,t+1} = S_{i,t} + \epsilon_{i,t}$. This model would predict 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.

¹²Sutton (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).

¹³In this case, indeed, $C_n[g_{i,t}] = 0$ for $n > 2$. We remind the reader that the second cumulant is actually the variance of the distribution.

¹⁴The 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). In Appendix B we briefly discuss the goodness of our estimators using Monte Carlo experiments.

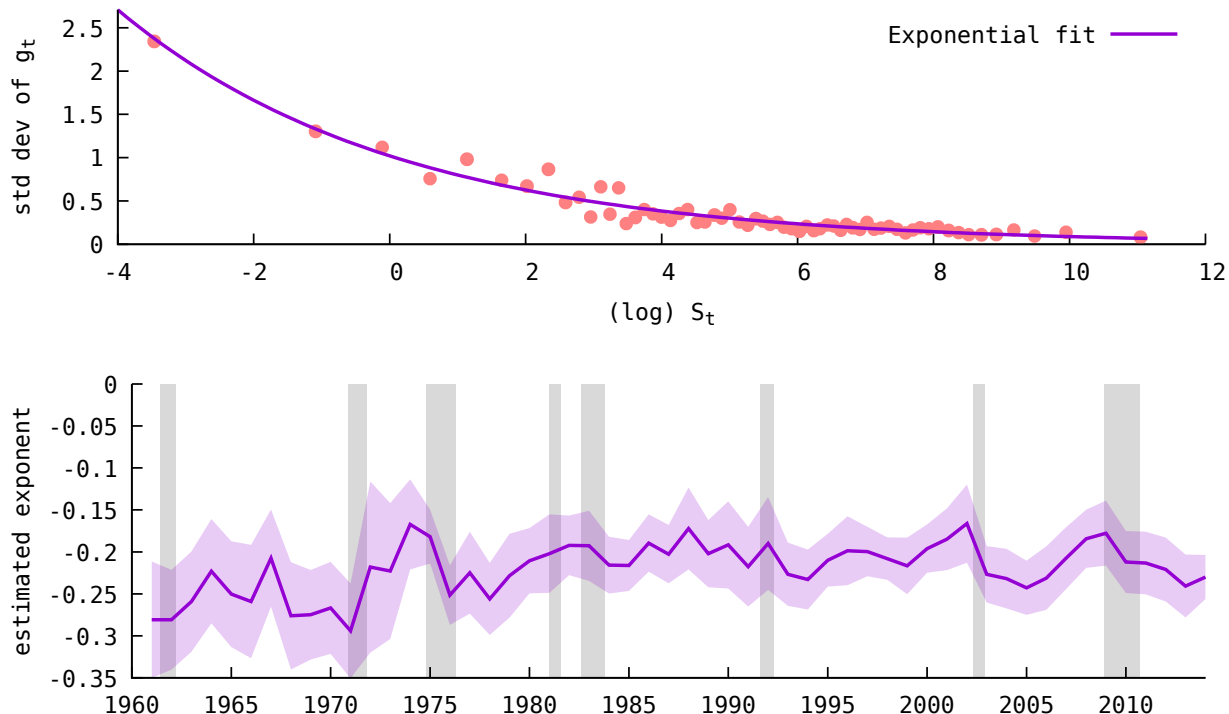


Figure 4: Standard deviation of growth rates σ_t as a function of initial (log) net sales $S_{i,t}$ together with an exponential fit in semi-log scale (**top panel**). Estimated exponent is $-0.23(0.01)$. Time evolution of the estimated exponent together with a 99% confidence band (**bottom panel**)

not have any empirical support.

First, as it is apparent from Figure 2 (right panel) and in line with Stanley et al. (1996) and Bottazzi and Secchi (2006a), the empirical density of $g_{i,t}$ presents tails substantially fatter than those of a Gaussian distribution also within the Compustat database.¹⁵ The existence of these fat tails is largely due to the fact that 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.

The dependence between volatility of growth and firm size is illustrated in Figure 4, where we report the standard deviation of growth rates in 2013 as a function of firm (log) size s_i .¹⁶ It is clear that the former declines with the 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;

¹⁵Results 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.

¹⁶We rank firms according to their size in a specific year, then we split them in equally populated 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.

Criscuolo et al., 2016).¹⁷ The bottom panel in the same figure reports the estimate of the scaling exponent for the time window 1960-2013. With a possible minor regime break in the late 70s, the exponent is characterized by a remarkable stability which confirms that the scaling property is not a peculiar feature of any particular year but it is rather persistent in time.

This evidence implies that in order to better proxy $E[G_t^{\text{COMP}}]$, we should include 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 $E[G_t^{\text{COMP}}]$. First, given the fat-tailed nature of the growth rate distribution, we need to retain all the cumulants because, 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.¹⁸ With these assumptions, the expression in (4) can be rewritten as a scaling adjusted central tendency index

$$\begin{aligned} H_t^2 \equiv E[G_t^{\text{COMP}}] &= \frac{\sum_i S_{i,t} e^{((\bar{S}/S_{i,t})^{2\beta_t} \bar{C}_{2,t} - C_{2,t})/2}}{\sum_i S_{i,t}} e^{C_{1,t} + C_{2,t}/2 + \sum_{n=3}^{\infty} C_{n,t}/n!} - 1 \\ &= \frac{\sum_i S_{i,t} e^{((\bar{S}/S_{i,t})^{2\beta_t} \bar{C}_{2,t} - C_{2,t})/2}}{\sum_i S_{i,t}} E[e^g] - 1, \end{aligned} \quad (6)$$

where we have factorized the size-dependent variance term and where $C_{2,t}$ stands for the (unconditional) variance of the log-growth rate distribution, computed at time t using the whole sample of firms. 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 of firms in equally populated size classes and we compute the log growth rate variance $\bar{C}_{2,t}$ of firms belonging to the size class including \bar{S} .¹⁹ Finally, we compute H^2 using the entire firm size distribution, summing across all firms, each weighed with its observed size $S_{i,t}$. Notice that the smallest firms, which are weighted less for their reduced size, have in fact an enhanced effect on H^2 due to their higher variance of growth implied by the scaling relation. The quantity $E[e^g]$ does in principle depends on all higher order cumulants of the firm growth rates distribution. However, once the volatility scaling is properly taken into account, the distribution of growth rates has definitely thinner tails and one can approximate $E[e^g]$ with $H^1 + 1$ (Bottazzi and Secchi, 2003a).²⁰

¹⁷Note that an exponential decay for the (log) firm size implies a power law decay for firm size.

¹⁸The 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.

¹⁹The 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 variance.

²⁰Higher moments will be partly analyzed in the next Section when we discuss the contribution of the skewness of the firm-level growth rates distribution to the aggregate dynamics.

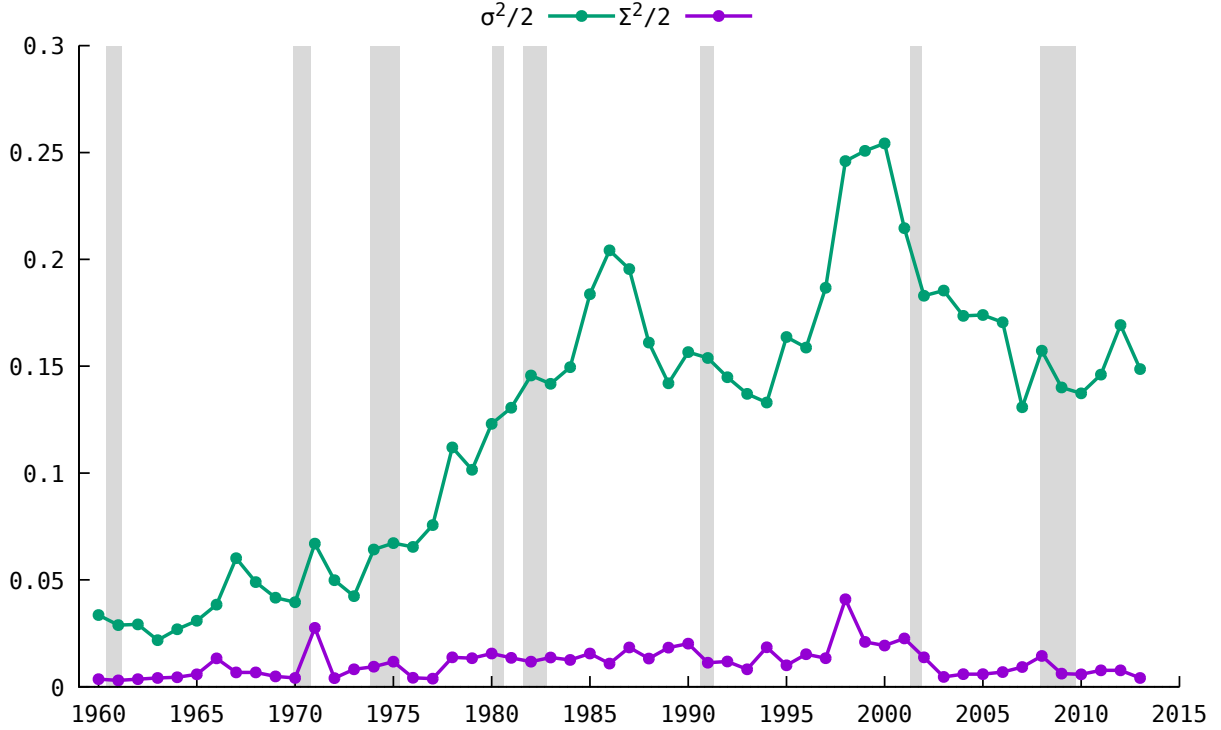


Figure 5: Time evolution over the period 1960-2013 of $\sigma_t^2/2$ and $\Sigma_t^2/2$. Shaded areas represent recessions according to the NBER business cycle dates.

The performance of H_t^2 in tracking the observed aggregate growth rate G_t^{COMP} can be judged once again from Figure 3. Two 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 Spearman 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 size and its relation with their growth rate becomes important.

To better remark the effect of the scaling of the variance, we reorganize the terms in (6) and truncate the distribution to the second order cumulant. We get

$$H_t^2 \sim e^{\mu t + \Sigma_t^2/2} - 1, \quad (7)$$

where

$$\Sigma^2 = 2 \log \left(\frac{\sum_i e^{s_{0,i} + \bar{\sigma}_g^2 e^{2\beta(s_i - \bar{s})/2}}}{\sum e^{s_{0,i}}} \right) \quad (8)$$

is an “effective” variance that takes into consideration the scaling of the growth rates variance with size. If the latter were not present, that is if $\beta_t = 0$, then it would be $\Sigma_t^2 = \sigma_t^2$ and one would get

back to the previous approximation, H_t^1 . Figure 5 reports the quantities $\sigma_t^2/2$ and $\Sigma_t^2/2$ for the different years in the database. These quantities represent the corrections to the average log-growth rate μ_t in H^1 and H^2 , respectively. As can be seen, for each year in the database, the effective variance is much lower than the variance of firm-level growth rates. This is the joint result of the scaling of the variance and of the presence of firms of hugely different sizes, that makes the impact of the scaling extremely relevant. Moreover, while the variance of growth at firm level increases significantly in the last years of the database, the effective variance remains rather constant. This is because, while in recent times firms are on average more volatile, the increase of the support of the size distribution, in the presence of a constant negative value of β , reduces the contribution of the firm-level volatility on the aggregate dynamics. In Appendix B we present a simulation exercise that clarify the effect of the scaling and of the shape of the size distribution on the performance of the indexes H^1 and H^2 .

Finally, note that the joint effect of the scaling of the variance of firms level growth-rates and a large support for the firm size distribution makes the effective variance Σ^2 so small that its correction with respect to μ_t in H_t^2 becomes negligible (cfr. Figure 5 and 3). 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 firm-level 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 , the variance of the reference size class $\sigma^2(\bar{S})$ and the firm size distribution.

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 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 again the right panel of Figure 2 which reports the probability density of

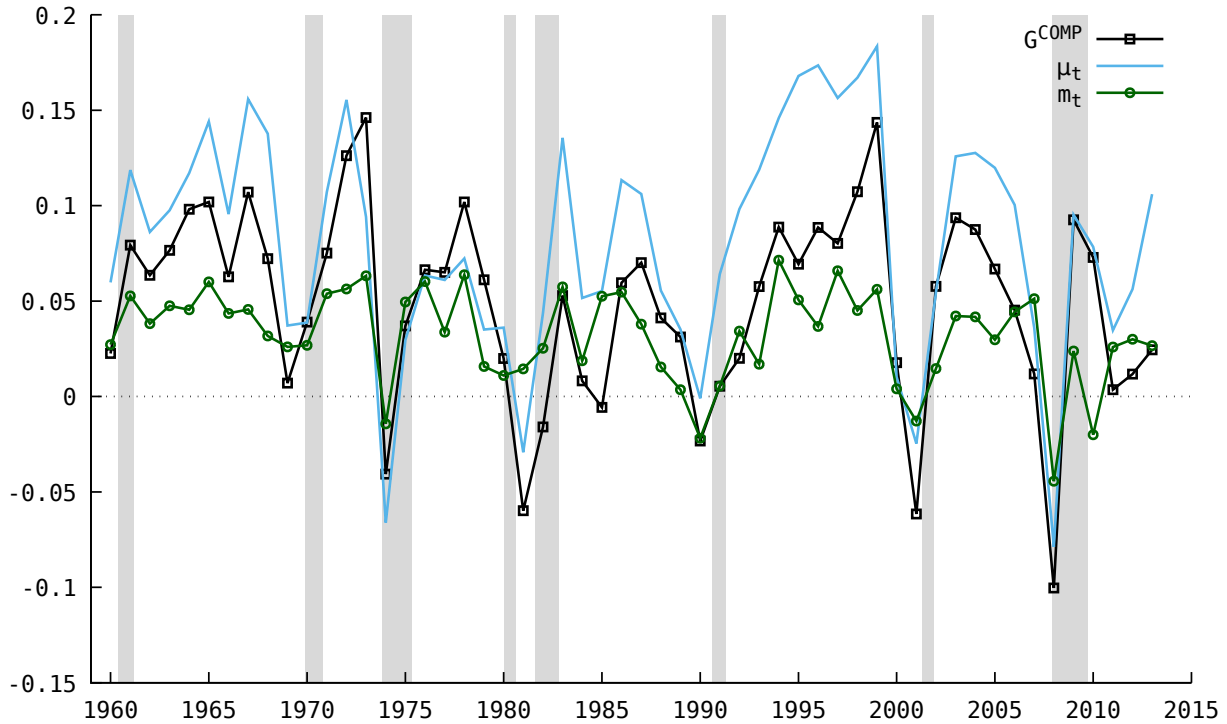


Figure 6: Time evolution over the period 1960-2013 of the aggregate Compustat growth rate G_t^{COMP} (black line with empty squares) and of the average μ_t (cyan line with no point) and of the modal m_t (dark-green line with empty circles) firms growth rate defined as the (log) difference of their Total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

the log growth rates of COMPUSTAT companies in 2013. Contrary to what one might conclude from 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 track 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 our understanding of the G_t^{COMP} dynamics by isolating in the index H_t^2 some measure of asymmetry. Given the strong similarity between H_t^2 and

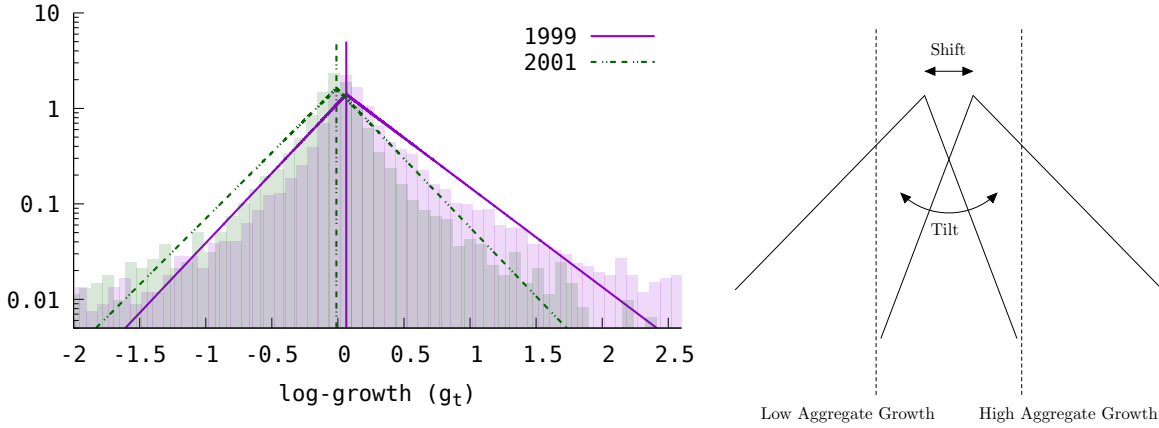


Figure 7: **Left:** Firm level growth rate distribution during an expansion (1999) and a contraction (2001) year. **Right:** The observed difference between the two years is a modest shift in the modal value but a relevant change in the tail behavior.

μ_t it seems natural to decompose the index in the sum of two component

$$H_t^2 \sim m_t + p_t, \tag{9}$$

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 (9), 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.

To this aim Figure 7 reports the firm (log) growth rate distribution in 1999 and 2001 highlighting the peculiar tent shape robustly observed in the literature. These two years have been chosen as they represent two highly different aggregate growth regimes: in 1999 the average (log) growth rate is about 0.18 while in 2001 it is -0.02 . On the contrary, the modes of the firm growth rates distributions, representing the most common, or typical, growth rate in each year, are closer to each other, being 0.06 in 1999 and 0.01 in 2001. Combining these two observation one can conclude that in moving from 1999 to 2001 the observed growth rate distribution has experienced a minor mode-shift leftward together with an important movement of firms from above to below the mode,

generating the observed change in the mean growth rate.

This distributional dynamics is schematically illustrated in the right panel of Figure 7. The observed separation between the mean and the mode in 1999 (high average growth) and 2001 (low average growth) is not consistent with a scenario where firms, while influenced by idiosyncratic factors, react homogeneously to the macroeconomic shocks hitting the economy. In this case, for the law of large numbers, we should have observed a simple shift of the distribution, with the mean and the mode moving together and ultimately sharing the same relation with the aggregate growth dynamics. Conversely, we observe a richer dynamics in which idiosyncratic individual shocks coexist with possible heterogeneous responses of individual companies to the aggregate shock: if some groups of firms over-react while other firms under-react, together with a possible common shift, we are likely to observe a change in the distribution of probability mass around the modal value breaking the symmetry and separating the average and the modal growth rates. The separate analysis of the dynamics of mean and mode provides a more flexible framework, as the shift and tilt movements of the micro distribution can be, in general, 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 richer description 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 (7) and the further decomposition in (9) 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

²¹The 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.

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^{COMP} and our statistics based on the micro log-growth rate distribution of Compustat firms we consider the following specification

$$G_t^{\text{COMP}} = \alpha + \sum_{\tau=0}^T \beta_{t-\tau}^m m_{t-\tau} + \sum_{\tau=0}^T \beta_{t-\tau}^p p_{t-\tau} + \epsilon_t \quad , \quad (10)$$

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) developed in Bickel and Frühwirth (2006) and briefly described in Appendix C.1. This specification is then compared with a model inspired by Holly et al. (2013)

$$G_t^{\text{COMP}} = \alpha + \sum_{\tau=0}^T \beta_{t-\tau}^m \mu_{t-\tau} + \sum_{\tau=0}^T \beta_{t-\tau}^p \gamma_{t-\tau} + \epsilon_t \quad , \quad (11)$$

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 quality of the fit of this simple model is good with an adjusted R-squared \bar{R}^2 of about 72%. This explanatory power is almost evenly distributed between m_t and p_t : as reported in column (1)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 (11). 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 residual correlation emerges between G_t^{COMP} and γ_t , see column 2(a). 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

²²One 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.

²³This 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.

²⁴This is a consideration which might well have implications going beyond the exercise presented here. For instance

Table 1: EXPLANATION AND PREDICTION - Compustat AGGREGATE

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

Next we turn our attention to the more general case by estimating (10) 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 than it. This double relation between G_t^{COMP} and the distribution of micro growth rates cannot be devised through the regression against 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 (10) and (11) 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 skewness is lower, with an \bar{R}^2 of 21%. As expected from the previous analysis, the skewness index has a relatively minor predictive power. None of these models is improved by including the lagged variable G_{t-1}^{COMP} which results in both cases non statistically significant.

Table 2: EXPLANATION AND PREDICTION - ROBUSTNESS CHECKS

	LAD Regression (1)	25% Trimming (2)	Parametric Mode and Tilt (3)	Parametric Mode and Tilt (4)	Parametric Mode and Tilt (5)	Parametric Mode and Tilt (6)	Including macro controls (7)	Including macro controls (8)
m_t	1.042*** (0.280)	1.201*** (0.060)	1.313*** (0.173)	1.266*** (0.184)			0.907*** (0.161)	
m_{t-1}					0.509* (0.295)	0.447 (0.317)		0.126 (0.368)
m_{t-2}					-0.057 (0.278)	-0.011 (0.294)		0.471* (0.261)
p_t	0.602*** (0.128)	1.072*** (0.104)	0.535*** (0.072)				0.624*** (0.105)	
p_{t-1}					0.427** (0.189)			0.464** (0.212)
p_{t-2}					-0.313* (0.160)			-0.422** (0.207)
$(a_r - a_t)_t$				0.794*** (0.121)		0.665** (0.300)		
$(a_r - a_t)_{t-1}$						-0.491* (0.270)		
$(a_r - a_t)_{t-2}$								
\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 (9) of the FGRD over the time span 1960-2013. $(a_r - a_t)_t$ 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.

Robustness checks

To check the robustness of our regression results we perform 3 sets of tests whose results are reported in Table 2.

A first possible concern has to do with our choice of not removing extreme growth events from the database. According to our analysis, see Appendix A.2, the largest variations in individual firm size observed in our database, represent legitimate events of the life of a business firm. However, the regression technique we use, the OLS, is known to be sensitive to outliers. To check the role played by extreme observations, we re-estimate our benchmark model in two ways: using Least Absolute Deviation (LAD) regression, a method which is more robust to the presence of extreme observations, and using OLS but trimming our individual firm growth rates at 25%. Using the LAD regression, instead of the OLS, does not seem to have any impact on our results, the estimated coefficients are unaffected.²⁵ Conversely, trimming the database has a significant impact on point estimates but does not change 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 . The value of these statistics might be very sensitive to the selected trimming threshold, a fact that confirms our skepticism in sample cleaning of this sort.

In our second test, we replace the non-parametric estimate of the mode of the distribution, based on the Half Sample Method (HSM) discussed in (Bickel and Frühwirth, 2006), with a parametric method based on the Asymmetric Power Exponential distribution (details are reported in Appendix C.2). We then re-estimate equation (10) using these parametric estimates with OLS, reporting the results in Columns (3)-(6) of Table 2. The point estimates are remarkably similar and all our conclusions remain the same.

A final possible concern regards the lack of some of the most common explanatory factors and predictors of aggregate growth in our regression models. 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, short term interest rate and 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)). Moreover, a regression model containing the additional explanatory factors alone is rather weak, with an \bar{R}^2 of 0.12.

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*** (0.081)				0.425*** (0.068)				63%				65%			
m_{t-1}		0.064 (0.123)		0.026 (0.097)		0.158 (0.109)		0.112 (0.084)		7%		2%	27%			7%
m_{t-2}		0.186* (0.099)		0.197** (0.089)		0.134 (0.090)		0.134 (0.080)		10%		6%	7%			4%
p_t	0.171*** (0.043)				0.137*** (0.036)				37%				35%			
p_{t-1}		0.181** (0.087)		0.241*** (0.075)		0.150* (0.077)		0.191*** (0.063)		31%		12%	33%			13%
p_{t-2}		-0.248*** (0.071)		-0.199** (0.075)		-0.184*** (0.056)		-0.123* (0.065)		52%		8%	33%			4%
μ_{t-1}			0.105* (0.057)				0.123** (0.051)				41%					62%
μ_{t-2}			-0.116* (0.066)				-0.091* (0.054)				36%					19%
γ_{t-1}			0.002 (0.002)				0.002 (0.002)				22%					18%
γ_{t-2}			0.000 (0.003)				0.000 (0.002)				1%					1%
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								

$m_t, p_t, \mu_t, \gamma_t$ represent the mode, the tilt as defined in equation (9), 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^{RFSDP} , 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.

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

With this aim we estimate equations (10) and (11) replacing the original dependent variable, G_t^{COMP} , with these macroeconomic statistics. Results are reported in Table 3. 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} as dependent variable (\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 the latter not being significant in practically any of our regressions.

Columns (2) and (3) report the result of the forecasting exercises, including only lagged values of the regressors. Three considerations seem relevant. Firstly, lagged values of m_t and p_t appear to have some 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.

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.

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

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

²⁷In 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. On the other hand, the correlation coefficient 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, so that the exercise remains meaningful.

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 growth rate. 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 the firm size distribution. These two distributional properties are embedded in a proposed synthetic index of central tendency, H^2 . Using H^2 , the tracking of aggregate fluctuations is improved by almost 100%, with a correlation between the index and the aggregate growth rate reaching 0.77. Probably more importantly, the derivation of H^2 highlights the role of the firm size distribution and the volatility scaling in shaping the observed aggregate fluctuations. A robust (negative) scaling jointly with the presence of firms of hugely different sizes have a depressing effect on the aggregate growth rate which becomes, consequently, similar to the average of the distribution of firm log growth rates.

We observe, however, that the agreement of our synthetic index with the aggregate growth rate is still far from being perfect. One important reason is the asymmetry of the distribution of firm growth rates, which H^2 does not take into account. The presence of this asymmetry has been largely ignored by the previous literature. To amend this shortcoming, we propose a parsimonious way to account for the observed asymmetry by decomposing the central tendency index using the mode of the firm log growth rates distribution and the distributional tilt, a residual term defined as the difference between the central tendency index and the mode. We show that a simple regression model based on mode and tilt 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, an approach previously adopted in the literature.

Remarkably the improved performance of our mode and tilt regressions is retained when we replace the aggregate growth rate of the Compustat sample with the GDP as the dependent variable. 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

$$X_t = \begin{cases} X_{4,t-1} + X_{1,t} + X_{2,t} + X_{3,t} & \text{if } t \geq 1976 \\ X_{3,t-1} + X_{4,t-1} + X_{1,t} + X_{2,t} & \text{if } t < 1976 \end{cases},$$

where $X_{q,t}$ represents the value of the variable X in the quarter $q \in (1, \dots, 4)$ 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.

²⁸See 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 .

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

$$AG_t^{\text{COMP}} = \frac{\sum_i(S_{i,t+1}) + \sum_k(S_{k,t+1})}{\sum_i(S_{i,t}) + \sum_j(S_{j,t})} - 1, \quad (12)$$

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.

²⁹Gabaix (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.

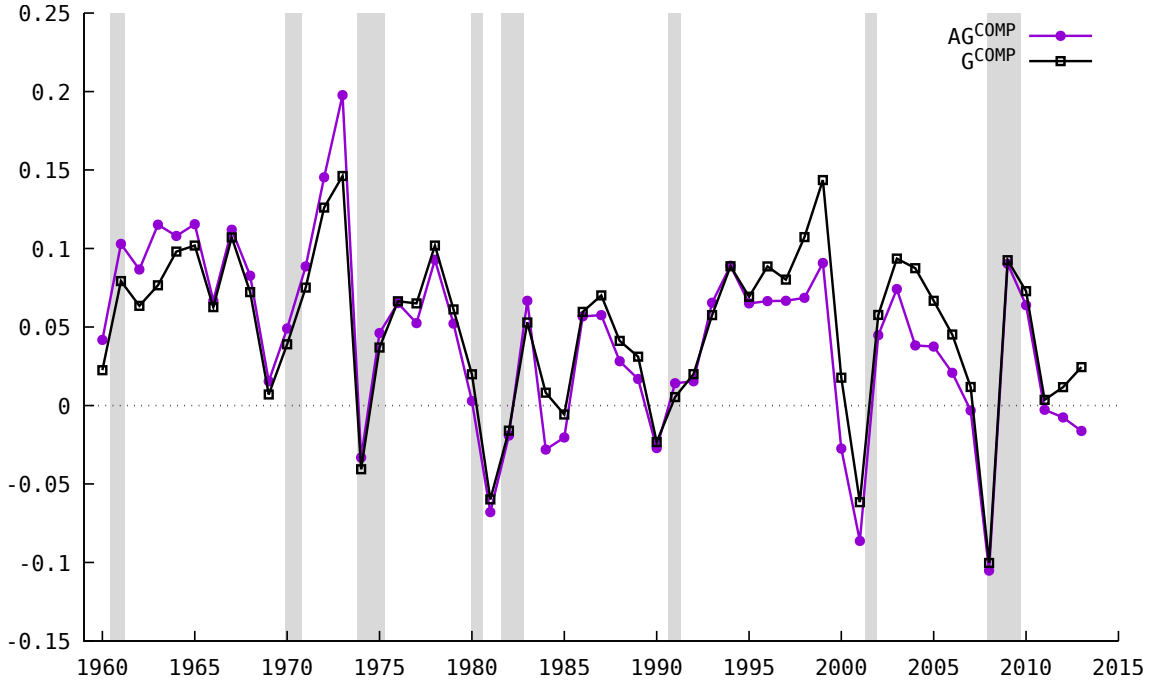


Figure 8: G_t^{COMP} and AG_t^{COMP} defined according to equation (2) and (12) respectively. Shaded areas represent recessions according to the NBER business cycle dates.

B Simulations

In this section we simulate the proportional random growth model used in Section 3 to check the performance of our indexes H^1 and H^2 in presence of the statistical fluctuations associated to finite samples. We also explore their behavior under different scenarios concerning the underlying firm size distribution.

Consider an economy with N firms and let $s_{i,0} = \log S_{i,0}$ be the initial (log) size of firm i assumed to be a random variable with mean μ_s and variance σ_s^2 . Let ϵ_i represent a random growth shock with 0 mean and variance σ^2 . Consistently with Section 3 we assume that the expected firm growth rate is independent of firm's size while its standard deviation scales exponentially with the (log) size with an exponent $\beta > 0$. Under these assumptions the proportional random growth model implies that the final (log) size reads

$$s_{i,1} = s_{i,0} + e^{-\beta(s_{i,0}-\mu_s)}\epsilon_{i,0} + \mu, \quad (13)$$

where μ is the expected log-growth rate. Notice that according to the specification in (13), the value σ represents the standard deviation of the log-growth rate of the firm with average size. Using $s_{i,0}$ and $s_{i,1}$ we can finally compute the indexes G , μ , H^1 and H^2 defined above and study their behavior. The timeline of the simulation is as follows:

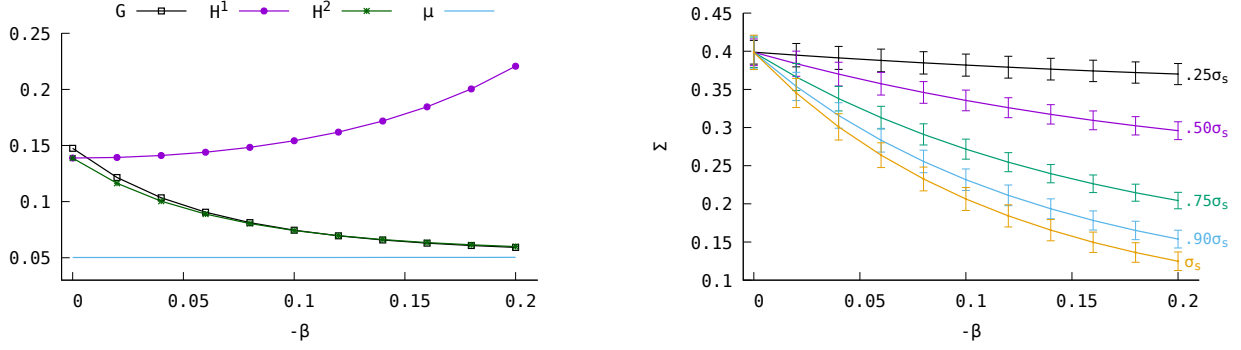


Figure 9: Expected values of the different indexes as a function of β (left panel). Effective standard deviation Σ as a function of β (right panel). Expected values and standard deviations are obtained from 100 independent realizations.

1. draw the initial size $s_{i,0}$ from the distribution f_s ;
2. draw growth shock $\epsilon_{i,0}$ from the distribution f_ϵ ;
3. compute the final size $s_{i,1}$ using equation (13);
4. compute G , μ , H^1 and H^2 .

Notice that the initial size of the firm $s_{i,0}$ is generated, at each iteration, in the simulation independently and from the same distribution. This implies that the size of the firm does not actually evolve. This is done in order to replicate the “unbalanced” nature of the empirical analysis and to avoid a diffusion behavior which is not observed in the data. To perform our exercise we set the number of firms N and the number of iterations T equal to 5000 and 100 respectively.³⁰ The rest of parameters are calibrated using the average over the last 20 years of the Compustat data set: $\mu_s = 5$, $\sigma_s = 2.8$, $\mu = 0.05$ and $\sigma = 0.4$. To run the simulation it remains to specify f_s and f_ϵ .

In line with the empirical evidence available we begin by assuming that f_s is Normal (that is firm size $S_{i,0}$ is Log-normally distributed) while f_ϵ is a fat tailed Laplace distribution and we explore the behavior of the model for $\beta \in [0, 0.2]$. Results of this simulation exercise are summarized in Figure 9. In the left-panel we report values of G , H^1 , H^2 and μ for different values of the scaling parameter β . When there is no scaling ($\beta = 0$) H^1 equals H^2 , as expected, and they both proxy well the net aggregate growth rate G . On the contrary μ looks like a poor proxy for G . Increasing β has two main effects. First, H^1 rapidly parts ways with G while H^2 moves even closer to G becoming almost identical to the aggregate net growth rate even for mild values of β . Second, both G and H^2 decrease when β increases with the indirect effect of making μ a better approximation of G when the scaling coefficient is close to -0.2. To sum up even in presence of finite samples the index H^2 provides a reliable estimate of the expected value of G , irrespective of the value of β . G and μ are similar only when it happens that the value of β is such that the expected log-growth rate μ is close to H^2 .

³⁰Expected values are computed over these 100 different realizations of the process.

To better understand the mechanism behind these results we move to the right panel of the same Figure 9. Here we report the behavior of the average Σ as a function of β and for different values of σ_s , the standard deviation of f_s . We start by looking at the lowest curve, that associated with the baseline value of σ_s , that is 2.8. When $\beta = 0$, as expected the effective quantity Σ is equal to σ and they are both 0.4. When β increases, in absolute value, the depressing effect of β on Σ is apparent. This is due to the fact that the largest firms, those that contribute the most to the aggregate growth rate, are characterized by a lower growth variance the higher the value of β . For $\beta = -0.2$ Σ is just above one quarter of the actual standard deviation σ . When we reduce the standard deviation of f_s the effect of a higher β is mitigated. Indeed a lower σ_s implies that larger firms in the economy tend to be smaller and with a higher variance of their growth rates implying a milder reduction in Σ as apparent in Figure 9. When the firm (log) size distribution is characterized by a standard deviation 0.25σ , increasing β does not have any significant effect on Σ .

C 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.

C.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.³¹ 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 ($n = 2^m$).³² 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 $j_1 \leq j_2 \leq n/4 + 1$;
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

³¹This procedure is implemented in the R package 'modeest' <https://cran.r-project.org/web/packages/modeest/modeest.pdf>.

³²They also generalize this algorithm to allow n to be any positive integer. See the details in the appendix of their paper.

the arbitrary selection of a “bandwidth” , which is often problematic. Indeed if the bandwidth is too large, then the mode cannot be precisely located, leading to a high bias. If, on the contrary, the bandwidth 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.

C.2 The Asymmetric Exponential Power distribution (AEP)

The AEP distribution introduced 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 $C = a_l A_0(b_l) + a_r A_0(b_r)$, $A_k(x) = x^{\frac{k+1}{x}-1} \Gamma\left(\frac{k+1}{x}\right)$ 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

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

Thus according to definition (9) the parametric distributional tilt is given by

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

The AEP parameters are estimated using Maximum Likelihood implemented in the package “Subbotools”.³³

³³The latest version of *Subbotools* can be found at <http://cafim.sssup.it/~giulio/software/subbotools/>. And Bottazzi (2004) provides the detailed manual.