

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

Long-Term Firm Growth: An Empirical Analysis of US Manufacturers 1959-2015

Giovanni Dos ^a
Marco Grazzi ^b
Daniele Moschella ^a
Gary Pisano ^c
Federico Tamagni ^a

^a Institute of Economics & EMbeDS, Scuola Superiore Sant'Anna, Pisa, Italy.

^b Department of Economic Policy, Università Cattolica del Sacro Cuore, Milan, Italy.

^c Harvard Business School, Harvard University, Cambridge, US.

2019/13

May 2019

ISSN(ONLINE) 2284-0400

Long-Term Firm Growth: An Empirical Analysis of US Manufacturers 1959-2015 *

Giovanni Dosi¹, Marco Grazzi², Daniele Moschella¹, Gary Pisano³, and
Federico Tamagni¹

¹Institute of Economics & EMbeDS, Scuola Superiore Sant'Anna, Pisa,
Italy

²Department of Economic Policy, Università Cattolica del Sacro Cuore,
Milan, Italy

³Harvard Business School, Harvard University, Cambridge, US

May 16, 2019

Abstract

Firm growth is an essential feature of market economies, shaping together macroeconomic performance and the evolution of industry structures. As a potential indicator of organizational “fitness” within a competitive environment, firm growth is also a central concern to both the practice and theory of business strategy. Despite both its theoretical and practical importance, though, growth remains a poorly understood property of firms. While previous studies have documented the highly skewed nature of firm growth rates, we know far less about the persistence of growth rates over long-periods of time. For instance, do “fast growers” tend to maintain their relative growth rates advantages over long-periods or is superior growth a transitory phenomenon? Is, as predicted by evolutionary and capability based theories of the firm, the process of firm growth path-dependent or is it more akin to a random walk? The answers to these questions are central to building a robust theory of firm growth. This paper attempts to address this gap in our empirical knowledge of firm growth using a dataset that spans 50 years, which allows the abandonment of the assumption, common to all incumbent studies, that

*We thank Giulio Bottazzi, Andrea Roventini, Angelo Secchi, anonymous referees, and several participants at the OrgScience Winter Conference in Park City February 2016; ICC Conference, Berkeley December 2016; the 11th EMAEE conference, Strasbourg June 2017; and seminars at EPFL, University of Florence for their insightful comments. We also thank Wayne Gray and the NBER-CES Manufacturing Industry Database for providing historical data on US sectoral deflators. This work has been partly supported by the European Commission under the H2020, GROWINPRO, Grant Agreement 822781). We acknowledge the financial support of the Harvard Business School Division of Research, and the assistance of Harvard Business Baker Library in assembling the data. The usual disclaimers apply.

the stochastic paths of all firms stem from the same generating process. These exploratory results indicate that growth rate persistence is there and may be even substantial for some firms, but it is rare. We also study the links between the micro-properties of firm growth within sectors and the patterns of aggregate growth of these same sectors. Indeed, we find circumstantial but widespread evidence that heterogeneity across firms correlates with industry dynamism.

Keywords: Firm growth; Persistence; Industrial Dynamics; Firm heterogeneity

JEL classification: C81, D21, D22, D24, L11, L21, L25, L62, L63

1 Introduction

Firm growth is an essential feature of market economies, shaping not only macroeconomic performance (such as GDP or employment dynamics) but also the evolution of industries. Why some enterprises might grow faster than others is of concern not only to economists interested in fundamental theoretical questions, but also to policy makers concerned with employment and antitrust, and to managers and investors given its prominence as a criterion for capital allocation in financial markets. Despite both the theoretical and practical importance, though, the understanding of firm growth and of its main drivers is surprisingly poor. In fact, since the publication of Edith Penrose's landmark book, *The Theory of the Growth of the Firm* (Penrose, 1959), six decades ago, little progress has been made in constructing a robust theory of firm growth. That contribution paved the way for the emergence of several related streams of literature claiming the central role of heterogeneous firm-specific characteristics as drivers of firm growth. Yet, not much dialogue has been in place between these streams and the empirical facts on size-growth dynamics emerging from the large literature spurred by the seminal contribution by Robert Gibrat (Gibrat, 1931). One has still a relatively limited understanding of the underlying organizational drivers of firm growth and, as a result, of the industry growth.

The starting point in building a theory of firm growth (like any theory) is the identification of empirical regularities and puzzles. A solid theory of firm growth needs to be able to explain firm growth patterns over long periods of time and, thus, it must be a theory of persistence of growth or lack of it. The empirical findings emerged within the Gibrat's Law literature over the past decades provide a fairly consistent picture about the properties of firm growth rates (from the classic Ijiri and Simon, 1977 to Evans, 1987b,a; Dunne et al., 1988; Stanley et al., 1996; Bottazzi and Secchi, 2003 among others; see also the discussion in Dosi, 2007). However, long-term persistent differences in growth rates are much less well characterized empirically. That is, we do not know empirically whether firms (like say Apple, Google, or Facebook) that have been relatively fast growers over a given period are more likely (than other firms) to be fast growers also in subsequent time periods. Putting another way, we do not have an agreed upon answer to the very basic question: Are firm growth rates and growth rates differences persistent over long periods of time?

The issue is not trivial. Evolutionary and capability-based theories of firm predict some degree of path dependence in firm growth rates based on the notion of quasi-stable differences in firm-specific capabilities and know-how. A general lack of persistence in growth rates would be a serious blow to such theories of the firm. Moreover, lack of persistence would also undercut much of the rationale for the field of strategic management and other bodies of management focused on helping practitioners identify drivers of long-term sustainable advantage. Without some stylized facts about the question of persistence, it is hard to image the construction of a robust.

To address this gap, this paper explores the growth processes of a large sample of US manufacturing firms over more than fifty years of observations and their relationships with some features of industry evolution.

Previous empirical studies on the persistence of growth (see e.g. Coad et al., 2014;

Capasso et al., 2013; Hölzl, 2014; Daunfeldt and Halvarsson, 2015; Bianchini et al., 2017; Moschella et al., 2019), provides limited evidence of persistence. The primary means of identifying persistence has been to estimate autoregressive coefficients (the presence of a statistically significant one being indicative of persistence). However, since the datasets used in previous studies have relatively short time horizons, they have generally been forced to pool observations across firms – an approach requiring the restrictive assumption that the data generating process on growth rates is the same for all firms and in turn only able to capture the average degree of persistence in the sample. There is some attempt to also look at distributional dynamics, via transition probability matrices (TPM) of growth, but always over relatively short sample time windows.

By taking advantage of an unusually long time series of firm level performance (50 years), we are able to estimate firm-specific autocorrelation coefficients (and thus dispense with any assumptions regarding firm homogeneity). And indeed we find evidence that firm-specific growth persistence does characterize some firms, but is rare. We also explore persistence in terms of relative firm growth (using rankings) and examine properties of TPM to assess whether firm growth rankings remain relatively constant over the long period of time available to us. This analysis, while quite different from estimating firm-specific autocorrelation structures, yields similar evidence on the relatively low degree of persistence in growth patterns.

A second major difference between our approach and that used in prior studies is that we explicitly examine the link between firm level and industry level growth patterns. A few studies address the common properties or the sectoral specificities in the distribution of sizes and growth rates (see Bottazzi and Secchi, 2003; Bottazzi et al., 2007a) but none to our knowledge has made explicit attempt to understand how industry and firm growth rates are related.¹ For instance, does participating in a high growth industry (e.g. semiconductors) increase the odds of a given firm to also enjoying high rates of growth? Does high industry growth correlate at all with growth persistence at the firm level? How are industry growth rates associated with the heterogeneity of firm growth rates in that industry? Connecting industry and firm level growth dynamics would seem to be an essential task for a theory of the growth of the firm. Our findings suggest a richer relationship between firm and industry growth rates than previously found in the literature. In particular, we find that higher industry growth tends to associate with a higher variance in firm growth rates. Moreover, high industry growth associates with higher persistence in growth rates of some firms. Posed different, high industry growth rates are like races with both a larger dispersion between “front runner” and the middle of the pack and the presence of some “long distance runners”.

In the following we shall present the basic benchmark model, together with some “stylized facts” stemming from the incumbent literature germane to our analysis (Section 2), followed by a brief description of the data we address (Section 3). Section 4 shall present pictures of the distributions of firms growth rates, while Section 5 focuses on their persistence properties. Section 6 broadens the analysis to the relationships between micro- and industry-level performances. Section 7 concludes with a discussion of

¹A theoretical effort in this direction is Knudsen et al. (2017), but it is not matched with actual empirical data.

the limits of our approach, and the implications for future research.

2 Long-run Firm Growth: Theoretical and Empirical Questions

Let us begin with the primitive question: what does it mean for a firm to “grow”? This is obviously an important issue to clarify if one wants to develop a theory and yield empirical insights about the phenomenon called “firm growth”. Firm growth is a *process by which organizations pursue market opportunities and the acquisition and accumulation of the resources required to exploit those opportunities*. A few elements of our definition are worth highlighting. First, while we often think of (and indeed measure) growth as an outcome (e.g. greater revenue), we define it as a process to emphasize its evolutionary dynamic character. Growth is a dynamic process unfolding over time, not an instantaneous response to market opportunities or market signals. Firm growth is necessarily a dis-equilibrating process. But our definition also highlights the fact that firm growth entails qualitative changes in the firm itself. Firms are complex entities composed of processes, routines, formal and informal know-how, “rules” and norms (cultures) for behavior, and governing structures. Growth means more than just getting “bigger”; growth involves changes to the internal “machinery” of the firm itself.

Our definition moreover highlights both the supply-side and demand-side drivers of growth. On the demand side, greater market opportunities are a pre-requisite for growth. Some firms grow by finding (or creating) new demand in their existing markets. Apple, for instance, created and exploited massive new demand in mobile communications devices. But demand side growth, along the lines suggested by Penrose, may also involve finding new markets (e.g. Apple offering music subscriptions and payment services). Growth, though, is also a function of supply-side factors. New markets often come from major technological breakthroughs. There would have been no market for transistors without the invention of the transistor, or no market for biotech drugs without DNA decoding. And even microwave ovens would not be there without the discovery of the water-heating properties of microwaves. However, potential opportunities coming from the supply side, even when they stem from public domain, are far from open to all firms because of differences in capabilities. Apple has capabilities and resources like brand (and a customer base) that make it feasible for them to become a payment provider - but its capabilities to enter the car industry are more limited (and its capabilities to enter, say, the drug industry, even less so). Again as suggested by Penrose, demand side opportunities and supply side constraints on growth are connected. The demand side opportunities open to the firm depend on its supply side capabilities and resources. We can not theoretically talk about firm growth and long-run firm growth without referencing the respective roles that both demand side and supply side forces play.

Why Should We Ever Expect Persistent Growth?

The fact that all firms grow or decline at each time step is not a theoretical puzzle. The big challenge concerns indeed the links between what is inside the technological and organizational “black box”, on the one hand, and the available quantitative evidence on growth patterns on the other.

The ubiquitous and persistent heterogeneities across firms in their organizational forms, technological capabilities, and strategies are generally accepted to drive large and persistent differences in competitiveness. Since such differential competitiveness tends to persist over time, one would plausibly expect that they in turn leads to persistence also in firm growth patterns. If present, persistence in growth must be due to firm-specific endogenous forces like innovation, advertising, increases in productivity, and so forth. The question is to what extent this characterizes many firms or just some of them. In fact, even if drivers may appear clear theoretically, growth persistence is not such a straightforward story. Let’s take the example of a company that commercializes a very successful innovation (say, Apple commercializing the iPhone). Not surprisingly, such an innovation would lead to an increase in demand for the innovator relative to rivals (as it did for Apple). But why should such demand growth persist? Imitation and rivalry should presumably act as equilibrating forces dampening the relative growth rate advantage over time. Even in the absence of imitation and rivalry, the positive growth effects of an innovation (or any other demand inducing factor) should dampen over time. At some point, demand saturates and demand growth reverts back to long-term industry averages. Thus, persistence of growth also requires persistence in the firm-specific endogenous drivers of demand. A *single* innovation (e.g. the iPhone) does not explain persistence - but a *sequence* of innovations (e.g. iPhone 3, 3G, 4, 4C, Apple Watch, Apple Pay, etc.) might. A one time improvement in productivity can not explain long-term persistent growth - but years of continuous productivity improvement might. Persistence thus requires some kind of quasi-stable firm-specific capability or attributes that sustain a flow over time of innovations, productivity improvements, or other demand-inducing events.

Revisiting stylized facts on firm growth

It is handy to start the analysis of firm growth processes by setting a sort of “strawman” which also happens to be a classic in the literature, namely the so called *Gibrat’s Law* (cf. Gibrat, 1931; Simon and Bonini, 1958; Steindl, 1965; Ijiri and Simon, 1977; Sutton, 1997; Dosi, 2007). Let

$$s_i(t+1) = \alpha + \theta_i s_i(t) + \varepsilon_i(t) \tag{1}$$

where $s_i(\cdot)$ are the log sizes of firm i at time t , $t+1$ and α is the sector-wide (both nominal and real) component of growth.

Gibrat’s law in its strong form suggests that:

- (a) $\theta_i = 1$ for every i

(b) $\varepsilon_i(t)$ is an independent and identically distributed random variable with zero mean

Hypothesis (a) states the “law of proportionate effect”: growth is a *multiplicative* process independent on initial condition, with no systematic scale effects. Note that were one to find $\theta_i > 1$ one ought to observe a persistent tendency toward monopoly. Conversely, θ_i significantly smaller than one would be evidence corroborating regression-to-the-mean, and, indirectly, witness for some underlying “optimal size” long-run attractor.

Hypothesis (a) is indeed the object of most empirical inquiries which typically look at an average estimate of θ across samples of firms followed over relatively short time horizons. The findings provide mixed support to the hypothesis:

- (i) most often, smaller firms - on average - grow faster (under the *caveat* that one generally considers small surviving firms);
- (ii) otherwise, no strikingly robust relationship appears between size and average rates of growth, as in fact $\theta=1$ provides a reasonably good approximation for the growth process across medium-large firms (cf. Mansfield, 1962; Hall, 1987; Kumar, 1985; Bottazzi et al., 2007a; Bottazzi and Secchi, 2006b; Coad, 2009; Lotti et al., 2003, among others).
- (ii) The relationship between size and growth is modulated by the age of firms themselves - broadly speaking, with age exerting negative effects on growth rates, but positive effects on survival probabilities, at least after some post-infancy threshold (see Evans, 1987a).²

Notice that such pieces of evidence may easily be consistent with evolutionary theories of industrial change. Indeed an evolutionary interpretation would be rather at odds with a notion of convergence to some invariant “optimal” size, with decreasing returns above it. Conversely, it is rather agnostic on the precise specification of non-decreasing returns. In particular, it does not have any difficulty in accepting a world characterized by nearly constant returns to scale, (i.e. by values of θ_i in Equation 1 on average not too far from one) jointly with drivers of firm growth on average uncorrelated with size itself.

Further precious clues on the basic characteristics of the processes of market competition and corporate growth are offered by the statistical properties of the error term ($\varepsilon_i(t)$ in Equation 1). In fact, if one were to find corroboration to hypothesis (b) according to which - to recall - growth would be driven by small “atomless” uncorrelated shocks, this would come as a serious blow to all theories claiming the importance of heterogeneous firm-specific characteristics as drivers of their performance - including of course their growth -, from Penrosian interpretation based on differential managerial abilities, to capability-based theory of the firm emphasizing the diversity in the knowledge and in the strategies which firms embody (within a vast literature see Winter, 1987; Teece et al., 1997; Dosi et al., 2000; Pisano, 2017). Posed differently, one should expect structure and over-time dependence in growth if one believes that some firms have unique capabilities or organizational attributes that can sustain advantages over time. If, in contrast, there

²Moreover, the relationship between size and growth appears to be influenced by the stage of development of particular industries along their life cycles: see Geroski and Mazzucato (2002).

is nothing more to firms than a set of relatively easy to replicate contractual agreements and production techniques, we should not expect growth to replicate over time. Gibrat’s law in its strong form should hold.

The available empirical studies, once again mostly based on samples of firms observed over relatively short periods of time, do reveal some more structure than postulated by Gibrat’s Law in its strong form. Indeed, *first*, the distribution of growth rates is invariably fat-tailed and well approximated at least by a Laplace distribution (see, among the many others, Stanley et al., 1996; Bottazzi and Secchi, 2006a) or by even fatter-tailed distribution functions.³ Such tail behaviour inconsistent with small Gaussian shocks is an extremely robust stylized fact, holding across countries, level of aggregation (broad manufacturing as well as for well defined sectors) and irrespective of the proxy for size that is employed (sales, employees, value added, etc). The economic implication is that the growth process of firms is far lumpier than what is assumed by standard model of growth and, further, this process induces some correlation within the industry, as obviously market shares must add up to one.

A *second* piece of evidence is centered around the concept of persistence of growth and complement the picture above. As mentioned in the introduction, most available studies have tried to estimate the average degree of autocorrelation in cross sections, with mixed results depending on the country, sector and time periods under study. Findings range from the view that growth is indeed a random walk advanced in Geroski (2002), to the case of long-lasting autocorrelation (up to the 7th lag) found in Bottazzi et al. (2001a) across the world-wide pharmaceutical industry, to the negative serial correlation found in Goddard et al. (2002) on Japanese quoted firms and in Bottazzi et al. (2007b) and Bottazzi et al. (2011) for Italian and French manufacturing.⁴ More recent empirical research tried to go beyond “average effects” and exploited quantile auto-regressions or transition probability matrices (TPM) to examine the degree of persistence in different quantiles of the growth rates distribution, focusing in particular on high-growth firms in the top tail. Yet, most of these studies cast doubt on the very existence of persistent high-growth. Most high-growth firms do not replicate their high-growth performance over time and are in fact “one-hit wonders” (Hölzl, 2014; Daunfeldt and Halvarsson, 2015). Moreover, it is unclear what characteristic is crucial for the performance of the very few firms that succeed to consistently remain in the high-growth quantiles: most of supposedly strong drivers of growth –such as productivity, profitability, financial structure, innovation and innovation persistence– do not correlate with high-growth persistence (see Bianchini et al., 2017; Moschella et al., 2019; Guarascio and Tamagni, 2019), while size has contrasting effects on high-growth in different studies (small high-growth firms display negative autocorrelation in Coad 2007 and Coad and Hölzl 2009, while persistent out-performers are more often present among micro firms in Capasso et al. 2013).

There is much less attention, if any, on the persistence of growth in the history of each single firms. In this context, persistence can be interpreted in at least two distinct,

³See recent findings in Williams et al. (2015) suggesting that the heavier-tailed Cauchy distribution may provide the best approximation.

⁴Findings on service firms provide a similarly mixed picture, as in Vennet (2001) on banking companies across OECD countries and Goddard et al. (2004) on US financial services.

although related ways. On the one hand, and similarly in spirit to the TPM analysis, one can focus on the growth performance of the firm *relative* to other firms in the industry, and try to assess the degree of stability in the intra-distributional rankings of each firm. On the other hand, one can take the complementary notion of persistence in *absolute* terms, looking at the degree of autocorrelation in the firm-specific time series of growth rates and thus estimate one (or more, depending on the lag structure) autocorrelation coefficients separately for each firm. So far this approach has been empirically unfeasible, mostly because the time-series of firm-level panel datasets are rather short (generally below 20 years), thus precluding this sort of estimates. On the contrary, as in this work we can resort to more than fifty years of observations, we are able to estimate firm-level autoregressive coefficients.

3 Data and Variables

In our empirical analysis we mostly resort to Compustat, covering all firms listed on the US stock exchange, which is available to us starting from 1959. The unusually long availability of firm-level data is crucial to allowing us to investigate the pattern of growth of firms in the very long run, its characteristics in terms of persistence and the relation between some properties emerging in the industries and sectoral growth over the long term.

In addition to Compustat we also employ the the NBER-CES Manufacturing Industry Database which covers the period 1958-2011, in order to integrate the dataset with industry-level variables.

Given the nature of the empirical analysis with its focus on persistence and long-term trend, one needs to construct a sample of firms with a sufficiently large number of observations over time, possibly without any breaks in the time series. To this purpose, we limit our analysis to firms that have been continuously present in Compustat for at least 30 consecutive years. This means that from a total of more than 32,000 firms that appeared in Compustat records, our preferred sample to study persistence is restricted to around 2,500 firms. We will focus our analysis on manufacturing firms, from SIC 20 to SIC 39, and thus the number of companies further shrinks to around 1,100.

Among different proxies of firm size used in the literature, we take sales as the object of our analysis, as this is a measure that directly relates to the success (or failure) on the product market. In line with most of the previous literature, firm growth rate is computed as the log-difference of firm size between two consecutive years

$$g_{i,t} = \ln Size_{i,t} - \ln Size_{i,t-1} \quad (2)$$

where *Size* stands for normalised size obtained by removing from firm-level sales (*SIZE*) their annual sectoral average as follows:

$$\ln Size_{i,t} = \ln SIZE_{i,t} - \frac{1}{N} \sum_i \ln SIZE_{i,t}$$

The normalization accounts for common trends, such as inflation and business cycles

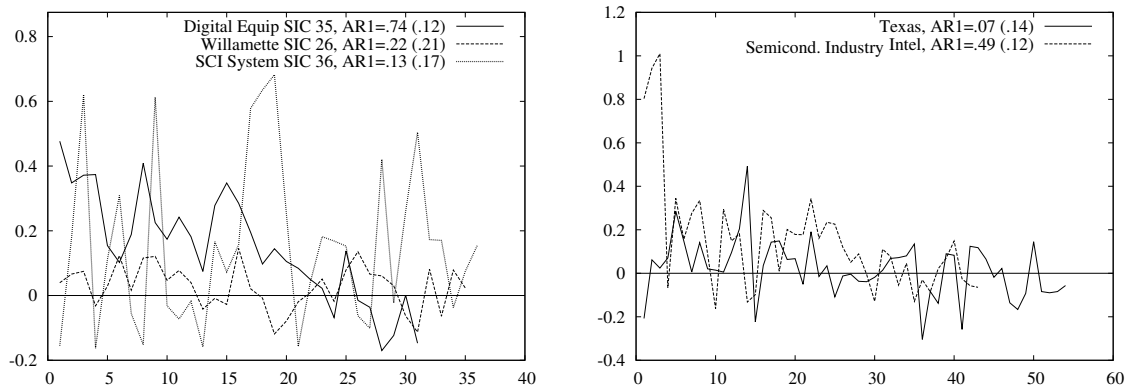


Figure 1: Time-series of yearly sales growth rates of selected firms, across different sectors (*left*) and within the semiconductor sector (*right*). Reported AR(1) coefficients (standard errors in parentheses) are from OLS estimation

effects in sectoral demand.⁵

Before moving to the more systematic analysis of the growth process of the next Section, we report here two impressionistic pieces of evidence on the inter-temporal dynamics that characterizes the growth patterns of firms. First, in Figure 1 we depict the time-series patterns of yearly sales growth rates of selected leading firms, across different sectors (see left plot) and within the semiconductor sector, one of the highest growth sector over the last sixty years (right plot). The profiles of firm growth are rather heterogeneous even within the same sector. Moreover, they are quite lumpy and erratic, with big jumps often followed by a slow reversion to the mean thereafter.

Further evidence on lumpiness of growth is presented in Figure 2. We rank the absolute growth episodes experienced by each firm from the highest to the lowest, and then report the mean and the median share of each rank in the total cumulated growth experienced by firms over the observed 30-year period.⁶ The results show that the largest growth event over the history of any one firms accounts for about 40% on average and about 30% in median of total sales growth experienced over the period (see left plot). By way of comparison, we also show the same ranks analysis for employment growth (right plot), which displays an even greater degree of lumpiness: the median highest absolute growth share is more than 50%.

4 Distributional analysis

We start the analysis of long-run micro properties by examining the distributions of growth rates at different levels of sectoral and temporal aggregation as well as at different points in time.

⁵Annual sales are deflated with the corresponding price index at the two-digit level. A standard limitation that we share with most empirical works is that we are unable to distinguish between organic or internal growth *versus* external growth.

⁶Absolute growth shares are, for each firm-year pair, the ratios of current (deflated) absolute sales growth to the total (deflated) absolute sales growth over 30 years. Notice also that in order to compute mean and median over the same number of observations per firm we use a balanced panel here.




Figure 2: Absolute growth shares by rank in sales (*left*) and employment (*right*), from 1975 to 2004. Shares on the vertical axis and ranks on the horizontal axis.

In order to provide an accurate analysis of the distributions of firm growth rates, we will resort to the Asymmetric Exponential Power (AEP) distributions, which is an extension of the Subbotin-Exponential Power distribution (Subbotin, 1923) introduced by Bottazzi and Secchi (2011) to properly account for asymmetries between the left and the right tail. Parametric (maximum likelihood) estimates of shape parameters allow to assess leptokurtosis, with the Gaussian distribution arising as a special case within this larger family.⁷

Figure 3 presents the distributions of sales growth rates for the aggregate manufacturing, in three different years (*left*), and for growth rates computed over varying time horizons (*right*).⁸ In general, the growth rates distribution displays the familiar “tent-shape” which has been found in several dataset (see, among others, Stanley et al., 1996; Bottazzi and Secchi, 2006a; Coad, 2007; Bottazzi et al., 2011) confirming the existence of a large share of ‘lumpy’ growth events, of both positive and negative sign. The estimates of the shape parameters b of the Subbotin distribution reveal that, if anything, growth rates became even more fat-tailed in more recent years: the estimated b was around 1 in 1963 and is around 0.6 in 2013 (recall that a Gaussian distribution would imply a parameter of 2). Moreover, the results for growth over different time horizons suggest that

⁷The AEP density presents the following functional form (Bottazzi and Secchi, 2011)

$$f_{\text{AEP}}(x; \mathbf{p}) = \frac{1}{C} e^{-\left(\frac{1}{b_l} \left|\frac{x-m}{a_l}\right|^{b_l} \theta(m-x) + \frac{1}{b_r} \left|\frac{x-m}{a_r}\right|^{b_r} \theta(x-m)\right)} \quad (3)$$

where $\mathbf{p} = (b_l, b_r, a_l, a_r, m)$, $\theta(x)$ is the Heaviside theta function and where the normalization constant reads $C = a_l A_0(b_l) + a_r A_0(b_r)$ with

$$A_k(x) = x^{\frac{k+1}{x}-1} \Gamma\left(\frac{k+1}{x}\right). \quad (4)$$

The two positive shape parameters b_r and b_l , describe the tail behavior in the upper and lower tail, respectively; two positive scale parameters a_r and a_l , associated with the distribution width above and below the modal value and one location parameter m , representing the mode. The standard Exponential Power distribution (Subbotin, 1923) arises under symmetry of tail, i.e. for $a_l = a_r$ and $b_l = b_r$.

⁸The plot on the left is performed on all available observations, whether in the one on the right, requiring a longer time horizon, we only include firms with at least 30 years of observations in their time series.

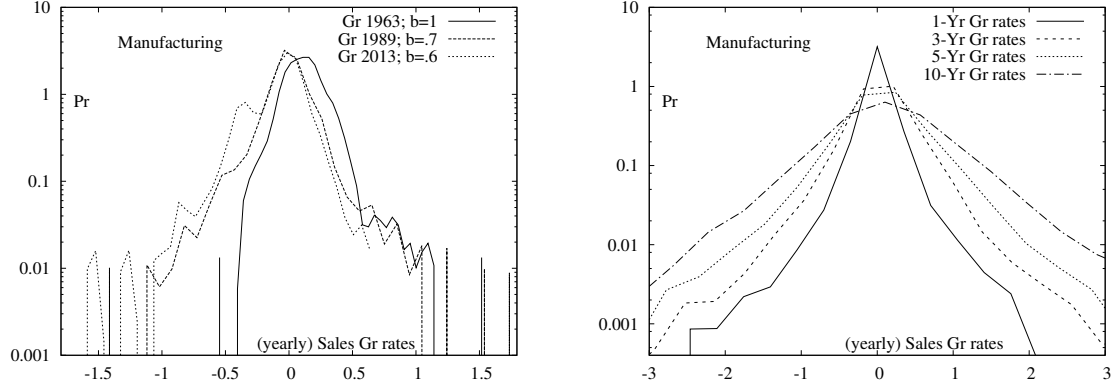


Figure 3: Empirical density of annual sales growth rates in three years (*left*) and over different time intervals (*right*), aggregate manufacturing.

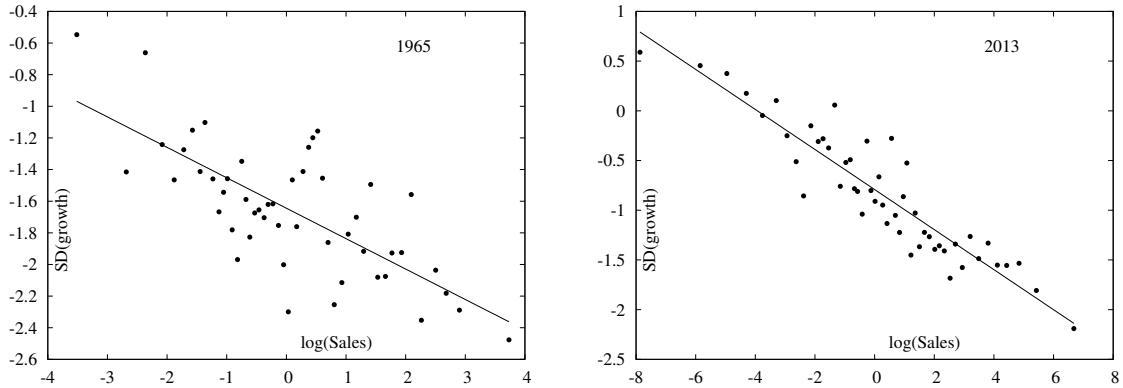


Figure 4: Binned relationship between the standard deviation of growth rates (on a log scale) and the size of the firm, together with a linear fit. Coefficients (and standard errors) are $-0.19(0.03)$ for 1965 and $-0.20(0.01)$ for 2013

aggregation over time maintains the fatness of the tail. Estimates of the b parameters (not reported) go from around 0.75 for the yearly growth rates, only up to 1.1 for the annual growth rates cumulated over a 10-year time interval.

In Figure 4 we investigate the relation between the standard deviation of growth rates and firm size. We split the normalized (log) size $\ln Size_{i,t}$ into 50 bins (quantiles) for each year and then we fit a linear relation between the log of the standard deviation and the mean log size within each bin. Results are reported for two different years (1965 and 2013) by way of example but consistent findings emerged for the entire sample period. They testify a significant negative relation (linear in logs, and exponential in levels) between size and the standard deviation of growth rates. The estimated slope is around -0.20 in both years, and the same value is in fact the average value of the coefficients estimated across all the years in the sample. This is very close to the size-variance scaling parameters estimated in previous studies, for both the US and other countries (Amaral et al., 1997; Bottazzi et al., 2001b; Bottazzi and Secchi, 2006b; Calvino et al., 2018).

We next move from the analysis of aggregate manufacturing to the distributional properties of growth rates within sectors. Figure 5 shows results using data for firms observed over 30 years for a set of selected sectors, by way of example of the findings

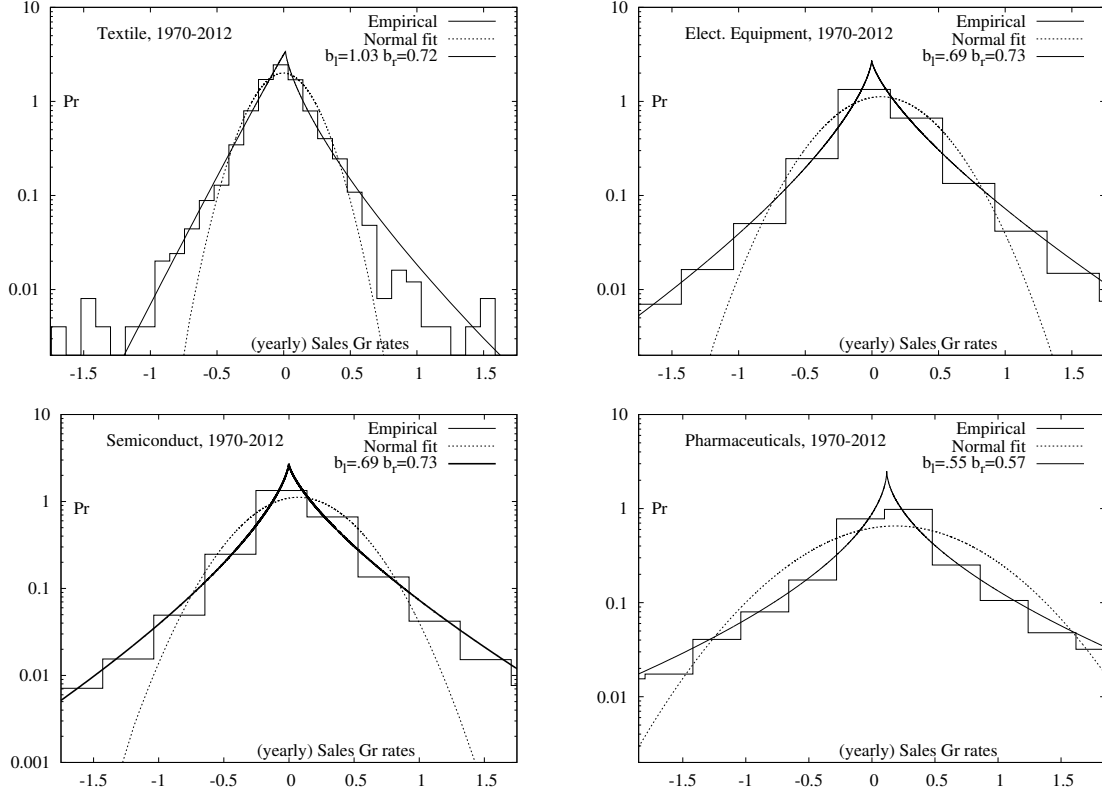


Figure 5: Empirical density of sales growth rates for four sectors, together with the normal and the AEP fits. Pooled data 1971-2011. Probabilities on the y-axis are on a log scale

consistently obtained also in other industries. Even within sector, the departure from normality of the empirical distribution is apparent, both for the tail behavior and for the asymmetry of the two tails. This confirms the stylised fact that heterogeneity, which in the plot gets revealed in the width of the support and in the lumpy nature of the growth process, is maintained also when increasing the level of disaggregation, thus recalling, as originally put forth by Griliches and Mairesse (1999), Mandelbrot’s fractal phenomenon: “the observed variability-heterogeneity does not really decline as we cut our data finer and finer. There is a sense in which different bakeries are just as much different from each others as the steel industry is from the machinery industry.” Granted this general finding, there is anyway room for second-order sectoral variation. It is indeed interesting to note that high-tech/high-growth sectors, like Semiconductors and Electrical Equipment, display fatter tail than low-tech/low-growth sectors like Textile.

While the above pieces of evidence allow us to reproduce a set of stylised facts already investigated in the literature over datasets shorter in time, the long firm level time-series available here allows to add new empirical analysis, concerning the patterns of growth of each single firm. In particular, we can even further disaggregate the analysis of growth rates distributions, moving from sectoral distributions to the study of the distribution of growth rates experienced by each and every firm, at least for those firms for which we have data over many years.

We focus on firms surviving at least 55 years, which compresses the sample to 225

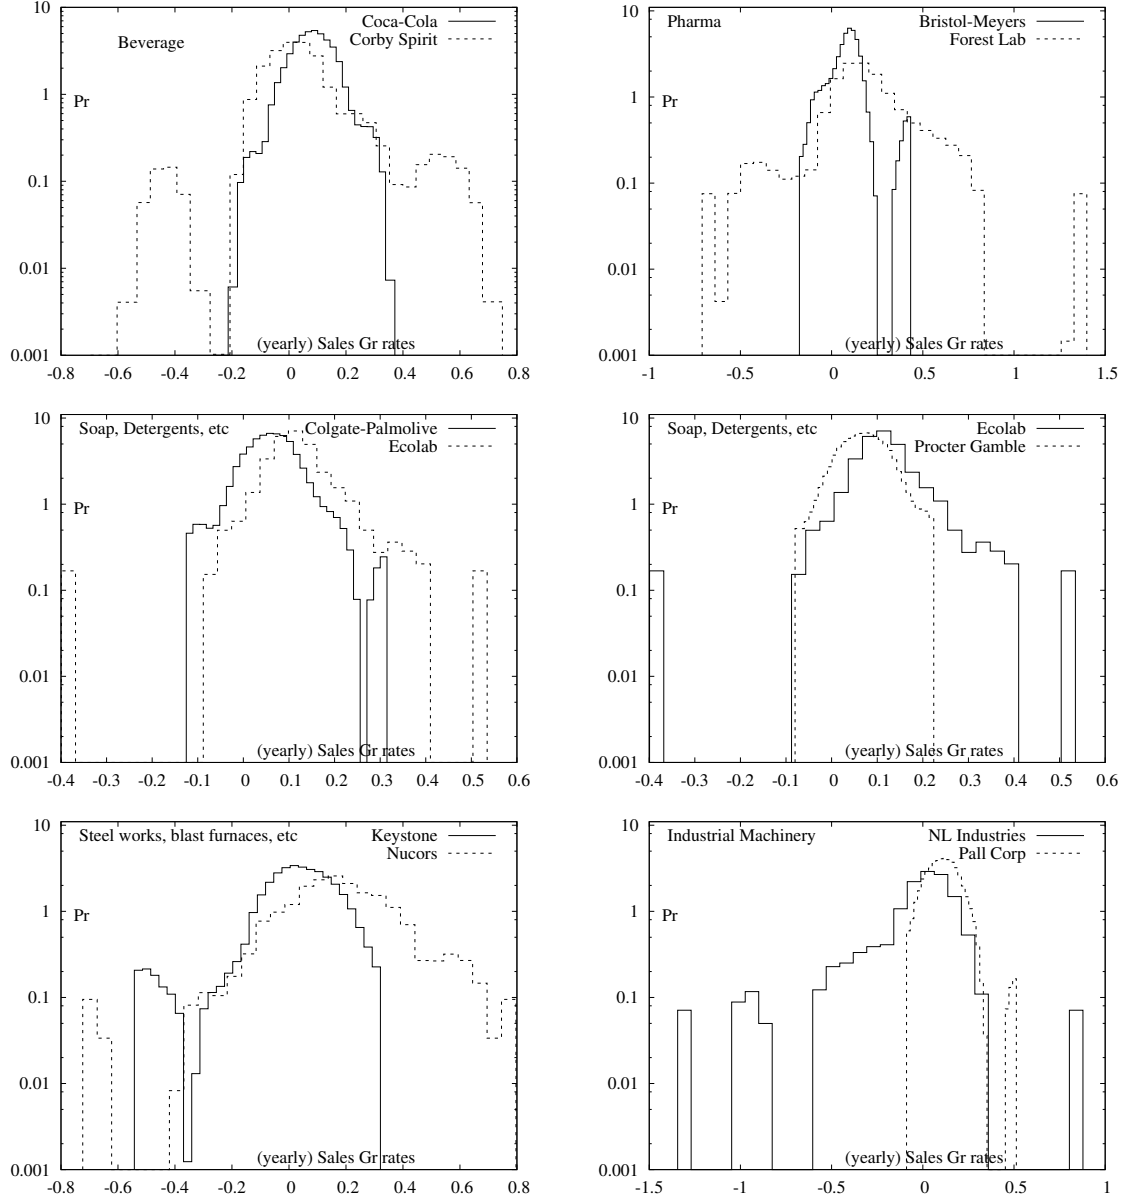


Figure 6: Empirical density of firm-specific sales growth rates. Each plot displays two firms within the same 3 digit SIC code. Only firms with at least 55 years of observation are considered.

firms, with at most 8 companies in the same 3 digit sector. In Figure 6 we report the single-firm growth rates distribution for two companies within the same 3 digit SIC code. This shows that Gaussianity is the exception also at the firm-level: fitting the AEP on the empirical distribution of a single firms' growth rates yields only 4 cases (out of the 225) in which both the left and right shape parameters are consistent with a Gaussian distribution.

Overall, the analysis of properties of growth rates distributions, at different levels of aggregation, already hints at the emergence of some regularities. *First*, we confirm the stylised fact already established on shorter in time datasets that fat-tails are ubiquitous. Thus, even within the long run perspective we undertake here, we corroborate the ex-

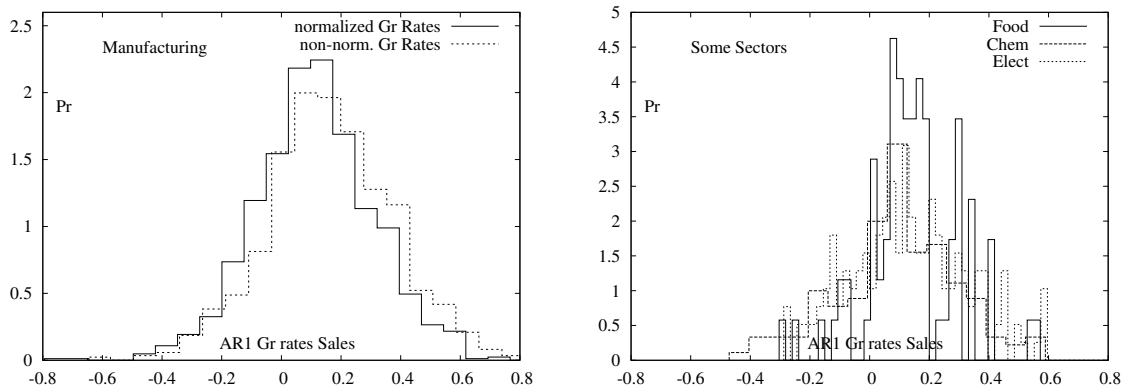


Figure 7: Empirical density of AR1 coefficients of sales growth rates for aggregate manufacturing (*left*) and for three sectors (*right*). The figure for aggregate manufacturing reports also the distribution of AR1 coefficients for sales growth rates without normalization

istence of much more structure in the growth process than what postulated according to a pure “white noise” process. *Second*, and this is germane to our firm-level analysis allowed by long time series, fat-tails are not only invariant to different levels of sectoral aggregation (total manufacturing vs. single sectors), but they are still there even when considering the distribution of growth rates of a single firm. *Third*, and again novel, the analysis of growth rates distributions by each firm reveals very different patterns within industries, hardly compatible with the conjecture of a common data generating process.

Granted all such structure in the growth process, what is the evidence about persistence? We tackle the question in the next section.

5 Persistence of growth

We first focus on what we call *absolute* persistence, by which we refer to the firm-specific autocorrelation structure in the growth process. More in detail, we estimate an AR(1) model separately for each firm which is present in the dataset for at least 30 years.

To summarize the results, Figure 7 shows the distribution of firm-level AR(1) coefficients, both for aggregate manufacturing and for three selected sectors. The support of the distribution is wide, from relatively large and negative autoregressive coefficients, to some outstanding firms displaying relatively high coefficients, like Digital Equipment and Intel (their AR(1) coefficients are 0.74 and 0.49, respectively, see Figure 1). The positive side of supports seems more populated, though, as we found that about 26% of firms display an autoregressive coefficient larger than 0.3. In fact, around 25% of the AR(1) coefficients are statistically different from 0 at the 5% confidence level: when considering only these significant coefficients, the proportion of values larger than 0.3 is around 93%. Of course, a positive AR coefficient simply signals the existence of some persistence, without telling whether such persistence is over a sequence of positive or negative growth events.

These findings on the autoregressive structure of firm-specific growth patterns com-

Table 1: Transition probability matrix between time t and $t + 1$, distribution of yearly sales growth rates

Aggregate Manufacturing		$t+1$				Food and Beverage		$t+1$			
		1	2	3	4			1	2	3	4
t	1	34.07	24.51	20.42	21.00	t	1	36.28	25.49	19.49	18.74
	2	24.63	30.11	27.31	17.95		2	24.50	32.56	26.51	16.43
	3	20.72	27.01	29.29	22.98		3	19.82	24.51	32.98	22.69
	4	21.19	18.81	23.39	36.61		4	21.44	16.90	22.54	39.12
	Shorrocks index=0.90 Bartholomew's index=0.36				Shorrocks index=0.86 Bartholomew's index=0.34						
Chemicals (no Pharma)		$t+1$				Pharmaceuticals		$t+1$			
		1	2	3	4			1	2	3	4
t	1	38.72	23.37	16.27	21.65	t	1	40.30	21.29	19.01	19.39
	2	24.45	35.27	26.89	13.39		2	26.23	36.48	24.59	12.70
	3	15.83	24.89	37.27	22.02		3	17.62	27.20	36.02	19.16
	4	22.83	15.93	21.08	40.16		4	19.57	14.04	23.83	42.55
	Shorrocks index=0.83 Bartholomew's index=0.33				Shorrocks index=0.82 Bartholomew's index=0.32						
Electrical Equipment (no Semicond.)		$t+1$				Semiconductors		$t+1$			
		1	2	3	4			1	2	3	4
t	1	35.19	24.89	19.66	20.26	t	1	34.29	25.71	20.71	19.29
	2	24.76	31.80	25.11	18.33		2	28.36	29.48	27.99	14.18
	3	20.89	26.54	32.71	19.86		3	23.53	26.10	27.57	22.79
	4	21.01	17.44	23.45	38.10		4	19.38	17.83	25.58	37.21
	Shorrocks index=0.87 Bartholomew's index=0.35				Shorrocks index=0.91 Bartholomew's index=0.35						

Note. Transitions across quartiles of the yearly sales growth rates distribution, 1971-2011, ranging from the lowest (1) to the highest (4).

plement the descriptive evidence on the lumpiness of growth emerged from growth distributions and growth ranks. Growth is characterized by big lumps –and the tent-shaped distribution further supports this conjecture– which tend to stretch over a few consecutive years, generating some persistence at least for some firms, as suggested by the positive AR(1). High persistence is however rare.

We then move to the analysis of transition probabilities which allow to examine the degree of persistence in *relative* growth performance of firms.

Tables 1 and 2 report, respectively, the 1-year and 3-year transitions across quartiles, for aggregate manufacturing and within selected sectors. In general, there is some structure in the process, although persistence is not high. The frequencies estimated in the different cells, ranging between 12.7 and 42.6, point in fact to some departures from an hypothetical “seemingly random” data generating process where the probabilities of staying in the same quartile and of jumping across all the other quartiles are all the same (and equal to 0.25). The probabilities on the main diagonal are in fact always higher than the 25% benchmark, and they are highest at the two extreme of the distribution: firms in the top and in the bottom quartile display the highest probabilities of remaining in the same class. Of course, much larger values (say around 0.7-0.8) would be revealing of strong persistence but, still, there is some more structure than a random process would imply. Off-diagonal values do not reveal a clear pattern: if anything longer jumps seems to be less frequent than jumping to a quartile closer to the initial one. Similar results emerge both for aggregate manufacturing and disaggregating by sector. However, a look at the sectoral properties hints at the possibility that persistence might go together with

Table 2: Transition probability matrix between time t and $t + 3$, distribution of yearly sales growth rates

Aggregate Manufacturing		$t+3$				Food and Beverage		$t+3$			
		1	2	3	4			1	2	3	4
t	1	29.59	23.18	22.43	24.81	t	1	32.54	23.13	21.69	22.65
	2	23.82	29.24	26.95	19.99		2	24.50	29.30	27.48	18.71
	3	22.99	27.05	28.22	21.74		3	21.61	25.48	31.94	20.97
	4	24.22	21.85	23.55	30.39		4	23.05	21.89	21.23	33.83
	Shorrocks index=0.94 Bartholomew's index=0.39				Shorrocks index=0.91 Bartholomew's index=0.37						
Chemicals (no Pharma)		$t+3$				Pharmaceuticals		$t+3$			
		1	2	3	4			1	2	3	4
t	1	31.45	23.46	21.87	23.22	t	1	34.94	22.49	22.49	20.08
	2	22.57	29.93	28.30	19.20		2	28.38	28.82	25.33	17.47
	3	19.42	29.67	30.16	20.75		3	15.16	29.92	36.89	18.03
	4	27.97	17.85	22.47	31.71		4	26.24	18.55	19.91	35.29
	Shorrocks index=0.92 Bartholomew's index=0.38				Shorrocks index=0.88 Bartholomew's index=0.35						
Electrical Equipment (no Semicond.)		$t+3$				Semiconductors		$t+3$			
		1	2	3	4			1	2	3	4
t	1	30.04	23.63	21.98	24.36	t	1	32.20	23.86	23.11	20.83
	2	25.05	29.98	27.09	17.88		2	25.40	27.78	28.57	18.25
	3	22.49	28.61	27.42	21.48		3	23.14	26.67	27.45	22.75
	4	24.52	20.46	25.07	29.95		4	26.14	22.82	22.82	28.22
	Shorrocks index=0.94 Bartholomew's index=0.38				Shorrocks index=0.95 Bartholomew's index=0.39						

Note. Transition across quartiles of the yearly sales growth rates distribution, 1971-2011, ranging from the lowest (1) to the highest (4).

knowledge specificities: see the case of the drug industry which emerges as the most persistent in terms of the ranking of growth rates (relatively higher frequencies on the main diagonal).

In order to summarize the degree of *overall* persistence revealed by the matrices, we compute two indexes of intra-distributional mobility. The Shorrocks' index captures the degree of persistence on the main diagonal: it takes value 0 if the TPM displays "full persistence" (i.e., if all firms remain in their initial quartile and the TPM is thus the identity matrix) and goes to $k/(k-1)$, with k being the number of states, if there is "full mobility" (i.e., if the matrix has all zeroes on the diagonal).⁹ The results, in the range of 0.8-0.9, confirm the presence of some more persistence than in a pure random case, but very low.

The Bartholomew index provides a more complete assessment, by taking into account also off-diagonal transitions, with higher weights assigned to longer jumps. As the Shorrocks' measure, it also takes value 0 in case of "full persistence" (all firms remains in their initial quartile), but it has an upper bound at 1 for "full mobility" (in this case to be interpreted as all firms making the longest possible jump compared to their initial class).¹⁰ The values obtained in the different TPM, all around 0.35, again point to rel-

⁹Formally, for a matrix A , the Shorrocks (1978) index $M_s(A)$ is computed as $(k - \text{trace}(A))/(k - 1)$. In our case, with 4 states, the upper bound for "full mobility" is equal to $4/3$, while the index value would be 1 for a benchmark "seemingly random" process with 0.25 in all cells of the matrix.

¹⁰Given a matrix B , the index (Bartholomew, 1973) is given by $M_B = \frac{1}{k-1} \sum_i \sum_j p_{i,j} |i-j|$ where $p_{i,j}$ is the value of the probability in row i , column j ; and $p_{i,\cdot}$ is the marginal distribution of class i in the year t of observation; $i-j$ denotes the distance between the two class (for a further discussion of

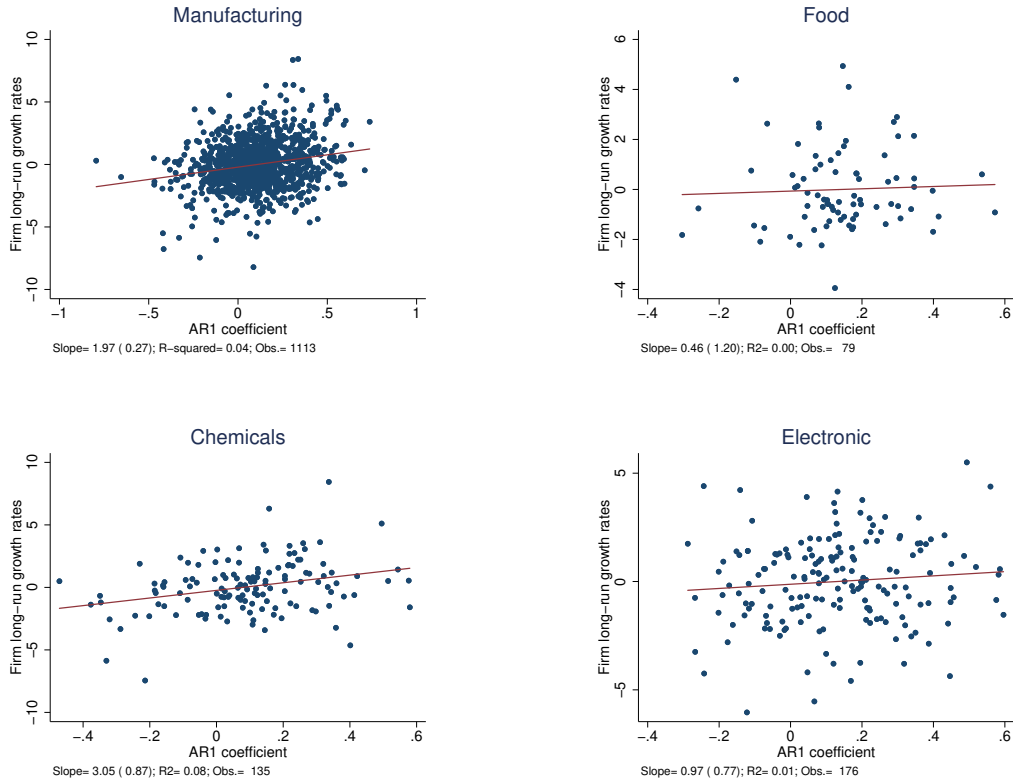


Figure 8: Scatter-plots of the relationship between AR1 coefficients and firm long-run sales growth rates. Below each plot, we report the slope of the linear regression line (with the associated standard error), the R^2 and the number of observations

atively low persistence, although they confirm the existence of some more structure than in a pure random process.

Taken together, these pieces of evidence on persistence, both in *absolute* and *relative* terms, are quite convincing.

Another aspect of micro-level persistence of growth processes that we can investigate is the relation between firm-level autocorrelation and firm-level growth performance in the long run. Figure 8 reports four scatter-plots of the relationship between AR(1) coefficients and long-run sales growth, computed for each firm as the log-difference between the (demeaned) size in the initial and final year over the entire period. We show results pooling across entire manufacturing, and for three selected sectors (Food, Chemicals, Electronics). They reveal a positive correlation, which is apparent in the plot for aggregate manufacturing and in the Chemicals sector. Indeed, the correlation is positive in all but one sector (Apparel) available in the data. However, the “strength” of this nexus is generally rather low when there at all. Across the whole sample, the mean correlation coefficient between AR(1) coefficients and long-run growth rates is around 0.2, but there is a quite puzzling heterogeneity across the different sectors: the correlation is negligible in a traditional sector like Food, but also in a high-tech sector like Electronics, while

mobility index computed starting from TPM, see among the others, Cantó and Ruiz, 2015). A “seemingly random” process with 0.25 probability in all cells would imply a value of about 0.42).




Figure 9: Manufacturing growth rates over the period 1961-2011, computed from the NBER-CES industry database (*blue line*) and from Compustat firms (*red line*).

it is significantly higher in sectors like Furniture (correlation coefficient 0.42) but also Instruments (correlation coefficient equal to 0.36).

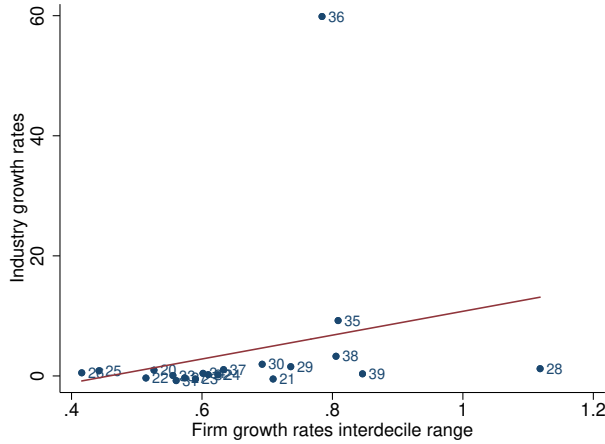
6 Micro properties and industry performance

How does the statistical properties of growth emerging at the micro level within sectors relate to aggregate patterns of growth at the industry level? As already mentioned, we rely upon the NBER-CES database to measure industry variables.¹¹ In Figure 9, we verify that the trend in aggregate output reported by NBER-CES is well in tune with the dynamics one gets by pooling together observations from Compustat companies. The plot indeed shows that even on a rather long time period the two aggregate time series are quite similar, corroborated by a correlation of around 0.9.

As a first exercise linking micro-properties to sectoral dynamics, we correlate sectoral long-run growth with within-sector heterogeneity of firm-level growth rates. In Figure 10, we take cumulated growth of sectors over the period 1970-2011 and proxy for within-sector heterogeneity through the interdecile range of growth rates across firms active in each sector.¹² A mild positive relationship emerges: sectors displaying a wider dispersion in the distribution of firm-level growth rates also report higher aggregate sectoral growth. This is much coherent with a picture in which (a) out-performers significantly contribute to shape the industry evolution, and, more generally, (b) the most dynamic industries witnesses also the co-existence of very good and bad performers (for a related discussion, refer to Dosi, 2007). Incidentally, note that sector 36 (“Computer Equipment and Electronic and Other Electrical Equipment and Components, Except Computer Equipment”) clearly displays cumulated growth much larger than those of other industries and the same holds

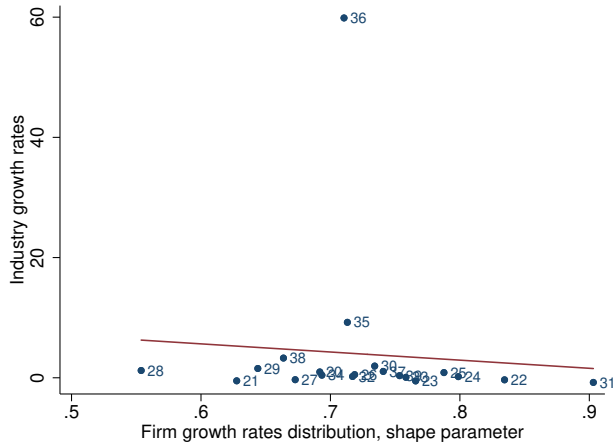
¹¹The dataset is described in Bartelsman and Gray (1996).

¹²Values at the top and bottom deciles are computed after pooling together observations from all years within any one sector.



(a) Slope=1.51 (6.00); Obs.=20

Figure 10: Cumulated industry sales growth rate over the period 1970-2011 (NBER-CES data) *vs.* interdecile range of the yearly firm sales growth rates (Compustat, pooled 1970-2011). Industry classification is SIC 1987. LAD slope (standard error) and number of observations are reported below the graph.




(a) Slope=-7.53 (5.60); Obs.=20

Figure 11: Cumulated industry sales growth rate over the period 1970-2011 (from NBER-CES data) *vs.* shape parameter b of the Subbotin distribution fit of yearly firm sales growth rates (Compustat firms, pooled 1970-2011). LAD slope (standard error) and number of observations are reported below the graph.

–though to a lesser extent– also for sector 35 (“Industrial and Commercial Machinery”). The representation of results is somewhat distorted by these two outliers, not only here but also in the rest of the section. We keep their figures in all the exercises along the section, but results are robust to dropping them from the analysis (see the Appendix).¹³

¹³All the fitted lines presented in this Section are estimated using a LAD estimator, which is more



(a) Slope=-3.18 (4.51); Obs.=20

(b) Slope=-2.78 (4.21); Obs.=20

Figure 12: Cumulated industry sales growth rate over the period 1970-2011 (NBER-CES data) *vs.* shape parameters b_l and b_r of the AEP fit of yearly firm sales growth rates (Compustat firms, pooled over 1970-2011). LAD slope (standard error) and number of observations are reported below the graph

Next, we explore whether industry long-run growth correlates with the degree of fat-tailness in firm-level growth rates. This provides a first piece of evidence about whether large, extreme events of growth hitting individual firms play a role in aggregate dynamics. In Figure 11 we proxy for fat-tail behavior through the Subbotin shape parameter b and relate it to cumulated industry growth over the years 1970-2011. We observe a mildly negative relation, suggesting that a more lumpy growth process at the firm level (lower b) is associated to more dynamic industry performances. The result is confirmed also in Figure 12, where we account for possible asymmetries between the right and the left tail, and consider the AEP left and right shape parameters b_l and b_r : again, sectors with relatively more extreme (positive or negative) firm-growth events are characterised by higher long-run industry growth.

To further investigate to what extent extreme firm-level growth events “drive” aggregate performance, we focus on the relationship between industry growth rates and the growth rates of the largest firms, revealing of granularity in the growth process. In fact, when firm size distributions are skewed and fat-tailed then idiosyncratic shocks to large firms bear a significant impact upon overall industry growth.

In order to check that, we exploit the methodology suggested by Gabaix (2011) and measure the contribution of top (larger) firms via sectoral “granular residuals”, defined for each year t as:

$$\Gamma_{j,t} = \sum_{i=1}^K \frac{SIZE_{i,j,t-1}}{SALES_{j,t-1}} (g_{i,j,t} - \bar{g}_{j,t}) \quad (5)$$

where j indicates the sector, $SIZE_{i,j}$ is firm-level sales of firm i in sector j (from Compustat data), $SALES_j$ are aggregate sectoral sales (from the NBER-CES data), $g_{i,j}$ is

robust than OLS to the presence of outliers. In the Appendix we show that OLS estimates without outliers are indeed very similar.

Table 3: Explanatory power of the granular residual with lag

SIC 2-digits	$\Gamma_{j,t}$	$\Gamma_{j,t-1}$	R^2	$Adj.R^2$
20	0.053	-0.677	0.051	0.010
22	-1.246	-2.109*	0.104	0.065
23	-0.560	0.033	0.014	-0.028
24	0.417	-0.408	0.010	-0.033
25	0.001	-0.987**	0.110	0.071
26	-0.253	0.005	0.010	-0.033
27	10.428***	0.904	0.175	0.139
28	-0.175	-0.899	0.129	0.091
29	-0.069	0.033	0.021	-0.021
30	-1.641*	-0.077	0.089	0.049
31	-0.897**	1.050***	0.188	0.152
32	0.737	0.868	0.024	-0.019
33	-0.334	-0.043	0.001	-0.043
34	-6.403***	2.245	0.176	0.140
35	0.988	-1.549	0.016	-0.027
36	0.297	4.632	0.164	0.128
37	0.502**	0.463*	0.159	0.122
38	-0.589*	0.140	0.074	0.034
39	0.226	-0.339	0.005	-0.038


Note. OLS regression of sectoral sales growth on the granular residual Γ computed over the top 4 largest firms in terms of previous-year sales (see equation (5)). Regression exploit data over the period 1962-2011. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

growth of firm i in sector j and \bar{g}_j is the average growth rate across all firms in the sector, and K is the number of top firms we consider in our granular residual.

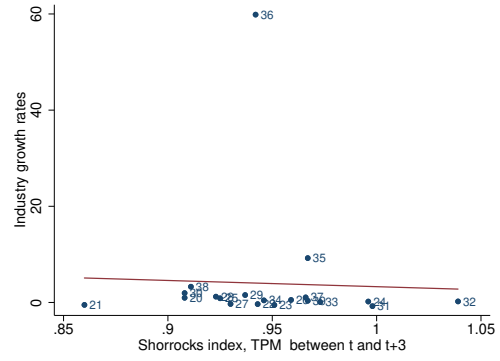
We set $K = 4$, and then regress the granular residual obtained in the different industries and year against sectoral sales growth. Table 3 report results from OLS regressions, taking contemporaneous and lagged residuals.

In general, we observe very low value of the $Adj.R^2$, meaning that the explanatory power of granular residuals with respect to sectoral growth is very poor. In one sector, Transportation (sector 37), characterized by a “mature” oligopolistic structure, granularity is positively correlated with industry growth with significant coefficients on both the contemporaneous and lagged value of the granular residual and a relatively high $Adj.R^2$ (around 12%) compared to most other sectors. In four sectors (27, 31, 34 and 36), the $Adj.R^2$ is also between 12% and 15% but with coefficients that are either negative or not significant.

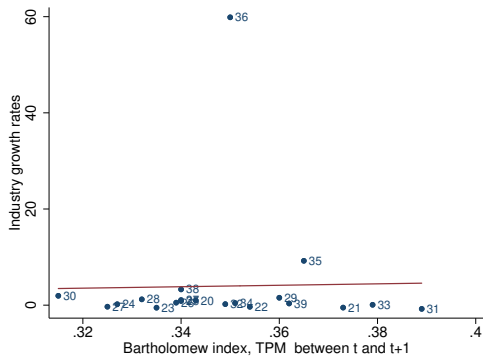
Lastly, we correlate aggregate sectoral growth with micro-level persistence, taking the sectoral values of the two mobility indexes computed above to summarize properties of TPM across quartiles. In Figure 13 we show the results for the indexes computed over 1-year and 3-year transitions. Recalling that both indexes have a lower bound of zero for “full persistence”, the mild negative correlation that we observe suggests that sectoral performance is higher in sectors where persistence is higher. Putting it another way, sectors that have grown relatively more in the long-run are seemingly those wherein




(a) Slope=-14.75 (12.34); Obs.=20



(b) Slope=-8.70 (11.26); Obs.=20



(c) Slope=-23.07 (35.88); Obs.=20



(d) Slope=-13.62 (22.17); Obs.=20

Figure 13: Cumulated industry sales growth rate over the period 1970-2011 (NBER-CES data) *vs.* Shorrocks and Bartholomew index of mobility from sectoral Transition Probability Matrix. Figure (a) and (c) consider TPM between t and $t + 1$, Figure (b) and (d) consider TPM between t and $t + 3$. LAD slope (standard error) and number of observations are reported below each graph

there has been less reshuffling of firm across the growth distribution.

7 Conclusions

If Gibrat (1931)'s hypothesis in its strong form and, much more recently Geroski (2002), were fully right, then a large number of theories stressing persistent differences in firm-level specificities as a key driver of persistence in firm growth and industrial dynamics would be undermined, and management analysis basically hopeless to identify what makes a firm grow and outcompete others in the long run. Notwithstanding this study offers partial and to some extent rudimentary evidence –mostly due to the low number of firms that one can track over a long enough period of time– we provide evidence that this is not the case: the growth process of firms has much more structure than what would be postulated on the grounds of a purely random process driven by independently identically distributed shocks.

First, our evidence corroborates the extremely robust fat-tail property of growth rates. That is, negative and positive growth events are lumpy and bear the mark of some underlying correlating mechanisms concerning both growth opportunities and market competition (the greater growth of some firms means less growth for other ones). This already known stylised fact is confirmed by our novel analysis of firm-specific growth rates time series.

Second, our long time series allows to study the stochastic paths characterizing each single firm which can well stem from different generating processes. Indeed, we find that a good share of firms display positive autocorrelation in their growth rates, although there are clear firm-specific patterns revealed by a good deal of heterogeneity on autocorrelation coefficients, ranging from very low to some specific quite high values. Interestingly, some of the larger values emerge for the technological leaders in two of the most dynamic industries, namely semiconductors and electrical equipment (including computers). Along the same line, we find relatively low persistence of relative growth from estimates of intra-distributional dynamics, again suggesting that high persistence is there, but rare and accruing some particular firms more than standing out as a widespread property.

Third, we uncover some correlations between firm-level statistical properties and industry-level patterns of growth. In fact, one observes (i) some correlation between industry growth and firm autocorrelation and (ii) a good correlation between dispersion of growth rates, on the one hand, and overall industry growth. This is indeed a sign that high growth industries are fueled by underlying industrial dynamics involving highly heterogeneous firms.

On the contrary, and *fourth*, there is hardly any evidence of a more-than-proportional contribution of the biggest firms to sectoral growth (i.e. of a “granular residual”). Rather, the high tide lift all boats irrespectively of their size.

There is an overall message coming from these findings. As firms are persistently heterogeneous in the growth patterns, the search for the underlying determinants and for the sources of sustained competitive advantages ought to continue to be central to the efforts of both industrial economics and management studies.

References

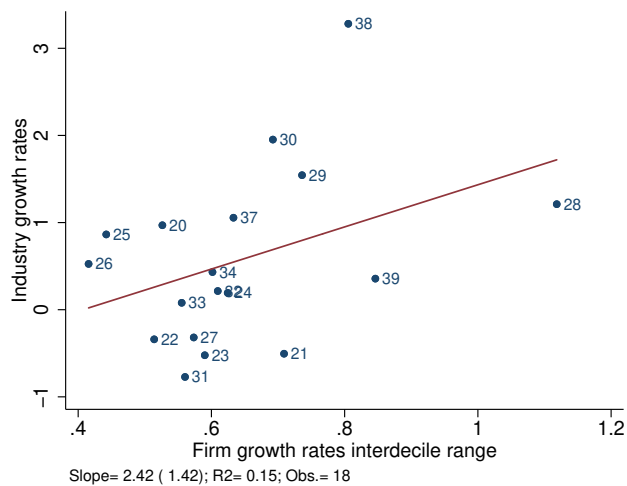
- AMARAL, L. A. N., S. V. BULDYREV, S. HAVLIN, P. MAASS, M. A. SALINGER, H. E. STANLEY, AND M. H. STANLEY (1997): “Scaling behavior in economics: The problem of quantifying company growth,” *Physica A*, 244, 1–24.
- BARTELSMAN, E. J. AND W. GRAY (1996): “The NBER Manufacturing Productivity Database,” NBER Technical Working Papers 0205, National Bureau of Economic Research, Inc.
- BARTHOLOMEW, D. J. (1973): *Stochastic models for the social process*, London: John Wiley & Sons, 2nd edition.
- BIANCHINI, S., G. BOTTAZZI, AND F. TAMAGNI (2017): “What does (not) characterize persistent corporate high-growth?” *Small Business Economics*, 48, 633–656.
- BOTTAZZI, G., E. CEFIS, G. DOSI, AND A. SECCHI (2007a): “Invariances and Diversities in the Evolution of Italian Manufacturing Industry,” *Small Business Economics*, 29, 137–159.
- (2007b): “Invariances and Diversities in the Evolution of Italian Manufacturing Industry,” *Small Business Economics*, 29, 137–159.
- BOTTAZZI, G., A. COAD, N. JACOBY, AND A. SECCHI (2011): “Corporate growth and industrial dynamics: evidence from French manufacturing,” *Applied Economics*, 43, 103–116.
- BOTTAZZI, G., G. DOSI, M. LIPPI, F. PAMMOLLI, AND M. RICCABONI (2001a): “Innovation and corporate growth in the evolution of the drug industry,” *International Journal of Industrial Organization*, 19, 1161–1187.
- BOTTAZZI, G., G. DOSI, AND G. ROCCHETTI (2001b): “Modes of Knowledge Accumulation, Entry Regimes and Patterns of Industrial Evolution,” *Industrial and Corporate Change*, 10, 609–38.
- BOTTAZZI, G. AND A. SECCHI (2003): “Common Properties and Sectoral Specificities in the Dynamics of U.S. Manufacturing Companies,” *Review of Industrial Organization*, 23, 217–232.
- (2006a): “Explaining the Distribution of Firms Growth Rates,” *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.
- CALVINO, F., C. CRISCUOLO, C. MENON, AND A. SECCHI (2018): “Growth volatility and size: A firm-level study,” *Journal of Economic Dynamics and Control*, 90, 390–407.

- CANTÓ, O. AND D. O. RUIZ (2015): *The Contribution of Income Mobility to Economic Insecurity in the US and Spain during the Great Recession*, Emerald, vol. 23 of *Research on Economic Inequality*, 109–152.
- CAPASSO, M., E. CEFIS, AND K. FRENKEN (2013): “On the existence of persistently outperforming firms,” *Industrial and Corporate Change*, forthcoming.
- COAD, A. (2007): “A Closer Look at Serial Growth Rate Correlation,” *Review of Industrial Organization*, 31, 69–82.
- (2009): *The Growth of Firms: A Survey of Theories and Empirical Evidence*, New Perspectives on the Modern Corporation, Edward Elgar, Cheltenham, UK.
- COAD, A., S.-O. DAUNFELDT, W. HÖLZL, D. JOHANSSON, AND P. NIGHTINGALE (2014): “High-growth firms: introduction to the special section,” *Industrial and Corporate Change*, 23, 91–112.
- COAD, A. AND W. HÖLZL (2009): “On the autocorrelation of growth rates,” *Journal of Industry, Competition and Trade*, 9, 139–166.
- DAUNFELDT, S.-O. AND D. HALVARSSON (2015): “Are high-growth firms one-hit wonders? Evidence from Sweden,” *Small Business Economics*, 44, 361–383.
- DOSI, G. (2007): “Statistical Regularities in the Evolution of Industries. A Guide through some Evidence and Challenges for the Theory,” in *Perspectives on Innovation*, ed. by F. Malerba and S. Brusoni, Cambridge University Press.
- DOSI, G., R. R. NELSON, AND S. WINTER (2000): “Introduction,” in *The nature and dynamics of organizational capabilities*, ed. by G. Dosi, R. R. Nelson, and S. Winter, Oxford University Press: Oxford.
- DUNNE, T., M. J. ROBERTS, AND L. SAMUELSON (1988): “Patterns of Firm Entry and Exit in U.S. Manufacturing Industries,” *Rand Journal of Economics*, 19, 495–515.
- EVANS, D. S. (1987a): “The Relationship between Firm Growth, Size, and Age: Estimates for 100 Manufacturing Industries,” *Journal of Industrial Economics*, 35, 567–81.
- (1987b): “Tests of Alternative Theories of Firm Growth,” *The Journal of Political Economy*, 95, 657–674.
- GABAIX, X. (2011): “The Granular Origins of Aggregate Fluctuations,” *Econometrica*, 79, 733–772.
- GEROSKI, P. A. (2002): “The Growth of Firms in Theory and in Practice,” in *Competence, Governance, and Entrepreneurship - Advances in Economic Strategy Research*, ed. by N. Foss and V. Mahnke, Oxford University Press: Oxford and New York.
- GEROSKI, P. A. AND M. MAZZUCATO (2002): “Learning and the sources of corporate growth,” *Industrial and Corporate Change*, 11, 623–644.

- GIBRAT, R. (1931): *Les inégalités économiques*, Librairie du Recueil Sirey, Paris.
- GODDARD, J., P. MOLYNEUX, AND J. WILSON (2004): “Dynamics of Growth and Profitability in Banking,” *Journal of Money, Credit, and Banking*, 36, 1069–1090.
- GODDARD, J., J. WILSON, AND P. BLANDON (2002): “Panel tests of Gibrats Law for Japanese manufacturing,” *International Journal of Industrial Organization*, 20, 415 – 433.
- GRILICHES, Z. AND J. MAIRESSE (1999): “Production Functions: The Search for Identification,” in *Econometrics and Economic Theory in the Twentieth Century: the Ragnar Frisch Centennial Symposium*, ed. by S. Steiner, Cambridge University Press: Cambridge.
- GUARASCIO, D. AND F. TAMAGNI (2019): “Persistence of innovation and patterns of firm growth,” *Research Policy*, 48, 1493–1512.
- HALL, B. H. (1987): “The Relationship Between Firm Size and Firm Growth in the Us Manufacturing Sector,” *Journal of Industrial Economics*, 35, 583–606.
- HÖLZL, W. (2014): “Persistence, survival, and growth: a closer look at 20 years of fast-growing firms in Austria,” *Industrial and corporate change*, 23, 199–231.
- IJIRI, Y. AND H. A. SIMON (1977): *Skew distributions and the sizes of business firms*, North-Holland, Amsterdam.
- KNUDSEN, T., D. A. LEVINTHAL, AND S. G. WINTER (2017): “Systematic differences and random rates: Reconciling Gibrats law with firm differences,” *Strategy Science*, 2, 111–120.
- KUMAR, M. S. (1985): “Growth, Acquisition Activity and Firm Size: Evidence from the United Kingdom,” *Journal of Industrial Economics*, 33, 327–338.
- LOTTI, F., E. SANTARELLI, AND M. VIVARELLI (2003): “Does Gibrat’s Law hold among young, small firms?” *Journal of Evolutionary Economics*, 13, 213–235.
- MANSFIELD, E. (1962): “Entry, Gibrat’s Law, Innovation, and the Growth of Firms,” *The American Economic Review*, 52, 1023–1051.
- MOSCHELLA, D., F. TAMAGNI, AND X. YU (2019): “Persistent high-growth firms in China’s manufacturing,” *Small Business Economics*, 52, 573–594.
- PENROSE, E. T. (1959): *The theory of the growth of the firm*, Oxford: Blackwell, 3rd ed.
- PISANO, G. P. (2017): “Toward a prescriptive theory of dynamic capabilities: connecting strategic choice, learning, and competition,” *Industrial and Corporate Change*, 26, 747–762.
- SHORROCKS, A. F. (1978): “The Measurement of Mobility,” *Econometrica*, 46, 1013–24.

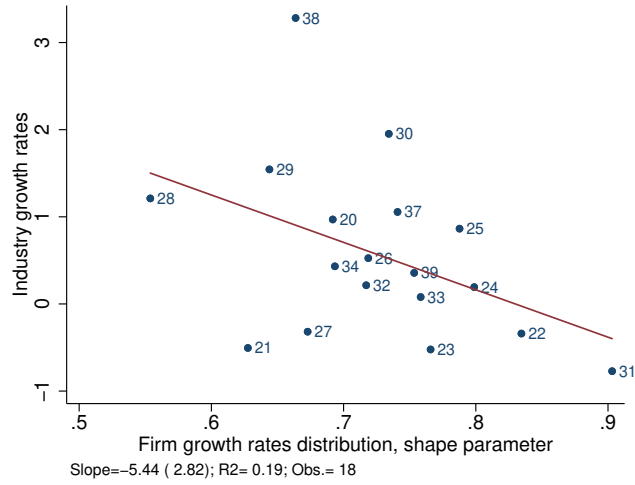
- SIMON, H. A. AND P. BONINI, CHARLES (1958): “The Size Distribution of Business Firms,” *American Economic Review*, 48, 607–617.
- STANLEY, M., L. AMARAL, S. BULDYREV, S. HAVLIN, H. LESCHHORN, P. MAASS, M. SALINGER, AND H. STANLEY (1996): “Scaling behaviour in the growth of companies,” *Nature*, 379, 804–806.
- STEINDL, J. (1965): *Random Processes and the Growth of Firms*, London: Charles Griffin & Company.
- SUBBOTIN, M. F. (1923): “On the Law of Frequency of Errors,” *Matematicheskii Sbornik*, 31, 296–301.
- SUTTON, J. (1997): “Gibrat’s Legacy,” *Journal of Economic Literature*, 35, 40–59.
- TEECE, D. J., G. PISANO, AND A. SHUEN (1997): “Dynamic Capabilities and Strategic Management,” *Strategic Management Journal*, 18, 509–533.
- VENNET, R. V. (2001): “The law of proportionate effect and OECD bank sectors,” *Applied Economics*, 33, 539–546.
- WILLIAMS, M. A., B. P. PINTO, AND D. PARK (2015): “Global evidence on the distribution of firm growth rates,” *Physica A: Statistical Mechanics and its Applications*, 432, 102–107.
- WINTER, S. G. (1987): “Knowledge and Competences as Strategic Assets,” in *The Competitive Challenge: Strategy and Organization for Industrial Innovation and Renewal*, ed. by D. J. Teece, Ballinger: Cambridge, MA.

Appendix A - Plots without sectors 35 and 36




(a)

Figure 14: Cumulated industry sales growth rate over the period 1970-2011 (from NBER-CES data) *vs.* interdecile range of the yearly firm sales growth rates (Compustat firms, pooled 1970-2011) . Sectors 35 and 36 excluded. Industry classification is SIC 1987. Slope (standard error), R^2 and number of observations are reported below the graph




(a)

Figure 15: Cumulated industry sales growth rate over the period 1970-2011 (from NBER-CES data) *vs.* shape parameter b of the Subbotin distribution fit of yearly firm sales growth rates (Compustat firms, pooled 1970-2011). Sectors 35 and 36 excluded. Slope (standard error), R^2 and number of observations are reported below the graph



(a)



(b)

Figure 16: Cumulated industry sales growth rate over the period 1970-2011 (from NBER-CES data) *vs.* shape parameter b_l and b_r of the AEP fit of yearly firm sales growth rates (Compustat firms, pooled 1970-2011). Sectors 35 and 36 excluded. Slope (standard error), R^2 and number of observations are reported below the graph

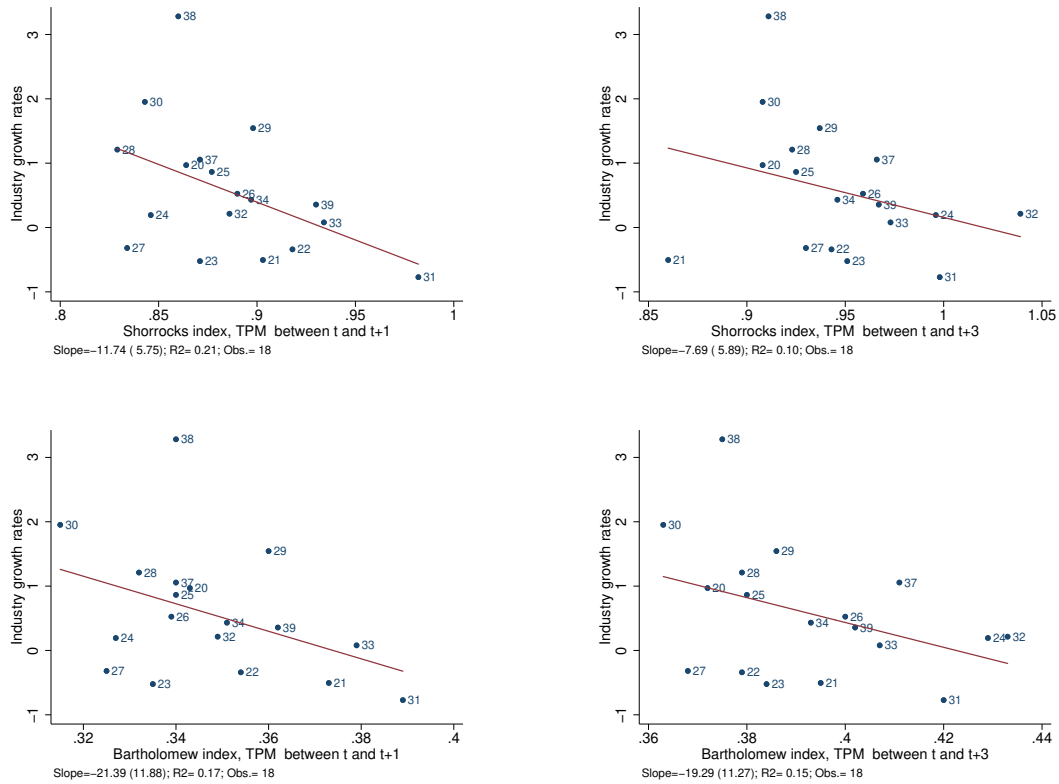


Figure 17: Industry sales growth rates over the period 1970-2011 *vs.* Shorrocks and Bartholomew index of mobility calculated from sectoral Transition Probability Matrix. Figure (a) and (c) consider TPM between t and $t + 1$, Figure (b) and (d) consider TPM between t and $t + 3$. Both indexes are increasing in mobility, meaning that for high value of the index there is less persistence. Sectors 35 and 36 excluded. Slope (standard error), R^2 and number of observations are reported below each graph

Appendix B - Sectors

Table 4: Sectors

SIC 2-digits	Description
20	Food And Kindred Products
21	Tobacco Products
22	Textile Mill Products
23	Apparel And Other Finished Products Made From Fabrics And Similar Materials
24	Lumber And Wood Products, Except Furniture
25	Furniture And Fixtures
26	Paper And Allied Products
27	Printing, Publishing, And Allied Industries
28	Chemicals And Allied Products
2834	Pharmaceutical preparations
29	Petroleum Refining And Related Industries
30	Rubber And Miscellaneous Plastics Products
31	Leather And Leather Products
32	Stone, Clay, Glass, And Concrete Products
33	Primary Metal Industries
34	Fabricated Metal Products, Except Machinery And Transportation Equipment
35	Industrial And Commercial Machinery And Computer Equipment
357	Computer and office equipment
36	Electronic And Other Electrical Equipment And Components, Except Computer Equipment
3674	Semiconductors and related devices
37	Transportation Equipment
38	Measuring, Analyzing, And Controlling Instruments; Photographic, Medical And Optical Goods; Watches And Clocks
39	Miscellaneous Manufacturing Industries