Industrial and Corporate Change, Volume 11, Number 4, pp. 705–723

Corporate growth and industrial structures: some evidence from the Italian manufacturing industry

Giulio Bottazzi, Elena Cefis and Giovanni Dosi

This work analyses the properties of corporate growth in a large longitudinal sample of Italian manufacturing firms. In particular, it focuses on the statistical properties of growth rates and on the influence of proxies for relative efficiency upon relative growth. In line with previous work, the emergence of 'fat tails' in growth rate distributions and the idiosyncratic nature of autocorrelation co-efficients confirm the existence of a structure in the growth process richer than the one normally assumed by the 'Gibrat Law' hypothesis and suggest the presence of firm-specific drivers of growth. At the same time, there is a remarkable puzzle concerning the absence of any negative relationship between size and growth variance and only weak influences of relative efficiencies upon growth dynamics.

1. Introduction

In this work we report the preliminary results of an investigation on industrial dynamics based on a decade of micro-longitudinal data from four Italian industries—pharmaceuticals, primary metals, machine tools and textiles—chosen as representative of quite diverse production technologies and learning modes. Here we begin addressing two sets of issues concerning (i) the shape of the size distributions and their possible inter-sectoral differences, and (ii) the characteristics of growth dynamics.

A classic reference, when dealing with the statistical properties of firm growth, is the so-called 'Law of Proportionate Effect' (or 'Gibrat's Law') (Gibrat, 1931), entailing processes of stochastic growth uncorrelated with size and basically driven by several small idiosyncratic events. If $x_i(t)$ stands for the logarithm of firm i size[1] at time t, according to this law its size at time $t+1$ reads[2]

[1]Where size can be measured with respect to some 'extensive' variable such as total sales, value added or number of employees. In the Gibrat literature the problem of what variables to choose has seldom been discussed and, when done, it has been mainly to affirm the irrelevance of such a choice (Stanley *et al.*, 1997, among others). As we shall see below, this might be a misleading assumption.

[2]For discussions, following the pioneering Ijiri and Simon (1977), cf., among others, Brock and Evans (1986), Boeri (1989), Sutton (1997), Geroski (2000), Dosi *et al.* (1995), Marsili (2001) and Cefis *et al.* (2001).

$$x_i(t+1) = a + bx_i(t) + \varepsilon_i(t) \quad (1)$$

where a is an industry-wide drift, b is an auto-regressive component and ε are random variables independent from x. Equation (1), under the restriction of $b = 1$, can be considered as a sort of 'null hypothesis' regarding firm dynamics. Note that it is also an hypothesis that makes evolutionary economists rather uncomfortable, in that it seems at odds both with several pieces of microeconomic evidence highlighting long-standing differences in technological and organizational competences across firms, and also with a notion of a competitive process systematically selecting within such a population of heterogeneous firms.[3]

Moreover, the overall evidence for 'Gibrat's Law' from the literature, often based on inadequate data, is rather mixed, and the analyses often simply amount to testing the statistical acceptability of the $b = 1$ restriction via the estimation of a first-order autoregressive process AR(1) on the whole panel data (for reviews and discussions, see Dosi *et al.*, 1995; Geroski, 2000; Sutton, 1997). As we argue at greater length elsewhere (Dosi *et al.*, 1995; Bottazzi *et al.*, 2001; Bottazzi, 2000; Cefis *et al.*, 2001) such an approach is likely to fall well short of identifying the possible underlying structures in the growth process. An alternative, proposed by Bottazzi *et al.* (2001), involves testing for: (i) 'fat tails' in the distribution of growth shocks with (relatively rare) 'spurs of growth'; (ii) possible autocorrelations of growth rates over time; and (iii) firm-specificities in growth patterns that are persistent over time (as suggested by Cefis *et al.*, 2001). Such properties do indeed emerge in the case of the world's leading pharmaceutical firms (Bottazzi *et al.*, 2001; Cefis *et al.*, 2001). However, an obvious issue regards the degrees of generality of such findings. Are the foregoing properties dependent upon the particular features of learning and competition of the drugs industry or, conversely, are they rather general characteristics of industrial dynamics? And, even if the latter hypothesis held true, to what extent are such characteristics influenced by industry-specific factors? We shall address these questions in the following.

In addition, size as such might not be the best variable upon which to condition growth events. Rather, it is much more in tune with an evolutionary idea of heterogeneity-cum-market selection[4] to search for proxies of relative degrees of firm 'competitiveness' and investigate their impact on firm growth profiles. This is what we shall also do below, using labor productivities as proxies for production efficiencies.

In Section 2 we briefly describe the database and the variables under scrutiny. Section 3 discusses the evidence on size distributions, the distribution of growth shocks and their possible autocorrelation. Section 4 considers the relationship between firm

[3]Incidentally, note that violations of Gibrat-type process of growth based on i.i.d. shocks are also implied by equilibrium models of industrial dynamics such as Jovanovic (1982) and Ericson and Pakes (1995): cf. Pakes and Ericson (1998).

[4]Within a rapidly expanding evolutionary literature on industrial dynamics let us just mention three of the 'seeding classics', namely Winter (1971, 1984), and Nelson and Winter (1982).

size and growth variances. In Section 5, we analyze relative labor productivities and their dynamics, while in Section 6 we study their relationships with growth profiles.

2. The database

This research draws upon the MICRO.1 databank developed by the Italian Statistical Office (ISTAT).[5] MICRO.1 contains longitudinal data over around a decade on a panel of several thousand Italian firms that employ 20 or more people; for statistical consistency we utilize the period 1989–1996.

In this work we are exclusively interested in the process of *internal* growth, as opposed to the growth due to mergers, acquisitions and divestments. In order to control for the latter we build 'super-firms' which account throughout the period for the union of the entities which undertake such changes. So, for example, if two firms merged at some time, we consider them merged throughout the whole period. Conversely, if a firm is spun off from another one, we 're-merge' them starting from the separation period.

Moreover, since the panel is open—due to entry, exit, fluctuations around the 20-employee threshold and variability in response rates—we consider only those firms that are present both at the beginning and at the end of our window of observation.

Firms are classified according to their sector of principal activity.[6] For the analysis that follows, as already mentioned, we have chosen pharmaceuticals,[7] primary metals,[8] machine tools[9] and textiles,[10] which can be reasonably taken as representative of Pavitt's taxonomic classes identified as 'science-based', 'scale-intensive', 'specialized supply' and 'supplier dominated', respectively (Pavitt, 1984). Here we consider Pavitt's categories as an attempt to classify different industrial sectors according to the diverse modes of generation and exploitation of novel opportunities of product and process innovation (cf. also Dosi, 1988; Marsili, 2001). In turn, diverse regimes of technological learning might well influence growth dynamics.

The statistical variables we consider here are the total number of employees $L_i(t)$, sales $S_i(t)$ and value added $V_i(t)$ of 'super-firm' i at time $t \in [1, \ldots, 8]$, together with labor productivity defined as $\Pi_i(t) = V_i(t)/L_i(t)$.

It is often convenient to analyze the normalized logarithm of those variables. For instance, regarding the number of employees we define

[5]The database has been made available to our team under the mandatory condition of censorship of any individual information.

[6]The Italian ATECO.3 classification closely matches the ISIC one.

[7]ATECO.3: 24.4 pharmaceuticals; 97 observations.

[8]ATECO.3: 27.1 ferrous and non-ferrous metals; 67 observations.

[9]ATECO.3: 29.4 machine tools; 114 observations.

[10]ATECO.3: 17.2 textiles; 171 observations.

$$l_i(t) = \log[L_i(t)] - < \log[L_i(t)] >_i \tag{2}$$

where $<\cdot>_i$ stands for the average over all the firms at a given time. Analogously, we define 'rescaled' log sales $s_i(t)$ and log value added $v_i(t)$. These variables are characterized by stationary distributions[11] and allow us to treat the growth process on these normalized quantities as a stationary one. Let us denote the various growth rates as

$$g_i^x(t) = x_i(t+1) = x_i(t) \tag{3}$$

where x takes the values l, s and v respectively for the number of employees, sales and value added.

Note that through this 'rescaling' procedure one washes away common trend effects due to both inflationary dynamics and real (i.e. constant price) expansion/contraction of the industry as a whole (including those captured by a in equation 1).

3. Size distributions and corporate growth

3.1 Size distributions

Due to the relatively low number of observations, it is safer to plot the distribution function rather than the probability densities. It is also handy to refer to a 'symmetric transformation' of such a distribution function. In what follows (for purposes of clarity) we will use the 'symmetrized' version of the distribution function $F(x)$ defined according to

$$F_s(x) = \begin{cases} F(x) & F(x) < 0.5 \\ 1 - F(x) & F(x) > 0.5 \end{cases} \tag{4}$$

(Under this convention, in what follows we drop the subscript s for convenience.)

In Figures 1 and 2 we show the upper and lower tails of the ('symmetrized') probability density. These figures show important differences both in tails and supports. The latter variable suggests different 'spreads' of sizes in different industries. Concerning the former, notice that higher tails means more concentrated industries, i.e. the share of the total industry possessed by the top fraction of firm population is higher when the probability density is asymptotically higher.

First, notice the striking intersectoral differences in our proxy for concentration, well in tune with an established 'stylized fact' from industrial economics. Interestingly, the representatives of 'science-based' and 'scale-intensive' sectors (pharmaceuticals and primary metals respectively) are more concentrated and display smaller support, i.e. relatively lower size asymmetries. The converse holds for 'specialized suppliers' and 'supplier-dominated' sectors (machine tools and textiles).

[11]We have checked the stationarity hypothesis using Kolmogorov–Smirnov tests and we find robust evidence supporting it: the significance is always >0.96.

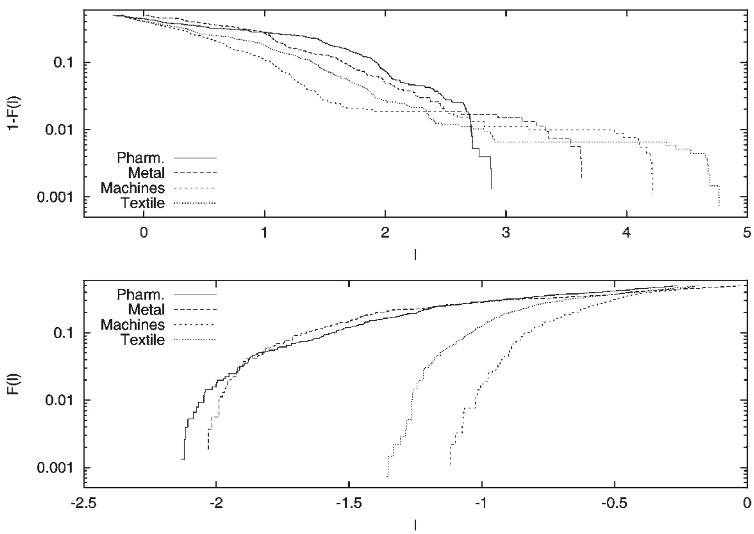

Figure 1 Upper (top) and lower (bottom) half of the size distribution function in terms of number of employees in the four sectors (computed using the whole database time horizon).

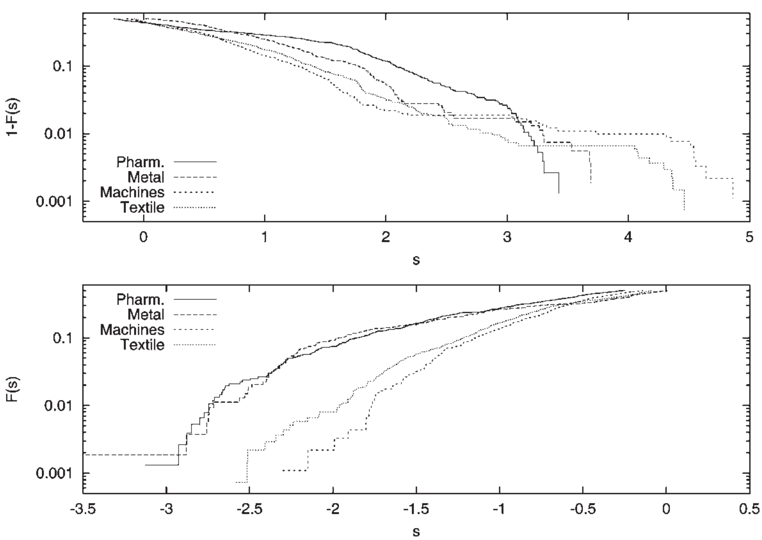

Figure 2 Upper (top) and lower (bottom) half of the size distribution function in terms of sales in the four sectors (computed using the whole database time horizon).

Second, the upper tails also show (with the exception of pharmaceuticals) some large gaps which can be intuitively interpreted as a sort of 'barrier' separating different segments of the industry, i.e. a core part from a fringe one. Indeed the large width of these gaps, compared to the average size of growth shocks (see also below), implies that the large majority of micro-dynamics develops separately inside the different segments, with rare crossing events.

Third, the lower tails appear to be more homogeneous across sectors (but we are also less confident about making any inference on a tail 'artificially bent' by a sampling threshold, further burdened by proportionally more frequent missing observations due to rather noisy response rates by smaller firms).

Fourth, the slope of the upper tails for all sectors tends to fall in the middle-to-high size range so that the curve takes a convex shape (this is in fact analogous to what happens on Italian data in Pareto fit to the top firms[12]) and the power-like behavior (i.e. a linear behavior in log terms) seems to be interrupted by a sudden decrease.

Incidentally notice also some of the further interpretative questions inspired by this evidence: among them, to what extent are these patterns influenced by the institutional specificities of the Italian case? And, conversely, how robust are inter-sectoral differences which hold for the same sectors across different countries?

3.2 Growth dynamics

In order to characterize the growth process, let us begin by checking if any relationship between size and growth is present in our data. Interestingly, both the growth means and growth variances do not display any relationship with size.[13] This circumstantial evidence for the weaker form of the 'Law of Proportionate Effect', prescribing the lack of any relationship between growth and size, appears at work here.

However, consider as a benchmark for the dynamics a 'stronger' Gibrat hypothesis, whereby growth shocks should be well described by a log-normal distribution,[14] and compare it with the actual distribution of g^l and g^s shown in Figures 3 and 4 respectively.

The plots clearly show how a log-normal distribution dramatically underestimates the 'fatness' of the observed tails.

Let us try then to fit the data using a more fat-tailed distribution. In order to do that we use the Subbotin family of distributions (Subbotin, 1923) with density of the form:[15]

---

[12]Cf. Dosi *et al.* (2000).

[13]The lack of relationship concerning growth means is a robust result which has been found many times elsewhere (cf. the evidence discussed in Sutton, 1997; Geroski, 2000) (see also Section 4). Conversely, the presence of a negative relationship between size and growth variances, which constitutes a quite typical feature of industrial data (cf. the discussion in Bottazzi, 2000), is not displayed by our data.

[14]This is indeed a straightforward conjecture, under the Central Limit Theorem, once the idea of growth as a sequence of random shocks is accepted on every time scale.

[15]This in turn generalizes on a similar procedure used by Stanley *et al.* (1997) where a Laplace distribution is used. The relevance of our generalization can be checked by looking at the fitted b exponents that are constantly lower then unity (i.e. the property of Laplace distributions: see also below).

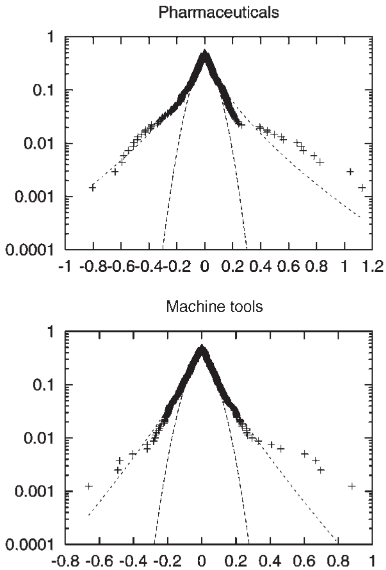

Figure 3 Probability distributions for the (log) labor growth g^l in the four sectors. Broken lines show fits for a normal distribution (lower one) and a Subbotin distribution (upper one). (For the parameters of the latter, see Table 1.)

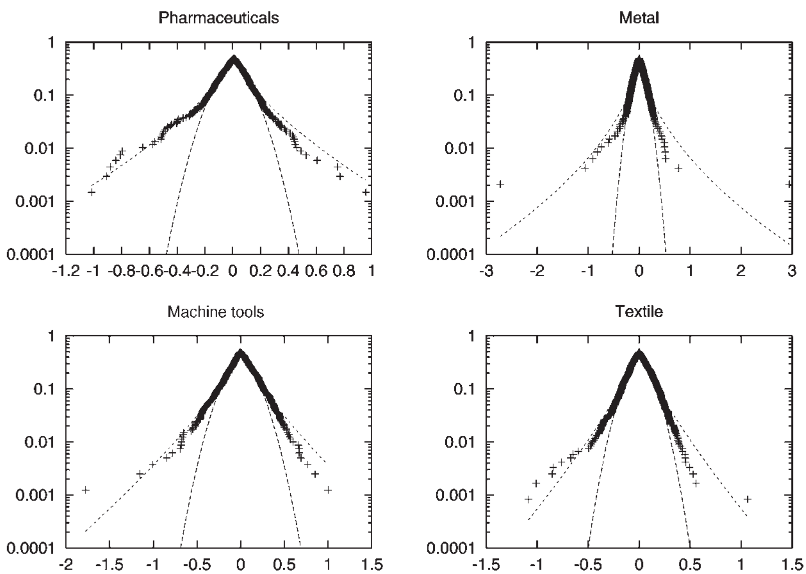

Figure 4 Probability distributions for the (log) sales growth g^s in the four sectors. Again, a normal (lower) and a Subbotin exponential (upper) fits are also shown. (For the parameters of the latter, see Table 2.)

$$f(x) = \frac{1}{2} \frac{\beta \alpha^{1/\beta}}{\Gamma(1/\beta)} e^{-\alpha |x|^{\beta}} \tag{5}$$

Here $\Gamma(x)$ is the gamma function and β represents a 'shape parameter', shaping the distribution tails: the distribution is leptokurtic for $\beta < 2$ and platikurtic for $\beta > 2$. The lower is β, the fatter are the tails. The 'scale' parameter α describes the central width of the distribution. The $2l$-th central moment of the Subbotin distribution reads

$$m_{2l} = \alpha^{\frac{-2l}{\beta}} \frac{\Gamma[(2l+1)/\beta]}{\Gamma(1/\beta)} \tag{6}$$

implying that the rescaled central moments (such as the kurtosis) do not depend on the parameter α. For $\beta = 2$ this distribution reduces to a Gaussian, and for $\beta = 1$ to a symmetric exponential, i.e. Laplace, distribution.

To fit observed data we use the associated probability distribution function:

$$F(x) = \begin{cases} \frac{1}{2}\left[1 - P\left(1/\beta, \alpha |x|^{\beta}\right)\right] & x < 0 \\ \frac{1}{2}\left[1 + P\left(1/\beta, \alpha x^{\beta}\right)\right] & x > 0 \end{cases} \tag{7}$$

where $P(a,x)$ stands for the incomplete gamma function:

$$P(a,x) = \frac{1}{\Gamma(a)} \int_0^x e^{-t} t^{a-1} \mathrm{d}t \tag{8}$$

In general the distribution (7) provides a good description of the observed frequencies over a wide range of values. Interestingly, the major drawback comes from a remarkable asymmetry of the growth distribution between positive and negative parts, both for the sales and the employees variables, at least in some of our sectors.

In Tables 1 and 2 we report the result of a least-square fitting procedure of (7) to the observed frequencies.

3.3 Autocorrelation

Another major question concerns the presence of autocorrelation in the growth dynamics of firms. Hence, we compute for each variable (employees, sales and value added), the histogram of the autocorrelation coefficients for all the firms in a given sector. The mean of this distribution represents the sampled autocorrelation computed using all the firms from the panel. Under the assumption that different firm histories were to represent different realizations of the same random process, this should indeed be the best estimate of the autocorrelation in the overall growth process. In fact, as reported in Tables 3 and 4, the means are in general close to zero and, even when their differences from zero are marginally significant, their small values (around 0.01)

Table 1 Parameters from least-squares fitting of the g^l distribution with a Subbotin function

	Pharmaceuticals	Metals	Machine tools	Textiles
α	8.27	9.03	10.99	10.37
β	0.58	0.39	0.77	0.49
σ^2	0.0234	0.02461	0.01161	0.01039

Table 2 Parameters from least-squares fitting of the g^s distribution with a Subbotin function

	Pharmaceuticals	Metals	Machine tools	Textiles
α	6.83	6.62	5.7	7.55
β	0.62	0.46	0.73	0.76
σ^2	0.0429	0.0743	0.0658	0.0307

Table 3 Mean and SD of the distribution of autocorrelation coefficients for labor growth for the four sectors. The significance of the Kolmogorov–Smirnov comparison test between the observed distributions and the distributions obtained with randomly resampled (bootstrapped) growth shocks is also shown

	Pharmaceuticals	Metals	Machine tools	Textiles
Mean	0.0789	0.093	0.095	0.123
σ	0.320	0.327	0.306	0.351
Significance (*P*)	0.0085	0.00034	9.511010^{-5}	4.6710^{-8}

Table 4 Mean and SD of the distribution of the autocorrelation coefficients for sales growth in the four sectors (cf. Table 3)

	Pharmaceuticals	Metals	Machine tools	Textiles
Mean	0.085	–0.016	–0.066	–0.124
σ	0.327	0.284	0.305	0.351
Significance (P)	0.0042	0.815	0.029	4.6710^{-8}

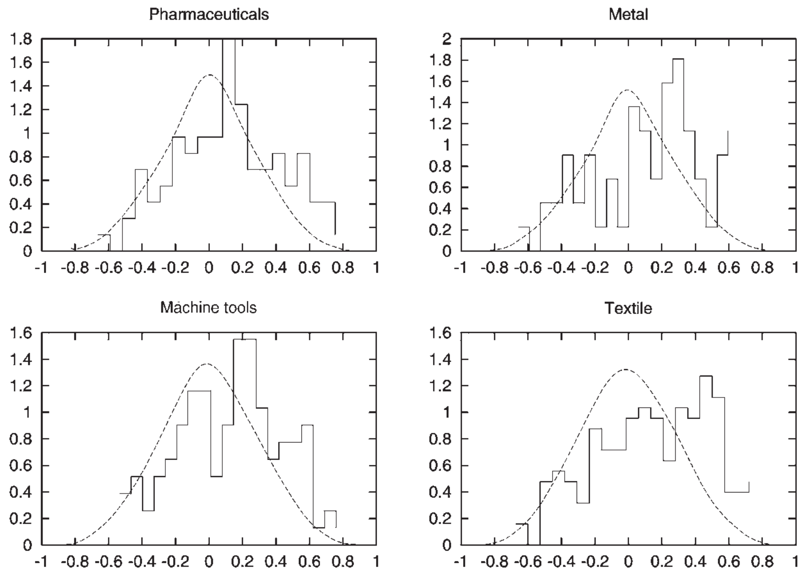

Figure 5 Observed frequency for the autocorrelation coefficient of g^l growth (steps function) and the associated density distribution computed using 10 000 bootstrapped time series (dotted line). See Table 3 for the mean values and the results of a Kolmogorov–Smirnov comparison between observed and bootstrapped distributions.

cannot, *prima facie*, suggest any remarkable, persistent difference in firms' growth profiles, at least over the relatively short time horizon that characterizes our database.

However, the assumption of 'identity' among firms turns out to be rather questionable. In order to check to what extent different firms' dynamics can be treated as the outcome of the same underlying process, one may compare the observed distributions of frequencies, shown in Figures 5 and 6 for the labor and sales variables respectively, with the distributions obtained from a dataset made of 'artificial firms' histories' that satisfy this identity requirement by construction. These histories can be obtained by a 'bootstrap sampling', i.e. by randomly extracting 'growth rates' from the set of all the observed growth rates.

If one then computes again the autocorrelation distribution on this 'artificial dataset', a different shape is obtained (cf. Figures 5 and 6). The difference between the two is revealed by performing a Kolmogorov–Smirnov test comparison between the 'artificial' and the observed distributions, and by looking at the obtained significance of P-values (i.e. the probability the observed differences between the distributions might be simply a matter of chance). As can be seen in Tables 3 and 4, the P-value is in many cases so low that it leads to a clear rejection of the 'identity' hypothesis between the growth processes of different firms.

Here, again, the evidence is circumstantial, but it is surprisingly well in tune with the findings from Cefis *et al.* (2001) hinting at powerful idiosyncratic patterns of growth.

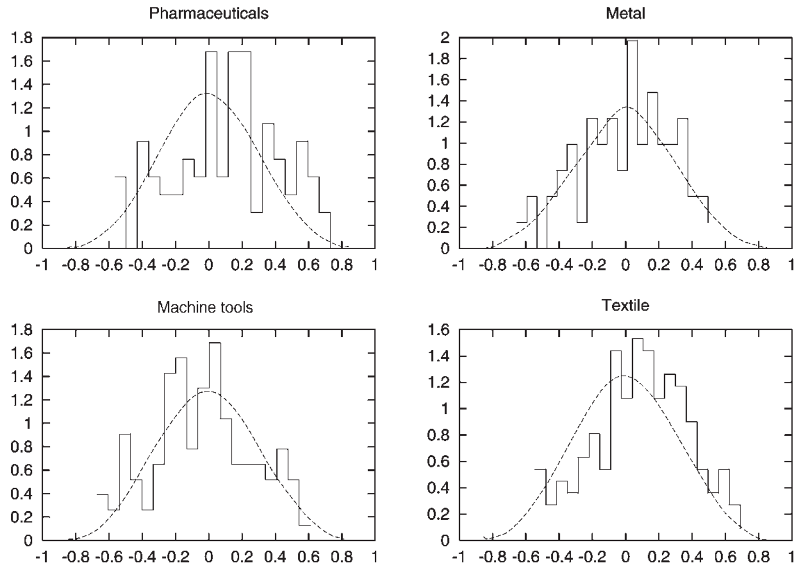

Figure 6 Observed frequency for the autocorrelation coefficient of g^s growth (steps function) and the associated density distribution computed using 10 000 bootstrapped time series (dotted line). See Table 4 for the mean values and the results of a Kolmogorov–Smirnov comparison between observed and bootstrapped distributions.

The promising conjecture is that growth dynamics are persistently asymmetric across firms, that firm-specific processes display a long memory, and that, together, we are still at a preliminary stage in identifying the underlying (technological and organizational) conditioning factors.

4. Corporate sizes and growth variances

As already mentioned, there is robust evidence to suggest that variance of growth rates falls with corporate size (cf. e.g. Evans, 1987a,b; Hall, 1987). Our data on a subset of Italian industries covering a comprehensive sample of firms with more than 20 employees conflict with such a common wisdom. As shown in Figure 7 for the number of employees, no such pattern appears.[16]

Bottazzi (2000) proposed an explanation for the negative variance–size relationship grounded in diversification patterns (for a similar interpretation on US manufacturing industry, see Stanley *et al.*, 1997).

In brief, Bottazzi (2000) and Bottazzi *et al.* (2001) show that the number of markets in which a firm diversifies bears a (less than proportional) relation to size and that the underlying dynamics is a (plausibly, competence-driven) branching process. In turn

[16]The analysis using sales gives very similar results.

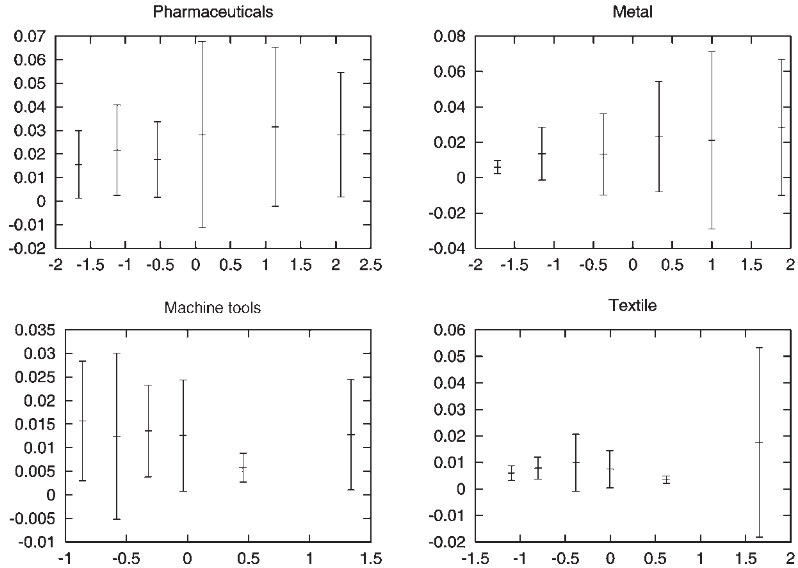

Figure 7 The average rate of growth of the number of employees g^l for equipopulated bins of firms partitioned according to the number of employees l. The error bars correspond to 3SD.

diversification across (uncorrelated) markets fully explains the observed coefficients of the negative relation between growth variance and size.

5. Labor productivities

Recall the definition, given in Section 2, of Π (the value added per employee) and of the rescaled (log) variable π, as such a proxy for relative labor productivities. Figure 8 presents the distributions of such quantities by sector. (Given the stationarity of the distributions over the considered time period, we pool all yearly observations together.)

First, note that an implication of the observed stationarity is the lack of any reduction in the distribution variances over time, suggesting a persistence in the micro-heterogeneity, prominently shown by the wideness of the distribution supports. Indeed, this strongly corroborates a central evolutionary conjecture on the inertial reproduction over time of *diverse capabilities* and related diverse performances, favored also by the general difficulties of ('boundedly rational') economic agents in learning new technological and organizational practices, and even in identifying the notionally most promising ones.

Second, again, fat tails reveal the systematic presence of a relatively large number of 'outperformers' and 'underperformers'—as compared to a normal distribution benchmark.

Finally, as suggested by the positive skewness of the distributions, the observations concerning the highest productivities are further away from the distribution averages, as compared to the lowest productivity observations.

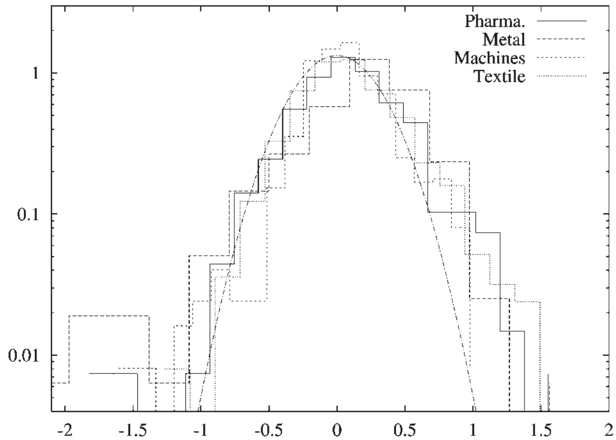

Figure 8 Probability densities of the productivity π for the four sectors under analysis (the normal distribution is plotted as visual guide).

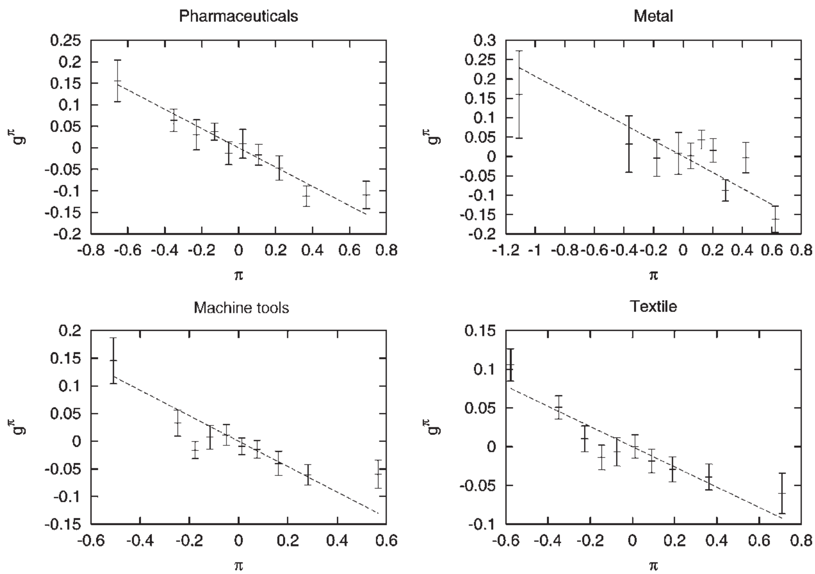

Figure 9 Regression of the productivity growth g^{π} on the average productivity π. The data are distributed according to the latter in 10 equipopulated bins. The regression parameters are reported in Table 5.

Given such distributions, what are the dynamics of productivity over time?

In Figure 9 we show the average productivity growth for different productivity bins.[17] An inverse relationship emerges, where more productive firms are on average

[17]'Bin' stands for a quantile in the distribution of the population in the variable at hand.

Table 5 The slope *r* and the asymptotic standard error as obtained with an OLS linear regression of the productivity growth versus actual productivity

	Pharmaceuticals	Metals	Machine tools	Textiles
r	–0.013	–0.23	–0.2	–0.22
σ_r	0.015	0.023	0.031	0.025
r/σ_r (%)	12	10	15	11

doomed to see their relative productivities decreasing relative to the industry average the following year. This is of course consistent with some process of learning and imitation amongst firms leading to the rapid diffusion of capabilities through the industry: the 'catching up' abilities of the technological followers appears to wash away relatively quickly positions of (temporary) leadership in production efficiency. However, no systematic reversion to the mean tendency emerges: distributions of relative productivities are stable over time. This basic evidence is corroborated by statistics (not shown here) concerning growth rates of productivity over the whole period against initial relative productivities: mild evidence of some catching-up tendency can be seen, but not enough to yield an increasing uniformity among firms.

6. Relative efficiency and corporate growth

A fundamental hypothesis of evolutionary theories is that differential levels of 'competitiveness' (or 'fitness') systematically affect relative growth rates of micro-entities.

Let us plot employees and sales growth for different bins over the π distribution. Results for the employees are shown in Figure 10. Here, a clear positive relationship appears, supported also by exercise of linear regression of the employment growth versus relative productivities, whose estimated coefficients are reported in Table 6. The positive relationship appears quite robust and, one should note, rather homogeneous for the different sectors.

Such a dynamic is well in tune with a 'replicator-type' process of market selection whereby, in probability, firms with above-average productivity tend to expand and those below average tend to shrink. However, in our data, we observe that this relationship disappears when considering firm growth as measured in terms of sales or value added rather than employees.[18] Even more puzzling, such a relation tends to disappear when considering long-run relationships (i.e. growth measured on larger time spans) between relative productivities on the one hand, and any growth indicator (e.g. employment, value added or sales), on the other.

[18]This is indeed a puzzle that we intend to explore in future works.

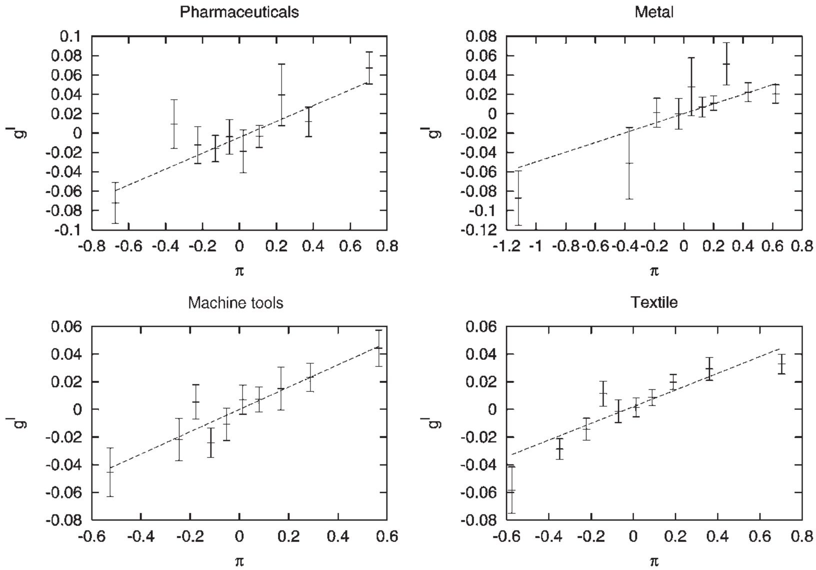

Figure 10 Regression of employee growth g^l against (relative) productivity π. The data are distributed according to the latter in 10 equipopulated bins. The regression parameters are reported in Table 6.

Table 6 The slope r and the asymptotic standard error as obtained with an OLS linear regression of employee growth versus labor productivity

	Pharmaceuticals	Metals	Machine tools	Textiles
r	0.081	0.049	0.08	0.06
σ_r	0.015	0.012	0.012	0.01
r/σ_r (%)	18	25	15	16

Finally, a prominent phenomenon highlighted by all our evidence is the role of *outliers*, i.e. by the presence of few remarkable *outperformers* and few remarkable *underperformers* which systematically appear and have non-negligible impact on the sector dynamics. This applies to cumulative productivity growth, to systematic growth in proxies of size (labor and value added), and also to the relationship between far-from-average productivity growth and far-from-average growth proxies. One is tempted here to think that most 'near-average' differences in our admittedly very noisy proxy for 'competitiveness' also pick up many factors of non-price competition, and, together, many roughly *neutral drifts* in technologies, organizational arrangements and strategies. Together, a few 'hopeful monsters'—in a biological analogy—stand out above the noise involved in our accounting proxies and drive systemic changes in productivity and market shares. At the opposite extreme, market selection seems to operate quite gently, if at all, *vis-à-vis* most 'near-average' agents. Its role, it seems, is mainly to cut out the very worse performers.

7. Conclusions

As already mentioned, this work is a preliminary study within a wider search of the statistical regularities of industrial dynamics. As such, it suggests both relatively robust insights into the nature of the underlying evolutionary process and also some intriguing challenges. In these conclusions let us mainly focus on the latter.

Let us start from the puzzling property of our data, which, in tune with a lot of the evidence reviewed in Geroski (2000), lack any strong autocorrelation in the growth process.

This is particularly puzzling since also in our data one finds abundant evidence of systematic heterogeneity across firms. First, as discussed above, we find at least circumstantial evidence of differences across firms in the generating processes of growth shocks. Second, and even more importantly, our data display striking persistent differences across firms in production efficiencies.

Why shouldn't these asymmetries in efficiency be reflected in more systematic selection processes autocorrelated over time?

Part of the answer might rest in the differences in the time scales at which productivity shocks arrive *vis-à-vis* the time scale at which market adjustments take place. After all, we have in the real world asynchronous processes of adjustments in production technologies, prices and market shares which might be badly reflected by an 'artificial' sampling over one-year time periods. (This is also akin the hypothesis put forward by Geroski, 2000.)

However, we are not convinced that this is by any means the whole story. A lot of evidence from the literature suggests that profits tends to be asymmetrically distributed and that such asymmetries are persistent over time. In future works we intend to check whether these properties apply also to our data and whether they are systematically correlated with asymmetries in efficiency. If that were the case, one would also have to draw far-reaching implications regarding the patterns of competition.

First, one might be forced to conclude that asymmetric efficiencies do not translate so much in systematic 'replicator-type' dynamics in the relative sizes of output but primarily in differential abilities to generate profits (and possibly affect relative sizes in the longer term only through the impact of profitability upon investment).

Second, an equally challenging implication of our evidence is that selection dynamics are primarily driven by outliers.

While qualitative evidence suggests that 'near-average' performances map into 'near-average' growth, some striking outliers systematically appear on both efficiency and growth indicators. It might well be that selection operates mostly in the long run, and mostly through the upper and lower distribution tails.

Relatedly, the dynamics of both the efficiency distributions and the revealed growth rates distributions suggest symmetry-breaking system behavior whereby outliers are the main drivers of long-term changes.

Another puzzle regards the evidence stemming from our data of any lack of relationship between growth variance and size—contrary to a lot of previous evidence from the literature (Sutton, 1997), and contrary also to our findings on the international pharmaceutical industry (Bottazzi, 2000; Bottazzi *et al.*, 2001). The lack of such a relationship in our Italian data might be interpreted on the grounds of different, possibly complementary, phenomena.

First, it might well be that diversification plays a relatively weaker role in Italian firms. Second, it could be the even when diversification occurs, it affects lines of business whose demand profiles tend to be highly correlated. Third, it could be a statistical artifact stemming from the ways 'firms' are defined, mainly for fiscal reasons.[19] Whatever the reason, the determinants of the variance in growth profiles is yet another challenging issue ahead.

All together, the foregoing evidence adds elements to the interpretation of the patterns of industrial evolution, with their generic invariances and their inter-sectoral differences. One of the apparent invariances regards the structure of the growth process, with 'fat-tailed' distributions of shocks—confirming the findings of Bottazzi *et al.* (2001). At the same time, the parameterizations of such distributions do depend on the specific sectors.

Concerning the underlying determinants of growth itself, the lack of robust correlations between proxies of efficiency and firms' growth continue to remain a puzzle for evolutionary analysts. Perhaps one should identify better proxies for the 'competitiveness' of each firms; or, maybe, markets do not work too well as selection devices at least on the time scale of our observations; or, the determinants of growth have highly idiosyncratic components that can only be captured through detailed firm by firm investigations. Whatever the answer, the issue stands as a major challenge facing evolutionary theory.

Acknowledgements

This work has been part of a research supported by the Italian Ministry of University and Research (MURST), by the S. Anna School of Advanced Studies, Pisa, and by the University of Bergamo (grant ex 60% no. 60CEFI01, Department of Economics). It has been accomplished with the precious collaboration of the Italian Statistical Office (ISTAT) and in particular of Roberto Monducci, whose competence and dedication

[19]Note that this could well be the case if the diversification of business groups has mostly occurred through the formation of formally separate legal entities (cf. the discussions, unfortunately in Italian, in Balconi, 1996; Barca, 1997).

made the idea of a joint research enterprise possible. Comments by several participants to the 'Nelson and Winter Conference', DRUID, Aalborg, June 20001—in particular Steve Klepper, Mariana Mazzucato, Luigi Orsenigo and Massimo Riccaboni—and by two anonymous referees have helped in shaping the present version of this work.

Address for correspondence

G. Dosi: S. Anna School of Advanced Studies, Piazza Martiri della Libertà 33, Pisa, Italy; gdosi@sssup.it.

References

Balconi, M. (1996), 'Sistemi informativi e crescita delle imprese. Un' interpretazione del recente sviluppo di gruppi medio-grandi in Italia e una rivisitazione teorica di crescita e dimensione', *Economia e Politica Industriale*, **92**, 35–78.

Barca, F. (ed.) (1997), *Storia del Capitalismo Italiano dal dopoguerra ad oggi*. Donzelli: Roma.

Boeri, T. (1989), 'Does firm size matter?', *Giornale degli Economisti e Annali di Economia*, **48**, 477–795.

Bottazzi, G. (2000), 'Firm diversification and the law of proportionate effect', LEM working paper, S. Anna School of Advanced Studies (http://lem.sssup.it).

Bottazzi, G., G. Dosi, M. Lippi, F. Pammolli and M. Riccaboni (2001), 'Innovation and corporate growth in the evolution of the drug industry,' *International Journal of Industrial Organization*, **19**, 1161–1187.

Brock, W. A. and D. Evans (1986), *The Economics of Small Businesses: Their Role and Regulation in the US Economy*. Holmes and Meier: New York.

Cefis, E., M. Ciccarelli and L. Orsenigo (2001), 'The growth of firms: from Gibrat's legacy to Gibrat's fallacy,' EPRIS working paper, University of Siena, Italy.

Dosi, G. (1988), 'Sources, procedures and microeconomic effect of innovation,' *Journal of Economic Literature*, **26**, 1120–1171.

Dosi, G., O. Marsili, L. Orsenigo and R. Salvatore (1995), 'Learning, market selection and the evolution of industrial structures,' *Small Business Economics*, **7**, 411–436.

Dosi, G., M. Riccaboni and R. Varaldo (2000), 'La dinamica dei rapporti tra grande e piccole imprese: alcune riflessioni tra teoria e analisi storica,' LEM working paper, S. Anna School of Advanced Studies (http://lem.sssup.it).

Ericson, R. and A. Pakes (1995), 'Markov-perfect industry dynamics: a framework for empirical work,' *Review of Economic Studies*, **62**, 53–82.

Evans, D. S. (1987a), 'The relationship between firm growth, size, and age: estimates for 100 manufacturing industries,' *Journal of Industrial Economics*, **35**, 567–581.

Evans, D. S. (1987b), 'Tests of alternative theories of firm growth,' *Journal of Political Economy*, **95**, 657–674.

Geroski, P. A. (2000), 'The growth of firms in theory and practice,' in N. Foss and V. Mahnke (eds), *New Directions in Economic Strategy Research*. Oxford University Press: Oxford.

Gibrat, R. (1931), *Les inégalités économiques*. Librairie du Recueil Sirey: Paris.

Hall, B. (1987), 'The relationship between firm size and firm growth in the US manufacturing sector,' *Journal of Industrial Economics*, **35**, 583–606.

Ijiri, Y. and H. A. Simon (1977), *Skew Distributions and the Sizes of Business Firms*. North Holland: Amsterdam.

Jovanovic, B. (1982), 'Selection and evolution in industry,' *Econometrica*, **50**, 649–670.

Marsili, O. (2001), *The Anatomy and Evolution of Industries*. Edward Elgar: Cheltenham.

Nelson, R. and S. Winter (1982), *An Evolutionary Theory of Economic Change*. The Belknap Press of Harvard University Press: Cambridge, MA.

Pavitt, K. (1984), 'Sectoral patterns of technical changes: towards a taxonomy and a theory,' *Research Policy*, **13**, 343–375.

Pakes, A. and R. Ericson (1998), 'Empirical implications of alternative models of firm dynamics,' *Journal of Economic Theory*, **79**, 1–45.

Stanley, M. H. R., L. A. Nunes Amaral, S. V. Buldyrev, S. Havlin, H. Leschhorn, P. Maass, M. A. Salinger and H. E. Stanley (1997), 'Scaling behavior in economics: empirical results and modelling of company growth,' *Nature*, **319**.

Subbotin, M. T. (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.

Winter, S. (1971), 'Satisficing, selection, and the innovating remnant,' *Quarterly Journal of Economics*, **85**, 237–261.

Winter, S. (1984), 'Schumpeterian competition in alternative technological regimes,' *Journal of Economic Behavior and Organization*, **5**, 287–320.