

---

# On the existence of persistently outperforming firms

---

Marco Capasso\*, Elena Cefis\*\*<sup>‡</sup> and Koen Frenken<sup>†</sup>

---

This study analyses persistence in growth rates for a data set of manufacturing firms of all sizes. Previous quantile autoregressions of firm growth rates show that extreme growth events are likely to be negatively correlated over time, thus questioning the existence of persistent outperformers. By supplementing the quantile regression analyses with transition probability matrices, our study shows that “bouncing” firms co-exist with persistent outperformers. This result is shown to be robust for all size classes as well as for the exclusion of firms who experienced acquisitions or spin-offs.

**JEL classification:** C14, L11, L25, L26, L60.

## 1. Introduction

This study analyses the growth rate distributions of the entire population of manufacturing firms in the Netherlands, in the period 1994–2004. In particular, we examine whether there is persistence in growth rates in terms of firms that persistently outperform or underperform.

A vast economics and management literature emphasizes that there are persistent asymmetries among firms in terms of size, innovation, and productivity. The existence of heterogeneous firms with different capabilities results in persistent heterogeneous firm performance (especially in relation to corporate growth and

---

\*Marco Capasso, Department of Economics, School of Business and Economics, Maastricht University, Tongersestraat 53, 6211 LM, Maastricht, The Netherlands. e-mail: marco.capasso@gmail.com

\*\*Elena Cefis, Department of Management, Economics and Quantitative Methods, University of Bergamo, via dei Caniana 2, 24127 Bergamo, Italy and LEM, Sant’Anna School of Advanced Studies, Pisa, Italy. e-mail: elena.cefis@unibg.it

<sup>†</sup>Koen Frenken, School of Innovation Sciences, Eindhoven University of Technology, P.O. Box 513, NL-5600MB Eindhoven, The Netherlands. e-mail: k.frenken@tue.nl

<sup>‡</sup>Main author for correspondence.

profitability). There are firms that persistently perform better or worse than others—or, in other words, there are persistent outperformers and underperformers.

As Dosi and Nelson (2009: 45) point out, the possible explanations for differences in corporate performance include differences in the ability to innovate and different production efficiencies. The persistently higher innovative capability and/or persistently greater efficiency should translate into persistently higher firm performance, and particularly persistently higher profitability and/or corporate growth.

On one side, substantial research effort has been devoted to examining profit persistence, and several empirical studies show that firms display persistent heterogeneity in levels of profitability (Mueller, 1986; Cubbin and Geroski, 1987; Mueller, 1990; Goddard and Wilson, 1999; Glen *et al.*, 2001; Cefis, 2003; Gschwandtner, 2004; Dosi, 2007), that is, profits do not seem to converge to a common rate of return. However, it is difficult to say to what extent the observed persistence in profitability heterogeneity reflects the persistence of differential “efficiency” levels that are not eroded away by the competitive process. Analogously, the way in which different profitability levels can translate into different employment growth paths is difficult to predict.

On the other side, when a firm has good innovative capabilities and/or high production efficiency it might be expected that it would grow with some persistence and that firms that perform persistently worse than their competitors should have persistent negative or null relative growth rates if the market selection mechanism allows them to survive. There is robust evidence that corporate growth rates result in distributions with fat tails (fat not only with respect to a Gaussian distribution, but also to a Laplace distribution) regardless of level of sectoral aggregation, country, and measures of size used (see among others for the United States Stanley *et al.*, 1996 and Bottazzi and Secchi, 2003; for Italy Bottazzi *et al.*, 2002 and Bottazzi *et al.*, 2007; for France Bottazzi *et al.*, 2011; for Austria Coad and Hözl, 2009; and for Denmark Reichstein and Jensen, 2005).

These studies show that growth rate distributions are often stable over time and, even if they cannot be considered as different draws from a same theoretical distribution, their shape remains constant over time, always displaying fat tails where outperformers and underperformers are concentrated. The stability of the distributions, or the stability of the shape of the distribution, may be given either by the fact that there is an intra-distributional mobility of firms that substitutes firms formerly in the tails with new ones, in every period, or by the fact that firms remain in the same positions as time goes by.

It is this aspect that this article addresses by analyzing the growth rate distributions of the *entire* population of firms to investigate whether there are firms that grow persistently (positively or negatively) in relative terms, and whether there is intra-distributional mobility among firms as time passes.

Investigating the persistence in growth rates has theoretical as well as policy implications. Among the studies that have investigated the validity of Gibrat’s Law, there are some that focus on the persistence of growth. In fact, finding persistence in

growth rates is one way to prove that Gibrat's Law does not hold, as Gibrat's Law implies independence (both across firms and along time) of growth events. From a policy point of view, persistence in growth rate is important to assess the role of the fast-growing firms (the so-called Gazelles) to evaluate whether their growth, and the resulting job creation, is a long-term phenomenon or their fast growth is intended to reverse in the subsequent periods.

The contributions of this work are fourfold: first, we use the entire population of manufacturing firms (without any employee number threshold), not assuming that firm growth rate distributions are symmetric, but investigating especially the asymmetric tails. Also, we shed more light on the different patterns among manufacturing firms conditional on firm size classes, which goes beyond most of the previous studies, including the works by Coad (2007), Bottazzi *et al.* (2007), and Bottazzi *et al.* (2011), which exclude micro firms from the analysis. A notable exception comes from the recent article by Coad and Hözl (2009: 143) whose study, in the words of the authors, "seeks to complement the existing literature on growth rate autocorrelation by focusing on the dynamics of micro firms."

Second, we account for exit and entry in the firm populations, and analyze persistence in growth rates based on an unbalanced panel and not just on surviving firms as Bottazzi *et al.* (2002, 2007) and Coad (2007) among others did. Considering only surviving firms may introduce a bias since we exclude negative growth rates for those firms that eventually exit the market, and the usually greater growth rates of newborn firms.

Third, we account for acquisitions and spin-offs, which means we can disentangle internal growth from external growth (the latter being due to acquisition activity). To the best of our knowledge, there are no other studies that consider acquisitions and spin-offs directly.<sup>1</sup>

Fourth, we do not only apply quantile regressions as in Coad (2007), Coad and Hözl (2009), and Reichstein *et al.* (2010), but we also qualify the results obtained in the quantile regression by estimating Transition Probability Matrices (TPM). These matrices allow us to detect whether "bouncing" firms (experiencing alternately highly positive and highly negative growth rates) coexist with persistent performers (experiencing persistent highly positive or highly negative growth).

The main result of this study questions the conclusions from previous studies that extreme growth events are negatively correlated over time. By supplementing the

---

<sup>1</sup>In Bottazzi *et al.* (2002, 2007) mergers and acquisitions (M&A) were considered indirectly, by considering any pair of firms that was merged or acquired during the time span of the study, as a unique "super-firm" along the whole span. In the present study, while firms that merged in a given year must always be excluded by the corresponding yearly subset of the pooled cross-section (since our data set does not enable us to track merged firms over time), episodes of acquisition and spin-off are explicitly considered, first by including in and then excluding from the cross-section of a given year all the firms that have experienced such episodes in that year.

quantile regression analyses with TPM, our study shows that bouncing firms coexist with persistent outperformers. This result is shown to be robust once firms who experienced acquisitions or spin-off are excluded. Furthermore, in differentiating among size classes, we find that the coexistence of bouncing and persistent outperformers is especially pronounced in micro firms.

The article is organized as follows. Section 2 reviews the literature and relates our contributions to previous work. Section 3 describes the data and the variables in our analysis and Section 4 discusses the methodology. Section 5 presents the results of the graphical analysis and of the regression analysis. Both sets of results are compared with those obtained in the previous literature. Section 6 makes a careful analysis of the quantile regression results and raises some doubts about the previous interpretations of the quantile regression results in the literature. Section 7 warns the reader about potential pitfalls within our regression analysis. Section 8 estimates the TPM and discusses the results. The analysis of the internal growth versus external growth is carried out in Section 9, and Section 10 concludes.

## 2. Literature review

Evolutionary economists emphasize that persistent asymmetries in firm corporate performance should be expected given the idiosyncratic routines, capabilities, and competencies of firms and their different learning processes. In particular, permanent differences in the ability to innovate and in production efficiencies are seen as two decisive factors that generate long-lasting differences in firm performance.

With regard to the ability to innovate, there is a growing empirical literature showing that firms have different capabilities to innovate, resulting in a small number of firms in each sector accounting for the majority of the innovative output, within a large group of minor innovators, and an even larger group of non-innovators (Geroski *et al.*, 1997; Cefis and Orsenigo, 2001). Furthermore, these different abilities to innovate across firms are persistent over time: there are “systematic” innovators that are able to innovate continuously (as opposed to “occasional” innovators that may produce an innovation once in a while, but with no continuity) and persistent non-innovators (Malerba and Orsenigo, 1996; Bottazzi *et al.*, 2001; Cefis, 2003; Peters, 2009; Raymond *et al.* 2010).

In terms of heterogeneity in production efficiency, there is robust evidence that production efficiency differs widely and persistently across firms and across plants (Baily and Chakrabarty, 1985; Rumelt, 1991; Baldwin and Rafiquzzaman, 1995; Jensen and McGuckin, 1997; Bartelsman and Doms, 2000; Dosi, 2007). These asymmetries can be found at relatively high levels of sectoral aggregation and do not depend on differences in relative factor intensities (Syverson, 2004; Bottazzi *et al.*, 2007; Dosi, 2007).

The persistently higher innovative capability and/or production efficiency should translate into permanent higher firm performance, and particularly into persistently higher corporate growth. The first studies to focus on persistence in growth rates of firms size regress current growth rates on lagged growth rates and find some persistence (Singh and Whittington, 1975; Dunne and Hughes, 1994, among others). Chesher (1979) was the first to introduce a first-order autoregressive structure in the error terms (the autoregressive coefficient was found to be significant and positive), aimed mainly at obtaining unbiased estimates of the autoregression coefficient of size (usually analyzed to test the validity of Gibrat's Law).

The empirical evidence on the persistence of growth rates is mixed and depends on the subsample analyzed and the methodology used. When subsamples that are homogeneous with respect to some factors (same sector and same age in Lotti *et al.* 2001; same technological regime in Almus and Nerlinger, 2000; etc.) are considered, we find generally nil or weak persistence in growth rates. Fotopoulos and Giotopoulos (2010) divide a sample of Greek manufacturing firms according to firm age and size and find that there is persistence only in certain groups, namely micro, small, and young firms. As Reichstein *et al.* (2010) point out, the methodology also plays a role in determining whether firm growth rates show persistence. They claim that "Gaussian statistics are unfit for studying firm growth" and that the results from studies that rely on them are misleading since they focus on central moments in the distribution. Even for studies that assume an exponential-like growth rate distribution, such as Stanley *et al.* (1996), Bottazzi *et al.* (2002, 2007), Bottazzi and Secchi (2003), but symmetric distributions (in the case of Bottazzi *et al.* due to data limitations: only firms with >19 employees) that concentrate on the central moment of the conditional distribution (in particular, the median), persistence is underestimated.

Lotti *et al.* (2003), Coad (2007), Coad and Höltz (2009), and Reichstein *et al.* (2010) apply quantile regressions to investigate the heterogeneity of firm growth, taking into consideration the entire conditional distribution of firm size or firm growth, and focusing on different percentiles separately. Lotti *et al.* (2003) test whether Gibrat's Law holds for new entrants in a given industry by running a quantile autoregression of firm size, and find that departures from Gibrat's Law may be due to the higher growth rates experienced by small entrant firms immediately after start-up. Coad (2007), using a quantile autoregression of firm growth rates, shows that small fast-growing firms are likely to display negative growth autocorrelation, while larger firms will often achieve smoother growth patterns. Coad and Höltz (2009) confirm that autocorrelation dynamics vary with firm size, and extend the previous findings to the case of micro firms, whose growth seems to be characterized by lumpy adjustment patterns. Reichstein *et al.* (2010) use quantile regression techniques to link firm growth rates to industrial and regional characteristics. Firm size again is found to exert a moderating influence on growth, but industry concentration and regional specialization seem to put a lower bound on the growth rate of large firms.

Our article aims to analyze persistence in growth rates, especially in the tails of the distributions, to assess whether there are persistent outperformers and persistent underperformers, supplementing the quantile regression analyses with TPM, on an unbalanced panel of data that comprises firms of all sizes and controlling for firm entry and exit (including M&A and spin-offs).

### 3. Data and variables

The data in this article were collected by the Centraal Bureau voor de Statistiek (CBS) and stem from the Business Register (BR) of enterprises. The BR database includes the entire population of firms registered for fiscal purposes in the Netherlands in the year considered. The database reports detailed information on the sector of the company, at the five-digit SBI (Standaard Bedrijfsindeling—the Dutch standard industry classification) level, number of employees and dates of entry and exit of firms in the market. The definition of entry and exit excludes changes in the firm's sector of activity; when this occurs, the firm is regarded as a continuing firm. Furthermore, the BR provides information on different types of exit, distinguishing between exits due to failure, to merger, to acquisition, and to radical restructuring. Our observation period covers the years 1994–2004. For each year, we selected all the manufacturing firms present in the BR. The population includes firms with zero employees, referred to as self-employment.

In this study, the main variables of interest are firm size, as proxied by employment, and its growth rate. Following Bottazzi *et al.* (2001), Bottazzi *et al.* (2007), and Bottazzi *et al.* (2011), we chose to analyze the behavior, for every firm, of the logarithm of firm size (where size is measured by the number of employees plus one, to avoid logarithms of zero) relative to the industry average, i.e. of the variable

$$s_i(t) = \ln(S_i(t)/\bar{S}_J(t)) \quad (1)$$

where  $\bar{S}_J(t)$  represents the geometric mean of size across all firms that at each time  $t$  belong to the same industry  $J$  to which firm  $i$  belongs. This is equivalent to

$$s_i(t) = \ln(S_i(t)) - \frac{1}{N_J} \sum_{j \in J} \ln(S_j(t)) \quad (2)$$

The relevant industry  $J$  for each firm is defined at the three-digit SBI level. The new rescaled log variables represent deviations from the industry average (at the three-digit level) for each firm. The use of the proportion of size  $s_i(t)$  as our basic variable, rather than plain size  $S_i(t)$ , or its logarithmic form, has important advantages:

- (i) The new variable controls for size differentials across industries at the three-digit level;
- (ii) It controls for differences in sectoral growth rates;

- (iii) It removes possible common shocks and more general common factors in the economy such as business cycle and inflation<sup>2</sup>;
- (iv) It has the advantage that it can be used to characterize distributions whenever the number of firms changes over time, and therefore provides an easy way to compare distributions with different numbers of observations.

Following most of the previous studies, we define the growth rate of firm size as the difference in the logarithmic proportional size between two consecutive years, namely,

$$g_i(t) = s_i(t) - s_i(t - 1) \quad (3)$$

and consequently the growth rate at time  $t - 1$  is  $g_i(t - 1) = s_i(t - 1) - s_i(t - 2)$ . We thus consider short-term growth rates. Notably, alternative definitions of growth rates have been used by Davis *et al.* (1996), Davis *et al.* (2006), and Hölzl (2009), which reduce the problems related to firm exit, and in particular the disadvantages deriving from the zero-size threshold posing a lower bound to logarithmic growth rates in a balanced panel. Our decision of defining growth rates as differences in size logarithms is dictated by the necessity of offering a clear comparison between our results and the results obtained in the recent literature on firm growth. Notice also that the procedure for normalizing size described previously automatically normalizes growth rates with respect to the average industrial growth rate.

As the focus of the article is on size, all analyses are conducted on the whole sample as well as on four different subsamples, obtained by dividing firms according to their size. The subsamples include the following:

- MICRO firms, <20 employees;
- SMALL firms, 20–99 employees;
- MEDIUM firms, 100–199 employees;
- LARGE firms, >200 employees.

The choice to analyze four size subsamples is to enable comparison with previous studies that usually do not include micro firms (Bottazzi *et al.*, 2002, 2007; Bottazzi and Secchi (2003); Coad, 2007), and generally focused mainly on medium–large firms. In general, this choice in the past has been dictated by the data availability. Since in our case data on the entire population are available, it is interesting to compare the different subsamples to provide some new empirical evidence on a part of the firm population that has been neglected. Between-group differences in the number of employees distribution are shown in Table 1 (pooled data of the 11 cross-sectional waves between 1994 and 2004 included).

---

<sup>2</sup>Given the economy's reliance on exports, the Dutch business cycle follows more or less the world's business cycle. In the period considered here, GDP growth was highest in 1999 (with 4.7%) and was lowest in 2002 (0.1%). The use of the proportion of size  $s_i(t)$  allows us to avoid problems linked to the pro-cycle nature of firm growth as described in Higson *et al.* (2004).

**Table 1** Descriptive statistics of the firms' number of employees, 1994–2004 (pooled)

	Micro firms	Small firms	Medium firms	Large firms	All firms
Sample size	604,294	64,458	9381	8315	686,448
Mean	3.20	40.89	138.43	636.77	16.26
Variance	16.19	392.31	721.68	2,043,773.00	29,889.42
Skewness	2.01	1.10	0.40	18.31	134.86
Kurtosis	6.57	3.23	2.21	427.02	26,050.61
Minimum	0	20	100	200	0
Quantile					
0.1	0	21	105	220	0
0.2	1	24	110	246	1
0.3	1	27	119	278	1
0.4	1	30	127	325	1
0.5	1	35	137	353	2
0.6	2	38	148	425	3
0.7	3	46	150	541	5
0.8	5	58	161	725	10
0.9	9	75	179	1155	24
Maximum	19	99	199	37,726	37,726

#### 4. Methodology

To analyse persistence and intra-distributional mobility of firm growth rates, we apply parametric and nonparametric methods. First, we verify whether growth rate distributions are stable over time, and whether there are some particularities in their shape. We perform a graphical analysis by comparing histograms of relative frequencies and compute Cramér von Mises test statistics (Anderson, 1962) to check whether distributions change along time. Given the presence of growth rate autocorrelations, the  $P$ -values of the test are not reliable; however, the test statistics can still provide an idea about the stability over time of the growth rate distributions, in the arbitrary assumption that the bias caused by the autocorrelations is stable over time. Second, we study the autoregressive structure of growth rates by means of the following quantile regression (Koenker and Bassett, 1978):

$$g_i(t) = \alpha_\theta + \beta_\theta g_i(t-1) + u_i(t) \quad (4)$$

with

$$\text{Quant}(g_i(t)|g_i(t-1)) = \alpha_\theta + \beta_\theta g_i(t-1)$$

where  $\alpha_\theta$  and  $\beta_\theta$  are the parameters to be estimated,  $u$  is the error term of the model, and  $\text{Quant}(g_i(t)|g_i(t-1))$  denotes the  $\theta^{\text{th}}$  conditional quantile of  $g_i(t)$  given  $g_i(t-1)$ .



As we allow for the effect of size by running the regression not only on the whole firm population but also separately on different size classes, there is no control for size directly in the regression. The parameter estimation procedure, for any given conditional quantile  $\theta$  between zero and one, solves the following minimization problem:

$$\min_{\alpha_\theta, \beta_\theta} \left\{ \sum_{i,t: g_i(t) \geq \alpha_\theta + \beta_\theta g_i(t-1)} \theta |g_i(t) - \alpha_\theta - \beta_\theta g_i(t-1)| + \sum_{i,t: g_i(t) < \alpha_\theta + \beta_\theta g_i(t-1)} (1 - \theta) |g_i(t) - \alpha_\theta - \beta_\theta g_i(t-1)| \right\}$$

There are two reasons for this choice. On the one hand, growth rate distributions display a “tent” shape, which departs from the assumption of error normality in favor of more heavy-tailed distributions. On the other hand, following the intuition suggested by Coad (2007) for firm growth autocorrelation and later used to relate firm growth to industry characteristics (Reichstein *et al.*, 2010), we prefer not to restrict attention to the mean or the median of the conditional distribution of the dependent variable (as in Ordinary Least Squares (OLS) and Least Absolute Deviation regressions), and instead retrieve a different growth autoregression parameter  $\beta_\theta$  for each quantile  $\theta$  of the conditional distribution. In our study, the regression is computed for a pooled cross-section of firms, such that the growth rates between years  $t - 1$  and  $t$  are regressed on the growth rates between years  $t - 2$  and  $t - 1$  for the same firms, considering all the years  $t$  between 1996 and 2004 simultaneously.

Third, we estimate slightly modified TPM on the deciles of the growth rate distributions to derive some descriptive statistics on firm intra-distributional mobility. Notice that, by using the firm distribution deciles and then by focussing only on the relative positions of firms within the distribution, our TPM analysis has the advantage of being less vulnerable to the risks associated to pooling the data when the growth distribution narrows its shape over time. For example, even when the highest growth rates in a given year are much higher or much lower than the highest growth rates in a different year, a firm having the highest growth rate in both years will not change decile, and will still be classified as a persistent outperformer.

The transition probability matrix  $\mathbf{P}$  is the matrix with  $p_{kh}$  as the elements measuring the probability of moving from decile  $k$  to decile  $h$  in one period (Hoel *et al.*, 1987). This probability is relatively high (low) when the corresponding value in the transition matrix is higher (lower) than 0.1. If for any given growth rate observed at time  $t - 1$  the probability of moving to a particular growth decile at time  $t$  were uniformly distributed across the 10 possible target deciles, then each of the 10 deciles at time  $t$  would have exactly 1/10 probability of including the new growth rate. We compute 1-year period TPM, that is, matrices where the input set (vertical axis) considers the growth  $g_i(t - 1)$ , and the target set (horizontal axis) considers the growth  $g_i(t)$ .

A final remark on methodological issues is required, namely in relation to the balancing procedure applied to the growth rate distributions. In fact, the panel is balanced for each couple of years, that is, for each year  $t$  we need firms to be present in year  $t$  and in year  $t-1$ . In addition, size classes are defined at year  $t-1$ . For example, the group of micro firms includes all firms with  $<20$  employees at year  $t-1$ , while at time  $t$  they may have  $>20$  employees. When we compute the growth quantile autoregressions and build the transition matrices, for each year  $t$  we consider the size of the firms present in the database at time  $t-2$ ,  $t-1$  and  $t$  (excluding exits that have occurred in each 3-year time span). The size classes are then defined at year  $t-1$  to be consistent with the subsampling procedure applied earlier when plotting the growth rate distributions. The quantile regressions and the TPM are used to explain the growth rates  $g_i(t)$  of the group defined at  $t-1$ , and therefore, for each time  $t$ , the dependent variables of the quantile regressions and the target sets of the TPM must be the growth rates  $g_i(t)$  of the group defined at  $t-1$ .

This choice is thus a consequence of the role played by the different methodologies applied in this article. This is the reason why we introduce a modification in computing the TPM, that is, we normalize each cell probability by the column sum and not the row sum, which is normally used when estimating TPM. Therefore, the 1-year transition probability is defined as follows:

$$p_{kh} = P(X_{t-1} = k | X_t = h) \quad (5)$$

where  $t = 1995, \dots, 2004$ .

The TPM  $\mathbf{P}$  is the matrix with  $p_{kh}$  as elements measuring the probability of moving to state  $h$  from state  $k$  in 1 year. In other words, each cell in our TPM shows the probability that firms have to start from the  $k$ -th decile at time  $t-1$  given they end up in the  $h$ -th decile at time  $t$ .

However, in practical terms the use of the column sum instead of the row sum exerts only a minor influence on the estimated probabilities. The use of deciles on both the row and the column implies that, except for the first and the final years of the considered time span, the firms belonging to the same  $k$ -th decile of growth rates in a given year  $t$  will contribute in equal terms to the  $k$ -th row sum and to the  $k$ -th column sum, since they belong for year  $t$  to the target set and for year  $t+1$  to the input set.

## 5. The results

### 5.1 Growth rate distributions

Figure 1 represents the empirical distribution of firm growth rates for all the years included in the data set, obtained by using the growth rates computed as in (3) on

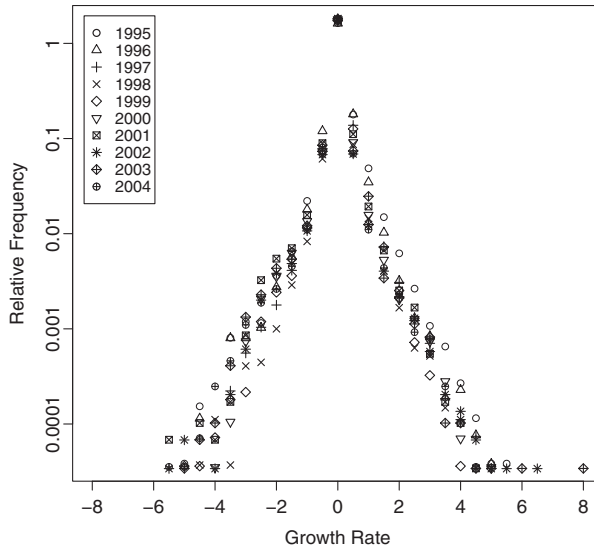


Figure 1 Growth rates distribution, 1995–2004.

the horizontal axis and the corresponding relative frequencies<sup>3</sup> on the vertical axis, and excluding firms that exited in each 2-year time span. Note that the scale of the vertical axis is logarithmic. The growth rate distributions look tent-shaped and approximately symmetric regardless of the year. The body of the distribution recalls the Laplace distribution found by Stanley *et al.* (1996) for firm growth rates, while the tails look slightly fatter than expected, suggesting the possibility of a better fitting by an exponential power (Subbotin) distribution with a shape parameter slightly  $< 1$  (see Bottazzi *et al.*, 2002).

Comparing the first and last year growth distributions (Figure 2), we can see a decrease in relative frequencies in the tails, especially in the right tail. This decrease is more evident in Figure 3, where the evolution over time of the 10th and the 90th percentiles show the shrink in the distribution toward its central part (mainly due to an increase of self-employed firms in the most recent years). Table 2 shows that the two-sample Cramér von Mises test statistic, computed for all the years considered to compare the distribution of growth rates in year  $t$  with the distribution of growth rates at time  $t - 1$ , is very high and leads us to think that the hypothesis that each pair of relative frequency distributions represents the double realization of a unique data-generating process should be rejected. However, the decrease in the statistic over time points to a tendency for less unstable distributions in more recent years.

<sup>3</sup>Relative frequencies have been computed by using the function “hist” of the software package R (option “density”), and indicate the number of firms having a given growth rate, weighted by the total number of firms in the sample and by the width of the bin used on the horizontal axis.




Figure 2 Growth rates distribution for the years 1995 and 2004.

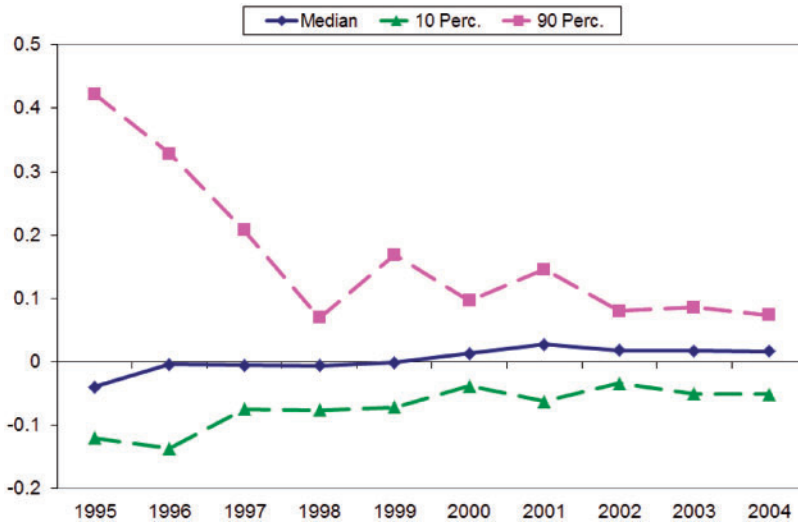



Figure 3 Evolution over time of median, 10th percentile and 90th percentile of the growth rate distribution.

**Table 2** Cramér von Mises statistic for comparing the growth rate distribution in year  $t$  with the growth rate distribution in year  $t - 1$

Year $t$	Value
1996	1175.081
1997	52.54619
1998	94.13955
1999	62.71444
2000	762.5912
2001	959.2009
2002	614.7111
2003	186.7675
2004	101.5324

An important departure from what Bottazzi *et al.* (2002, 2007) find, is the width of the distribution, which appears to be much larger in our case. A possible explanation of this different width (or scale) might be that the firms considered by Bottazzi *et al.* (2002, 2007, 2011) had >19 employees in all the years under consideration. Therefore, the numerous changes in size in the shift from micro to small or larger firms were not observable in their data set. If we disaggregate according to size, as shown in Figure 4, it is evident that the growth rate variance decreases as firm size increases, a phenomenon already described by Hymer and Pashigian (1962), Axtell (2001), and Sutton (2002). Moreover, the growth distribution is shown to lose its symmetry in favor of an increase in the right tail for firms with <20 employees, and an increase of the left tail for all the other firms. The growth rate distributions of small, medium, and large firms thus appear left-skewed and bimodal, with the slope of the left tail appearing much flatter than the slope of the right tail. For micro firms, on the other hand, the distribution appears left-truncated, with the slopes of the two tails looking similar. The apparent symmetry in Figure 1 is a result of the combination of the left-skewness observed for small, medium, and large firms, related to episodes of poor performance experienced by larger firms, and the left-truncation observed for micro firms, related to the fact that firms cannot have less than zero employees and therefore exiting firms cause the truncation.

The negative relation between size and growth (now meant as growth level, not growth variance) has been found in several studies (e.g. Evans, 1987; Hall, 1987). The doubts raised by the possibility that this relation could only be due to the exclusion from the database of exiting firms were dismissed by Dunne and Hughes (1994) and Harhoff *et al.* (1998), who confirm that smaller companies grow faster even after controlling for sample attrition. The growth rate distributions in our data set allow



**Figure 4** Growth rates distribution, 1995–2004: micro firms (top left); small firms (top right); medium firms (bottom left); large firms (bottom right).

us clearly to distinguish between the sample attrition effect, shown by the left truncation of the distribution for micro firms, and the real economic phenomenon of a negative effect of size on growth, shown by the shrinking in the right tail as we move from micro to small, medium, and large firms. In other words, the negative influence of size found by regression studies in the literature can be described in terms of the asymmetric Subbotin distribution terms used by Bottazzi and Secchi (2003), not as a negative effect of size on the location parameter (the distribution always appears centered on zero), but as a negative effect of size on the scale parameter of the right part of the distribution.

It could be argued that in a balanced panel the scale of the distribution also can be biased by the exiting firms, and therefore even the positive distributions of micro and small firms are not comparable, as the frequencies of the right tail must be considered in relation not only to the overall number of frequencies in the whole “truncated tent” but also to the number of firms that exited the database.

To show that the change in scale in the right part of the distribution is not simply an outcome of sample attrition, we rescaled the distributions with respect to the total number of firms existing at time  $t - 1$ , including firms that no longer exist at time  $t$ . Whatever negative growth rate we assign to these firms, the scale of the right (i.e. the positive growth rate) part of the distribution is directly comparable across size classes.


Therefore, after pooling all the 10 cross-sectional waves, we divide the frequencies of growth rates of micro and small firms (shown in disaggregated form in the top panels of Figure 4) by the sum of all the firms existing in the first of each pair of years for which we computed growth rates. In other words, we divide the frequencies of the growth rates between years  $t - 1$  and  $t$  for micro and small firms that survive in both years by the number of firms existing at time  $t - 1$  including firms that exited in time  $t$ . The resulting relative frequencies that correspond to positive (normalized) growth rates are plotted on a log scale in Figure 5 (relative frequencies equal to zero are represented as zeros on the log scale). The plot shows that, even after controlling for the number of exit firms, extreme events of positive growth are clearly relatively much more frequent for micro firms than for larger firms, and can still be encompassed in a Laplace/Subbotin framework where the tent (on its right side) simply has a lower scale for higher average size.

Following the “regression fallacy” argument by Davis *et al.* (1996), it can be argued that the changing asymmetry of the growth rate distribution, when moving from small to large firms, could be explained by a technical artifact, rather than a regression-to-the-mean law. Indeed, firms that are classified as small could be firms that were previously large and, after experiencing a transitory negative growth event, are coming back to their original status (same goes for small firms that have become large). As the focus of our study is on persistence of growth rates, the following analysis includes an appreciation of the relative amount of transitory events, as opposed to stable paths of growth, that take place within the economy, and thus goes to the root of the “regression fallacy” argument.


## 5.2 Regression analysis

We next turn to the quantile regressions to detect possible patterns of autocorrelation in firm growth. In particular, for the pooled cross-section of firms that existed between 1994 and 2004, obtained after balancing the panel at each 3-year time span, we regress the growth rate for each firm at time  $t$  on the growth rate for the same firm at time  $t - 1$ .

As expected, the autoregression parameter  $\beta_\theta$  is very close to zero at the median of the conditional distribution (i.e. for  $\theta$  close to 0.5), for the whole sample as well as for the different size classes (see Figures 6 and 7). However, obtaining a value of the autoregression parameter that is close to zero is in this case also depending on the robustness of the median regression (as opposed to an OLS regression), which does




**Figure 5** Density plot of normalized growth rates, 1995–2004: Zoom on the right (positive) part of the distribution; frequencies expressed in relation to the overall initial sample size (including exits); zeros in the plot represent zero density.



**Figure 6** Quantile autoregression of growth—whole sample.





**Figure 7** Quantile autoregression of growth—micro firms (top left); small firms (top right); medium firms (bottom left); large firms (bottom right).

not consider the firms that lie in the tails of the conditional distribution. The subsamples shown in Figure 7 are different in terms of their behavior at quantiles that do not correspond to the median.

Given the high percentage of micro firms within our database, the whole sample and the subsample of micro firms show many similarities. In particular, the coefficient is close to zero for most quantiles, which could be expected given that the discrete nature of the employment variable prevents micro firms from experiencing frequent growth events, but becomes negative for the extreme quantiles. In other words, the micro-firms that at a given year experience extreme events of positive growth are on average firms that in the previous year experienced events of lower growth. For the negative case, the absolute value of the size decrease at a given year is not necessarily strong: it can also be mild, being linked on average to previous episodes of even milder negative growth. The negative relation can be detected on a larger range of quantile values when we consider firms with >19 and <100 employees (small firms). However, both the coefficient values and the width of the confidence intervals at high quantiles look larger, indicating that above-average growth at time  $t$  cannot be attributed completely to a rebound after poor performance at time  $t - 1$ .

The results for micro and small firms would seem to confirm the results in the literature (see e.g. Garnsey and Heffernan, 2005), while the results for medium and large firms, depicted in the bottom panels of Figure 7, are less similar to those found in previous studies. We find there is a positive parameter for quantiles between 0.1 and 0.5 in the subsample of medium firms and for quantiles between 0.5 and 0.9 in the subsample of large firms, while the parameter remains negative for the extreme quantiles. In other words, medium firms situated in medium–high quantiles seem to benefit from previous episodes of growth, but this is not the case for extreme events.

The fact that firms belonging to extremely high quantiles have a negative coefficient even when these are large firms, is robust to the non-normalization of growth rates to average sectoral growth, and defines a difference with respect to the work done by Coad (2007) and Coad and Hözl (2009), who find a positive coefficient for large firms. However, these studies introduced firm size directly into the regression as an independent variable, while here we study a “pure” autocorrelation of growth rates and consider firm size differences only by subsampling the data into the four groups previously described. Moreover, the two above-mentioned works consider size at  $t-1$  where growth rates between  $t-2$  and  $t-1$  and between  $t-3$  and  $t-2$  are among the explanatory variables, which is consistent with the objective of studying the direct effect of size on growth, but which makes the interpretation of results on the autocorrelation of growth rates even more difficult.

## 6. Discussion of the regression results

In Section 5 we analyzed the behavior of firms situated in different positions in the conditional growth rate distribution at time  $t$ , i.e. conditional to growth rates at time  $t-1$ , and found that, when considering the whole sample, firms with very high or very low conditional growth rates at time  $t$  are characterized by a negative autoregression parameter, while this coefficient is negligible for firms that are close to the conditional median.

To understand the meaning of this result in unconditional terms, and thus to connect it with our research question about the persistence of firm performance, we need to go back to the definition of quantile regression. Since quantile regressions estimate a linear relation between the independent variable, in our case the growth rate at time  $t-1$ , and the quantiles of the conditional distribution of the dependent variable, in our case the growth rate at time  $t$ , then an interpretation of the previous results can be based on the way in which the conditional distribution of the growth rates at time  $t$  changes as the growth rate at time  $t-1$  increases, i.e. the way in which the growth of a given firm at time  $t-1$  affects the probability of growth of the same firm at time  $t$ . As Figures 1 and 4 have shown, the growth rate distributions for our database are characterized by the same tent shape already observed in the previous literature, and seem to be always centered around zero, i.e. having a median close to zero.

A simple example can show how extending this zero-median property to all the *conditional* distributions of growth rates at time  $t$ , and assuming also that the skewness of the same conditional distributions decreases as the conditioning growth rate at time  $t - 1$  increases, can explain the quantile regression results described in the previous section.

Suppose there are three groups of seven firms, where the firms of the first, second, and third group have experienced, respectively, a negative, zero, and positive growth rate at year  $t - 1$ . Suppose also that, at time  $t$ , three firms of the first group experience a high (low in absolute value) negative growth rate, one remains stable, and three firms experience a very high positive growth rate. This is equivalent to say that, conditioning on a negative growth rate at time  $t - 1$ , the conditional distribution of growth rates at time  $t$  has a median equal to zero and a positive skewness.


Suppose also that, at time  $t$ , three firms of the second group (i.e. the ones that were stable at time  $t - 1$ ) experience a moderate negative growth rate, one remains stable, and three firms experience a moderate positive growth rate. This is equivalent to say that, conditioning on a zero growth rate at time  $t - 1$ , the conditional distribution of growth rates at time  $t$  has median and skewness equal to zero.

Finally, suppose that, at time  $t$ , three firms of the third group (i.e. the ones that had a positive performance at time  $t - 1$ ) experience a very low (high in absolute value) negative growth rate, one remains stable, and three firms experience a low positive growth rate. This is equivalent to say that, conditioning on a zero growth rate at time  $t - 1$ , the conditional distribution of growth rates at time  $t$  has median equal to zero and negative skewness.

All the previous assumptions can be summarized by stating that the conditional distributions of growth rates at time  $t$  (conditioned on the growth rate at time  $t - 1$ ) have always median equal to zero, and their skewness decreases as the conditioning growth rate at time  $t - 1$  increases.

Such hypothetical scenario is described by Figure 8 in which growth rates at time  $t - 1$  are measured on the  $x$  horizontal axis, and growth rates at time  $t$  are measured on the  $y$  vertical axis. The three groups of firms are then represented as three groups of seven points, belonging to a bi-dimensional space defined by the  $x$  and  $y$  axes, that are situated, respectively, at a negative (say,  $-4$ ), zero, and positive (say,  $+4$ ) value on the  $x$  axis, but where within each group the  $y$  coordinates are different. The median of the groups is equal, but the highest  $y$ s in the first group are higher than the highest  $y$ s in the second and third groups, and the lowest  $y$ s in the first group are higher than the lowest  $y$ s in the second and third groups. This is tantamount to saying that the slope of the linear relation linking  $x$  and the high conditional quantiles of  $y$  (conditional to  $x$ ) is negative, the slope of the linear relation linking  $x$  and the low conditional quantiles of  $y$  (conditional to  $x$ ) is negative, and the slope of the linear relation linking  $x$  and the conditional median of  $y$  (conditional to  $x$ ) is flat.

Consider now the first and the fifth sextile (examples, respectively, of a low and a high quantile). In Figure 8, the points corresponding to the conditional first and fifth



**Figure 8** Theoretical example of negative slope on extreme conditional quantiles and zero slope on the conditional median.

sextiles, respectively (for each  $x$  we have 7 points, i.e. 6 intervals), are marked by a cross and a triangle, and have been fitted by quantile regressions estimated on the whole sample of 21 points at quantiles  $1/6$  and  $5/6$ . The slopes of the two lines is negative, which depends simply on the fact that the conditional distribution of  $y$  given  $x = -4$  is right skewed, while the conditional distribution of  $y$  given  $x = 4$  is left skewed. Note that, if the tails of all the conditional (and therefore also the unconditional) distributions are sufficiently fat, the  $y$  values may be very high even if the slope of the line obtained for the fifth conditional sextile is negative. It should be noted also that, in this scenario, the median of the conditional distribution of  $y$  remains fixed, but the mean is decreasing with  $x$ .

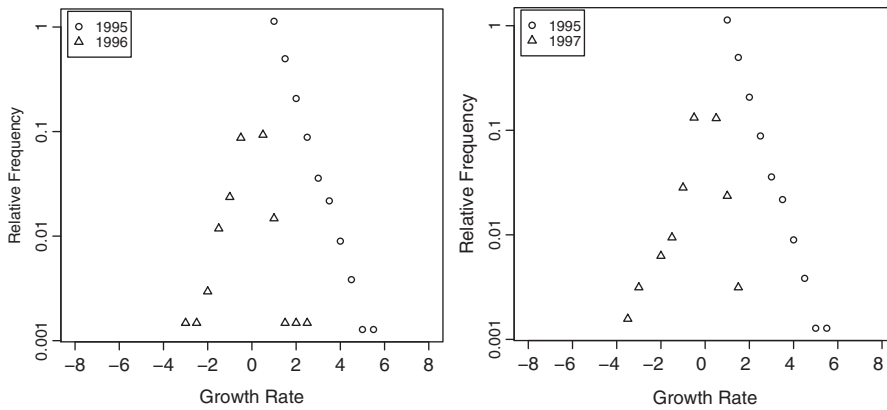
If we interpret the 21 points as 21 firms whose growth rates at time  $t-1$  are measured on the  $x$  axis and growth rates at time  $t$  are measured on the  $y$  axis, then, in presence of fat tailed growth rate distributions, a negative coefficient of the fifth conditional sextile would correspond to a subsample of firms (top-left corner of the figure) with very high growth rates at time  $t$  and negative growth rates at time  $t-1$ , and a subsample of firms (top-right corner of the figure) with lower but still high growth rates at time  $t$  and high growth rates at time  $t-1$ . Roughly speaking, for that given conditional quantile, the slope of the quantile regression line (i.e. the autoregression parameter) is negative, but the intercept is so high that it is still possible for a firm to show persistent growth. On the other hand, on average, the growth rate at time  $t$  is lower for firms that experienced high growth rates at time  $t-1$ . Summing up, the previous theoretical example, consistent with the empirical results on quantile autoregressions of growth rates, illustrates that, even if firms having a high

growth rate at time  $t - 1$  are on average more likely to experience low growth rates at time  $t$ , still their probability of replicating their good performance at time  $t$  (i.e. of being persistent outperformers) can be far from negligible.


If we return to the real data, in Figure 9 (left) we compare the frequency distribution of the growth rates of firms showing high growth (the top three percentiles of all firms) in 1995 with the distribution of the growth rates of the same firms in 1996 (left). It appears that being in the right tail of the distribution in year 1995 raises the probability of not being in the central body of the distribution in 1996. Moreover, it raises also the probability of being in the left tail (compare the triangles at the left of  $-2$  on the horizontal axis with the triangles at the right of  $2$ ). If we repeat the exercise for a 2-year lag, i.e. growth rate distributions for the same firms in 1997, we obtain the scenario depicted in the right part of Figure 9, which shows an even fatter left tail of the distribution. Thus, the firms belonging to the right tail in 1995 have a relatively higher probability of showing poor performance in 1997, as if a mean-reversion process were acting on the selected firms. In other words, experiencing high growth in a given year seems to increase the probability of experiencing negative growth in subsequent years. The simplest explanation for this could be that extreme (positive) growers move into a higher size class and therefore their growth rate distribution becomes more left-skewed.

The scatterplot in Figure 10 charts normalized growth rates at time  $t$  against normalized growth rates at time  $t - 1$  for all the years between 1996 and 2004. The star shape of the scatterplot is a direct consequence of the tent shapes of the unconditional distributions of the growth rates. If we concentrate on the points that, for each interval of the horizontal axis, have, respectively, the highest and the lowest value on the vertical axis, i.e. the firms that belong, respectively, to the highest and lowest quantiles of the conditional distribution of growth rates at time  $t$ , we get the same pattern as described in Figure 8: the two groups of points seem to define two lines with negative slopes, and the conditional distributions seem to be more and more left-skewed as the conditioning growth rate at time  $t - 1$  gets higher and higher. Notice that only the part of the graph that lies at the left of the vertical axis could be justified partially by the exclusion of exiting firms, which makes it difficult for a firm to experience two consecutive strong negative growth rates and still remain in our database.

As predicted, the negative coefficients obtained from the extreme quantile regressions do not preclude the existence of many firms in the top-right corner that are growing persistently and belong in both years to a high quantile of the unconditional growth rate distribution. Roughly speaking, the quantile regression, when run on the highest conditional quantiles, compares the highest part of the scatterplot in top-left part of the graph with the highest part of the scatterplot in the top-right part of the graph. The negative autoregression coefficient obtained in this study and in the studies by Coad (2007) and Coad and Hölzl (2009) tells us that the upper bound of the scatterplot is higher in the top-left corner than in the top-right corner, i.e. the



**Figure 9** Growth rates distribution for the years 1996 (left panel, triangles) and 1997 (right panel, triangles) of the fastest growing firms (top three growth percentiles) of year 1995, as compared with their distribution in 1995 (left and right panels, circles).



**Figure 10** Scatterplot of growth rates at time  $t$  against growth rates at time  $t - 1$ , 1995–2004.

upper bound of the scatterplot becomes lower and lower as we move from left to right of the picture. In economic terms, this means that the highest growth rates at year  $t$  observed among firms that had a poor performance at year  $t - 1$  are generally higher than the highest growth rates at year  $t$  observed among firms that already showed good performance at year  $t - 1$ .

As a consequence, at year  $t$ , the highest growth rates within the whole population generally are achieved by firms that showed poor performance at year  $t - 1$ , but there will be many firms that experience high growth rates (not the highest in the population, but higher than the average) in two consecutive years, i.e. there will be persistent winners. In particular, when the negative value of the autoregression parameter is not too high in absolute terms so that, roughly speaking, the upper bound of the scatterplot does not cross the horizontal axis (which applies to our data), then we can be certain that persistent winners exist. They correspond to the points in the top-right corner of the scatterplot and clearly experience episodes of very high growth in both years compared with the total population.

Of course, the visual inspection of the scatterplot overemphasizes the behavior of outliers, which are more visible because most observations are plotted on top of each other “within the star” close to the axes. In fact, most of the strong intra-distributional movements over time, that is, the big spurts from an unconditional quantile at time  $t - 1$  to a very distant quantile at time  $t$ , may correspond to very small changes in the absolute value of the growth rate, given the Laplace-like unconditional distribution of growth rates. Similarly, a low slope coefficient of the quantile regression for quantiles close to the median may be more meaningful in economic terms than a higher slope coefficient obtained for an extreme quantile. At the same time, the quantile regression is able to show only some of the firm heterogeneity by fitting with one line the points corresponding to a given conditional quantile, while more information on outperformers and underperformers could be retrieved by considering explicitly those points that, within that quantile, are far from the linear fit.

For these reasons, we need to apply another methodology that will directly infer growth persistence between  $t - 1$  and  $t$  among firms belonging to a given unconditional quantile at year  $t$ , and to include outlying firms in the resulting picture. Koenker and Hallock (2001) argue that a least squares estimation fit to subsets of the response variable according to its unconditional distribution is not possible because truncation of the dependent variable would lead to a sample selection problem (Heckman, 1979). As well as suggesting the (conditional) quantile regression, these authors suggest the alternative of a binary response model for the probability that the response variable exceeds some prespecified cutoff value.

Given that firm heterogeneity and outperformers are the focus of this study, a unique coefficient of a probit model linking the probability of belonging to a given unconditional growth quantile at time  $t$  (dependent variable) to growth at time  $t - 1$  (independent variable) might be insufficiently flexible for our purposes. Rather, for each couple of unconditional quantiles of the growth distribution at time  $t - 1$  and

time  $t$ , it would be useful to know the probability of a firm ending up in the first quantile at time  $t$  starting from the second quantile at time  $t-1$ . This can be achieved by means of probability transition matrices built on the growth rate distributions of time  $t$  and  $t-1$ .

## 7. Caveat about the regression results

Our study is based on a pooled cross-section, since we have discarded the idea of introducing panel data techniques to analyze our data set. Some readers could argue against our decision, especially for what concerns the quantile regression approach, for which our decision of pooling the data could look particularly arbitrary.

In the recent literature on industrial dynamics, several examples can be found in which quantile regressions have been used on pooled cross-sectional data, sometimes with the aid of yearly dummies or after cross-sectionally demeaning the data, a procedure that can avoid problems of macro trends but not the problem of firm-specific (possibly constant over time) characteristics. Examples can be found in firm growth rate autoregressions (Coad and Hölzl, 2009), vector autoregressions including firm growth among the variables (Coad and Rao, 2010), and firm size autoregressions (Bottazzi *et al.*, 2007).

However, the methodological research on quantile regressions is currently advancing in the direction of panel data techniques, thus providing the researcher with new methods that only a few years ago were not available (see for instance Koenker, 2004; Lamarche, 2010; Galvao, 2011) and thus giving the industrial researcher the choice of considering explicitly the panel nature of the data, or pooling the data in a unique cross-section, or only considering and comparing results obtained on different cross-sectional waves (as in Coad, 2007).

To explain why we prefer not to go in the direction of panel data techniques, it is worth considering the simpler case of an OLS estimation in the context of panel autoregressions. There are two possibilities: individual effects are present in the firm size autoregression model, or individual effects are present (also) in the firm growth autoregression model. In the former case, much of the estimation bias would be lost when log-differencing the data (what we implicitly do because we consider a growth, rather than size, autoregression model), although results would still be inconsistent (see Bond, 2002: 145). In the latter case, the problem is more difficult to face: following Bond (2002), we could rewrite our growth autoregression model (4) as

$$g_i(t) = \alpha + \beta g_i(t-1) + \eta_i + u_i(t) \quad (6)$$

where  $\eta_i$  is an unobserved individual-specific time-invariant effect (for sake of simplicity, now we do not differentiate  $\beta$  across quantiles).

Notice that, in the context of autoregressions, there is an obvious correlation between the individual-specific variable and the other independent variables (that



include the dependent lagged one), and therefore panel data techniques should deal with fixed, rather than random, effects.

If the data generating process was well proxied by model (6), using a different model that neglects the existence of individual effects (the whole problem can be seen as an omitted variable bias) would lead to an OLS estimation of  $\beta$  that is biased upward (Bond, 2002). In the economic context of our study, the bias can be explained in the following way.

Suppose there are firm-specific features that push some firms' growth constantly over time, and they are represented by  $\eta_i$  in equation (6). At each time, the growth of the firm depends on this push  $\eta_i$ , that is, firm-specific and constant over time, and on the growth of the same firm in the previous time (partly due to the same push). The autoregression parameter only refers to this second influence, and not to the constant push. If we now exclude  $\eta_i$  from the estimated model, as we have excluded it throughout our empirical analysis, our estimation would attribute to  $\beta$  also the effect of  $\eta_i$ , that is, the persistent growth of a firm whose  $\eta_i$  is positive and high would be considered as a domino effect over time of subsequent positive growth rates, each one caused by the previous one (as in a pure autoregressive model), rather than by a push that is firm-specific and influences the growth rates without being influenced by it. Roughly speaking, when not considering the existence of  $\eta_i$ , a firm that is persistently growing around time  $\tau$  looks like a firm that, because of an event occurred in a given year, say  $u_i(\tau - 5)$ , has entered a growth path that becomes self-sustained until a bad event, say  $u_i(\tau + 5)$ , occurs to stop the path. Instead, in the real data-generating process modeled as (6), the persistent path of positive growth rates is not self-sustained, nor is triggered by the event  $u_i(\tau - 5)$ , but is created by the firm-specific characteristics  $\eta_i$  that make the path of persistent growth resilient to, rather than dependent on, the events  $u_i(t)$ . In technical terms, all this line of reasoning can be summarized as an upward biased estimation of  $\beta$ .

We want to draw the attention on two points: the meaning of  $\beta$  and the meaning of  $u_i(t)$ .  $\beta$  does not have the same meaning in models (4) and (6). In particular, in Equation (4),  $\beta$  is much closer to an autocorrelation parameter, that is, to a parameter that describes, rather than explains, the growth persistence. A high  $\beta$  in (4) corresponds to a high growth rate autocorrelation, while in (6) it simply contributes to a high growth rate autocorrelation. Therefore, if we use the estimation of  $\beta$  as a device to detect some growth autocorrelation for some subsets of firms (in particular, of firms that correspond to some conditional quantiles of the empirical growth distribution), we are referring to the  $\beta$  of (4). We are detecting, rather than explaining, the autocorrelation. A full explanation of the autocorrelation would lead us first to use model (6), and then we could go a step further, introducing other explanatory variables that are firm-specific and possibly not constant over time, reducing more and more the importance of both  $\beta$  and  $\eta_i$  in the model. Notably, this is the path toward which the research on industrial dynamics is now moving (see Stam, 2010), and, if followed with a quantile regression approach, entails many more technical

challenges in the interpretation of empirical results (since conditional growth distributions become conditional on many variables, and so even the economic meaning of conditional quantiles becomes difficult to grasp). However, this is not the case of our study, which we confine to empirically detecting hidden phenomena of firm growth, thus leaving to economic theory the role of understanding them, without the claim of testing empirically the economic theory itself. In the same sense, the interpretation of the errors in (4) goes more toward summarizing many omitted variables, possibly firm-specific, and that we do not consider explicitly, rather than toward the description of events that can trigger alone a self-sustained persistent growth. As with all the omitted variable cases, and in particular for the cases in which omitted variables are constant over time in a panel context, the estimations are affected by the simplification of the model.

The fact that in model (4) the parameter  $\beta$  varies across different conditional quantiles makes the picture even more complicated. We mix up different years of observations in each decile, with consequences on the quality of our estimates that cannot be easily assessed (a problem that affects also the transition matrix analysis of the following section, where we consider unconditional quantiles). While hoping that the current gap between methodological advances in quantile regressions and empirical applications to industrial dynamics will soon be filled, we consider our study as a necessary step for broadening the view on firm growth quantile autoregressions, by showing how our quantile regression results can be connected with our other results obtained in this study using alternative approaches on the same data, and by shedding new light on the quantile regression results already obtained in the previous literature.

## 8. Transition matrices

As Table 3 illustrates, although intra-distribution mobility seems high, the values on the matrix diagonal show clear persistence in growth rates. Values  $>0.10$  occur in all columns. The diagonal values in the first two columns (0.12 and 0.21) can be characterized as persistent fast-declining firms and the values in the last two columns (0.22 and 0.11) as persistent fast-growing firms. So we have winners and losers. However, the high values in the lower left and upper right corners confirm the idea that firms that experience extreme growth behavior in a given year are also likely to experience extreme growth behavior in the following year but in the opposite direction, revealing a strong rebound effect. As only some of the outperformers are able to repeat their performance the next year, and the others are likely to move toward the low quantiles of the distribution, we can claim that it is not possible to generalize about any “hare and tortoise” or “snowball” effect: the only persistence over time common to the whole set of outperformers and underperformers is a higher tendency to experience extreme events (i.e. a much higher dynamics than the rest of the population).

**Table 3** One-year probability transition matrix of growth rates: all sizes

	Decile in year $t$									
	1	2	3	4	5	6	7	8	9	10
Decile in year $t - 1$										
1	0.12	0.13	0.09	0.08	0.08	0.07	0.08	0.08	0.14	0.15
2	0.09	0.21	0.11	0.06	0.10	0.05	0.09	0.06	0.11	0.12
3	0.08	0.10	0.20	0.15	0.11	0.07	0.12	0.05	0.03	0.09
4	0.08	0.08	0.19	0.23	0.07	0.11	0.03	0.06	0.07	0.08
5	0.08	0.08	0.08	0.10	0.13	0.15	0.08	0.17	0.04	0.09
6	0.07	0.04	0.07	0.11	0.19	0.17	0.08	0.14	0.06	0.08
7	0.08	0.10	0.04	0.08	0.07	0.11	0.20	0.17	0.08	0.08
8	0.09	0.06	0.08	0.07	0.08	0.15	0.16	0.12	0.10	0.09
9	0.15	0.11	0.06	0.03	0.08	0.05	0.09	0.08	0.22	0.12
10	0.16	0.11	0.08	0.08	0.08	0.08	0.08	0.08	0.13	0.11

To find out whether size plays a role in this growth behavior, we replicate the analysis at a more disaggregated level (see Tables 3–6, respectively) for each firm size group. For micro firms (Table 4), the results are similar to those obtained for the whole sample, except for a slightly higher “rebound” effect signaled by very high values at the bottom-left and top-right corners. Compared with non-micro firms, the rebound effect is highest for micro firms. At the same time, persistent Schumpeterian winners and losers are much more pronounced for micro firms than non-micro firms based on the high diagonal values in the first two and last two columns in the table. This result is in line with the notion of a Schumpeterian Mark I, or entrepreneurial regime versus a Schumpeter Mark II, or routinized regimes scenario (Winter, 1984; Malerba and Orsenigo, 1996; Breschi *et al.*, 2000). In the first regime, innovation is driven by the entry of micro firms operating in a technologically uncertain environment. Some firms catch on to a promising innovation path, while other firms bet on a technological dead end. Consequently, one would expect more “Schumpeterian” winners and losers expressed by the persistence of extreme growth events. In the second regime, innovation stems primarily from large firms operating in a technologically more stable environment, with less heterogeneous innovation rates among firms, spread across more products and over longer time horizons. This regime is consistent with less persistence in extreme growth events.

For small firms (Table 5), we still find higher than expected values on the main diagonal, but this persistence is much lower than for micro firms, and the top-left cell contains a value that is  $<0.1$ , i.e. it is difficult for small firms to experience two consecutive strong size decreases. Instead, the 0.16 probability in the bottom-left cell

**Table 4** Probability transition matrix of growth rates: micro firms

	Decile in year $t$									
	1	2	3	4	5	6	7	8	9	10
Decile in year $t - 1$										
1	0.12	0.11	0.10	0.08	0.07	0.06	0.08	0.07	0.13	0.17
2	0.12	0.24	0.07	0.07	0.11	0.05	0.05	0.10	0.07	0.12
3	0.08	0.08	0.24	0.15	0.11	0.07	0.10	0.07	0.02	0.08
4	0.07	0.11	0.18	0.20	0.07	0.11	0.04	0.03	0.11	0.08
5	0.08	0.08	0.08	0.10	0.13	0.17	0.08	0.15	0.05	0.09
6	0.06	0.05	0.10	0.12	0.15	0.17	0.09	0.13	0.07	0.08
7	0.07	0.10	0.02	0.08	0.11	0.11	0.16	0.19	0.08	0.07
8	0.09	0.05	0.06	0.10	0.09	0.15	0.17	0.08	0.13	0.08
9	0.14	0.08	0.05	0.02	0.08	0.04	0.16	0.09	0.20	0.13
10	0.16	0.10	0.08	0.09	0.08	0.08	0.08	0.08	0.14	0.11

**Table 5** Probability transition matrix of growth rates: small firms

	Decile in year $t$									
	1	2	3	4	5	6	7	8	9	10
Decile in year $t - 1$										
1	0.08	0.09	0.09	0.09	0.10	0.09	0.11	0.11	0.10	0.13
2	0.08	0.11	0.12	0.09	0.11	0.09	0.09	0.10	0.11	0.10
3	0.09	0.08	0.11	0.13	0.11	0.10	0.08	0.10	0.11	0.10
4	0.10	0.09	0.09	0.12	0.16	0.08	0.11	0.07	0.09	0.10
5	0.11	0.09	0.09	0.10	0.11	0.14	0.10	0.08	0.08	0.10
6	0.09	0.10	0.10	0.08	0.09	0.11	0.13	0.10	0.10	0.09
7	0.10	0.10	0.10	0.10	0.07	0.12	0.11	0.12	0.11	0.09
8	0.08	0.11	0.11	0.10	0.08	0.09	0.10	0.13	0.10	0.09
9	0.11	0.13	0.11	0.09	0.08	0.09	0.08	0.11	0.11	0.10
10	0.16	0.09	0.08	0.10	0.10	0.09	0.09	0.09	0.10	0.10

reveals that relatively many small firms experience a strong rebound effect from extreme positive growth at time  $t - 1$  to extreme negative growth at time  $t$ .

For medium firms (Table 6), the size decrease cannot be linked precisely to a particular autocorrelation pattern, as the left part of the table shows high values

**Table 6** Probability transition matrix of growth rates: medium firms

	Decile in year $t$									
	1	2	3	4	5	6	7	8	9	10
Decile in year $t - 1$										
1	0.12	0.10	0.10	0.09	0.10	0.08	0.10	0.11	0.09	0.12
2	0.09	0.13	0.09	0.11	0.12	0.08	0.08	0.10	0.09	0.10
3	0.10	0.10	0.08	0.13	0.13	0.09	0.08	0.09	0.10	0.11
4	0.10	0.12	0.09	0.11	0.11	0.12	0.09	0.08	0.08	0.10
5	0.11	0.08	0.11	0.11	0.10	0.13	0.12	0.08	0.07	0.10
6	0.09	0.09	0.12	0.09	0.09	0.12	0.12	0.10	0.10	0.09
7	0.09	0.11	0.09	0.09	0.08	0.12	0.11	0.11	0.11	0.08
8	0.09	0.11	0.11	0.08	0.09	0.09	0.10	0.11	0.13	0.10
9	0.08	0.10	0.12	0.10	0.07	0.07	0.10	0.12	0.13	0.11
10	0.12	0.08	0.09	0.09	0.11	0.11	0.10	0.10	0.11	0.08

around the main diagonal and in the bottom-left corner. This result was signaled by the quantile regression results in the bottom-left panel of Figure 7, where the very large confidence interval prevented the identification of a clear pattern of autocorrelation for medium-sized firms. However, the high values in the first and last cells in the first column in Table 6 add new information: medium firms that experience very bad performance at time  $t$  are likely to have emerged from another extreme event, positive for some and negative for others, at time  $t - 1$ .

Finally, for large firms (Table 7), it should be noted that the companies experiencing the highest-growth events at time  $t$  (corresponding to the last column of the matrix) have a particularly high probability of coming from the two highest-growth deciles (0.14 and 0.12) and the lowest-growth decile (0.13) at time  $t - 1$ . On the other hand, firms showing relatively bad performance at time  $t$  seem to be those that experienced episodes of high growth (high values in the bottom-left part of the table) rather than low growth (low values in the top-left part of the table). As shown by the quantile regressions, growth persistence is still important in the central part of the distribution (high values in the central part of the main diagonal in Table 7).

## 9. Acquisitions and spin-offs

The above analysis examines patterns of growth, but does not distinguish between different modes of growth. In this section we concentrate on acquisitions and/or spin-offs as sources of external growth as opposed to internal growth. While still pooling in a unique cross-section all the growth rates observed at each year between

**Table 7** Probability transition matrix of growth rates: large firms


	Decile in year $t$									
	1	2	3	4	5	6	7	8	9	10
Decile in year $t - 1$										
1	0.09	0.11	0.09	0.11	0.07	0.11	0.10	0.08	0.10	0.13
2	0.08	0.11	0.09	0.11	0.15	0.07	0.10	0.11	0.10	0.09
3	0.08	0.09	0.10	0.09	0.13	0.08	0.11	0.11	0.11	0.09
4	0.08	0.07	0.10	0.17	0.11	0.10	0.09	0.11	0.09	0.08
5	0.09	0.10	0.10	0.12	0.12	0.12	0.10	0.08	0.08	0.07
6	0.12	0.10	0.09	0.08	0.10	0.14	0.11	0.09	0.08	0.08
7	0.11	0.09	0.10	0.07	0.09	0.13	0.09	0.13	0.09	0.10
8	0.09	0.10	0.14	0.08	0.08	0.07	0.10	0.11	0.13	0.10
9	0.12	0.11	0.10	0.07	0.07	0.08	0.10	0.10	0.11	0.14
10	0.13	0.11	0.09	0.09	0.08	0.08	0.11	0.08	0.11	0.12

1995 and 2004, we now distinguish among firms that were involved in acquisitions and firms from which other firms were spun off in the corresponding period.

We can define as “*mutations*” acquisitions and spin-off events, and “*mutated firms*,” firms that during time  $t$  experienced acquisitions and/or spin-offs, since firm size is proxied by the number of employees at the *end* of each period.

In the same figure we plot the growth rate distribution of all the firms and the growth rate distribution of the firms who made the acquisitions, to provide the histogram in Figure 11 (the vertical axis now refers to the log absolute frequencies instead of relative frequencies, so that the reader can directly observe the weight of the acquiring firms over the whole sample, and all the cross-sectional waves have been pooled). The left tail of the acquiring firms’ distribution appears to be parallel to the left tail of the whole sample, which means that, given the log scale of the vertical axis, the proportion of firms making acquisitions remains constant for all levels of negative growth. For the right tail the picture is different: the tail of the acquiring firms appears fatter, signaling that the proportion of acquiring firms increases as the growth rate becomes higher. Figure 12 shows that the asymmetry is greater in the growth rate distributions of the firms that spun off other firms. The right tail is very steep while the left tail is very fat, signaling that only a few of the split firms experienced positive growth. If we merge the two graphs, we obtain Figure 13, which suggests that especially in the body of the distribution there may be firms that experienced both acquisitions and spin-offs.

We then estimated the TPM on the population of firms excluding those that experienced episodes of acquisitions or spin-off during each 3-year time span used




**Figure 11** Growth rates distribution, 1995–2004: all firms (light colour) vs. firms making acquisitions (dark colour).

to estimate the matrix (see Table 8). With the exception of the cells in the bottom-right corner, whose higher numbers mean a higher persistence of growth in the top quantiles, the TPM is similar to that obtained from pooling internal and external growth. On the other hand, moderate growth episodes appear less persistent, while intra-distributional mobility in low quantiles appears to be unchanged. In other words, extreme growth events based on external sources are not likely to be repeated over time, or at least not at the same magnitude, whereas the (relatively) extreme events that are based solely on internal growth may be persistent over time. Note also that excluding only firms that have made acquisitions (Table 9) or have spun off firms (Table 10) does not change the matrix significantly. Thus, our conclusions regarding the coexistence of bouncing firms and persistent outperformers are shown to be robust if we exclude firms involved in acquisitions or spin-offs.

## 10. Conclusions

Using data on the whole population of Dutch manufacturing firms for the period 1994–2004, we tested for the existence of firms that are able to sustain exceptionally




**Figure 12** Growth rates distribution, 1995–2004: all firms (light colour) vs. firms in the process of spin-off (dark colour).

good performance over time, and analyzed their behavior in the light of the stylized facts in the literature on industrial dynamics and in particular Gibrat's Law.

The data show that Dutch manufacturing firms display the typical tent-shaped growth rate distribution observed by Stanley *et al.* (1996). However, plotting the distributions for four different size groups shows that for large firms extreme events are mainly negative: extreme positive growth is rare. This may be due to the fact that although large firms may be able to exploit more business opportunities than smaller firms, they cannot overcome the barrier of market size. We can see also that the location of the distribution appears stable across different size subsamples, with the growth median always close to zero. Thus, firm size seems to have an effect only on the skewness in the growth distribution, which is lower for larger firms.

This leads to a different interpretation of the quantile regression results in the literature, which are partially confirmed by our data. The quantile autoregressions of firm growth rates show that extreme growth events are likely to be correlated negatively over time, although the negative autocorrelation is weaker for large firms. It may be difficult for firms over time to repeat very good or very bad performance as some sort of rebound effect tends to occur after an extreme event. A technical explanation for this may be that extreme growth events cause important changes in





**Figure 13** Growth rates distribution, 1995–2004: all firms (light colour) vs. firms making acquisitions or in the process of spin-off (dark colour).

firm size, and a positive extreme growth event for a small firm (less likely for a large firm) results in a medium–large size firm and thus a lower probability of another strong growth event. The same explanation holds for large firms experiencing high negative growth, which reduces their size and makes them more likely to experience an episode of positive growth.

The estimated TPM, however, showed that, even if a bouncing back from an extreme growth event is the rule for most firms, there are still firms that are able to repeat their good performance over time. On the one hand, this result confirms that very dynamic firms are usually very different from other firms in that they can show strong autocorrelation patterns (Coad, 2007). On the other hand, within the sample of very dynamic firms we can see that there are two, coexisting, well-defined subsets, a bigger one showing high negative autocorrelation and a smaller one showing high positive autocorrelation. If a firm, at a given year, experiences an extreme growth event, it is safe to say that the same firm is unlikely to be stable in the following year and can be expected to experience another extreme event; however, it is not possible *a priori* to predict the direction of such an event. Many of today's extreme growers will decrease in size tomorrow, but some will maintain positive extreme growth, while a few will stabilize at the newly achieved size.

**Table 8** Probability transition matrix of growth rates: whole sample, excluding firms that have been subject of acquisitions or spin-off episodes

		Decile in year $t$									
		1	2	3	4	5	6	7	8	9	10
Decile in year $t - 1$											
1		0.12	0.12	0.09	0.08	0.08	0.07	0.08	0.08	0.14	0.15
2		0.09	0.21	0.11	0.07	0.10	0.04	0.08	0.07	0.10	0.12
3		0.08	0.09	0.21	0.16	0.11	0.08	0.09	0.06	0.03	0.09
4		0.08	0.09	0.17	0.24	0.07	0.12	0.04	0.06	0.08	0.08
5		0.08	0.08	0.09	0.09	0.13	0.16	0.08	0.17	0.04	0.08
6		0.07	0.04	0.08	0.09	0.20	0.17	0.08	0.14	0.06	0.07
7		0.08	0.10	0.04	0.08	0.08	0.11	0.19	0.17	0.08	0.08
8		0.09	0.06	0.08	0.08	0.09	0.14	0.17	0.10	0.11	0.08
9		0.15	0.10	0.06	0.03	0.07	0.04	0.11	0.07	0.23	0.13
10		0.16	0.10	0.08	0.08	0.08	0.08	0.08	0.08	0.13	0.12

**Table 9** Probability transition matrix of growth rates: whole sample, excluding firms that have operated acquisitions

		Decile in year $t$									
		1	2	3	4	5	6	7	8	9	10
Decile in year $t - 1$											
1		0.12	0.12	0.09	0.08	0.08	0.07	0.08	0.08	0.14	0.15
2		0.09	0.21	0.12	0.06	0.10	0.04	0.08	0.07	0.10	0.12
3		0.08	0.10	0.20	0.16	0.11	0.08	0.09	0.06	0.03	0.09
4		0.08	0.08	0.17	0.24	0.06	0.12	0.04	0.06	0.08	0.08
5		0.08	0.09	0.08	0.09	0.12	0.16	0.08	0.17	0.04	0.08
6		0.07	0.04	0.08	0.09	0.20	0.17	0.08	0.14	0.06	0.07
7		0.08	0.10	0.04	0.08	0.08	0.11	0.19	0.17	0.08	0.08
8		0.09	0.06	0.08	0.08	0.09	0.14	0.17	0.10	0.11	0.08
9		0.15	0.10	0.06	0.03	0.07	0.04	0.11	0.07	0.23	0.13
10		0.16	0.10	0.08	0.08	0.08	0.08	0.08	0.08	0.13	0.12

**Table 10** Probability transition matrix of growth rates: whole sample, excluding firms that have been subject to spin-off

	Decile in year $t$									
	1	2	3	4	5	6	7	8	9	10
Decile in year $t - 1$										
1	0.12	0.13	0.09	0.08	0.08	0.07	0.08	0.08	0.14	0.15
2	0.09	0.21	0.11	0.06	0.10	0.05	0.08	0.06	0.11	0.12
3	0.08	0.09	0.20	0.15	0.11	0.07	0.12	0.05	0.03	0.09
4	0.08	0.08	0.19	0.23	0.07	0.11	0.03	0.06	0.07	0.08
5	0.08	0.08	0.08	0.10	0.13	0.15	0.08	0.17	0.05	0.09
6	0.07	0.04	0.07	0.11	0.19	0.17	0.08	0.13	0.06	0.08
7	0.07	0.10	0.04	0.08	0.07	0.11	0.20	0.17	0.08	0.08
8	0.09	0.06	0.08	0.07	0.09	0.15	0.15	0.12	0.10	0.09
9	0.15	0.10	0.06	0.03	0.08	0.05	0.10	0.08	0.22	0.13
10	0.16	0.11	0.08	0.08	0.08	0.08	0.08	0.08	0.13	0.11

The coexistence of bouncing firms with persistent outperformers shows up more clearly in the case of micro firms, thus providing support for the notion of a Schumpeterian Mark I, or entrepreneurial, regime, as opposed to Schumpeter Mark II, or routinized regime (Winter, 1984; Malerba and Orsenigo, 1996; Breschi *et al.*, 2000). It is also robust to the exclusion, from the analysis, of external growth, since we repeated the exercise after excluding firms that experienced episodes of acquisition or spin-off. On the other hand, our results do not completely support the thesis that “Gazelles” are the source of long-lasting employment. In fact, when we exclude micro firms (since Gazelles are usually defined as young small–medium firms, for example, with >10 or 20 employees depending on the authors, see Henrekson and Johansson, 2010) the probability that a positive growth event is followed by another positive one is small.

Our analysis was aimed at shedding light on the heterogeneity of persistent growth paths that can be detected within a firm population. The identification of the determinants of these different paths (such as firm age, innovative capabilities, export propensity, etc.) goes beyond our current purpose, but certainly deserves attention and further research. Some fruitful attempts in this direction have been made recently by Hölzl (2009), Hölzl and Friesenbichler (2010), and Coad and Rao (2010) who investigate the causal relationship between firm innovative capabilities and firm growth.

On the other hand, the recent progress in the methodologies concerning quantile regressions could provide an improvement in the analysis of firm heterogeneity, both

in summarizing the different paths of growth followed by different subsets of firms, and in identifying the different responses, in terms of growth, that different firms can have to changes in the values of the explanatory variables. One of the main limitations of our study consists in not having explored the panel nature of our data. The developments in quantile regressions for panel data, as in Lamarche (2010) and Galvao (2011), could lead to a refinement of our work through the adoption of new models, which face firm heterogeneity from a different perspective and are not subject to the estimation risks of a pooled cross-section analysis.

## Acknowledgements

The authors want to thank Andrew Bernard, Tom Brökel, Werner Hölzl, André Lorentz, Sandro Sapiro, Erik Stam, the participants at the DIME Workshop on “The dynamics of firm evolution: productivity, profitability and growth” held in Pisa on October 3–4, 2008, and the participants at the EMAEE Conference on “Evolution, behavior and organization” held in Jena on May 21–23, 2009. The empirical analysis for this research was carried out at the Centre for Research of Economic Microdata at Statistics Netherlands (CBS). The views expressed in this article are those of the authors and do not necessarily reflect the policies of Statistics Netherlands. The authors thank Gerhard Meinen and CBS on-site staff for their collaboration.

## Funding

This work was supported by Utrecht University [High Potential Grant (HIPO) to E.C. and K.F.]; and the University of Bergamo [grant ex 60%, n. 60CEFI10, Department of Management, Economics and Quantitative Methods, to E.C.].

## References

- Almus, M. and E. Nerlinger (2000), ‘Testing “Gibrat’s Law” for young firms - empirical results for West Germany,’ *Small Business Economics*, **15**(1), 1–12.
- Anderson, T. W. (1962), ‘On the distribution of the two-sample cramér-von mises criterion,’ *Annals of Mathematical Statistics*, **33**(3), 1148–1159.
- Axtell, R. (2001), ‘Zipf distribution of U.S. firm sizes,’ *Science*, **293**, 1818–1820.
- Baily, M. N. and A. K. Chakrabarty (1985), ‘Innovation and productivity an US industry,’ *Brookings papers on Economic Activity*, **2**, 609–632.
- Baldwin, J. R. and M. Rafiqzaman (1995), ‘Selection versus evolutionary adaptation, learning and post-entry performance,’ *International Journal of Industrial Organization*, **13**(4), 501–522.

- Bartelsman, E. J. and M. Doms (2000), 'Understanding productivity: lessons from longitudinal microdata,' *Journal of Economic Literature*, **38**(3), 569–594.
- Bond, S. R. (2002), 'Dynamic panel data models: a guide to micro data methods and practice,' *Portuguese Economic Journal*, **1**(2), 141–162.
- 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**(7), 1161–1187.
- Bottazzi, G., E. Cefis and G. Dosi (2002), 'Corporate growth and industrial structures: some evidence from the Italian manufacturing industry,' *Industrial and Corporate Change*, **11**(4), 705–723.
- Bottazzi, G. and A. Secchi (2003), 'Why are distributions of firm growth rates tent-shaped?' *Economic Letters*, **80**(3), 415–420.
- Bottazzi, G., E. Cefis, G. Dosi and A. Secchi (2007), 'Invariances and diversities in the patterns of industrial evolution: some evidence from Italian manufacturing industries,' *Small Business Economics*, **29**(1), 137–159.
- Bottazzi, G., A. Coad, N. Jacoby and A. Secchi (2011), 'Corporate growth and industrial dynamics: evidence from French manufacturing,' *Applied Economics*, **43**(1), 103–116.
- Breschi, S., F. L. Malerba and L. Orsenigo (2000), 'Technological regimes and schumpeterian patterns of innovation,' *Economic Journal*, **110**(463), 388–410.
- Cefis, E. and L. Orsenigo (2001), 'The persistence of innovative activities: a cross-countries and cross-sectors comparative analysis,' *Research Policy*, **30**(7), 1139–1158.
- Cefis, E. (2003), 'Is there any persistence in innovative activities?' *International Journal of Industrial Organization*, **21**(4), 489–515.
- Chesher, A. (1979), 'Testing the law of proportionate effect,' *Journal of Industrial Economics*, **27**(4), 403–411.
- Coad, A. (2007), 'A closer look at serial growth rate correlation,' *Review of Industrial Organization*, **31**(1), 69–82.
- Coad, A. and W. Hözl (2009), 'On the autocorrelation of growth rates,' *Journal of Industry, Competition and Trade*, **9**(2), 139–166.
- Coad, A. and R. Rao (2010), 'Firm growth and R&D expenditure,' *Economics of Innovation and New Technology*, **19**(2), 127–145.
- Cubbin, J. and P. Geroski (1987), 'The convergence of profits in the long run: inter-firm and inter-industry comparisons,' *Journal of Industrial Economics*, **35**(4), 427–442.
- Davis, S. J., J. Haltiwanger, R. Jarmin and J. Miranda (2006), 'Volatility and dispersion in business growth rates: publicly traded versus privately held firms,' *NBER Working Papers 12354*. National Bureau of Economic Research, Inc.
- Davis, S. J., J. Haltiwanger and S. Schuh (1996), 'Small business and job creation: dissecting the myth and reassessing the facts,' *Small Business Economics*, **8**(4), 297–315.

- Dosi, G. (2007), 'Statistical regularities in the evolution of industries: a guide through some evidence and challenges for the theory,' in F. Malerba and S. Brusoni (eds), *Perspectives on Innovation*. Cambridge University Press: Cambridge, UK.
- Dosi, G. and R. R. Nelson (2009), 'Technical change and industrial dynamics as evolutionary processes, LEM WP 2009/07, LEM, Scuola Superiore Sant'Anna: Pisa. Forthcoming,' in B. Hall and N. Rosenberg (eds), *Handbook of Innovation*. Elsevier: Amsterdam/New York.
- Dunne, P. and A. Hughes (1994), 'Age, size, growth and survival: UK companies in the 1980s,' *Journal of Industrial Economics*, **42**(2), 115–141.
- Evans, D. S. (1987), 'The relationship between firm growth, size and age: estimates for 100 manufacturing industries,' *Journal of Industrial Economics*, **35**(4), 567–581.
- Fotopoulos, G. and I. Giotopoulos (2010), 'Gibrat's law and persistence of growth in Greek manufacturing,' *Small Business Economics*, **35**(2), 191–202.
- Galvao, A. F. Jr (2011), 'Quantile regression for dynamic panel data with fixed effects,' *Journal of Econometrics*, **164**(1), 142–157.
- Garnsey, E. and P. Heffernan (2005), 'Growth setbacks in new firms,' *Futures*, **37**(7), 675–697.
- Geroski, P. and A. Jacquemin (1988), 'The persistence of profits: a European Comparison,' *Economic Journal*, **98**, 375–389.
- Geroski, P. A., J. Van Reenen and C. F. Walters (1997), 'How persistently do firms innovate?' *Research Policy*, **26**(1), 33–48.
- Glen, J., K. Lee and A. Singh (2001), 'Persistence of profitability and competition in emerging markets,' *Economics Letters*, **72**, 247–253.
- Goddard, J. A. and J. O. S. Wilson (1999), 'The persistence of profit: a new empirical interpretation,' *International Journal of Industrial Organization*, **17**, 663–687.
- Gschwandtner, A. (2004), 'Evolution of profit persistence in the US: evidence from four 20-years periods,' *Vienna Economics Papers*, **0410**. University of Vienna, Department of Economics.
- Hall, B. H. (1987), 'The relationship between firm size and firm growth in the US manufacturing sector,' *Journal of Industrial Economics*, **35**(4), 583–606.
- Harhoff, D., K. Stahl and M. Woywode (1998), 'Legal form, growth and exit of west german firms - empirical results for manufacturing, construction, trade and service industries,' *Journal of Industrial Economics*, **46**(4), 453–488.
- Heckman, J. J. (1979), 'Sample selection bias as a specification error,' *Econometrica*, **47**(1), 153–161.
- Henrekson, M. and D. Johansson (2010), 'Gazelles as job creators: a survey and interpretation of the evidence,' *Small Business Economics*, **35**(2), 227–244.
- Higson, C., S. Holly, P. Kattuman and S. Platis (2004), 'The business cycle, macroeconomic shocks and the cross-section: the growth of UK quoted companies,' *Economica*, **71**, 299–318.
- Hymer, S. and P. Pashigian (1962), 'Turnover of firms as a measure of market behavior,' *The Review of Economics and Statistics*, **44**(1), 82–87.

- Hoel, P. G., S. C. Port and A. C. J. Stone (1987), *Introduction to Stochastic Processes*. Waveland Press Inc.: Prospect Heights, IL.
- Hözl, W. (2009), 'Is the R&D behaviour of fast-growing SMEs different? Evidence from CIS III data for 16 countries,' *Small Business Economics*, **33**(1), 59–75.
- Hözl, W. and K. Friesenbichler (2010), 'High-growth firms, innovation and the distance to the frontier,' *Economics Bulletin*, **30**(2), 1016–1024.
- Jensen, B. J. and R. H. McGuckin (1997), 'Firm performance and evolution: empirical regularities in the U.S. microdata,' *Industrial and Corporate Change*, **6**(1), 25–47.
- Koenker, R. (2004), 'Quantile regression for longitudinal data,' *Journal of Multivariate Analysis*, **91**(1), 74–89.
- Koenker, R. and G. Bassett (1978), 'Regression quantiles,' *Econometrica*, **46**, 33–50.
- Koenker, R. and K. F. Hallock (2001), 'Quantile regression,' *Journal of Economic Perspectives*, **15**(4), 143–156.
- Lamarche, C. (2010), 'Robust penalized quantile regression estimation for panel data,' *Journal of Econometrics*, **157**(2), 396–408.
- Lotti, F., E. Santarelli and M. Vivarelli (2001), 'The relationship between size and growth: the case of Italian newborn firms,' *Applied Economics Letters*, **8**(7), 451–454.
- Lotti, F., E. Santarelli and M. Vivarelli (2003), 'Does Gibrat's law hold in the case of young, small firms?' *Journal of Evolutionary Economics*, **13**(3), 213–235.
- Malerba, F. and L. Orsenigo (1996), 'Schumpeterian patterns of innovation are technology-specific,' *Research Policy*, **25**, 451–478.
- Mueller, D. C. (1986), *Profits in the Long Run*. Cambridge University Press: Cambridge, England.
- Mueller, D. C. (1990), *The Dynamics of Company Profits: an International Comparison*. Cambridge University Press: Cambridge.
- Peters, B. (2009), 'Persistence of innovation: stylised facts and panel data evidence,' *Journal of Technology Transfer*, **34**, 226–243.
- Raymond, W., P. Mohnen, F. Palm and S. S. van der Loeff (2010), 'Persistence of Innovation in Dutch manufacturing: is it spurious?' *Review of Economics and Statistics*, **92**(3), 495–504.
- Reichstein, T., M. S. Dahl, B. Ebersberger and M. B. Jensen (2010), 'The devil dwells in the tails: a quantile regression approach to firm growth,' *Journal of Evolutionary Economics*, **20**(2), 219–231.
- Reichstein, T. and M. B. Jensen (2005), 'Firm size and firm growth rate distributions - the case of Denmark,' *Industrial and Corporate Change*, **14**(6), 1145–1166.
- Rumelt, R. P. (1991), 'How much does industry matter?' *Strategic Management Journal*, **12**, 167–185.
- Singh, A. and G. Whittington (1975), 'The size and growth of firms,' *Review of Economic Studies*, **42**(1), 15–26.

- Stam, E. (2010), 'Growth beyond Gibrat: firm growth processes and strategies,' *Small Business Economics*, **35**(2), 129–135.
- Stanley, M. H. R., L. A. N. Amaral, S. V. Buldyrev, S. Havlin, H. Leschorn, P. Maass, M. A. Salinger and H. E. Stanley (1996), 'Scaling behavior in the growth of companies,' *Nature*, **379**, 804–806.
- Sutton, J. (2002), 'The variance of firm growth rates: the 'scaling' puzzle,' *Physica A: Statistical Mechanics and its Applications*, **312**(3–4), 577–590.
- Syverson, C. (2004), 'Market structure and productivity: a concrete example,' *Journal of Political Economy*, **112**(6), 1181–1222.
- Winter, S. G. (1984), 'Schumpeterian competition in alternative technological regimes,' *Journal of Economic Behavior and Organization*, **5**, 137–158.