Economics Letters 164 (2018) 5–9

Contents lists available at ScienceDirect

Economics Letters

journal homepage: www.elsevier.com/locate/ecolet

Gibrat's law and quantile regressions: An application to firm growth[☆]

Roberta Distante[a], Ivan Petrella[b,c], Emiliano Santoro[d,*]

[a] A.P. Moller-Maersk, Strategy Department, Esplanaden 50, 1098, Copenhagen, Denmark
[b] CEPR, United Kingdom
[c] Warwick Business School, University of Warwick, Scarman Rd, CV4 7AL Coventry, United Kingdom
[d] Department of Economics, University of Copenhagen, Østerfarimagsgade 5, 1353, Copenhagen, Denmark

ARTICLE INFO

Article history:
Received 5 November 2017
Accepted 19 December 2017
Available online 27 December 2017

JEL classification:
D22
L11
C21

Keywords:
Firm growth
Size
Age
Conditional quantiles

ABSTRACT

The nexus between firm growth, size and age in U.S. manufacturing is examined through the lens of quantile regression models. This methodology allows us to overcome serious shortcomings entailed by linear regression models employed by much of the existing literature, unveiling a number of important properties. Size pushes both low and high performing firms towards the median rate of growth, while age is never advantageous, and more so as firms are relatively small and grow faster. These findings support theoretical generalizations of Gibrat's law that allow size to affect the variance of the growth process, but not its mean (Cordoba, 2008).

1. Introduction

The literature on industrial dynamics has devoted much attention to unveiling the nexus between firm size and growth. In this respect, the theoretical proposition known as Gibrat's law (Gibrat, 1931) – which predicts randomness of firm growth rates – has been widely tested. Linear econometric frameworks employed to validate this hypothesis have delivered mixed evidence (see Sutton, 1997 for a comprehensive review of the literature). While some studies have found a tendency for large firms to grow faster than small ones (Samuels, 1965; Singh and Whittington, 1975), others have appreciated a tendency of small firms to grow faster (Hall, 1987; Evans, 1987a, b; Dunne et al., 1989).[1] More recently, (Cordoba, 2008) has introduced a generalization of Gibrat's law that allows size to affect the variance of the growth process, but not necessarily its mean. This property is relevant to both models of economic growth featuring balanced-growth conditions, as well as short-run frameworks attempting to explain business cycles as phenomena emerging from idiosyncratic shocks to different production units (see, e.g., Carvalho and Grassi, 2015). This note shows how a consensus among these views may be reached, once firm heterogeneity is properly accounted for and firm growth is tracked over a long time span. It does so by re-examining the size-growth conundrum through the lens of conditional quantile regressions.

Empirical tests of Gibrat's law have typically relied on cross-section regressions or short-panel econometric techniques that impose homogeneity in the parameters across units and over time (Urga et al., 2003). On one hand, the first approach ignores the information contained in firm-specific time variation of growth rates. On the other hand, while considering information available for different periods of time, a major drawback of the second approach is to pool potentially heterogeneous firms as if their data were generated according to the same process. To overcome these problems, we examine firm growth by means of conditional quantile regressions (see Koenker and Gilbert Bassett Jr., 1978; Koenker, 2005), so as to allow factors such as size and age to exert different effects depending on the speed at which firms expand/contract. In fact, there is no reason to anticipate that the marginal effects of the covariates on the shape of the density are invariant over the domain of firm growth.

Quantile regressions have been implemented in the analysis of the determinants of firm size (e.g., Machado and Mata, 2000; Cabral and Mata, 2003) and growth (e.g., Coad, 2007; Coad and

☆ We would like to thank – without implicating – one anonymous referee, Andrea Caggese, Alex Coad, Sean Holly and Ron Smith for their insightful and constructive comments. All remaining errors are our own.

* Corresponding author.
E-mail addresses: roberta.distante@maersk.com (R. Distante), Ivan.Petrella@wbs.ac.uk (I. Petrella), emiliano.santoro@econ.ku.dk (E. Santoro).

[1] Consistent with the assumption of decreasing returns to scale, these works show that small firms tend to grow faster than large ones. This implies a mean reversion effect on firm size, which introduces an overall limit to the variance of the size distribution, as firm size converges in the long run towards an optimal level.

https://doi.org/10.1016/j.econlet.2017.12.028
0165-1765/

Rao-Nicholson, 2008; Reichstein et al., 2010). However, these studies have typically focused on short time windows, while a vast body of empirical evidence has shown how the density of firm growth displays marked variation over time, both at business cycle and lower frequencies (Higson et al., 2002, 2004; Holly et al., 2013). In light of this, we employ a long panel of COMPUSTAT data on manufacturing firms, accounting for the presence of time effects in the set of determinants of firm growth. In addition, we condition the quantiles of firm growth to firms' age. In this respect, Haltiwanger et al. (2013) have recently stressed the importance of controlling for age when examining the relationship between growth and size.[2]

We detect marked heterogeneity in the impact of firm characteristics on the growth process. Age is never advantageous to firm growth, and more so as firms are relatively small and grow faster. By contrast, size exerts a negative (positive) effect on firms that grow above (below) the median rate, with the marginal impact increasing in absolute terms as firms grow/decline faster. This implies a tendency to mean reversion, so that size differences between firms are transitory. This is an important finding, as it lends support to the generalization of Gibrat's law that allows size to affect the variance of the growth process, but not its mean (see Cordoba, 2008). Our results are robust to controlling for firm turnover – as implied by the analysis of different balanced panels – as well as to splitting the sample over a number of dimensions (e.g., distinguishing between durable and non-durable producing firms). Notably, we report a marked tendency for relatively large firms to display a faster pace of convergence to the mean size in the overall sample. This signals a weak degree of adaptation to an evolving competitive environment as firms grow larger, but not necessarily older.

The remainder of the paper is laid out as follows: Section 2 introduces the quantile regression framework; Section 3 presents the data and reports quantile-based evidence on the relationship between firm growth, size and age; Section 4 concludes.

2. Quantile regression analysis

Quantile regressions are especially useful when dealing with non-identically distributed data. In these situations one should expect to observe significant discrepancies in the estimated 'slopes' at different quantiles with respect a given set of covariates (Machado and Mata, 2000). Such discrepancies may reflect not just into location shifts, but also into scale shifts (i.e., changes in the degree of dispersion) and/or asymmetry reversals (i.e., changes in the sign of the skewness). Define the τth quantile of the distribution of a generic variable y, given a vector of covariates $\mathbf{x}$, as:

$$Q_\tau \left(y|\mathbf{x}\right) = \inf\left\{y|\, F\left(y|\,\mathbf{x}\right) \geq \tau\right\}, \quad \tau \in (0, 1)\,, \tag{1}$$

where $F\left(y|\,\mathbf{x}\right)$ denotes the conditional distribution function. A least squares estimator of the mean regression model would be concerned with the dependence of the conditional mean of y on the covariates. The quantile regression estimator tackles this issue at each quantile of the conditional distribution. In other words, instead of assuming that covariates shift only the location of the conditional distribution, quantile regression looks at the potential effects on the whole shape of the distribution. The statistical model we opt for specifies the τth conditional quantile of firm-level growth, g_{it}, as a linear function of the vector of covariates, $\mathbf{x}_{it}$:[3]

$$Q_\tau \left(g_{it}|\mathbf{x}_{it}\right) = \mathbf{x}_{it}'\boldsymbol{\beta}_\tau, \ \tau \in (0, 1)\,. \tag{2}$$

3. Data and model specification

As it has been noted by Urga et al. (2003), short panels are much more informative on high-frequency variations in corporate growth rates than they are on low-frequency variations. As a result, growth rates can appear more random than they really are and important long-run or secular variations in growth rates may be overlooked. Short panels may also erroneously lead one to reject the view that firms have natural life cycles or systematically evolve through number of stages (see Binder et al., 2005). Finally, it is important to recall that reasonably long panels ($T > 30$) may alleviate problems of autocorrelated residuals. To account for these issues, we employ annual accounting COMPUSTAT annual data on manufacturing firms over six decades (1950–2010).[4]

Real sales are taken as a proxy for firm size, which is denoted by s_{it}.[5] We then compute firm i's growth rate as $g_{it} \equiv \left(s_{it} - s_{it-1}\right)/[(s_{it} + s_{it-1})/2]$.[6] This definition is widely employed in the literature on industrial dynamics, as it shares some useful properties of log-differences and has the advantage of accommodating entry and exit (see Haltiwanger et al., 2013).[7] In line with the industrial dynamics tradition, the econometric framework includes firm-level ($t-1$) size and age in the vector of covariates. In addition, we include industry dummies at the 3-digit SIC code level – which account for the fact that firm growth, size and age distributions vary by industry – as well time dummies, which aim at controlling for the behavior of the distribution over time. The resulting framework generalizes the first order Galton–Markov model $g_{it} = \beta s_{it-1} + u_{it}$, where u_{it} is an error term, assumed to be *i.i.d.* across firms and over time. Note that $\beta < 0$ implies that small firms grow faster than bigger ones, while for $\beta > 0$ the opposite holds true. Gibrat's Law holds instead if the estimated parameter $\hat{\beta}$ is not significantly different from zero, so that growth turns out to be stochastic and independent of size.

Prior to looking the effects of firm size and age over the spectrum of firm growth, it is important to examine the behavior of the density over the time span we consider. Fig. 1 graphs the quantiles of firm growth. A first observation to be made is that different parts of the distribution do not follow the same time path, neither at relatively high nor lower frequencies. As documented by Comin and Philippon (2006) and Comin and Mulani (2006), the density has slowly become more sparse over time. Our evidence points to increasing dispersion as a phenomenon that primarily hinges on the evolution of the tails of the distribution, while the

[2] In fact, their analysis shows that the systematic inverse relationship between firm size and net growth rates highlighted in prior analyses is entirely attributable to most new firms being classified in small size classes. By contrast, once firm age is controlled for, they report no systematic relationship between firm size and growth.

[3] Ideally, one would prefer to implement quantile panel regressions, allowing for both firm-specific and time effects (see, e.g., Powell, 2010). However, this is computationally demanding, even in the presence of a limited number of covariates. In Distante et al. (2015) we show that quantile estimates are robust to the exclusion of firm-specific effects.

[4] Our data selection has privileged the time-dimension of the COMPUSTAT panel, along with its availability. On the downside, it might be argued that, in light of including only quoted companies, these data are biased towards relatively large firms. However, on a priori grounds there is no reason to believe that this property should be crucial in explaining our facts.

[5] Various measures – including the value of assets of a firm, employment and sales – have been traditionally used to proxy firm size. Where data have been available for the various measures the results have generally been invariant to the measure of size (see Evans, 1987a; Hall, 1987).

[6] We remove firms growing (declining) beyond a 100% rate. Replicating the analysis with growth rates defined as log-differences or under alternative cut-off intervals does not qualitatively affect the analysis.

[7] Along with being symmetric around zero and bounded between −2 (exit) and 2 (entry), this growth rate represents a second order approximation of the log-difference for growth rates.