Firms' Size Distributions and Industrial Dynamics: An EVT Approach to the Italian Case
(preliminary version)

Pasquale Cirillo*
DEP/IMQ - Università Bocconi Milano
4th August 2006

Abstract

This paper focuses on the analysis of firms' growth rates, dynamics and size distributions in Italy, which are known to be highly skewed with fat tails. At a first glance, such an analysis could be considered a rehash, but it is not. In fact, even if one can find a lot of papers on these topics, this work is characterized by two interesting novelties: the first one concerns the EVT (Extreme Value Theory) approach we use to study our distributions, their tails and evolution over time; while the second one regards the data we analyse, which come from the very complete, reliable and comprehensive Centrale dei Bilanci database, made up of a twenty-year panel of more than forty thousand manufacturing firms.

In particular, as far as the approach is concerned, we use a set of graphical, parametric, nonparametric and simulation methods, which allow us to exploit our data in order to get all the possible information.

The results we achieve are quite interesting: even if they generally confirm and strengthen well-known empirical evidences, they also propose some new developments.
JEL classification: C16, L11, L16

1. Introduction

Starting from the pioneering works of Gibrat (1931), many researchers have studied industrial dynamics, focusing their attention on firms' size distribution (FSD) and growth rates. The importance of these topics is due to their relevance for public policies and economic growth. The interested reader can find interesting and stimulating results in Simon(1955), Quandt (1966a-b), Ijiri et al. (1977), Stanley et al. (1996).

In particular, recent empirical evidences underline a FSD (in terms of capital) highly skewed to the right and characterized by fat tails, such as a Pareto, while growth rates show a laplacian behavior (Bottazzi et al., 2004). All this seem to deny the validity of Gibrat's law of proportionate effects (GB), according to which firms' growth is independent from their size. Statistically speaking, GB can be translated into a lognormal FSD and a Gaussian growth rates' distribution.

In this paper we try to undersand if these empirical evidences hold in Italy, using the comprehensive CEBI industrial database and an EVT approach (Gumbel, 1958). Our findings say that power law behavior in FSD and laplacian growth rates seem to hold in we consider the whole economy, while they may change if we enter into more specific levels.

*mail to pasquale.cirillo@phd.unibocconi.it.

The papers is organized as follows: Section 2 shortly describes the database we used for our studies; Section 3 is devoted to the analysis of the whole firms' size distribution; Section 4 focuses its attention on FSD tails; Section 5 shows our results as far as public and joint-stock companies (PJS) are concerned; Section 6 deals with growth rates; while Section 7 concludes.

2. The Database

All our empirical analysis are based on firm-level observations from the CEBI database[1], for the period 1983-2002. CEBI, formerly developed by the Bank of Italy, is now mantained by Centrale dei Bilanci Srl. It represents, at the moment, one of the most complete, reliable and comprehensive industrial datasets in Italy.

Thanks to several queries on the original database, we have collected a panel of more than forty thousands Italian manufacturing firms[2], all satisfying the following: (i) no missing capital data in each year; (ii) reliable data for capital, employees and costs. For each firm (and year), we have more than 20 variables, from age to financial ratios. Many of them are not used in this work, but they will be essential for future developments.

From our first panel we have also extracted a second sample made up of public and joint stock companies, that's to say those big companies that are legally required to publicly report their accounts. In this second panel, we have about two thousand firms, i.e. four times the firms Cabral et al. (2003), one of our benchmarks for this article, used for their analysis on Portuguese companies.

3. Firms' Size Distribution

In this section we analyse the whole size distribution of Italian manufacturing firm and its evolution over time. In particular, following Cabral et al. (2003), we will focus our attention on how FSD develops and behaves, finding out that, apart from an average growth of the economy that makes the mean value of the distribution increase, its shape is quite regular, suggesting that FSD is quite stable over time at economywide level. On the contrary, we will show that stability does not hold if we condider FSD by age group: indeed the distribution of firms changes over time as they get older. In particular the lognormal distribution seems to be a good proxy for old firms. All this is confirmed in Section 5, where we analyse PJS companies that, in general, represent the oldest firms on the market.

Such a duality is very interesting since it suggests that macro regularities over time (i.e. FSD stability) come from changes and evolutions at the micro level, as stated by Evolutionary and Agent-based Economics (for a reference, see Tesfatsion, 2002). Cabral et al. (2003), but also Gallegati et al. (2005), believe that the stability of the whole FSD is due to the firms' entry-exit process in the economy, that makes the market composition almost constant. As far as the FSD by age group, its evolution is rather due to the stronger financial constraints new firms have to face, whereas mature surviving firms are not (or less) financially constrained. Cirillo (2006) reproduces the distributional duality proposing an agent-based model in which financial constraints are not the only obstacles for firms, adding institutional and geographical limitations too.

Figure 1 shows the evolution of FSD over time considering all the firms in our panel. The figure is based on the so called Zipf's plot, that is a plot of the log of the ranks versus the log of the variable being analyzed. Such a tool is very useful in detecting if a distribution shows fat tails and, in particular, if the well-known Zipf's law(Zipf, 1932; Okuyama et al., 1999) on the power law behaviour of tails holds. As in figure 1, a clear linear trend in the plot rejects the hypothesis of lognormal behavior and points out the presence of regularly varying tails. A distribution function F is said to be regular varying with index α (Zipf's index) when

$$\lim_{u\to\infty} \frac{\overline{F}(ux)}{\overline{F}(u)} = x^{-\alpha}, \tag{1}$$

where u represents a proper threshold. Figure 1 also shows that FSD is highly skewed to the right, confirming the usual empirical evidences. To complete the analysis, in figure 2 we show the nonparametric kernel estimates of FSD.

The evolution of FSD over time shows that, apart from an obvious scale difference (economy grows indeed), the distribution shape is always the same, suggesting that all the different FSDs in the plot belong to the same family. All this states in the Italian case what Cabral et al. (2003) have discovered for the Portuguese firms.

In order to investigate which distribution can be considered a good proxy for our data, we have performed several graphical and analytical tests. First of all, as in Embrechts et al. (1997), we have studied the behaviour of the mean excess function (MEF). In fact, MEF is surely a good tool to explore data behaviour since it graphically characterises different distribution families. Figure 3 shows the evolution of the mean excess function versus threshold over time. Even in this case, all the FSD show a common behaviour: a clearly upwarding trend, that represents a strong signal for the presence of a Pareto-like distribution.

Figure 4 presents a quantile-quantile plot of FSD. We have chosen year 1988 in order to have a clearer plot, but all the years show a similar behaviour. For a first graphical result, we have compared our data with a theoretical Lomax distribution with shape parameter α equal to one. As one can see in the figure, Lomax, also known as Pareto-II, seems to be a good proxy, since most of the points in the plot lie on the diagonal. The choice of a Lomax distribution is mainly due to the high scale of our data and the hints given by figures 1 and 3 (see Kleiber et al., 2003).

Analytical tests confirm our graphical intuition. Kruskall-Wallis' statistics on

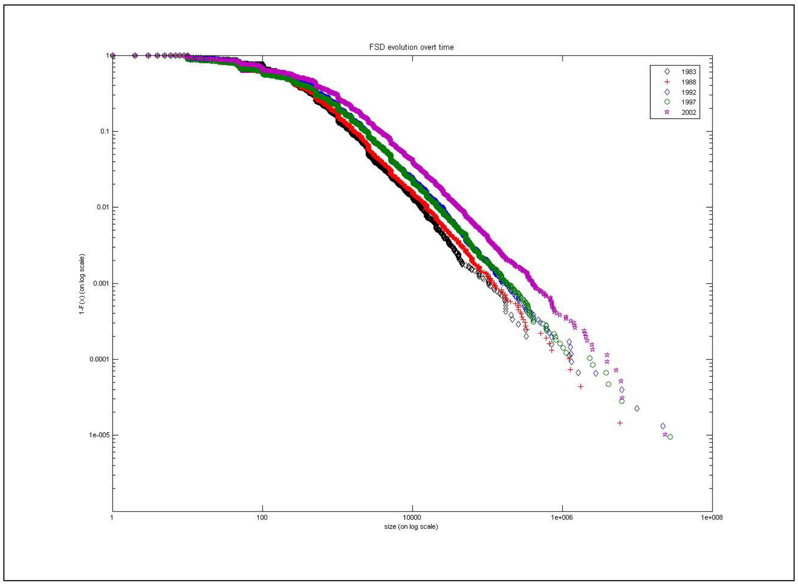

Figure 1: Evolution of FSD over time (Zipf's plot)

ranks does not reject the null of a Lomax distribution with a confidence interval of 95%. Kolmogorov-Smirnov test does the same, analysing the similarity between the theoretical and the empirical distributions thanks to the distance induced by the supremum metrics. For all these reasons Gibrat's idea of a lognormal FSD is not acceptable if we consider all the firms on the market, without any distinction.

As already set, our panel from CEBI database is quite complete and offers several useful variables for the analysis. One of these variables is firms' age. For every firm, indeed, we know the date/year of birth and its years of activity (an inactive firm can be considered as not aging). Following Cabral et al. (2003) our aim is to investigate if FSD differs among firms when we divide them into age groups. In particular, creating three different "cohorts" (1-6 years old, 7-12 years old and 12+ years old), we want to compare FSDs.

Figure 5 shows how FSDs by age group differ in 2002. Obviously, similar results hold for all the other years and even the pooled case, as for all the evidences we present in this paper. As one can see in the figure, the youngest firms (1-6 y.o.) are characterized by an highly skewed distribution with heavy tails, while the oldest ones present a simmetrical distribution, very similar to a Gaussian. Since size is expressed in logs, all this can be interpreted as follows: young and mature firms follow a Pareto-like distribution, old companies tend to a lognormal. Analytical

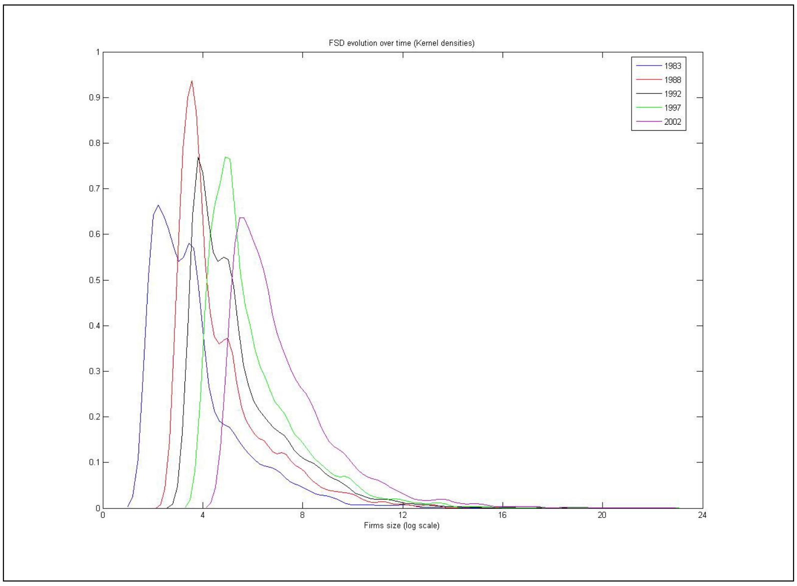

Figure 2: Evolution of FSD over time (kernel densities)

test such as Kolmogorov-Smirnov and Jarque-Bera confirm our graphical findings: the first one with a 90% c.i, the second with 95%. This behaviour can be seen as a partial validation of Gibrat's first results. Probably the lognormal distribution is not a good proxy for FSD if we consider all firms together, but it surely performs better if we only consider mature surviving firms, as if lognormal distribution were a limit distribution in firms' evolution. Maybe Gibrat's law of proportionate effects could be considered as a non-universal law, but simply a specific evidence. We will try to clarify this statement in section 6, analysing growth rates, the other fundamental component in Gibrat's law.

4. Tails

Lomax can be seen as a distribution coming from mixing two Pareto-I distributions, with similar scale but different shapes. In this section we focus our attention on tails, in order to get good estimates of the shape parameter α over time. In particular, we aim to show that, even if the distributional shape of FSD remains alost constant over time, there are some changes in the shape parameters α's, indicating that the number of big firms slightly increases over time. Anyway, these changes are not sufficient for the FSD to makeover.

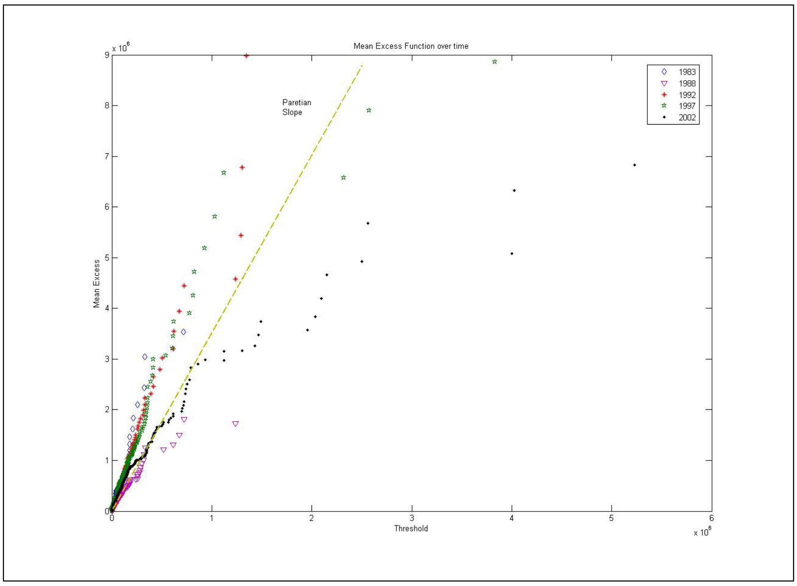

Figure 3: Mean excess function versus threshold over time.

Following the usual procedures in EVT, we have tested if the tails (top 15% firms) of FSD follow a GPD behaviour, where GPD stands for Generalised Pareto Distribution.

Starting from the well-known Fisher-Tippett Theorem, which deals with the convergence of maxima, the GPD distribution represents one of the most important limiting cases for observations over an high threshold.

Its functional form is:

$$H(x) = \begin{cases} 1 - (1 + \xi \frac{x}{\beta})^{-\frac{1}{\xi}} & if\ \xi \neq 0 \\ 1 - e^{-\frac{x}{\beta}} & if\ \xi = 0 \end{cases}, \tag{2}$$

where $\beta > 0$ and x is such that $1 + \xi x > 0$ and ξ is the shape parameter (tail index $\alpha = \frac{1}{\xi}$).

There are three different situations:

1. $\xi > 0 \rightarrow$ GPD distribution becomes the classical Pareto distribution;
2. $\xi = 0 \rightarrow$ GPD distribution converges to the exponential distribution;
3. $\xi < 0 \rightarrow$ GPD distribution is then known as Pareto II.

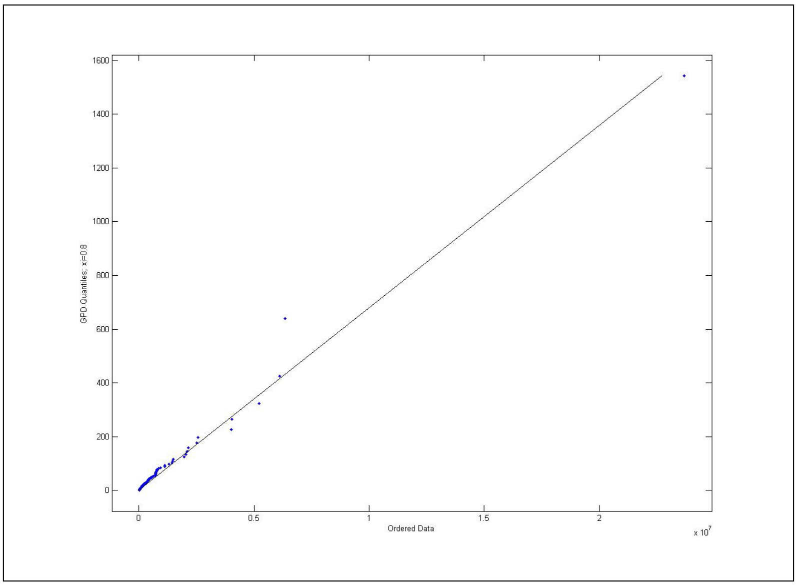

Figure 4: Quantile-quantile plot for FSD. Sample year: 1988.

Even in this case, we have performed both analytical and graphical studies to test our hypothesis.

Belonging to a Lomax is surely a strong signal for data to have GPD-like tails. Anyway the most rigorous way to assert this is to check for Von Mises' and extremal conditions (VME).

VME state that a distribution function F (with density f)belongs to the domain of attraction of an extreme type distribution (among which GPD) if

$$\lim_{x \uparrow x_T} \frac{(1+\gamma x) f(x)}{\overline{F}(x)} = c > 0, \tag{3}$$

where $\gamma \in \mathbb{R}$ and x_T represents the right endpoint of the distribution F. GPD is the only distribution for which VME hold with equality (Falk et al, 2003).

In order to have an indirect validation of VME, we have used the tests proposed by Dietrich et al. (2002) and Hüsler et al. (2006), concerning the convergence of quantiles. A first proof of GPD-like behaviour of tails has so been collected. In particular, after estimating $\widehat{\beta}$ and $\widehat{\xi}$ with the method of moments, we have compared the quantiles of distribution $H(x)$ with those presented in the tables by Drees et al. (2006), non rejecting the null.

As far as the graphical tools, we have compared tails, as in figure 6, and ecdf's,

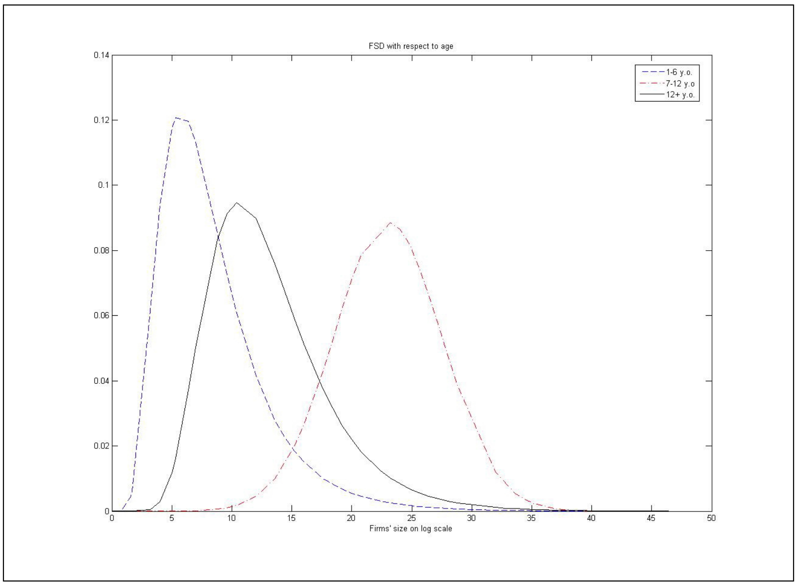

Figure 5: FSD by group age in 2002.

as in 7. Once again, the choice of a single year (1992) has been simply made to guarantee a major readability of the results. Anyway, in both cases it is possible to see a very good fitting of data. We can so assume they follow a GPD distribution.

Such an assumption is surely strengthened by the fact that the whole FSD seems to be a Lomax. Lomax, in fact, is one of the three specific distributions one can obtain from GPD.

Being in the field of GPD makes our data suitable to a well-known series of statistical tests to estimate the shape parameter α. In particular, instead of the classical regression on ranks, typically used in industrial studies, we can have more accurate estimates thanks to GPD simulation, Hill's semiparametric index and MLE.

GPD simulation is a very useful method for the analysis of tails. In fact, it's quite simple and gives very good estimates. Starting from the method of moments, one calculates $\hat{\beta}$ and $\hat{\xi}$ and then simulates a GPD$(\hat{\beta}, \hat{\xi})$. At this point the empirical quantiles on data are compared with those of the just simulated GPD. If the fitting is good, everything stops, otherwise, following a grid method, $\hat{\beta}$ and $\hat{\xi}$ change and the procedure is repeated. Once the point estimates are sufficiently good, one starts playing with data, excluding one-by-one all the biggest observations, that could be possible outliers. The result is a plot as in figure 8. The procedure stops

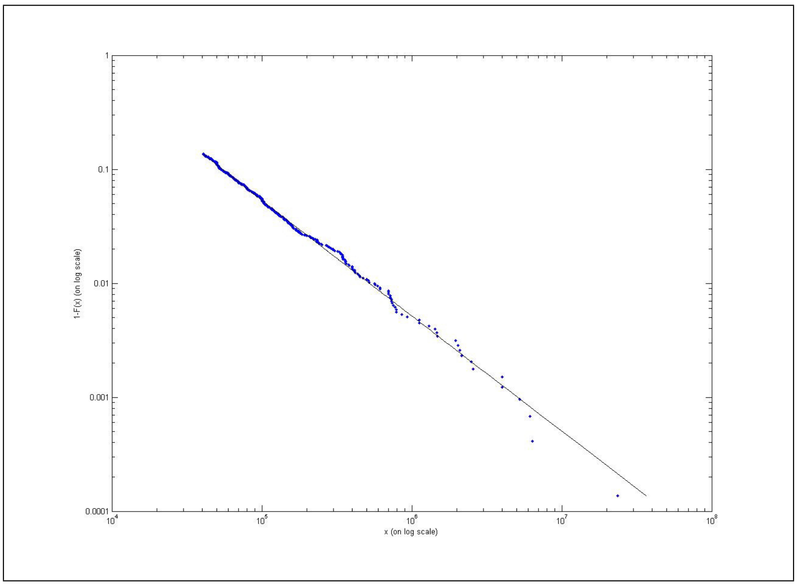

Figure 6: Comparison between the theoretical GPD tail and empirical data. Year 1992.

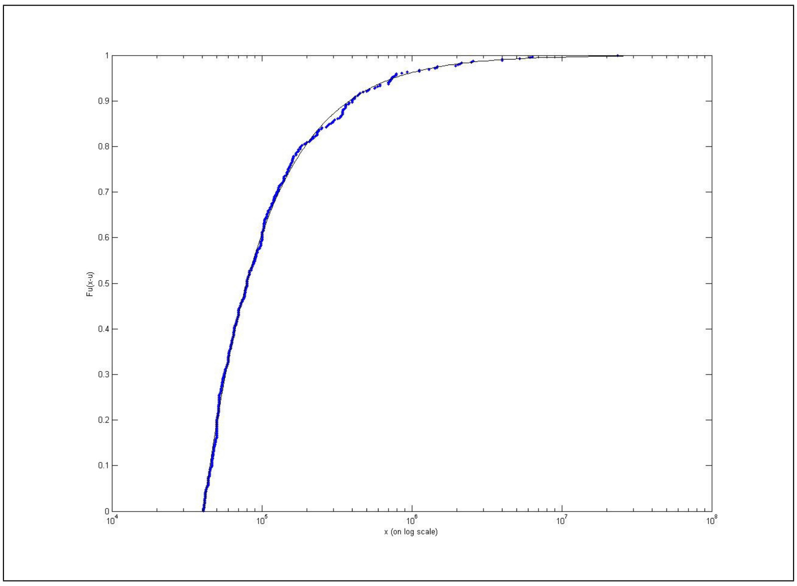

Figure 7: Comparison between the theoretical GPD and empirical data. Year 1992.

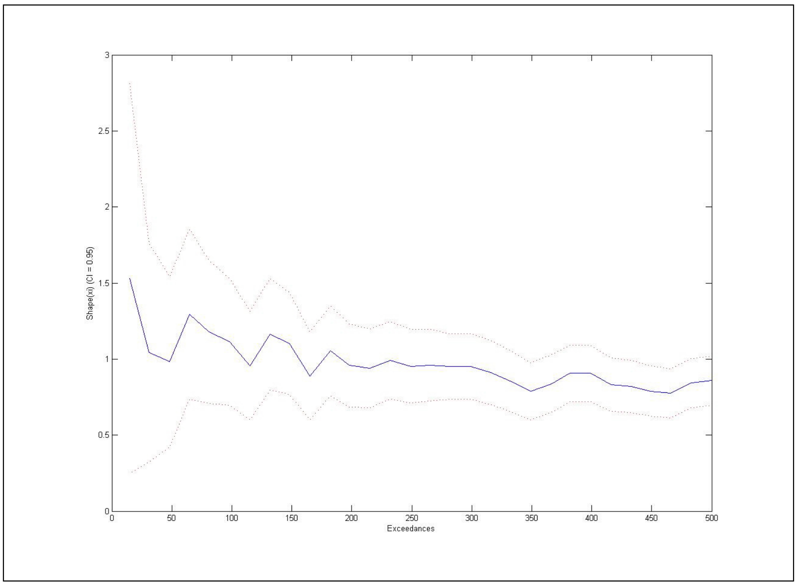

Figure 8: GPD shape parameter via simulation of 30 different models. Year 2002.

when the shape parameter $\hat{\alpha} = \frac{1}{\hat{\xi}}$ assumes a quite stable behaviour.

On the contrary, the well-known Hill's estimator $\bar{\xi}$, together with the Pickard's one, is the most used way to determine the shape parameter $\bar{\alpha} = \frac{1}{\bar{\xi}}$ of a distribution belonging to a Paretian family.

In particular

$$\bar{\xi} = \frac{1}{k-1} \sum_{i=1}^{k-1} \ln x_{i,N} - \ln x_{k,N} \quad for\ k \geq 2, \tag{4}$$

where k is the upper order statistics and N the sample size. Such an index is proved to have very good asymptotic properties. Even in this case, letting k change, it is possible to have a plot as in figure 9.

The interested reader can find a good introduction to both methods in Embrechts et al. (1997). No need to discuss MLE.

In table 1 we have collected the α estimates for the three methods we have used in different years. These values allow us to make some considerations about the evolution of FSD tails over time. First, one can underline that, considering the different procedures, α_{1983} is not very far from α_{2002} indicating that there has been only a little change in the tails' fatness. Moreover, apart from 1997 (the same

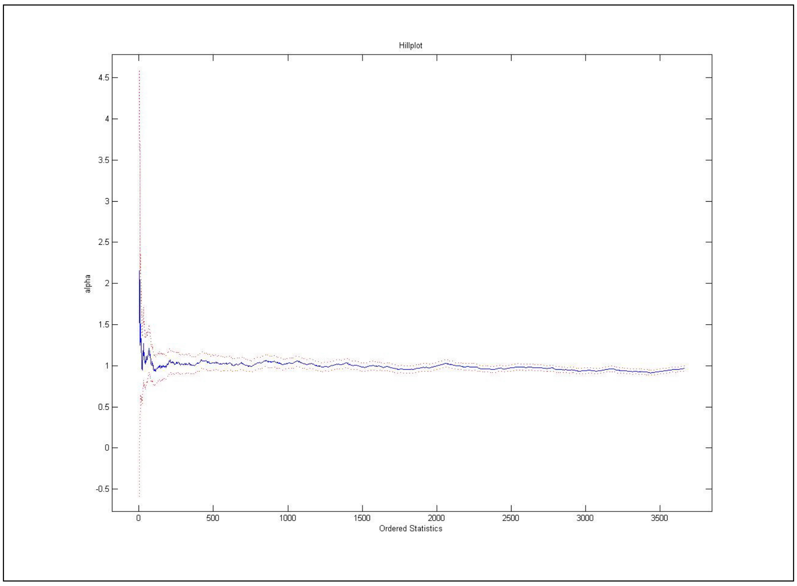

Figure 9: Hill's index plot. Year 2002.

happens in 1985 and 1991), the values of α monotonically decrease, suggesting that the tails of FSD is slightly increasing, probably because the quote of big firms in the economy has grown over time. In add, all the values are around the unity, perfectly fitting (i) regular variation, (ii) Zipf's law and (iii) the Paretian field (0,2], empirically determined in the main industrial literature (Quandt, 1966a).

Finally, a deeper analysis of data suggests that FSD is quite stable over time, apart from some obvious scale shifts. All this is consistent with the findings of Stanley et al. (1996), Amaral et al. (1996) and Cabral et al. (2003).

$\alpha(se)$	**Sims**	**Hill**	**MLE**
1983	1.0803 (0.5121)	1.0815 (0.4721)	1.0811 (0.4212)
1988	1.0632 (0.3823)	1.0657 (0.3933)	1.0671 (0.3937)
1992	1.0410 (0.4672)	1.0385 (0.4175)	1.0388 (0.4248)
1997	1.0412 (0.4922)	1.0392 (0.4296)	1.0395 (0.4301)
2002	0.9852 (0.3875)	0.9859 (0.3911)	0.9849 (0.3900)

Table 1: Shape parameter's estimates with different methods

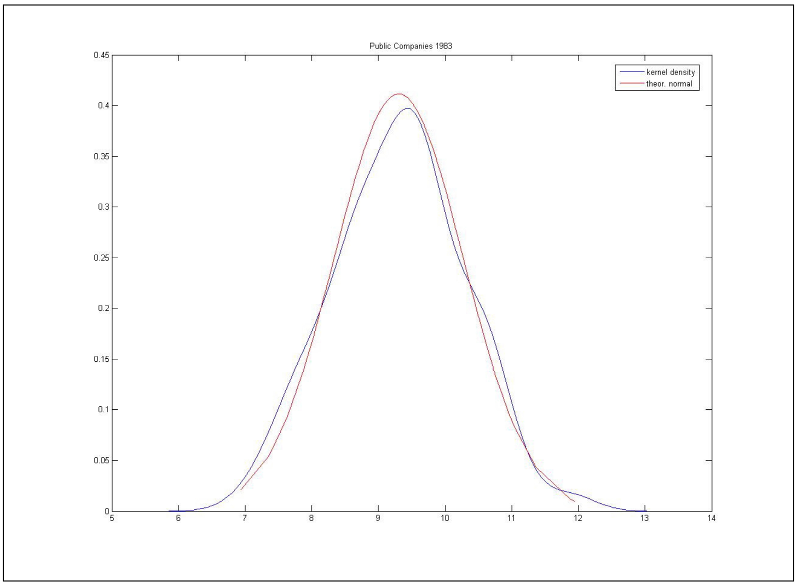

Figure 10: Comparison between the kernel estimate of PJS companies' distribution and a lognormal distribution. Year 1983.

5. What about PJS Companies?

The aim of this section is to investigate if the results of Cabral et al. (2003) for public and joint-stock companies hold for the Italian market. In their interesting work, the authors find out that PJS firms show particular features as far as their size distribution, that seems to be almost lognormal.

In fact, if we estimate their df's on a log scale we note that they are quite simmetric and similar to a Gaussian. Figures 10 and 11 indicate that such a pattern is present in different years (1983 and 2002 for exaple), i.e. in all years apart from 1991 that behaves negatively in many of our analysis. It should be useful to understand if this is due to a problem of data or it reproduces the Italian crisis of early 90's. Anyway, at least graphically, a lognormal distribution (a normal on logs) is probably a good proxy for data.

In particular, in the figures, one can observe the comparison between kernel densities estimates and theoretical Gaussians. The similarity is evident.

Obviously, we have analytically searched for normality on all log observations, using both Kolmogorov-Smirnov and Jarque-Bera tests. In all years, apart 1991 as already said, we cannot reject the null of gaussianity with a confidence interval

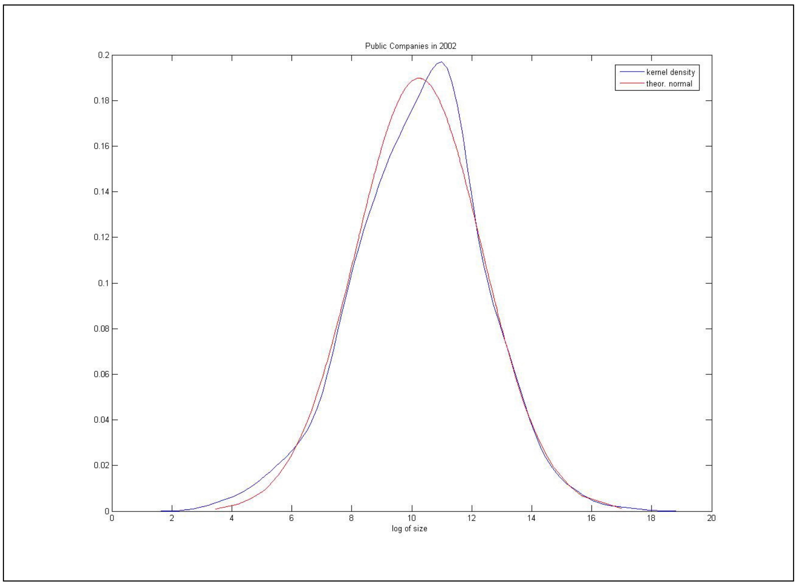

Figure 11: Comparison between the kernel estimate of PJS companies' distribution and a lognormal distribution. Year 2002.

of 90% (1989, 1992) or 95% (all the rest). Such a result is quite strong.

Cabral et al. (2003) think such a behaviour is due to the particular conditions of PJS. Public and joint-stock companies, in fact, are in general big old firms (not necessarily the largest ones) that have survived in the economy and, being mature, are not financially constrained. This allow them to differently develop. According to us, this explanation is surely true, even if we believe that financial constrains are not the only cause. Indeed we consider that geographical and institutional limitations play a strong role too, as stated in the main agent-based literature.

Finally, we cannot forget PJS are particular companies, subject to specific law requirements. We then have to control for sampling bias in order to validate our results. The lognormal distribution could arise from a selection process in which larger firms are picked up with an increasing probability, so that their distribution is close to a lognormal. Further analysis is so needed.

6. Growth Rates

As far as firms' growth rates are concerned, several studies (Axtell, 2001; Bottazzi et al., 2004; Hall, 1987) find a tent-shape behaviour. In particular, the

Pars	**1994 data**
μ	-0.0029 (0.0011)
a	0.0419 (0.0213)
b	1.0032 (0.4688)
$-loglik$	1.1921

Table 2: Estimated Subbotin's Parameters for growth rates in 1994

Laplace and Lévy distributions seem to provide the best fitting (Mandelbrot, 1960; Bottazzi and Secchi, 2003; Gabaix, 2004).

We have investigated if the empirical distributions of growth rates (in terms of capital) belong to the well-known Subbotin's family (Subbotin, 1923; Bottazzi et al., 2004), which represents a generalization of several particular cases, such as Laplace (see the red line in figure 12 for an example) and Gaussian distributions.

The functional form of Subbotin's family is:

$$f(x,a,b) = \frac{1}{2ab^{\frac{1}{b}}\Gamma\left(1+\frac{1}{b}\right)} e^{-\frac{1}{b}\left|\frac{x-\mu}{a}\right|^{b}}, \tag{5}$$

where μ is the mean, b and a two different shape parameters and Γ is the standard Gamma. If $b \to 1$ the Subbotin distribution becomes a Laplace, a Gaussian for $b \to 2$.

Using the maximum likelihood method[3], we have estimated the three Subbotin's parameters on our data for all years. Table 2 contains the results for year 1994.

At a first glace, one can observe that:

1. since b is very near to 1, growth rates' distribution (GRD) is in the field of attraction of the Laplacian case[4];

2. growth rates are centered on 0;

3. the value of a, the Laplacian shape parameter, indicates the fatness of tails.

As far as the evolution of GRD, in this case it is not possible to find a clear pattern. In fact, apart from a certain stability in Laplacianity, the tails of the distributions change over time, both increasing and decreasing. Probably one could better understand such a behaviour looking at the business cycles in the economy. We will work on it in future works.

However, the results we have obtained considering data by age group are very interesting. Figure 5 shows GRD for the three "cohorts" in 1997 and, apart from 1991-92, the same pattern is present in all years. We can note that right tails are greater than left ones in all the cases.

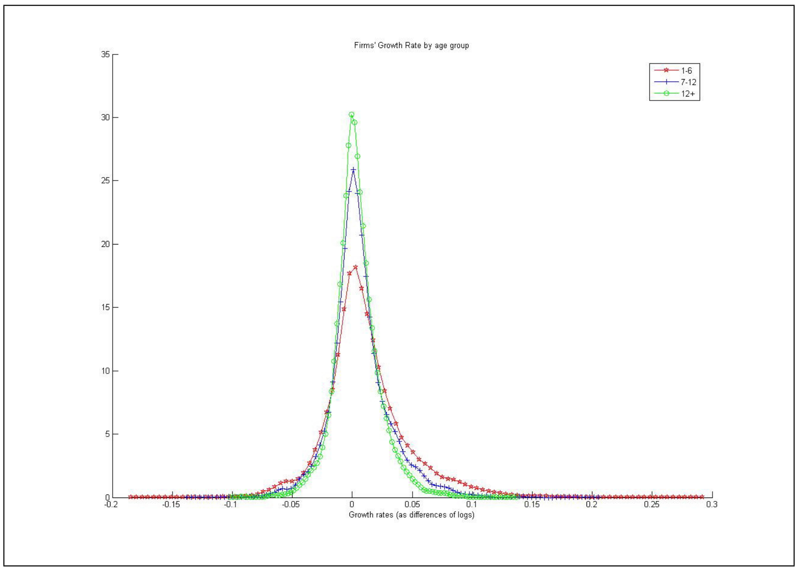

Figure 12: Growth rates' distribution by age group. Year 1997.

Pars (1997)	**1-6 y.o.**	**7-12 y.o.**	**12+ y.o.**
μ	0.0182 (0.0042)	0.0091 (0.0035)	-0.0213 (0.0097)
a	0.1279 (0.0213)	0.0916 (0.0213)	0.0623 (0.0213)
b	1.2152 (0.4688)	0.9772 (0.4688)	1.7056 (0.4688)
$-loglik$	2.0033	1.7315	2.1042

Table 3: Estimated Subbotin's Parameters for growth rates by age group in 1997

In table 3 we present the estimates for the three GRD by age group in 1997. The main features we can notice are:

1. the three distributions are almost centered on zero;

2. since 0.13>0.09>0.06, the distribution of youngest firms show fatter tails and this is evident also in figure 12;

3. the distributions of 1-6 and 7-12 years old firms are very close to a Laplace, being their β's near the unity;

4. the distribution of big firms, on the contrary, shows a β nearer to 2 than to 1, as if the distribution were tending to a Gaussian.

This last characteristic is quite interesting. Even if we cannot state that the GRD of big firms is normal (Jarque-Bera and Kolmogorov-Smirnov reject the null), we cannot deny it is quite different from the other cases ($\beta \to 2$): growth rates are more concentrated and extreme-type values are not present.

The most important aspect of such a behaviour is that it could support Gibrat's law, at least in a weaker sense, if we combine it with a lognormal FSD, as we have found in previous analysis. At the moment, this can only be a supposition, since further studies are required. We will work on it in the future.

Very similar results are available for PJS companies. We avoid to show results here in order not to be prolix. Moreover, an extended analysis of PJS firm is object of a forthcoming paper we are preparing.

7. Conclusions and future research

In this paper we have analysed several empirical evidences, finding out that they hold even in the Italian case. CEBI database, that we have used for our studies, shows to be a very good source of information and we will exploit it more in the future.

The main novelties of this paper are related to the behaviour of FSD and GRD if we split firms into age groups. In certain cases, in fact, Gibrat's law seems to hold, but more knowledge is needed. For example, we will follow the evolution of all the firms born in a certain year and we will make this for all the available years.

As far as the evolution of the whole FSD, we have found that there is a good stability of Pareto-like distributions over time and the same happens if we consider GRD and Laplace.

Very interesting findings concern PJS companies, that seems to behave in a different way with respect to other firms.

The approach we have used is quite new for industrial studies, since it focuses on EVT graphical and analytical methods. EVT gives the researchers very powerful tools to explore data, mainly if one concentrate on high (or low) magnitude observations, as those in distributions tails. In forthcoming papers, we will focus our attention on largest companies, trying to understand in what they differ from the others. Using particular tools such as the peaks over threshold method, even in the bayesian approach (Cirillo et al., 2006), we will try to see if the probability of becoming a big firm and survive over time has particular characteristics. Intuition says yes, but further analysis are required.

Obviously, beyond empirical studies, a good way to comprehend industrial dynamics is to model it. The aim is to construct models that reproduce reality. We have worked on these topics in past (Bianchi et al., 2005; Cirillo, 2006) and we will continue in the future.

Notes

[1]The author kindly thanks Bank of Italy for providing data to him and Prof. C. Bianchi (University of Pisa).

[2]This database is probably one of the biggest dataset ever used in Italy for Industrial studies. All this surely strengthens our results.

[3]The results are very similar, using the method of moments.

Estimation has been made thanks to a Matlab code we have written using the original C code by Bottazzi et al. (2002) as a constant benchmark.

[4]Some authors prefer a truncated Lévy distribution. The *querelle* is open. See Kleiber and Kotz (2003).

References

Amaral L., Buldyrev S., Havlin S., Leschhorn H., Maas P., Salinger M., Stanley E. and M. Stanley, Scaling Behavior in Economics: I. Empirical Results for Company Growth . *Journal de Physique*,7:621-33, 1997.

Axtell R., Axelrod R., Epstein J.M., Cohen M.D., Aligning Simulation Models: A Case Study and Results. *Computational and Mathematical Organization Theory*, Volume 1, N. 1, 1996.

Bianchi C., Cirillo P., Gallegati M., Vagliasindi P.A. Validating and Calibrating Agent-Based Models: a Case Study, submitted, 2005.

Bottazzi G., Secchi A., Explaining the distributions of firm growth rates. *LEM Rand Journal of Economics, forthcoming*, 2004.

Cabral L.M.B., Mata J., On the evolution of the Firm Size Distribution: Facts and Theory *The American Economic Review*, Volume 93, N. 4, 2003.

Cirillo P., Firms' Evolution and Constraints, submitted, 2006.

Cirillo P., Muliere P., A Bayesian Nonparametric Approach to the Analysis of Rare Events, preliminary working paper, 2006.

Dietrich D., de Haan L., Hüsler J., Testing Extreme Value Conditions. Extremes, N. 5, 71-85, 2002.

Drees H., de Haan L., Li D., Approximations to the tail empirical distribution function with application to testing extreme value conditions. *Journal of Statistical Planning and Inference*, forthcoming, 2006.

Embrechts P., Mikosch T., Kluppelberg C., *Modelling Extremal Events.* Springer Berlin and New York, 1997.

Gabaix X., Power laws and the origins of aggregate fluctuations. *MIT Economic Department: mimeo*, 2004.

Gallegati M., Delli Gatti D., Gaffeo E., Giulioni G., Kirman A., Palestrini A., Russo A., Complex Dynamics and Empirical Evidence forthcoming *Information Sciences*, 2005.

Gumbel, E.J., *Statistics of Extremes.* Columbia University Press, 1958.

Hall B.E., The relationship between firm size and growth. *Journal of Industrial Economics*, 35:583–606, 1987.

Hüsler J., Li D., Weak convergence of the empirical mean excess process with application to estimate the negative tail index. *Working paper IMSV Bern,* 2006.

Ijiri Y., Simon H.A., *Skew Distributions and the Size of Business Firms.* North Holland, Amsterdam, 1977.

Kleiber C., Kotz S., *Statistical Size Distributions in Economics and Actuarial Sciences.* Wiley, 2003.

Mandelbrot B., The pareto-lévy law and the distribution of income. *International Economic Review*, 1:79–106, 1960.

Okuyama K., Takayasu H., Takayasu M., Zipf's law in income distribution of companies. *Physica*, (A269):125–131, 1999.

Pareto V., *Cours d'économie politique.* Ed. Rouge, Lousanne, 1896.

Quandt R.E., On the size distribution of firms. *American Economic Review*, 56:416–432, 1966a.

Quandt R.E., Old and new methods of estimation and the pareto distribution. *Metrika*, 10:55–82, 1966b

Simon H. A., On a class of skew distribution functions. *Biometrika*, 42(3/4):425–440. 1955.

Stanley M., Amaral L., Buldyrev S., Havling S., Leshorn H., Maas P., Salinger M., Stanley E., Scaling behavior in the growth of companies. *Nature*, (379):804–806, 1996.

Subbotin M.T., The Law of Frequency of Error. *Matematicheskii Sbornik*, 31, 296–301, 1923.

Tesfatsion L., Agent-based computational economics. *ISU Economics Working Paper*, (1), 2002.

Zipf G.K., Selective studies and the principle of relative frequency in language. 1932.