Physica A 432 (2015) 102–107

Contents lists available at ScienceDirect

Physica A

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

Global evidence on the distribution of firm growth rates[☆]

Michael A. Williams, Brijesh P. Pinto [*], David Park

Competition Economics LLC, 2000 Powell Street, Suite 510, Emeryville, CA 94608, United States

HIGHLIGHTS

- We use global data to examine empirically the distribution of firm growth rates.
- We use the data to test eight theoretical distributions with EDF statistics.
- The consensus finding in the literature supports the Laplace distribution.
- We find firm growth rates are best fit by the heavier-tailed Cauchy distribution.

ARTICLE INFO

Article history:
Received 1 October 2014
Received in revised form 27 January 2015
Available online 9 March 2015

Keywords:
Firm growth rates
Theoretical distributions
Empirical distribution function tests

ABSTRACT

The consensus finding in the literature is that the distribution of firm growth rates is best approximated by the Laplace distribution, a particular case of the Subbotin, or exponential power, family of probability distributions. Using a richer database than prior studies and testing for more theoretical distributions, we find that the distribution of firm growth rates is best approximated by the heavier-tailed Cauchy distribution.

© 2015 Elsevier B.V. All rights reserved.

1. Introduction

Starting with the path-breaking research of Stanley et al. [1], a consensus finding has emerged that firm growth rates display a "tent-shaped" Laplace distribution.[1] This result has been obtained in several studies—for example, Nunes Amaral et al. [5], Bottazzi et al. [6], Bottazzi and Secchi [7], Buldyrev et al. [8], Alfarano and Milaković [9], Riccaboni et al. [10], Hölzl [11], and Erlingsson et al. [12]. We explore the distribution of firm growth rates using a richer database than used in these prior studies. Our global database has 13,342 firms in 57 industries from 43 countries over the period 1999–2010. These firms had total 2010 revenues of $38.5 trillion or 61% of world GDP.

Building on prior work on the distribution of firm growth rates, we present empirical distribution function, or EDF, tests for the log of the sales growth rate for eight distributions: Cauchy, exponential, gamma, Laplace, logistic, log-normal, normal, and Weibull. We find, in common with the consensus finding, that the Laplace distribution performs reasonably well in approximating the distribution of firm growth rates. However, we show that the distribution of firm growth rates is

☆ We thank Leslie Park for research assistance. We thank Christophe Dutang, Hugo Mialon, Simon Wilkie, and Wei Zhao for helpful comments.

* Corresponding author.
E-mail addresses: mwilliams@competitioneconomics.com (M.A. Williams), bpinto@competitioneconomics.com (B.P. Pinto), dpark@competitioneconomics.com (D. Park).

[1] A related "scaling law" is one that pertains to firm size. For example, Axtell [2] finds that firm size is characterized by the Zipf distribution. Palestrini [3] studies the theoretical relationship between the Laplace firm growth-rates distribution and the Pareto firm size distribution. A related study by Cabral and Mata [4] demonstrates, using data on Portuguese manufacturing firms, that firm size distribution has evolved toward a lognormal distribution.

http://dx.doi.org/10.1016/j.physa.2015.02.103
0378-4371/© 2015 Elsevier B.V. All rights reserved.

substantially better approximated by the heavier-tailed Cauchy distribution, a particular case of the Lévy α-stable family of densities.[2] Thus, our results agree, for example, with the empirical finding of Bottazzi et al. [14] that "[w]hilst the Laplace distribution of growth rates was repeatedly found in previous studies and appeared to be emerging as something of a 'stylized fact', we observe here that the growth rates of French firms are even fatter tailed than expected, a property which holds with disaggregation".

Understanding the empirical distribution of firm growth rates is important because the distribution places constraints on theoretical models of firm growth. In turn, models of firm growth are generally based on dynamic models of innovation and investment (see, for example, Refs. [8,6]).[3] In order to be empirically relevant, such models should generate firm growth rate density functions consistent with the empirical densities. To date, no theoretical models exist in which firm growth rates follow a Cauchy distribution. Thus, our results invite research on dynamic models of competition in which firms' growth rates follow a Cauchy distribution.

2. Data and calculation of firm growth rates

We obtain data on firms' growth rates using a unique, proprietary database that provides financial data for 13,342 firms spanning 57 different industries across 43 countries for the years 1999–2010.[4] To account for inflation, data on firm sales are deflated using country-specific consumer price indices (CPIs).[5] Let S_t^{ijk} denote the annual sales in year t of firm k located in country i and belonging to industry j. The annual sales data are reported to the nearest $1000 and, thus, are not binned. Normalized log sales, s_t^{ijk}, are defined as follows:

$$s_t^{ijk} = \log(S_t^{ijk}) - \frac{1}{|N_t^{ij}|}\sum_{i=1}^{|N_t^{ij}|}\log(S_t^{ijk}) \tag{1}$$

where N_t^{ij} denotes the set of firms in the sample determined by country i, industry j, and year t. The firm growth rate, r_t^{ijk}, is defined as follows:

$$r_t^{ijk} = s_t^{ijk} - s_{t-1}^{ijk}. \tag{2}$$

Defining the firm growth rate in this manner eliminates possible country-specific industrywide trends in the mean of annual sales (see Refs. [17,18]).

For each sample determined by particular values of i, j, and t, we test the hypothesis that the data $\{r_t^{ijk}\}_{k\in N_t^{ij}}$ are distributed according to a particular empirical distribution function F. In performing the EDF tests, we restrict our analysis to samples satisfying $|N_t^{ij}| \geq 10$. Imposing this restriction yields 63,568 observations of firm growth rates across 27 countries for a total of 9267 firms that are grouped into 2004 country/industry/year samples, with an average of approximately 32 firms per sample.

Summary statistics on the financial characteristics of these firms are presented in Table 1. Pooled across countries, industries, and years, the size distributions of firms, as measured by their capital or market value, are positively skewed. The mean firm capital is approximately six times larger than the median firm capital, and the mean firm market value is approximately seven times larger than the median firm market value.

3. EDF tests of theoretical distributions

To find the theoretical distribution that best approximates the empirical distributions of firm growth rates, we perform goodness-of-fit tests using EDF test statistics. Let F denote the empirical CDF and G denote the hypothesized distribution (e.g., the normal distribution). Let (X, δ) be any metric space of CDFs. Then, the distance $\delta(G, F)$ is an EDF test statistic. Commonly used EDF test statistics include the Anderson–Darling and Cramér–von Mises test statistics [19]. The Anderson–Darling test statistic is based on the parametric family $V_n = n\int_{-\infty}^{\infty}|G(x) - F(x)|^2\psi(x)\mathrm{d}G(x)$, where n is the number of observations. Setting $\psi(x) = \{F(x)(1 - F(x))\}^{-1}$ yields the Anderson–Darling test statistic. The Cramér–von Mises test statistic is derived by setting $\psi(x) = 1$ in the family $V_n = n\int_{-\infty}^{\infty}|G(x) - F(x)|^2\psi(x)\mathrm{d}G(x)$.

The parameters of each distribution are estimated using maximum-likelihood estimation [20]. Direct optimization of the log-likelihood is performed with the Nelder–Mead method. We first test the null hypothesis that the firm growth rates are normally distributed. As shown in Table 2, the null hypothesis is not rejected based on the Anderson–Darling

[2] The Lévy α-stable family ($\alpha \in (0, 2]$) includes the Gaussian, or normal, ($\alpha = 2$) and Cauchy ($\alpha = 1$) distributions as special cases. All non-Gaussian Lévy α-stable distributions (i.e., $\alpha < 2$) present heavy tails [13].

[3] To explain the consensus "tent-shape" finding, Bottazzi and Secchi [7] posit a mathematical framework based on the random assignment of business opportunities among firms. Manas [15] derives a mixed normal-Laplace distribution of firm growth rates using an alternative approach based on the direct stochastic modeling of firm sales. Metzig and Gordon [16] develop an agent-based model that yields a tent-shaped aggregate distribution of firm growth rates.

[4] Source: evaDimensions, www.evadimensions.com.

[5] CPIs are obtained from http://research.stlouisfed.org/.

Table 1
Summary statistics of global firm data (1999–2010).

Variable[a]	Median	Mean	Std. dev	10th percentile	90th percentile
Capital[b] ($mil)	344.09	2161.68	8615.62	38.81	3957.00
Market value[c] ($mil)	513.20	3509.85	14,340.53	50.22	6449.28
Sales[d] ($mil)	355.52	2350.77	10,156.84	27.94	4326.57

Notes:.
[a] Data on each variable are deflated using country-specific CPIs.
[b] Total operational assets, net of funding.
[c] Market value of the firm's net operational assets.
[d] Firm revenue.

Table 2
Percentage of country/industry/year samples not rejected.

Distribution	Anderson–Darling test	Cramér–von Mises test
Cauchy	97.55	98.25
Exponential	0.00	0.00
Gamma	44.56	45.51
Laplace	64.67	66.72
Logistic	56.89	60.83
Log-normal	42.17	44.86
Normal	42.37	44.36
Weibull	33.98	33.73
Number of samples	2004	2004
Time period	1999–2010	1999–2010

Note: statistical significance is determined at the 5% level.

test statistic in only 42.37% of the 2004 country/industry/year samples. Next, we test the null hypothesis that the firm growth rates are distributed according to the Cauchy, exponential, gamma, Laplace, logistic, log-normal, normal, and Weibull distributions.[6] At the 5% significance level, the null hypothesis that the firm growth rates are distributed according to the Cauchy distribution is not rejected in 97.55% of all country/industry/year samples, while the Laplace and logistic distributions are not rejected in 64.67% and 56.89% of all samples. Using the Cramér–von Mises test statistic, the Cauchy distribution is not rejected in 98.25% of all country/industry/year samples, while the Laplace and logistic distributions are not rejected in 66.72% and 60.83% of all samples. We conclude that the Cauchy distribution best fits the empirical distributions of firm growth rates in these samples.

Prior research has shown that the Laplace distribution fits the tails of firm growth rate distributions somewhat inaccurately (see, for example, Ref. [14]). As a result, *ad hoc* adjustments have been suggested to "fatten" the tails of the Laplace distribution to better fit empirical firm growth rate distributions. For example, Fu et al. [21], Buldyrev et al. [8], and Riccaboni et al. [10] use distributions that combine a Laplace central region with power-law tails. As demonstrated here, such *ad hoc* adjustments are not necessary with the Cauchy distribution. In only a small percentage of country/industry/year combinations do the data reject the null hypothesis of a Cauchy distribution.

Tables 3 and 4 show the EDF test results by country using the Anderson–Darling and Cramér–von Mises test statistics. The tables show that in each of the 27 countries, the Cauchy distribution best fits the empirical distributions of firm growth rates in the industry/year samples.

Schwartzkopf et al. [22] and Metzig and Gordon [16] explain fat-tailed firm growth rate distributions as being a collective phenomenon. This suggests examining the size distributions of firm growth rates across industries for given country/year samples. Table 5 shows the EDF test results using the Anderson–Darling and Cramér–von Mises test statistics for the 209 country/year samples in our data. The table shows that the Cauchy distribution best fits the empirical distributions of firm growth rates across industries in the country/year samples.

Table 2 shows that among the 2004 country/industry/year samples, the Laplace distribution is the next-best fit of the empirical distributions of firm growth rates. In this sense, our results are consistent with the consensus finding that the Laplace distribution performs reasonably well as a fit of the empirical distribution of firm growth rates. However, Tables 2, 3, 4, and 5 demonstrate clearly that the Cauchy distribution performs substantially better.

The Cauchy distribution is a stable, "fat-tailed" distribution with an infinite variance. This variance condition has an interesting economic implication—both very low and very high firm growth rates will occur much more frequently than predicted by the normal (Gaussian) distribution. For example, in the domain of observations three or more standard deviations away from the mean of the standard normal distribution, the Cauchy distribution with location 0 and scale 1 has more than 75 times the percentage of observations than the corresponding percentage for the standard normal distribution.

[6] The EDF tests for the Cauchy, logistic, and normal distributions are performed using unadjusted data on the firms' growth rates. Tests for the other distributions are performed after adding a constant to the growth rates so there are only positive values for that variable.

Table 3
Percentage of industry/year samples not rejected using the Anderson–Darling test.

Country	Number of samples	Cauchy	Exponential	Gamma	Laplace	Logistic	Log-normal	Normal	Weibull
Australia	57	100.00	0.00	42.11	61.40	49.12	40.35	38.60	29.82
Brazil	11	100.00	0.00	45.45	18.18	54.55	45.45	45.45	0.00
Canada	78	98.72	0.00	51.28	56.41	53.85	48.72	47.44	43.59
Chile	7	100.00	0.00	85.71	71.43	85.71	85.71	85.71	71.43
China	189	100.00	0.00	38.62	65.08	48.15	37.04	36.51	29.63
France	34	100.00	0.00	73.53	61.76	82.35	70.59	64.71	52.94
Germany	41	100.00	0.00	68.29	63.41	78.05	68.29	65.85	60.98
Greece	7	100.00	0.00	57.14	85.71	57.14	57.14	57.14	42.86
Hong Kong	28	100.00	0.00	39.29	50.00	42.86	39.29	32.14	35.71
India	84	100.00	0.00	53.57	77.38	69.05	50.00	51.19	35.71
Italy	20	100.00	0.00	75.00	60.00	80.00	75.00	65.00	55.00
Japan	307	96.42	0.00	45.93	68.40	60.59	42.35	43.32	27.04
Malaysia	90	100.00	0.00	47.78	65.56	64.44	46.67	47.78	46.67
Mexico	9	100.00	0.00	88.89	55.56	88.89	88.89	88.89	66.67
Norway	4	100.00	0.00	50.00	75.00	75.00	50.00	50.00	50.00
Poland	9	100.00	0.00	88.89	77.78	88.89	77.78	88.89	66.67
Saudi Arabia	4	100.00	0.00	75.00	0.00	75.00	50.00	50.00	50.00
Singapore	20	100.00	0.00	55.00	30.00	55.00	50.00	50.00	50.00
South Africa	12	100.00	0.00	75.00	41.67	75.00	75.00	75.00	50.00
South Korea	125	100.00	0.00	54.40	72.80	67.20	50.40	51.20	51.20
Spain	5	100.00	0.00	80.00	40.00	80.00	80.00	80.00	80.00
Sweden	7	100.00	0.00	85.71	57.14	85.71	85.71	85.71	85.71
Switzerland	18	100.00	0.00	61.11	33.33	66.67	61.11	61.11	61.11
Taiwan	105	96.19	0.00	60.00	74.29	76.19	53.33	56.19	51.43
Thailand	22	100.00	0.00	81.82	50.00	86.36	77.27	77.27	59.09
United Kingdom	139	100.00	0.00	64.03	63.31	72.66	61.15	60.43	54.68
United States	572	94.23	0.00	23.25	64.34	39.34	22.20	23.08	14.51

Note: Statistical significance is determined at the 5% level.

Table 4
Percentage of industry/year samples not rejected using the Cramér–von Mises test.

Country	Number of samples	Cauchy	Exponential	Gamma	Laplace	Logistic	Log-normal	Normal	Weibull
Australia	57	100.00	0.00	43.86	59.65	54.39	43.86	40.35	29.82
Brazil	11	100.00	0.00	45.45	27.27	54.55	45.45	45.45	0.00
Canada	78	98.72	0.00	50.00	57.69	65.38	48.72	48.72	43.59
Chile	7	100.00	0.00	85.71	71.43	85.71	85.71	85.71	71.43
China	189	100.00	0.00	38.10	70.90	53.97	37.57	37.04	31.22
France	34	100.00	0.00	73.53	61.76	82.35	70.59	67.65	52.94
Germany	41	100.00	0.00	68.29	65.85	78.05	68.29	68.29	56.10
Greece	7	100.00	0.00	57.14	85.71	57.14	57.14	57.14	42.86
Hong Kong	28	100.00	0.00	39.29	50.00	46.43	39.29	35.71	35.71
India	84	100.00	0.00	53.57	78.57	71.43	52.38	52.38	30.95
Italy	20	100.00	0.00	75.00	60.00	85.00	75.00	75.00	50.00
Japan	307	98.37	0.00	47.23	71.01	65.47	44.95	44.95	26.71
Malaysia	90	100.00	0.00	48.89	67.78	68.89	50.00	48.89	47.78
Mexico	9	100.00	0.00	88.89	55.56	77.78	88.89	88.89	66.67
Norway	4	100.00	0.00	50.00	75.00	75.00	50.00	50.00	50.00
Poland	9	100.00	0.00	88.89	66.67	88.89	88.89	88.89	77.78
Saudi Arabia	4	100.00	0.00	50.00	0.00	75.00	50.00	25.00	50.00
Singapore	20	100.00	0.00	55.00	30.00	60.00	55.00	50.00	50.00
South Africa	12	100.00	0.00	75.00	41.67	75.00	75.00	75.00	41.67
South Korea	125	100.00	0.00	56.00	75.20	72.00	54.40	56.00	49.60
Spain	5	100.00	0.00	80.00	40.00	80.00	80.00	80.00	80.00
Sweden	7	100.00	0.00	85.71	71.43	85.71	85.71	85.71	85.71
Switzerland	18	100.00	0.00	61.11	33.33	61.11	61.11	61.11	61.11
Taiwan	105	98.10	0.00	59.05	77.14	78.10	59.05	60.00	49.52
Thailand	22	100.00	0.00	81.82	50.00	86.36	81.82	77.27	59.09
United Kingdom	139	100.00	0.00	65.47	62.59	77.70	64.75	63.31	53.96
United States	572	95.28	0.00	25.52	66.43	42.66	25.52	25.17	15.38

Note: Statistical significance is determined at the 5% level.

At four or more standard deviations away from the mean, the Cauchy distribution has more than 2000 times the percentage of observations.

Probability–probability (P–P) plots and quantile–quantile (Q–Q) plots of the empirical and theoretical distributions provide additional goodness-of-fit measures and show further that the Cauchy distribution describes the distribution of firm growth rates. If $F(\cdot)$ and $G(\cdot)$ represent two distribution functions with the same support, then the P–P plot in functional form

Table 5
Percentage of country/year samples not rejected.

Distribution	Anderson–Darling test	Cramér–von Mises test
Cauchy	99.39	99.58
Exponential	0.00	0.00
Gamma	60.44	60.63
Laplace	54.33	55.46
Logistic	68.11	69.96
Log-normal	58.39	60.25
Normal	57.47	58.95
Weibull	49.09	48.07
Number of samples	209	209
Time period	1999–2010	1999–2010

Note: Statistical significance is determined at the 5% level.

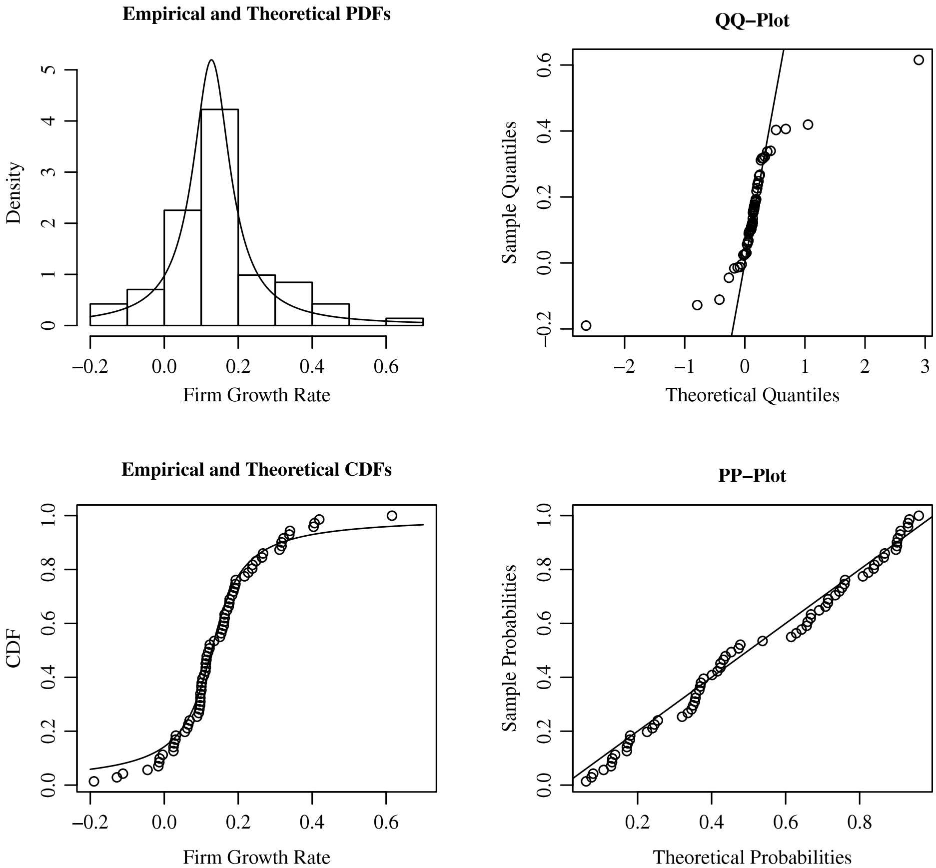

Fig. 1. P–P and Q–Q plots and the empirical and theoretical Cauchy densities and distributions for the US Chemicals industry in 2010.

equals $F(G^{-1}(p))$, where $0 < p < 1$. The Q–Q plot equals $F^{-1}(G(x))$ for all real x (see Ref. [23]). For example, Fig. 1 presents the P–P and Q–Q plots and the empirical and theoretical Cauchy densities and distributions for the US Chemicals industry in 2010. The Cauchy distribution is parameterized by the maximum-likelihood location and scale parameters. Both diagrams show clustering along the relevant 45-degree line, providing empirical evidence that the data are distributed according to the estimated Cauchy distribution.

4. Conclusions

EDF testing rejects the hypothesis that the size distribution of firm growth rates is Gaussian. We show that while the Laplace distribution does reasonably well, as found in prior studies, in approximating the distribution of firm growth rates, the Cauchy distribution performs substantially better. As Dosi et al. [24] conclude, this has an important economic implication: "this revealed structure in the stochastic process describing industrial evolution bears clear signs common to all complex system dynamics including, the fat-tailed distributions in the rates of changes of all the variable of interest.

That, in turn, is likely to witness for the existence of some underlying correlation mechanism, which makes the system self-organized in its growth process. In these respects, all the evidence on industrial change corroborates the exciting conjecture that evolutionary phenomena tend to undergo non-Gaussian lives influenced by persistent positive or negative interactions among agents within and across the relevant populations''.[7] Our results invite research on dynamic models of competition in which firms' growth rates follow a Cauchy distribution.

References

[1] M.H. Stanley, L.A. Amaral, S.V. Buldyrev, S. Havlin, H. Leschhorn, P. Maass, M.A. Salinger, H.E. Stanley, Scaling behaviour in the growth of companies, Nature 379 (6568) (1996) 804–806.

[2] R.L. Axtell, Zipf distribution of us firm sizes, Science 293 (5536) (2001) 1818–1820.

[3] A. Palestrini, Analysis of industrial dynamics: a note on the relationship between firms' size and growth rate, Econom. Lett. 94 (3) (2007) 367–371.

[4] L.M. Cabral, J. Mata, On the evolution of the firm size distribution: facts and theory, Amer. Econ. Rev. (2003) 1075–1090.

[5] L.A. Nunes Amaral, S.V. Buldyrev, S. Havlin, P. Maass, M.A. Salinger, H. Eugene Stanley, M.H. Stanley, Scaling behavior in economics: the problem of quantifying company growth, Physica A 244 (1) (1997) 1–24.

[6] G. Bottazzi, G. Dosi, M. Lippi, F. Pammolli, M. Riccaboni, Innovation and corporate growth in the evolution of the drug industry, Int. J. Ind. Organiz. 19 (7) (2001) 1161–1187.

[7] G. Bottazzi, A. Secchi, Why are distributions of firm growth rates tent-shaped? Econom. Lett. 80 (3) (2003) 415–420.

[8] S.V. Buldyrev, J. Growiec, F. Pammolli, M. Riccaboni, H.E. Stanley, The growth of business firms: facts and theory, J. Eur. Econ. Assoc. 5 (2–3) (2007) 574–584.

[9] S. Alfarano, M. Milaković, Does classical competition explain the statistical features of firm growth? Econom. Lett. 101 (3) (2008) 272–274.

[10] M. Riccaboni, J. Growiec, F. Pammolli, Innovation and Corporate Dynamics: a Theoretical Framework. Tech. Rep., University Library of Munich, Germany, 2011.

[11] W. Hölzl, Persistence, survival and growth: a closer look at 20 years of high growth firms and firm dynamics in Austria, in: Workshop on High-Growth Firms, Stockholm, Sweden, 2011.

[12] E.J. Erlingsson, S. Alfarano, M. Raberto, H. Stefánsson, On the distributional properties of size, profit and growth of Icelandic firms, J. Econ. Interact. Coord. 8 (1) (2013) 57–74.

[13] R. Ibragimov, Heavy-tailed densities, in: S.N. Durlauf, L.E. Blume (Eds.), The New Palgrave Dictionary of Economics, Palgrave Macmillan, Basingstoke, 2009.

[14] G. Bottazzi, A. Coad, N. Jacoby, A. Secchi, Corporate growth and industrial dynamics: evidence from French manufacturing, Appl. Econ. 43 (1) (2011) 103–116.

[15] A. Manas, French butchers don't do quantum physics, Econom. Lett. 103 (2) (2009) 101–106.

[16] C. Metzig, M.B. Gordon, A model for scaling in firms' size and growth rate distribution, Physica A 398 (2014) 264–279.

[17] G. Bottazzi, A. Secchi, Explaining the distribution of firm growth rates, Rand J. Econ. 37 (2) (2006) 235–256.

[18] G. Bottazzi, A. Secchi, Common properties and sectoral specificities in the dynamics of US manufacturing companies, Rev. Ind. Organ. 23 (3–4) (2003) 217–232.

[19] E.L. Lehmann, J.P. Romano, Testing Statistical Hypotheses, Springer, 2006.

[20] M.L. Delignette-Muller, R. Pouillot, J.-B. Denis, C. Dutang, fitdistrplus: help to fit of a parametric distribution to non-censored or censored data. R package version 1.0-2, 2014.

[21] D. Fu, F. Pammolli, S.V. Buldyrev, M. Riccaboni, K. Matia, K. Yamasaki, H.E. Stanley, The growth of business firms: theoretical framework and empirical evidence, Proc. Natl. Acad. Sci. USA 102 (52) (2005) 18801–18806.

[22] Y. Schwartzkopf, R.L. Axtell, J.D. Farmer, The cause of universality in growth distributions, 2010. arXiv:1004.5397.

[23] E.B. Holmgren, The PP plot as a method for comparing treatment effects, J. Amer. Statist. Assoc. 90 (429) (1995) 360–365.

[24] G. Dosi, S. Lechevalier, A. Secchi, Introduction: interfirm heterogeneity—nature, sources and consequences for industrial dynamics, Ind. Corp. Change 19 (6) (2010) 1867–1890.

[7] Dosi et al. [24, p. 1872].