Eur. Phys. J. B **57**, 205–211 (2007)
DOI: 10.1140/epjb/e2007-00153-0

THE EUROPEAN PHYSICAL JOURNAL B

How do output growth-rate distributions look like? Some cross-country, time-series evidence

G. Fagiolo[1,a], M. Napoletano[2,b], and A. Roventini[3,c]

[1] University of Verona, Italy and Sant'Anna School of Advanced Studies, Pisa, Italy
[2] Chair of Systems Design, ETH Zurich, 8032 Zurich, Switzerland and Sant'Anna School of Advanced Studies, Pisa, Italy
[3] University of Modena and Reggio Emilia, Italy and Sant'Anna School of Advanced Studies, Pisa, Italy

Received 31 August 2006 / Received in final form 19 December 2006
Published online 30 May 2007 –

Abstract. This paper investigates the statistical properties of within-country gross domestic product (GDP) and industrial production (IP) growth-rate distributions. Many empirical contributions have recently pointed out that cross-section growth rates of firms, industries and countries all follow Laplace distributions. In this work, we test whether also within-country, time-series GDP and IP growth rates can be approximated by tent-shaped distributions. We fit output growth rates with the exponential-power (Subbotin) family of densities, which includes as particular cases both Gaussian and Laplace distributions. We find that, for a large number of OECD (Organization for Economic Cooperation and Development) countries including the US, both GDP and IP growth rates are Laplace distributed. Moreover, we show that fat-tailed distributions robustly emerge even after controlling for outliers, autocorrelation and heteroscedasticity.

PACS. 89.65.Gh Economics; econophysics, financial markets, business and management – 89.90.+n Other topics in areas of applied and interdisciplinary physics – 02.60.Ed Interpolation; curve fitting

1 Introduction

In recent years, empirical cross-section growth-rate distributions of diverse economic entities (i.e., firms, industries and countries) have been extensively explored by both economists and physicists [1–10].

The main result of this stream of literature was that, no matter the level of aggregation, growth rates tend to *cross-sectionally* distribute according to densities that display tails fatter than those of a Gaussian distribution. From an economic point of view, this implies that growth patterns tend to be quite lumpy: large growth events, no matter if positive or negative, seem to be more frequent than what a Gaussian model would predict.

For example, at the microeconomic level, growth rates of US manufacturing firms (pooled across years) appear to distribute according to a Laplace [1,4]. This result robustly holds even if one disaggregates across industrial sectors and/or considers cross-section distributions in each given year [5,6]. Moreover, in some countries (e.g., France) firm growth rates display tails even fatter than those of a Laplace density [9]. Interestingly, similar findings are replicated at higher aggregation levels: both growth rates of industrial sectors [7,10] and countries [2,3,7] display tent-shaped patterns.

Existing studies have been focusing only on *cross-section* distributions. In this paper, on the contrary, we ask whether fat-tailed distributions also emerge *across time within a single country.* More precisely, for any given country, we consider gross domestic product (GDP) and industrial production (IP) time series and we test whether their growth-rate distributions can be well approximated by densities with tails fatter than the Gaussian ones.

Our analysis shows that in the US both GDP and IP growth rates distribute according to a Laplace. Similar results hold for a large sample of OECD (Organization for Economic Cooperation and Development) countries. Interestingly enough, this evidence resists to the removal of outliers, heteroscedasticity and autocorrelation from the original time series. Therefore, fat-tails emerges as a inherent property of output growth residuals, i.e. a fresh stylized fact of output dynamics.

Our work differs from previous, similar ones [2,3,7] in a few other respects. *First*, we depart from the common practice of using annual data to build output growth-rate distributions. We instead employ monthly and quarterly data. This allows us to get longer series and

[a] *Present address*: Sant'Anna School of Advanced Studies, Piazza Martiri della Libertà 33, 56127 Pisa, Italy;
e-mail: `giorgio.fagiolo@univr.it`
[b] e-mail: `mnapoletano@ethz.ch`
[c] e-mail: `aroventini@sssup.it`

better appreciate their business cycle features. *Second*, we fit output growth rates with the exponential-power (Subbotin) distribution [11], which encompasses Laplace and Gaussian distributions as special cases. This choice allows us to measure how far empirical growth-rate distributions are from the Normal benchmark. *Finally*, we check the robustness of our results to the presence of outliers, heteroscedasticity and autocorrelation in output growth-rate dynamics.

The paper is organized as follows. In Section 2 we describe our data and the methodology we employ in our analysis. Empirical results are presented in Section 3. Finally, Section 4 concludes.

2 Data and methodology

Our study employs two sources of (seasonally adjusted) data. As far as the US are concerned, we use data drawn from the FRED database. More specifically, we consider quarterly real GDP ranging from $1947Q1$ to $2005Q3$ (235 observations) and monthly IP from $1921M1$ to $2005M10$ (1018 observations). Analyses for the OECD sample of countries are performed by relying on monthly IP data from the "OECD Historical Indicators for Industry and Services" database ($1975M1 - 1998M12$, 288 observations).

The main object of our analysis are output growth rates $g(t)$, defined as:

$$g(t) = \frac{Y(t) - Y(t-1)}{Y(t-1)} \cong y(t) - y(t-1) = dy(t), \quad (1)$$

where $Y(t)$ is the output level (GDP or IP) at time t in a given country, $y(t) = \ln[\mathrm{Y(t)}]$ and d is the first-difference operator.

Let $T_n = \{t_1, ..., t_n\}$ be the time interval over which we observe growth rates. The distribution of growth rates is therefore defined as $G_{T_n} = \{g(t), t \in T_n\}$. We study the shape of G_{T_n} in each given country following a parametric approach. More precisely, we fit growth rates with the exponential-power (Subbotin) family of densities [32], whose functional form reads:

$$f(x) = \frac{1}{2ab^{\frac{1}{b}}\Gamma(1+\frac{1}{b})} e^{-\frac{1}{b}|\frac{x-m}{a}|^b}, \quad (2)$$

where $a > 0$, $b > 0$ and $\Gamma(\cdot)$ is the Gamma function. The Subbotin distribution is thus characterized by three parameters: a *location* parameter m, a *scale* parameter a and a *shape* parameter b. The location parameter controls for the mean of the distribution. Therefore it is equal to zero up to a normalization that removes the average growth rate. The scale parameter is proportional to the standard deviation.

The shape parameter is the crucial one for our aims, as it directly gives information about the fatness of the tails: the larger b, the thinner are the tails. Note that if $b = 1$ the distribution reduces to a Laplace, whereas for $b = 2$ we recover a Gaussian. Values of b smaller than one indicate super-Laplace tails (see Fig. 1 for an illustration).

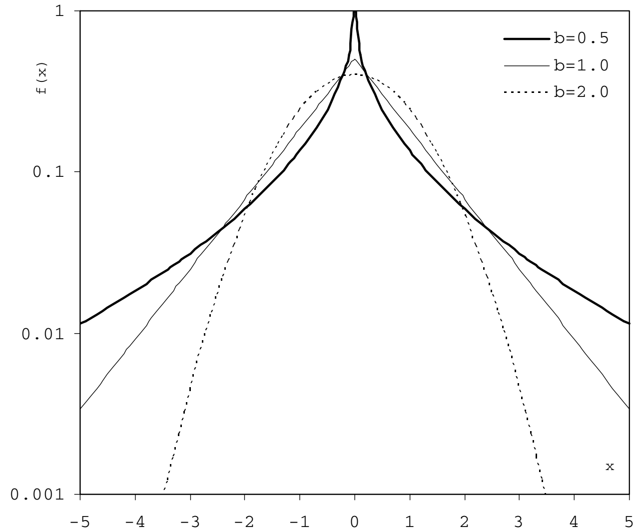

Fig. 1. The exponential-power (Subbotin) density for $m = 0$, $a = 1$ and different shape parameter values: (i) $b = 2$: Gaussian density; (ii) $b = 1$: Laplace density; (iii) $b = 0.5$: Subbotin with super-Laplace tails. Note: Log scale on the y-axis.

In our exercises, we fit empirical distributions G_{T_n} with the Subbotin density (2) by jointly estimating the three parameters by a standard maximum likelihood procedure (see [12] for details).

3 Empirical results

In this section we present our main empirical results. We begin with an analysis of US growth-rate distributions. Next, we extend our results to other OECD countries. Finally, we turn to a robustness analysis of growth residuals, where we take into account the effects of outliers, heteroscedasticity and autocorrelation.

3.1 Exploring US output growth-rate distributions

Let us start by some descriptive statistics on US output growth rates. Table 1 reports the first four moments of US time series. Standard deviations reveal that after World War II growth rates of industrial production and GDP have been characterized by similar volatility levels. The standard deviation of IP growth rates becomes higher if the series is extended back to 1921. Skewness is close to zero: -0.09 for GDP and $\approx$0.33 for IP. Notice that both the Jarque-Bera and Lilliefors normality tests reject the hypothesis that our series are normally distributed. Furthermore, the relatively high reported kurtosis values suggest that output growth-rate distributions display tails fatter than the Gaussian distribution. In order to better explore this evidence, we fit US output growth-rate distribution with the Subbotin density (see Eq. (2)).

Consider GDP first. In the first row of Table 2, we show the maximum-likelihood estimates of Subbotin parameters and their standard errors [33]. Estimates indicate that GDP growth rates seem to distribute according

Table 1. US output time series: summary statistics. P-values in parentheses.

Series	Mean	Std. Dev.	Skewness	Kurtosis	Jarque-Bera test	Lilliefors test
GDP	0.0084	0.0099	−0.0891	4.2816	15.4204 (0.0000)	0.0623 (0.0279)
IP (1921)	0.0031	0.0193	0.3495	14.3074	5411.7023 (0.0000)	0.1284 (0.0000)
IP (1947)	0.0028	0.0098	0.3295	8.1588	784.0958 (0.0000)	0.0822 (0.0000)

Table 2. US output growth-rate distribution: estimated subbotin parameters.

	Estimated parameters					
	$\widehat{b}$		$\widehat{a}$		$\widehat{m}$	
Series	Par.	Std. Err.	Par.	Std. Err.	Par.	Std. Err.
GDP	1.1771	0.1484	0.0078	0.0006	0.0082	0.0006
IP (1921)	0.6215	0.0331	0.0091	0.0004	0.0031	0.0002
IP (1947)	0.9940	0.0700	0.0068	0.0003	0.0030	0.0003

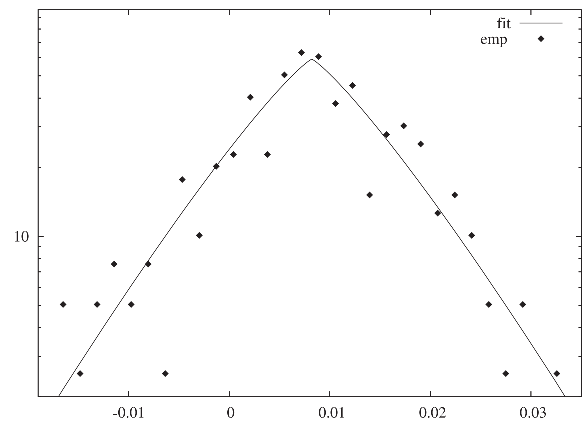

Fig. 2. Binned empirical densities of US GDP growth rates (emp) vs. Subbotin fit (fit).

to a Laplace: the shape parameter $\widehat{b}$ is equal to 1.18, very close to the theoretical Laplace value of one. Therefore, US output growth rates display tails fatter than a normal distribution. This can be also seen from Figure 2, where we plot the binned empirical density vis-à-vis the fitted one.

Next, we employ monthly industrial production (IP) as a proxy of US output [34]. Notice that, by focusing on IP growth, we can study a longer time span and thus improve our estimates by employing a larger number of observations. During the period 1921–2005, the IP growth-rate distribution displays tails much fatter than the Laplace distribution (see Fig. 3 and the 2nd row of Tab. 2), an outcome probably due to the turmoils of the Great Depression.

Moreover, in order to better compare IP growth-rate distribution with the GDP one, we also carry out an investigation on the post war period only (1947–2005). Notwithstanding this breakdown, our results remain unaltered. In the post-war period, the IP growth-rate distribution exhibits the typical "tent-shape" of the Laplace density (cf. Fig. 4). This outcome is confirmed by a $\widehat{b}$ very close to one (see the third row of Tab. 2). As pointed out

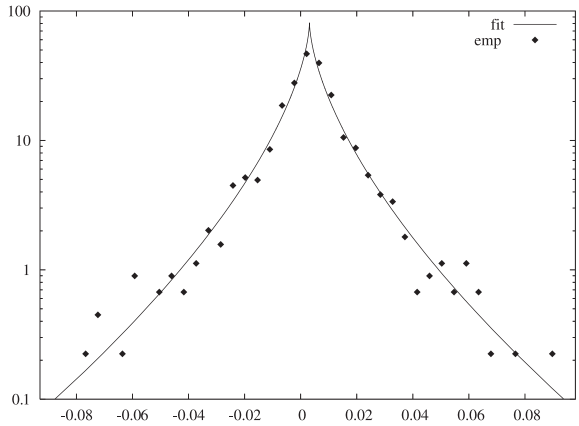

Fig. 3. Binned empirical densities of US IP growth rates vs. Subbotin fit. Time period: $1921M1-2005M10$.

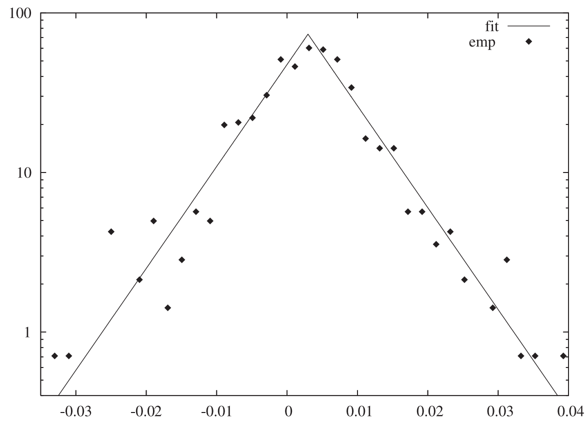

Fig. 4. Binned empirical densities of US IP growth rates vs. Subbotin fit. Time period: $1947M1-2005M10$.

by the lower standard error, the estimate of b is much more robust when we employ IP series instead of the GDP one.

To perform a more precise check, one might also compute the Cramér-Rao interval $[\widehat{b}-2\sigma(\widehat{b}), \widehat{b}+2\sigma(\widehat{b})]$, where $\sigma(\widehat{b})$ is the standard error of $\widehat{b}$ (Tab. 2, second column). A back-of-the-envelope computation shows that, for all the three growth-rate series, normality is always rejected. Moreover, one cannot reject the Laplace hypothesis for both GDP and IP-1947 series, whereas tails appear to be super-Laplace for IP-1921 [35].

Albeit Cramér-Rao intervals are valid only asymptotically, all the above results are confirmed by both standard goodness-of-fit (GoF) tests (e.g., Kolmogorov-Smirnov, Kuiper, Cramér-Von Mises, Quadratic Anderson-Darling) and likelihood ratio tests [36].

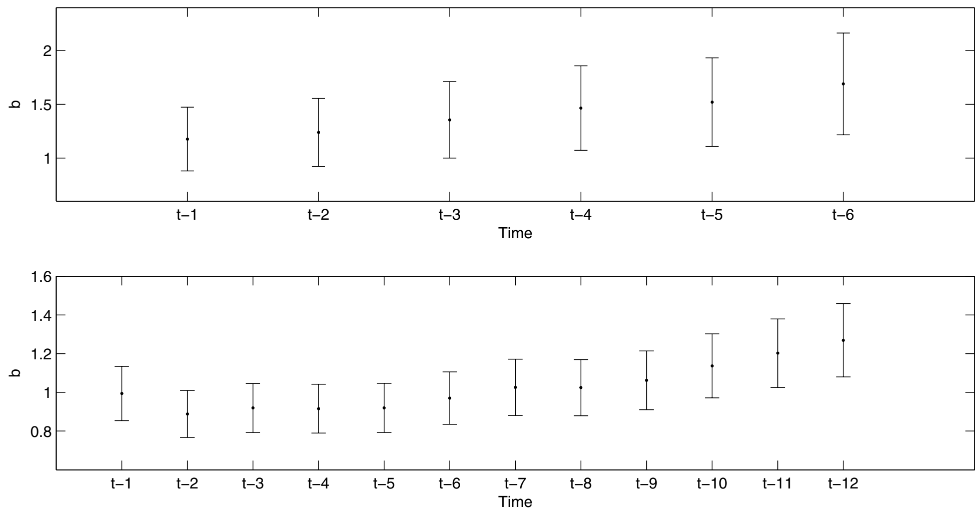

Fig. 5. US output growth rates: estimated subbotin b parameter for different time lags. Error bars (Cramér-rao bounds): $\pm 2\sigma(\widehat{b})$. Top panel: GDP data; bottom panel: IP data.

Finally, in line with [14] we inspect the distributions of output growth rates computed over longer lags. More precisely, we consider growth rates now defined as:

$$g_\tau(t) = \frac{Y(t) - Y(t-\tau)}{Y(t-\tau)} \cong y(t) - y(t-\tau) = dy_\tau(t), \quad (3)$$

where $\tau = 1, 2, ..., 6$ when GDP series is employed, and $\tau = 1, 2, ..., 12$ when IP series is under study. In line with [14], we find that the shape parameter estimated on GDP data becomes higher as τ increases (cf. the top panel of Fig. 5). When we consider IP series, the $\widehat{b}$ first falls and then starts rising (see the bottom panel of Fig. 5). Therefore, as the "growth lag" increases, tails become thinner (see [15] for similar evidence in the contest of stock returns). Nevertheless, estimated shape coefficients almost always remain statistically smaller than two.

3.2 Cross-country analyses

In the previous section we have provided evidence in favor of fat-tailed (Laplace) US output growth-rate distributions. We now perform a cross-country analysis in order to assess whether this regularity pertains to the US output only, or it might also be observed in other developed countries. Our analysis focuses (in addition to the US) on the following OECD countries: Canada, Japan, Austria, Belgium, Denmark, France, Germany, Italy, The Netherlands, Spain, Sweden, and the UK.

We start by analyzing the basic statistical properties of output growth-rate time series (cf. Tab. 3). In order to keep a sufficient time-series length, we restrict our study to industrial production series only. The standard deviations of the IP series range from 0.0073 (US) to 0.0404 (Japan). In half of the countries that we have analyzed, the distributions of IP growth rates seem to be slightly right-skewed, whereas in the other half they appear to be slightly left-skewed. The analysis of the kurtosis reveals that in every country of the sample the IP growth-rate distribution is more leptokurtic than the Normal distribution. Indeed, apart from Spain and Canada, standard normality tests reject the hypothesis that IP growth series are normally distributed.

Given this descriptive background, we turn to a country-by-country estimation of the Subbotin distributions. Estimated coefficients are reported in Table 4.

The results of the cross-country analysis confirm that output growth rates distribute according to a Laplace almost everywhere [37]. Excluding Canada, estimated "shape" coefficients are always close to 1. If one considers the Cramér-Rao interval $[\widehat{b} - 2\sigma(\widehat{b}), \widehat{b} + 2\sigma(\widehat{b})]$, the only country where output growth-rate distribution does not appear to be Laplace is Canada, whose $\widehat{b}$-interval lies above one [38].

3.3 Robustness checks: outliers, heteroscedasticity, and autocorrelation

The foregoing discussion has pointed out that within-country output growth-rate distributions are markedly non-Gaussian. The evidence in favor of Laplace (or super-Laplace) densities robustly arises in the majority of OECD countries, it does not depend on the way we measure output (GDP or IP), and it emerges also at frequencies more amenable for the study of business cycles dynamics (i.e. quarterly and monthly). Notice also that our analysis does not show any clear evidence in favor of asymmetric Laplace (or Subbotin) growth-rate distributions [39]. Hence, almost all OECD countries seem to exhibit (with a probability higher than we would expect) large, positive growth events with the same likelihood of large, negative ones.

This "fresh" stylized fact on output dynamics must be however scrutinized vis-à-vis a number of robustness

Table 3. Cross-country analysis of IP time series: summary statistics. P-values in parentheses.

Series	Mean	Std. Dev.	Skewness	Kurtosis	Jarque-Bera test	Lilliefors test
Canada	0.0021	0.0113	−0.2317	3.5631	5.9848 (0.0400)	0.0391 (0.3556)
USA	0.0026	0.0073	−0.1505	4.6337	31.6281 (0.0000)	0.0705 (0.0000)
Japan	0.0027	0.0404	−0.2250	4.6895	35.0981 (0.0000)	0.0944 (0.0000)
Austria	0.0024	0.0253	0.1707	5.7806	90.8554 (0.0000)	0.0565 (0.0279)
Belgium	0.0013	0.0401	−0.5689	5.9446	115.6987 (0.0000)	0.0884 (0.0000)
Denmark	0.0025	0.0340	0.1214	7.2748	213.3210 (0.0000)	0.0958 (0.0000)
France	0.0013	0.0130	0.1525	3.7251	6.9217 (0.0300)	0.0740 (0.0007)
Germany	0.0015	0.0212	0.0098	9.2312	453.1891 (0.0000)	0.0875 (0.0000)
Italy	0.0017	0.0321	0.0453	5.8380	93.3429 (0.0000)	0.0692 (0.0021)
The Netherlands	0.0015	0.0285	−0.0350	6.5731	148.3145 (0.0000)	0.0741 (0.0007)
Spain	0.0017	0.0401	0.2559	4.0067	14.5026 (0.0000)	0.0469 (0.1310)
Sweden	0.0016	0.0302	−0.2955	37.0700	13627.2129 (0.0000)	0.1153 (0.0000)
UK	0.0012	0.0140	−0.1631	8.4090	342.3813 (0.0000)	0.0712 (0.0013)

Table 4. Cross-country analysis of IP time series: estimated subbotin parameters.

	Estimated parameters					
	$\widehat{b}$		$\widehat{a}$		$\widehat{m}$	
Country	Par.	Std. Err.	Par.	Std. Err.	Par.	Std. Err.
Canada	1.6452	0.2047	0.0104	0.0007	0.0020	0.0010
USA	1.2980	0.1516	0.0060	0.0004	0.0031	0.0004
Japan	0.8491	0.0901	0.0259	0.0020	0.0021	0.0014
Austria	1.2499	0.1446	0.0204	0.0014	0.0010	0.0014
Belgium	1.0202	0.1125	0.0284	0.0021	0.0011	0.0017
Denmark	0.8063	0.0847	0.0215	0.0017	0.0000	0.0012
France	1.2623	0.1464	0.0106	0.0008	0.0010	0.0007
Germany	0.9768	0.1067	0.0144	0.0011	0.0024	0.0008
Italy	1.0778	0.1204	0.0237	0.0017	0.0010	0.0015
The Netherlands	1.2133	0.1393	0.0223	0.0016	0.0019	0.0015
Spain	1.4583	0.1755	0.0352	0.0024	0.0021	0.0029
Sweden	0.8826	0.0944	0.0168	0.0013	0.0010	0.0009
UK	1.0972	0.1230	0.0103	0.0008	0.0019	0.0006

checks. More precisely, the above results can be biased by two classes of problems. First, the very presence of fat-tails in the distribution of country-level growth rates might simply be due to the presence of outliers. Thus, one should remove such outliers from the series and check whether fat tails are still there. Second, our within-country analysis relies on pooling together growth-rate observations over time. Strictly speaking, the observations contained in G_{T_n} should come from i.i.d. random variables. In other words, we should verify that fat tails do not characterize growth rates only, but they are a robust feature of growth residuals (also known as "innovations"). To do so, one might remove the possible presence of any structure in growth-rate time series due to autocorrelation and heteroscedasticity, and then fit a Subbotin density to the residuals.

Our robustness analyses seem to strongly support the conclusion that fat-tails still characterize our series also after having controlled for outliers, autocorrelation and heteroscedasticity. More precisely, in the first row of Table 5 we have reported the estimates of the Subbotin parameters in the case of US GDP, after having removed the most common types of outliers [16]. The estimate for the shape parameter ($\widehat{b}$) still remains close to one, thus reinforcing evidence in favor of Laplace fat-tails.

Moreover, in order to remove any structure from the growth-rate process, we have fitted a battery of ARMA specifications to the growth-rate time series obtained after cleaning-up outliers and we have selected the best model through the standard Box and Jenkins's procedure. In Table 5, second row, we report — for the case of US GDP — our Subbotin estimates for the distribution of residuals of the best ARMA model, which turned out to be an AR(1) without drift. In particular this model was able to produce serially uncorrelated residuals. Indeed, none of the Ljung-Box tests performed on the residuals of this AR(1) model (at 9, 18 and 27 lags, respectively) rejected the null hypothesis of absence of autocorrelation in the residuals (both at 5% and 1% significance levels). However, the best fit for the distribution of the AR(1) residuals is still a Subbotin very close to a Laplace ($\widehat{b} = 1.2696$). Similar results hold also for the IP growth-rate series and are sufficiently robust across our sample of OECD countries. Finally, we ran (at 10, 15 and 20 lags, respectively) standard Ljung-Box and Engle's ARCH heteroscedasticity tests (on the squared residuals of our best ARMA

Table 5. US GDP growth-rate distribution: robustness checks.

	Estimated parameters					
	$\widehat{b}$		$\widehat{a}$		$\widehat{m}$	
After removing	Par.	Std. Err.	Par.	Std. Err.	Par.	Std. Err.
Outliers only	1.2308	0.1568	0.0073	0.0006	0.0000	0.0006
Outliers and autocorrelation	1.2696	0.1628	0.0071	0.0006	0.0000	0.0006

Note: Estimated Subbotin parameters after having removed outliers only (first row) and after having removed both outliers and autocorrelation (second row) from the original output growth-rate series. Outlier removal performed using TRAMO [17]. Autocorrelation removal performed fitting an ARMA model to outlier-free residuals. Best ARMA model: AR(1) w/o drift.

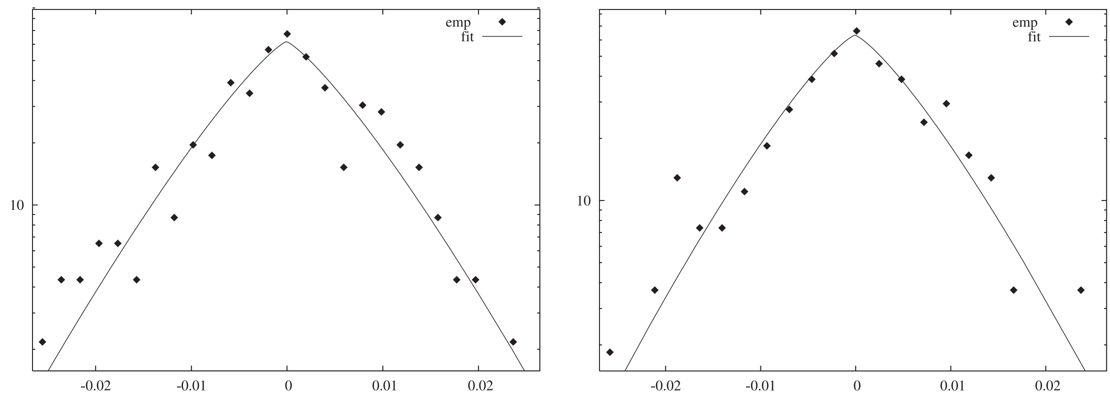

Fig. 6. Controlling for outliers and autocorrelation in US output growth rates. Binned Empirical Densities (emp) vs. Subbotin Fit (fit). Left: residuals after removing outliers only. Right: residuals after removing outliers and autocorrelation. Outlier removal performed using TRAMO [17]. Autocorrelation removal performed fitting an ARMA model to outlier-free residuals. Best ARMA model: AR(1) w/o drift.

models) without detecting evidence in favor of non-stationary variance over time. Indeed, for US GDP growth residuals both tests did not reject (at the 5% significance level) the null hypothesis of absence of heteroscedasticity.

As Figure 6 shows for US GDP, fat-tailed Laplace densities seem therefore to robustly emerge even after one washes away from the growth process both outliers and autocorrelation structure (i.e., when one considers growth residuals as *the* object of analysis).

4 Concluding remarks

In this paper we have investigated the statistical properties of GDP and IP growth-rate time-series distributions by employing quarterly and monthly data from a sample of OECD (Organization for Economic Cooperation and Development) countries.

We find that in the US, as well as in almost all other developed countries of our sample, output growth-rate time series distribute according to a symmetric Laplace density. This implies that the growth dynamics of aggregate output is *lumpy*, being considerably driven by "big events", either positive or negative. We have checked this result against a number of possible sources of bias. We find that lumpiness appears to be a very property of the data generation process governing aggregate output growth, as it appears to be robust to the removal of both outliers and autocorrelation.

At a very general and rather broad level, the robust emergence of fat-tailed distributions for within-country time series of growth rates and residuals can be interpreted as a fresh, new stylized fact on output dynamics, to be added to the long list of its other known statistical properties [40].

From a more empirical perspective, our results (together with the already mentioned cross-section ones [1–10]) ought to be be interpreted together with recent findings against log-normality for the cross-section distributions of firm and country size [24–28], and on power-law scaling in cross-country per-capita GDP distributions [29]. This joint empirical evidence seems to suggest that in economics the room for normally-distributed shocks and growth processes obeying the "Law of large numbers" and the "Central limit theorem" is much more limited than economists were used to believe. In other words, the general hint coming from this stream of literature is in favor of an increasingly "non-Gaussian" economics and econometrics. A consequence of this suggestion is that we should be very careful in using econometric testing procedures that are heavily sensible to normality of residuals [41]. On the contrary, testing procedures that are robust to non-Gaussian errors and/or tests based on Subbotin- or Laplace-distributed errors should be employed when necessary.

Finally, country-level, non-Gaussian growth-rate distributions (both within-country and cross-section) might have an important implication on the underlying generating processes. Suppose to interpret the country-level growth rate in a certain time period as the result of the aggregation of microeconomic (firm-level) growth shocks across all firms and industries in the same time period. The emergence of within-country non-Gaussian growth distributions strongly militates against the idea that country growth shocks are simply the result of aggregation of independent microeconomic shocks over time. Therefore, some strong correlating mechanism linking in a similar way at every level of aggregation the units to be aggregated seems to be in place. This interpretation is in line with the one proposed by [2,4,7] who envisage the widespread presence of fat tails as an indicator of the overall "complexity" of any growth process, mainly due to the strong inner inter-relatedness of the economic organizations under study.

Thanks to Giulio Bottazzi, Carolina Castaldi, Giovanni Dosi, Marco Lippi, Sandro Sapio, Angelo Secchi, Victor M. Yakovenko, and one anonymous referee for their stimulating and helpful comments. All usual disclaimers apply.

References

1. M.H. Stanley, L.A.N. Amaral, S.V. Buldyrev, S. Havlin, H. Leschhorn, P. Maass, M.A. Salinger, H.E. Stanley, Nature **379**, 804 (1996)
2. D. Canning, L.A.N. Amaral, Y. Lee, M. Meyer, H.E. Stanley, Econ. Lett. **60**, 335 (1998)
3. Y. Lee, L.A.N. Amaral, D. Canning, M. Meyer, H.E. Stanley, Phys. Rev. Lett. **81**, 3275 (1998)
4. L.A.N. Amaral, S.V. Buldyrev, S. Havlin, M.A. Salinger, H.E. Stanley, M.H. Stanley, Physica A **244**, 1 (1997)
5. G. Bottazzi, A. Secchi, Rev. Ind. Org. **23**, 217 (2003)
6. G. Bottazzi, A. Secchi, Econ. Lett. **80**, 415 (2003)
7. C. Castaldi, G. Dosi, Working Paper 2004/18, Laboratory of Economics and Management (LEM), Sant'Anna School of Advanced Studies, Pisa, Italy (2004)
8. D. Fu, F. Pammolli, S. Buldyrev, M. Riccaboni, K. Matia, K. Yamasaki, H. Stanley, Proceedings of the Nationaly Academy of Science **102**, 18801 (2005)
9. G. Bottazzi, A. Coad, N. Jacoby, A. Secchi, Working Paper 2005/21, Laboratory of Economics and Management (LEM), Sant'Anna School of Advanced Studies, Pisa, Italy (2005)
10. S. Sapio, G. Thoma, Working Paper 2006/09, Laboratory of Economics and Management (LEM), Sant'Anna School of Advanced Studies, Pisa, Italy (2006)
11. M. Subbotin, Matematicheskii Sbornik **31**, 296 (1923)
12. G. Bottazzi, Working Paper 2004/14, Laboratory of Economics and Management (LEM), Sant'Anna School of Advanced Studies, Pisa, Italy (2004)
13. G. Fagiolo, M. Napoletano, A. Roventini, Tech. Rep. 2006-23, Laboratory of Economics and Management, Sant'Anna School of Advanced Studies (2006), available in the RepEc/Ideas on-line archive at the URL: `http://ideas.repec.org/p/ssa/lemwps/2006-23.html`
14. G. Bottazzi, A. Secchi, RAND Journal of Economics **Forthcoming** (2006)
15. A. Silva, R. Prange, V. Yakovenko, Physica A **344**, 227 (2004)
16. C. Darné, O. Diebolt, J. Mon. Econ. **51**, 1449 (2004)
17. V. Goḿez, A. Maravall, in *A Course in Time Series Analysis*, edited by D. Peña, G. Tiao, R. Tsay (Wiley, New York, 2001)
18. C.R. Nelson, C.O. Plosser, J. Mon. Econ. **10**, (1982)
19. J. Cochrane, J. Polit. Econ. **96**, 893 (1988)
20. J. Cochrane, Quarterly J. Econ. **109**, 241 (1994)
21. O. Blanchard, D. Quah, Am. Econ. Rev. **79**, 655 (1989)
22. R.G. Rudebusch, Am. Econ. Rev. **83** (1993)
23. C.J. Murray, C.R. Nelson, J. Mon. Econ. **46**, (2000)
24. L. Cabral, J. Mata, Am. Econ. Rev. **93**, 1075 (2003)
25. G. Dosi, Working Paper 2005/17, Laboratory of Economics and Management (LEM), Sant'Anna School of Advanced Studies, Pisa, Italy (2005)
26. R. Axtell, CSED Working Paper 44, The Brookings Institution, Washington, DC (2006)
27. D. Quah, Econ. J. **106**, 1045 (1996)
28. D. Quah, J. Econ. Growth **2**, 27 (1997)
29. C. Di Guilmi, E. Gaffeo, M. Gallegati, Econ. Bull. **15**, 1 (2003)
30. J. Sutton, J. Econ. Lit. **35**, 40 (1997)
31. R. Barro, X. Sala-i Martin, J. Polit. Econ. **100**, 223 (1992)
32. More on fitting Subbotin distributions to economic data is in [5,6]
33. The period of strong growth experienced by the US economy after World War II is probably responsible for the positive location parameter $\hat{m}$ (0.0082), which implies a positive sample average growth rate
34. IP tracks very closely GDP in almost all countries. More precisely, the GDP-IP cross-correlation profile mimics from time $t-6$ to time $t+6$ the GDP auto-correlation profile
35. Cramér-Rao bounds are also graphically reported in Figure 5 at lag $t-1$
36. See [13] for details. Our exercises also show that fat tails emerge not only for the EP density, but also if one fits alternative families of heavy-tailed distributions (e.g. Cauchy, Student-t, Levy-stable). However, such alternative fat- and medium-tailed densities perform worse than the EP in fitting our data according to GoF tests
37. Again, all these results are confirmed by both GoF and likelihood ratio tests, see [13] for details
38. Another exception is Denmark, where the upper bound of its $\hat{b}$-interval is slightly below one
39. As shown in [13], likelihood ratio tests reject the hypothesis that data come from asymmetric EP densities
40. Two rather undisputed stylized facts of output dynamics — at least for the US — are: (i) GNP growth is positively autocorrelated over short horizons and has a weak and possibly insignificant negative autocorrelation over longer horizons; (ii) GNP appears to have an important trend-reverting component [18–23]
41. Such as Gibrat-like regressions for the dependence of firm growth on size [30] and cross-section country growth-rate analyses [31]