Macroeconomic Dynamics, **12**, 2008, 425–443. Printed in the United States of America.
doi: 10.1017/S1365100508070223

LÉVY-STABLE PRODUCTIVITY SHOCKS

EDOARDO GAFFEO
University of Trento

In this paper, we analyze the distribution of TFP growth rates at the four-digit sectoral level for the United States. We find that, contrary to the usual assumption employed in the literature on business cycles theory, technological shocks are not normally distributed. Instead, a Lévy-stable distribution with a divergent variance returns a better fit to the data.

Keywords: TFP, Lévy-Stable Distributions, Business Cycles

1. INTRODUCTION

This paper deals with productivity shocks at a sectoral level, and how they relate to aggregate business fluctuations. A well-known argument in multisector real business cycle models [see, e.g., Long and Plosser (1983)] is that as the number of sectors or industries considered in the analysis becomes large, aggregate volatility must tend to zero very quickly. This result, which follows directly from the Law of Large Numbers (LLN), rests on the hypothesis that each sector is periodically buffeted with idiosyncratic, identically and independently distributed shocks to Total Factor Productivity (TFP). As negative and positive shocks tend to cancel out each other, in an economy composed of N sectors—each one of approximately size $1/N$ of GDP—aggregate volatility must converge to zero at a rate $N^{1/2}$ [Lucas (1981)]. Furthermore, such a *curse of aggregation* is so compelling to offset, under rather general conditions, any shock-propagation effects due to factor demand linkages among industries [Dupor (1999)]. Hence, for a multisector business cycle model to be able to replicate aggregate fluctuations with a degree of volatility in line with that observed in real data, one has necessarily to appeal to aggregate shocks.

Several mechanisms recently have been proposed for allowing small idiosyncratic shocks to cause large macroeconomic fluctuations, all of them based on the idea of forcing the rate at which the LLN applies to slow down.[1] Horvath (1998, 2000), for instance, shows that if the input-output matrix is characterized by few full rows and many sparse columns, so that few sectors are key inputs for all the others, the LLN applies at a rate proportional to the rate of increase in the number

Many thanks to an anonymous referee, Marco Bee, Giuseppe Espa, and seminar participants at the Universities of Essex and Politecnica delle Marche for their comments. Address correspondence to: Edoardo Gaffeo, Department of Economics and CEEL, University of Trento, Via Inama 5, I-38100 Trento, Italy; e-mail: edoardo.gaffeo@economia.unitn.it.

 1365-1005/08

of such key input sectors. Given that real data for the U.S. input-output matrix suggest that out of N sectors, only $N^{1/2}$ are broadly used as inputs, a sizable portion of aggregate volatility may be caused by shocks to selected, individual sectors. Gabaix (2005), in turn, proves that if the firms' size distribution follows a Zipf's Law, an empirical regularity recently confirmed for the United States by Axtell (2001), an economy composed of N firms will display aggregate fluctuations with size proportional to $1/\ln N$, rather than $1/N^{1/2}$.

At the heart of these lines of reasoning lays the assumption that sectoral shocks' probabilistic processes possess finite moments. In this paper, we dispute the appropriateness of such an assumption. In particular, we show that the probability density function of TFP growth rates for four-digit manufacturing industries plausibly belongs to the class of Lévy-stable distributions. Introduced and characterized back in the 1920s by Paul Lévy (1925) and first sold to economists by Benoit Mandelbrot as the core of the "Stable Paretian hypothesis" for the distribution of personal incomes and commodities' price dynamics (1960, 1963), stable[2] distributions can accommodate a rich set of features such as heavy tails, skewness, leptokurtosis, and infinite second and higher moments. Furthermore, when the associated index of stability α is comprised between 0 and 1, they do not even possess the mean. Far from representing exotic objects, the class of Lévy-stable distributions naturally emerges as the domain of attraction for sums of independently and identically distributed random variables, as soon as the assumption that the variables to be summed possess finite variance is dropped.[3] In such a case, the so-called Generalized Central Limit Theorem applies. The Gaussian distribution is a special case of the Lévy-stable family when $\alpha = 2$, and it is the only one for which the second moment exists. Although the variance of a non-Gaussian Levy-stable distribution is infinite, the sample variance is of course finite for any realization of the process but the probability to observe extreme events—in our case, extreme variations in sectoral TFP growth—is much higher than what we should expect if the data were normally distributed. This causes the characteristic scale of aggregate fluctuations to be of order $N^{1/\alpha}$, with $1<\alpha<2$, that is larger than the Gaussian case $N^{1/2}$.

The remainder of this paper is organized as follows. In Section 2, we motivate our empirical analysis and briefly introduce the class of Lévy-stable distributions. In Section 3, we assess the fit of Lévy-stable distributions as a statistical model for the TFP growth rates of U.S. manufacturing sectors. Section 4 contains some concluding comments.

2. MOTIVATION

To motivate the empirical analysis contained in Section 3, let us start from the growth accounting model proposed by Hulten (1978),[4] who showed that the rate of increase in GDP, g_{GDP}, caused by Hicks-neutral i.i.d. productivity shocks τ to N industries is equal to

$$g_{\text{GDP}} = \sum_{i=1}^{N} \frac{S_i}{Y} \tau_i, \qquad \textbf{(1)}$$

where S_i indicates industry i's final sales, while Y represents GDP. Suppose, as it is usually done in real business cycle (RBC) theory, that sectoral productivity shocks have a common, finite variance σ_τ^2. If, for the sake of simplicity, we let each industry to account for $1/N$ of GDP, from (1) it follows that

$$\sigma_{\text{GDP}} = \frac{\sigma_\tau}{\sqrt{N}}. \quad (2)$$

As disaggregation becomes finer—or, alternatively, as the number of industries N grows large—one should expect the aggregate standard deviation to be vanishingly small. If we take an estimate of industries' volatility on the average at 6%, and consider a economy with 450 industries,[5] we end up with an aggregate volatility of 0.15%, one order of magnitude smaller than the empirically measured one for the United States. An obvious way to overcome such a *curse of aggregation* consists in substituting idiosyncratic shocks with aggregate shocks, a strategy largely employed in the RBC literature in which business cycles stylized facts are usually replicated by means of one-sector, representative agent models.

Alternatively, individual sectoral shocks may be consistent with aggregate fluctuations as soon as some mechanism capable to slow down the rate at which the LLN applies can be found. A natural candidate consists in relaxing the assumption that individual units are of equal size. Gabaix (2005), for instance, builds a model in which the firms' size distribution is Pareto to show that individual firm-level shocks can cause sizable aggregate fluctuations even as one considers an economy with $N = 10^6$ firms. Horvath (1998, 2000), in turn, demonstrates that if some sectors are important input suppliers while some other sectors are not, individual shocks to the former ones are a major source of aggregate volatility.

Admittedly, far less attention—if none at all—has been so far devoted to the distributional properties of idiosyncratic shocks σ_τ. The fundamental and generally undisputed assumption is that the (log)increments (i.e., growth rates) of TFP are identically and independently Gaussian distributed. Indeed, if one considers sectoral TFP changes as the sum of many small and independent variations in the production possibility sets of various profit centers [Hansen and Prescott (1993)], the Central Limit Theorem suggests the Gaussian as a natural candidate.[6]

However, even a cursory look at the data suggests us to be skeptical about such an inference. In Figure 1, we report the time series of TFP growth rates for a randomly selected small sample of four-digit U.S. industries,[7] plus the ± 3 standard deviations interval around the sample mean, which, if the data are Gaussian distributed, contains 99.7% of the probability mass. A remarkable fact emerges neatly: in spite of being composed of a limited number of observations (38), three out of four series display at least one outlier, which, under the Gaussian assumption, would have a negligible probability to occur. In other words, a casual visual inspection of the data indicates that time series with presumably independent increments tend to exhibit abrupt discontinuities or *extreme events*.[8]

Extreme events as those depicted earlier are attributable to the prominent weight exerted by large firms in shaping the dynamics of industries at the four-digit level.

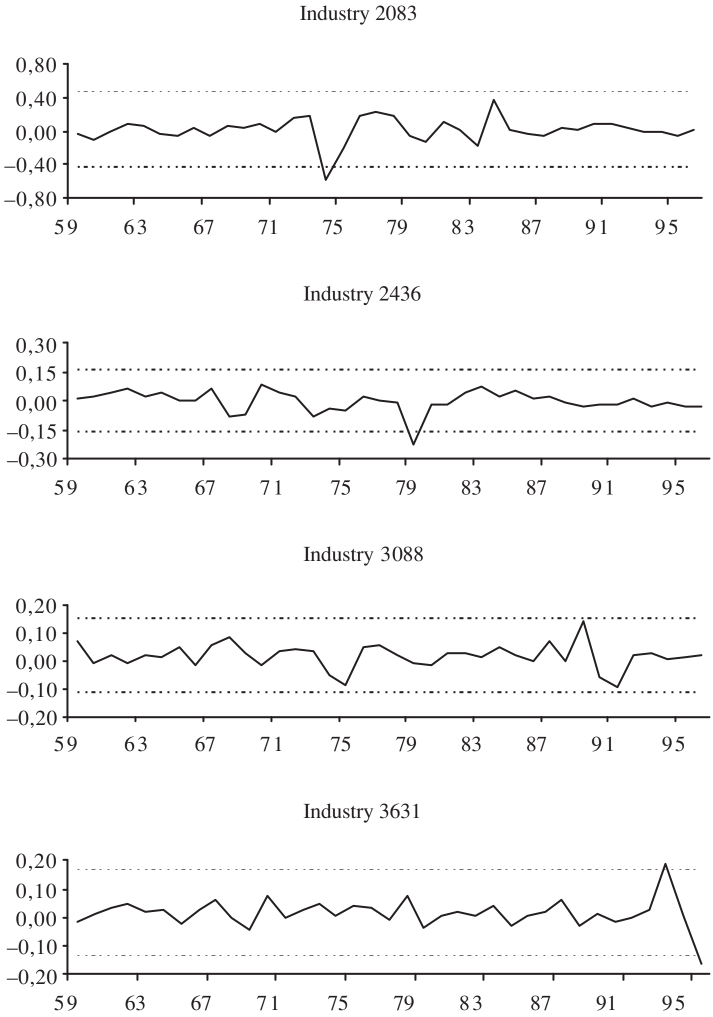

FIGURE 1. TFP growth rates time series for randomly selected four-digit industries. Dotted lines represent the ±3 standard deviations interval around the mean.

As a matter of example, the two extreme episodes experienced between 1994 and 1996 by the Household Cooking Equipment industry (SIC 3631) can be almost entirely ascribed to a major restructuring (with a cut of 2,000 jobs) and a later sudden increase in labor and material costs at Whirpool, which in those years shared about 80% of the U.S. market with other three main competitors (General Electric, AB Electrolux, and Maytag).[9]

Admittedly, the issue of very large productivity movements over short periods of time have been documented and discussed in the literature on several occasions.

Section 4 of Prescott (1998) constitutes an outstanding example. The main points, which seem to have been unnoticed by previous authors, however, are that such a phenomenon suggests that TFP shocks might well be characterized by non-Gaussian heavy-tailed probability density functions, and that this could have remarkable implications for macroeconomic theory.

Heavy tails are quite popular in finance nowadays, where the increasing availability of high-frequency data has allowed scholars to appreciate that, although the microstructure may be very different, some simple statistical properties of price fluctuations in foreign exchange, stock, and futures markets are very general. In particular, returns' distributions in financial markets tend to be characterized by sample leptokurtosis, cluster volatility, and unstable variance [Pagan (1996)]. All of these features, which were well known at least from the early 1960s, are inconsistent with Gaussian distributed processes, and led Benoit Mandelbrot to develop the "Stable Paretian hypothesis" as an alternative [Mandelbrot (1960, 1963)].[10] The core of Mandelbrot's working hypothesis rests on a family of distributions known as Lévy-stable distributions [Lévy (1925)]. Provided that Lévy-stable distributions do not in general possess a closed-form solution for their density, they can be expressed conveniently in terms of their characteristic function. Among the many different available parameterizations, we choose the $S_0(\alpha, \beta, \gamma, \delta)$ parameterization proposed by Nolan (2007),[11] according to which the characteristic function of Y is given by

$$E\exp(itY) = \begin{cases} \exp\left\{-\gamma^\alpha |t|^\alpha \left[1 + i\beta\left(\tan\frac{\pi\alpha}{2}\right)(\text{sign } t)\left((\gamma|t|)^{1-\alpha} - 1\right)\right] + i\delta t\right\} & \text{if } \alpha \neq 1 \\ \exp\left\{-\gamma|t|\left[1 + i\beta\frac{2}{\pi}(\text{sign } t)\left(\ln|t| + \ln\gamma\right)\right] + i\delta t\right\} & \text{if } \alpha = 1 \end{cases}. \quad (3)$$

A major advantage of the functional form (3) is that the four parameters have intuitive interpretations. The characteristic exponent or index of stability α, which has a range $0 < \alpha \leq 2$, measures the probability weight in the upper and lower tails of the distribution. In general, the pth moment of a stable random variable is finite if and only if $p < \alpha$. Thus, for $1 \leq \alpha \leq 2$, a Lévy-stable process possesses a mean equal to the location parameter δ (which in turn indicates the center of the distribution), but it has infinite variance, whereas if $\alpha < 1$, even the mean of the distribution does not exist. β, defined on the interval $-1 \leq \beta \geq 1$, measures the asymmetry of the distribution, with its sign indicating the direction of skewness. Finally, the scale parameter γ, which must be positive, expands or contracts the distribution around the location parameter δ. The Lévy-stable distribution function nests several well-known distributions, such as the Gaussian $N(\mu, \sigma^2)$ (when $\alpha = 2$, $\beta = 0$, $\gamma = \sigma^2/2$ and $\delta = \mu$), the Cauchy ($\alpha = 1$ and $\beta = 0$) and the Lévy-Smirnov ($\alpha = 0.5$ and $\beta = \pm 1$).

Although other statistical models, such as student-t distributions, mixture models, or time-varying variance models, have proved to be successful in capturing important features of the data such as leptokurtosis, skewness, or instable variance, all of them are somehow arbitrary, as they are not limiting distributions. On the contrary, Lévy-stable distributions represent an attractor in the functional space of probability density functions, in that the Generalized Central Limit Theorem [Gnedenko and Kolmogorov (1954)] states that the only possible limiting distribution for sums of independently and identically distributed random variables belongs to the Lévy-stable family. It follows that the conventional Central Limit Theorem is just a special case of this—a special case that applies whenever one imposes the condition that each of the constituent random variables has a finite variance. If for theoretical reasons one wants to appeal to some kind of LLN, Lévy-stable distributions are the most natural candidate. In particular, Lévy-stable distributions are stable under convolution. Simply stated, if we sum N i.i.d. Lévy-distributed variables with characteristic exponent α, the renormalized sum T_N:

$$\mathrm{T}_N = \frac{\sum_{i=1}^{N} \tau_i}{N^{1/\alpha}}, \tag{4}$$

is also Lévy-stable with characteristic exponent α. In addition to other interesting properties,[12] non-Gaussian Lévy-stable distributions (i.e., for $\alpha < 2$) are characterized by tails that are asymptotically Pareto distributed with exponent α.

The paradigm shift toward the Lévy-stable distribution for productivity shocks has far-reaching implications for the "sectoral *vs.* aggregate shocks" debate.[13] To simplify the analysis as much as possible, let us suppose that TFP growth rates τ are Lévy-stable *iid* random variables $S_0(\alpha, 0, \delta, 0)$, with $1 < \alpha < 2$.[14] Making use of (1), if we let each industry to account for $1/N$ of aggregate GDP, the sample variance of GDP growth rates is given by

$$\tilde{\sigma}^2_{\mathrm{GDP}} = \frac{\sum_{i=1}^{N} \tau_i^2}{N^2}. \tag{5}$$

By the property of invariance under convolution (4), we can write $N^{2/\alpha}\tilde{T} = \sum_{i=1}^{N} \tau_i^2$, where $\tilde{T}$ is a Lévy-stable distributed random variable. Thus

$$\tilde{\sigma}_{\mathrm{GDP}} = \frac{\tilde{T}^{\frac{1}{2}}}{N^{\left(\frac{\alpha-1}{\alpha}\right)}}. \tag{6}$$

It follows that if TFP shocks are non-Gaussian Lévy-stable distributed, aggregate fluctuations decays with N at the rate $\frac{\alpha-1}{\alpha}$, that is much more slowly than $N^{-\frac{1}{2}}$ as would be implied by Gaussian distributed shocks.[15]

3. EMPIRICAL ANALYSIS

We conduct our empirical analysis with data obtained from the NBER-CES Manufacturing Productivity (MP) database.[16] The MP contains annual information on 459 manufacturing industries as defined by the 1987 Standard Industrial Classification system. Data, which span from 1959 through 1996, have been deflated at 1987

TABLE 1. Statistical properties of TFP growth rate distribution, full sample

Attribute		Percentile	
Mean	0.0063	1st	−0.1712
Max	0.6178	10th	−0.0583
Min	−0.6215	25th	−0.0243
St. Dev.	0.0635	50th	0.0068
Skewness	−0.0539	75th	0.0369
Kurtosis	13.413	90th	0.0702
Observations	17442	99th	0.1838

prices [Bartelsmann and Grey (1996)]. Thus, the full sample comprises 17,442 data points. The TFP growth rates have been constructed using the Törnqvist productivity growth index based on a five-factor production function:

$$\tau = \dot{y} - \sum_i \omega_i \dot{x}_i, \qquad i = K, N, L, M, E, \tag{7}$$

where τ is TFP growth rate, y is the output quantity, ω_i is the average cost share of input i over the current and previous period, x_i $(i = K, N, L, M, E)$ is the quantity of input i (capital, production worker hours, nonproduction workers, nonenergy materials, and energy, respectively) and a $(\dot{})$ denotes a log first-difference. Table 1 reports attributes and quantiles of the TFP growth rates distribution.

The skewness statistic indicates that the distribution of productivity shocks is, for any practical purpose, symmetric. The kurtosis statistics, on its part, signals that the distribution is significantly leptokurtic. Figure 2, in which the empirical density of TFP growth rates (solid line) is overlaid with the fitted Gaussian density (thick dashed line) shows that the data depart sensibly from normality.[17] The empirically density is sharply peaked, and its tails are much more heavier than predicted by the fitted Gaussian density function. In addition to visual inspection, both parametric and nonparametric statistical tests reject the Gaussian null hypothesis for the TFP growth rates data at the 1% marginal significance level.[18]

As an alternative statistical model, we consider the Lévy-stable distribution. Parameter estimates have been obtained by recurring to three alternative methods:[19] (1) the quantile method proposed by McCulloch (1986); (2) the empirical characteristic function technique (ECF) in the version proposed by Kogon and Williams (1998); and (3) the maximum likelihood method (ML) described in Nolan (1997).[20] Results are reported in Table 2.

If the data are stably distributed, parameter estimates from the quantile method, the empirical characteristic function method, and maximum likelihood should not differ sensibly, given that all the three procedures are consistent [Nolan (1999)]. Indeed, this informal diagnostic seems to support the idea that the TFP growth rates are Lévy-stable distributed. If we take as the most accurate estimate the

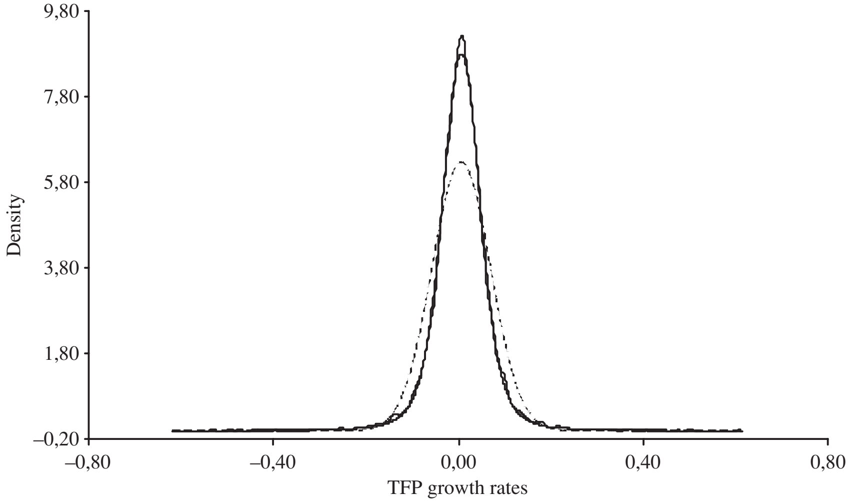

FIGURE 2. Empirical (solid line), fitted Lévy-stable (thin dashed line) and fitted Gaussian (thick dashed line) density functions of TFP growth rates, full sample.

TABLE 2. TFP growth rate distribution parameter estimates, full sample; for maximum likelihood estimates, the 95% confidence bounds are reported

Method	α	β	γ	δ
Quantile	1.5609	−0.0095	0.0316	0.0069
ECF	1.7035	−0.0509	0.0330	0.0069
ML	1.6324 ± 0.0232	−0.0282 ± 0.0563	0.0323 ± 0.0005	0.0068 ± 0.0008

one obtained by maximum likelihood, we find that the characteristic exponent is 1.63, whereas we cannot reject the hypothesis that the skewness parameter is null, signaling that the probability weights on both the upper and the lower tails are higher than the Gaussian case and approximately equal.

This intuition is confirmed by a further look at Figure 2, which clearly shows that the fitted stable density (thin dashed line)—contrary to the Gaussian one—does a good job in capturing the sharp peakedness of TFP data. Furthermore, one can appreciate from Figure 3, which offers an expanded view of both tails, that the non-Gaussian Lévy-stable distribution turns out to be a far better model for extreme outcomes as well. In fact, the Gaussian density does not assign any probability mass to positive and negative productivity shocks in excess of ±20%.[21]

Given that a four-parameter model as the Lévy-stable provides a better fit than a two-parameter model as the Gaussian almost by definition, the next step consists in assessing whether the Lévy-stable can be considered a good description of the data per se. To accomplish this task, we return to several exploratory data analysis

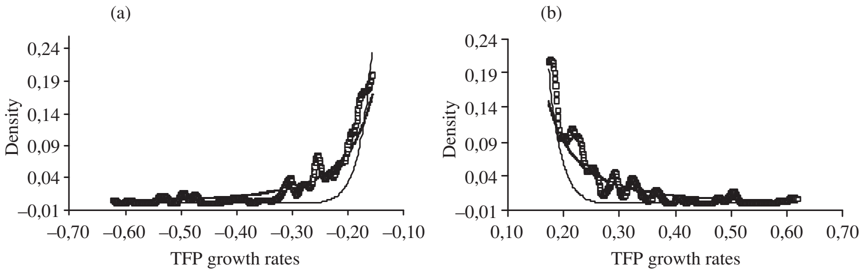

FIGURE 3. Expanded lower (panel a) and upper (panel b) tails of the empirical (squares), fitted Lévy-stable (dashed) and fitted Gaussian (solid) density functions, as shown in Figure 2.

and regression techniques. All of the results presented in what follows refer to maximum likelihood estimation.

The first diagnostic exercise consists in plotting the density functions for real data and the Lévy-stable and Gaussian models on semilogarithmic paper. Such a transformation, by accentuating larger observations, allows us to explore more deeply the behavior on the tails. The bulk of the divergent variance argument according to which aggregate volatility is proportional to $N^{-\left(\frac{\alpha-1}{\alpha}\right)}$ rests on the tails of the empirical distribution being asymptotically Paretian. A good fit of empirical data to a non-Gaussian Lévy-stable distribution in the central regime could well be consistent with finite variance as soon as observations higher than a threshold $|l|$ decay faster than a power law [Mantegna and Stanley (1994)]. As shown in Figure 4,

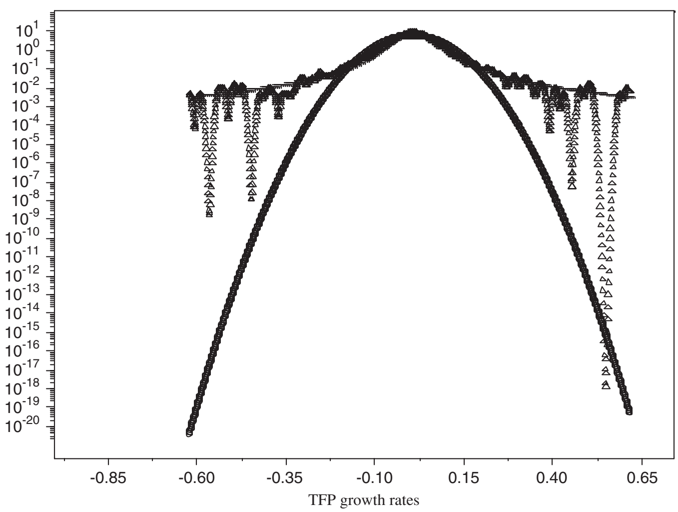

FIGURE 4. Semilog plot of the empirical (triangles), fitted Lévy-stable (thin line), and fitted Gaussian (thick line) density functions of TFP growth rates, full sample.

https://doi.org/10.1017/S1365100508070223 Published online by Cambridge University Press

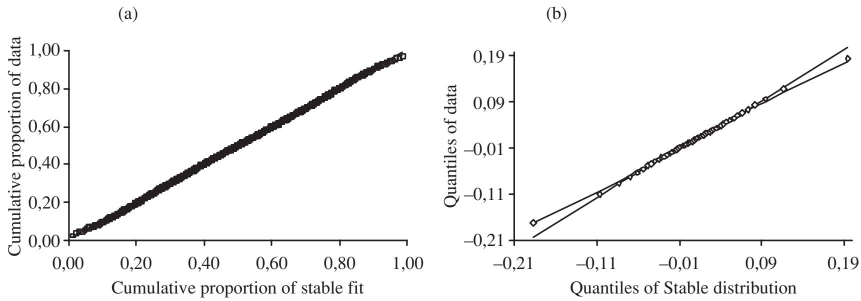

FIGURE 5. Variance-stabilized p-p plot (panel a) and thinned q-q plot (panel b) for Lévy-stable fitted TFP growth rates, full sample. In panel b diamonds are quantiles, whereas the two curves define the 95% confidence interval.

regardless of large fluctuations due to sparseness of data for large (in absolute value) TFP shocks, both extreme tails behave in good accordance with the Paretian hypothesis, thus suggesting the appropriateness of the infinite variance stable model.

Second, we follow Nolan (1999) in assessing whether the data are consistent with the hypothesis of stability by means of diagrammatic diagnostic tests. The variance-stabilized probability plot presented in panel a of Figure 5, in which a transformation is applied to make the variance in the p-p plot uniform,[22] signals that the stable fit is extremely good even for extreme values. Furthermore, panel b of Figure 5 presents a thinned q-q plot, where a point at every 4.5% of the data is given (diamonds),[23] which shows that TFP growth rates are consistent with the random variation of a Lévy-stable distribution and fall within the corresponding 95% confidence bands (lines). All in all, both diagnostics strongly support the hypothesis that productivity shocks at the four-digit level for the U.S. economy are Lévy-stable distributed.

As a motivation for a third diagnostic exercise, recall that by definition a random vector $\tau = (\tau_1, \ldots, \tau_N)$ is said to be Lévy-stable if all components τ_i are α-stable with a common α. The testing idea consists in partitioning the full sample into 20 subsamples corresponding to two-digit industries, in order to have a sufficiently high number of degrees of freedom. Estimates of the index of stability α for the full sample and for all subsamples, or at least for a significant number of them, should then be consistent. Both the Bera-Jarque and the Kolmogorov-Smirnov tests reject Normality for each two-digit industries group at the 1% marginal significance level except for Petroleum and Coal Products, for which the null hypothesis of normality is rejected at the 5% significance level. Lévy-stable estimates obtained with the quantile, the empirical characteristic function, and the maximum likelihood methods are listed in Table 3. In each case, the stable hypothesis seems to be confirmed by the value assumed by estimated parameters. Furthermore, the goodness of fit of the Lévy-stable model is confirmed by the q-q plots shown in Figure 6, from which it clearly emerges that the two-digit TFP growth rates data always fall within the

TABLE 3. TFP growth rates distribution parameter estimates, two-digit groups

Method	α	β	γ	δ
Industry 20				
Quantile	1.3689	−0.1168	0.0329	0.0069
ECF	1.5219	−0.0284	0.0347	0.0085
ML	1.4208 ± 0.0713	−0.0869 ± 0.1226	0.0333 ± 0.0017	0.0057 ± 0.0026
Industry 21				
Quantile	1.4555	−0.0707	0.0283	−0.0021
ECF	1.6655	−0.3739	0.0307	−0.0026
ML	1.5517 ± 0.2520	−0.0353 ± 0.5235	0.0238 ± 0.0050	−0.0034 ± 0.0083
Industry 22				
Quantile	1.5656	−0.1540	0.0293	0.0113
ECF	1.8095	−0.4415	0.0313	0.0108
ML	1.7160 ± 0.0987	−0.2338 ± 0.2888	0.0307 ± 0.0019	0.0113 ± 0.0036
Industry 23				
Quantile	1.5448	−0.0151	0.0323	0.0049
ECF	1.7676	−0.2273	0.0340	0.0038
ML	1.6548 ± 0.0885	−0.0226 ± 0.2266	0.0329 ± 0.0018	0.0058 ± 0.0033
Industry 24				
Quantile	1.5675	0.0761	0.0315	0.0027
ECF	1.7693	0.1878	0.0330	00022
ML	1.6650 ± 0.1190	0.0412 ± 0.3109	0.0322 ± 0.0024	0.0038 ± 0.0044
Industry 25				
Quantile	1.6696	−0.2020	0.0272	0.0032
ECF	1.8371	−0.5107	0.0279	0.0022
ML	1.8547 ± 0.1101	−0.0647 ± 0.6593	0.0281 ± 0.0021	000227 ± 0.0042
Industry 26				
Quantile	1.6535	−0.0619	0.0278	0.0044
ECF	1.6925	−0.3499	0.0275	0.0042
ML	1.6217 ± 0.1203	−0.1594 ± 0.2785	0.0267 ± 0.0021	0.0039 ± 0.0036
Industry 27				
Quantile	1.5939	−0.0036	0.0241	0.0014
ECF	1.8406	0.2039	0.0262	0.0025
ML	1.7702 ± 0.1214	0.0914 ± 0.4422	0.0259 ± 0.0020	0.0005 ± 0.0038
Industry 28				
Quantile	1.5922	−0.1312	0.0372	0.0133
ECF	1.7621	−0.1305	0.0390	0.0122
ML	1.6615 ± 0.0909	−0.1297 ± 0.2323	0.0380 ± 0.0022	0.0144 ± 0.0039
Industry 29				
Quantile	1.6149	−0.3042	0.0326	0.0094
ECF	1.7804	−1.0000	0.0345	0.0089
ML	1.6627 ± 0.2161	−0.3928 ± 0.5316	0.0334 ± 0.0046	0.0078 ± 0.0084
Industry 30				
Quantile	1.4814	−0.5087	0.0206	0.0205
ECF	1.8428	−1.0000	0.0233	0.0228
ML	1.7248 ± 0.1158	−0.6518 ± 0.2974	0.0228 ± 0.0017	0.0167 ± 0.0033

TABLE 3. (Continued.)

Method	α	β	γ	δ
Industry 31				
Quantile	1.4534	−0.0703	0.0348	0.0002
ECF	1.7092	−0.4116	0.0383	−0.0017
ML	1.5601 ± 0.1510	−0.1710 ± 0.1510	0.0366 ± 0.0036	0.0016 ± 0.0062
Industry 32				
Quantile	1.5830	−0.0627	0.0326	0.0084
ECF	1.7531	−0.0217	0.0342	0.0079
ML	1.7003 ± 0.0943	−0.1819 ± 0.2612	0.039 ± 0.0020	0.0096 ± 0.0037
Industry 33				
Quantile	1.6534	−0.1389	0.0369	0.0060
ECF	1.8060	−0.1600	0.0385	0.0058
ML	1.7279 ± 0.0922	−0.1673 ± 0.2850	0.0378 ± 0.0022	0.0066 ± 0.0041
Industry 34				
Quantile	1.6532	0.1550	0.0304	0.0025
ECF	1.7119	0.1742	0.0308	0.0018
ML	1.6793 ± 0.0792	0.0397 ± 0.2136	0.0305 ± 0.0015	0.0027 ± 0.0028
Industry 35				
Quantile	1.4936	0.2254	0.0315	0.0030
ECF	1.6454	0.2217	0.0332	0.0035
ML	1.6108 ± 0.0692	0.3144 ± 0.1519	0.0330 ± 0.0015	0.0021 ± 0.0026
Industry 36				
Quantile	1.6905	−0.1613	0.0337	0.0162
ECF	1.7283	−0.1256	0.0337	0.0174
ML	1.6667 ± 0.0807	−0.0217 ± 0.2122	0.0331 ± 0.0017	0.0144 ± 0.0030
Industry 37				
Quantile	1.5310	−0.0319	0.0309	0.0065
ECF	1.6906	0.0562	0.0330	0.0066
ML	1.6574 ± 0.1157	−0.0893 ± 0.2950	0.0325 ± 0.0024	0.0073 ± 0.0043
Industry 38				
Quantile	1.8398	0.4933	0.0341	0.0018
ECF	1.8561	0.1073	0.0348	−0.0011
ML	1.7858 ± 0.1087	0.0916 ± 0.4174	0.0343 ± 0.0024	0.0048 ± 0.0046
Industry 39				
Quantile	1.4864	−0.0555	0.0297	0.0106
ECF	1.6477	−0.1247	0.0309	0.0114
ML	1.5690 ± 0.1184	−0.0830 ± 0.2508	0.0301 ± 0.0023	0.0099 ± 0.0040

95% confidence bands associated with maximum likelihood. Interestingly enough, the maximum likelihood estimate of the characteristic exponent for the full sample ($\alpha = 1.63$) is fully consistent—in that such a value is inside the 95% confidence interval—with the corresponding estimates for 17 out of 20 subsamples.

It must be pointed out that very large observations on the tails of the unconditional pooled distribution may be due to large aggregate disturbances hitting

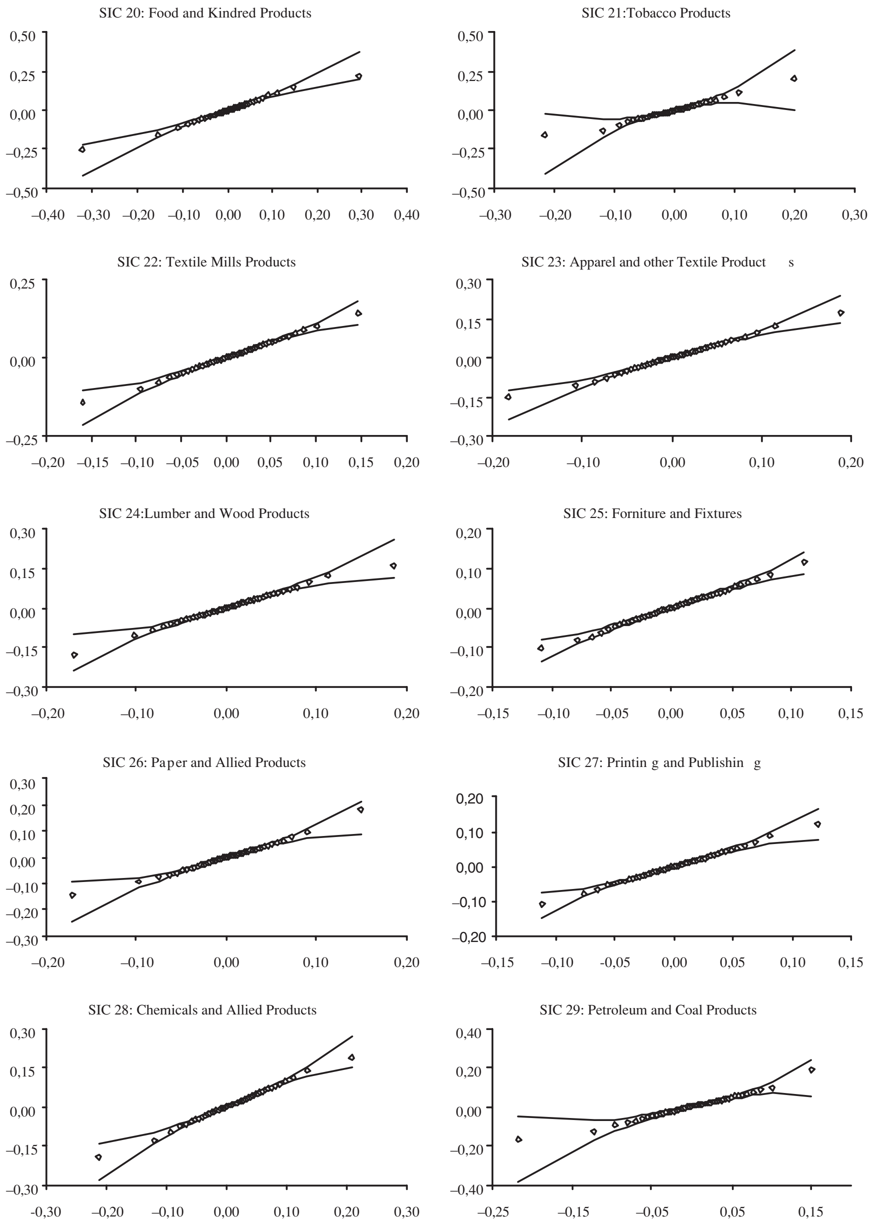

FIGURE 6. q-q plots for Lévy-stable fitted TFP growth rates, two-digit groups.

sectors in particular years. The severe oil shocks of the 1970s, for instance, had likely caused large negative TFP movements for almost all industries. Instead of being a signature of stable TFP shocks at a sectoral level, the right and left tails in

https://doi.org/10.1017/S1365100508070223 Published online by Cambridge University Press

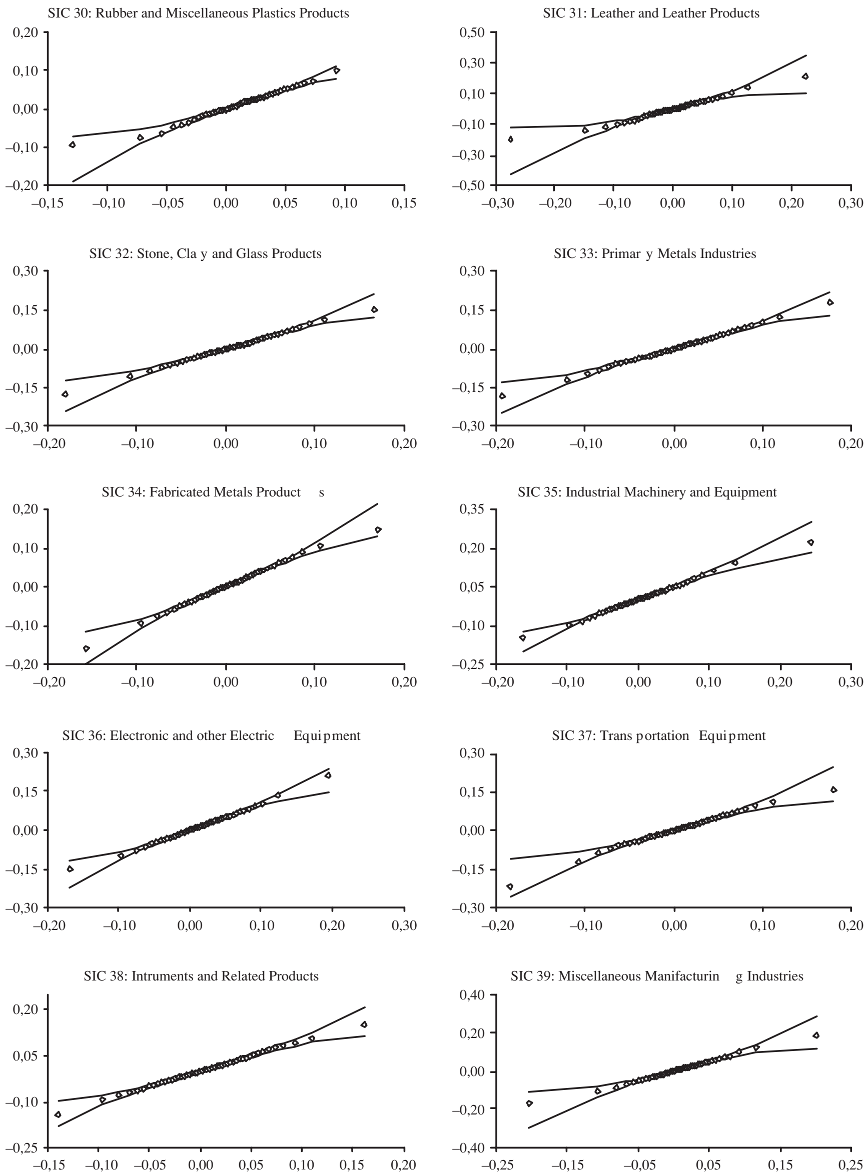

FIGURE 6. Continued.

the distribution reported in Figure 2 could therefore be the outcome of aggregate shocks. Thus, as a final exercise, we control for possible effects of common shocks.

https://doi.org/10.1017/S1365100508070223 Published online by Cambridge University Press

The standard approach entails the use of common factor models, according to which TFP growth rates corresponding to sector i, τ_{it}, can be decomposed as

$$\tau_{it} = \lambda_i' f_t + u_{it}, \quad t = 1, \dots, T;\ i = 1, \dots, N \tag{8}$$

where $f_t = [f_{1t}, \dots, f_{rt}]'$ is a vector of r common factors, while $\lambda_t = [\lambda_{it}, \dots, \lambda_{ir}]'$ is a vector of i's responses to common components. Finally, the vector $u_t = [u_{1t}, \dots, u_{Nt}]'$ comprises N idiosyncratic components. The decomposition (8) between common shocks and idiosyncratic disturbances is acceptable as soon as $E(f_t u_t') = 0$. The null hypothesis we want to test is that u_t is non-Gaussian Lévy-stable distributed.

Unfortunately, in our case, the standard common factors approach is not useful, as the principal component estimator of the loading matrix $\Lambda = [\lambda_{1\cdot}, \dots, \lambda_{N\cdot}]'$ requires the matrix $\Sigma = E(u_t u_t')$ to be finite. This amounts to impose a restriction on the stochastic process for u_t, which hopelessly prevents us from testing our null hypothesis.

As an alternative line of attack, we make use of the maximum likelihood estimator for linear regressions with stable disturbances developed by McColloch (1998). Such an ML estimator, in particular, allows jointly estimates of regression coefficients and of the index of stability characterizing the errors distribution, under the assumption that the random disturbances follow a symmetric stable distribution with median zero, $S_0(\alpha, 0, \gamma, 0)$.

The empirical strategy we choose is simply to run, for each year from 1959 to 1996, a cross-sectional equation of the type

$$\tau_i = f_1 + u_i, \quad i = 1, \dots, 459, \tag{9}$$

under the assumption that in each year a unique common shock hits all sectors simultaneously. The idiosyncratic disturbances we get from each linear regression can be immediately interpreted as the annual distribution of sectoral shocks netted from common influences. If the estimated characteristic exponent of the errors distribution turns out to be significantly lower than 2, we cannot reject the null hypothesis that for that year the idiosyncratic component of TFP sectoral shocks is non-Gaussian stable. If, on the contrary, α is equal to 2, we should reject the null in favor of the alternative of Gaussian idiosyncratic errors.

Table 4 reports the estimated index of stability of the errors distribution obtained from the 38 yearly cross-section regressions. The characteristic exponents are in all cases significantly lower than 2. These results confirm that, even after having controlled for aggregate shocks, we cannot reject the hypothesis that TFP growth rates follow a non-Gaussian stable distribution with infinite variance.

TABLE 4. Characteristic exponents of the errors distribution from cross-section linear regressions with stable disturbances

	α (s.e.)		α (s.e.)
1959	1.4967 (0.0715)	1978	1.5308 (0.0747)
1960	1.4935 (0.0717)	1979	1.5136 (0.0748)
1961	1.5686 (0.0739)	1980	1.4947 (0.0738)
1962	1.6281 (0.0800)	1981	1.5538 (0.0700)
1963	1.4500 (0.0704)	1982	1.5679 (0.0693)
1964	1.5908 (0.0735)	1983	1.4483 (0.0719)
1965	1.7071 (0.0827)	1984	1.5629 (0.0770)
1966	1.6389 (0.0757)	1985	1.5278 (0.0711)
1967	1.8781 (0.0666)	1986	1.5027 (0.0709)
1968	1.9031 (0.0594)	1987	1.5022 (0.0727)
1969	1.6302 (0.0742)	1988	1.6202 (0.0750)
1970	1.7838 (0.0701)	1989	1.5373 (0.0700)
1971	1.7503 (0.0713)	1990	1.6273 (0.0729)
1972	1.4675 (0.0768)	1991	1.5569 (0.0716)
1973	1.5068 (0.0795)	1992	1.6260 (0.0724)
1974	1.5160 (0.0746)	1993	1.5823 (0.0742)
1975	1.5342 (0.0772)	1994	1.5327 (0.0691)
1976	1.4383 (0.0717)	1995	1.5707 (0.0795)
1977	1.5102 (0.0748)	1996	1.5799 (0.0743)

4. CONCLUSIONS

The extension from one-sector to multisector models represents a challenge for the theory of business cycles. If shocks to economic activity are idiosyncratic, that is, at a firm or at a sectoral level, the Law of Large Numbers implies that aggregate volatility should be negligible, as positive shocks tend to be offset by negative shocks. The standard assumptions are that idiosyncratic shocks are given by changes of the Solow residual or, put differently, of the total factor productivity, and that such changes can be modeled as Gaussian increments. In other words, the growth rates of TFP are random variables with finite mean and variance.

In this paper, we dispute this last assumption, suggesting a statistical model for TFP growth rates alternative to the Gaussian. In particular, our proposal consists in modeling TFP shocks as Lévy-distributed random variables, which are characterized by a divergent variance. This assumption can find a motivation in that the family of Lévy-stable distributions is an attractor in the functional space of probability density functions for sums of i.i.d. random variables, so that one can appeal to the generalized version of the Central Limit Theorem for its application.

Lévy-stable estimations, obtained by means of the quantile, the empirical characteristic function, and the maximum likelihood techniques, respectively, return a good fit, both for the whole sample as derived from the NBER-CES Manufacturing

Productivity database for industries at the four-digit level, and for groups at the two-digit levels. Our results also seem to be robust after controlling for possible cross-sectional dependence due to aggregate shocks.

NOTES

1. In addition to the ones discussed in the main text, other mechanisms that may slow down convergence implied by the LLN include the following: (1) local interactions and nonconvexities leading to self-organized criticality [Scheinkman and Woodford (1994)]; (2) imperfect competition and strategic complementarities [Murphy et al. (1989)]; and (3) large multiplier effects [Jovanovic (1987)].

2. These distributions are called stable because they are invariant under convolution. In other words, sums of stable distributed random variables are stable distributed.

3. See, for example, Samorodnitsky and Taqqu (1994).

4. See also Gabaix (2005).

5. As we will see later, these figures are consistent with real data for four-digit U.S. manufacturing industries.

6. For an introductory exposition of the degree of heterogeneity across establishments in the distribution of output and productivity growth rates in the U.S. economy, see Haltiwanger (1997).

7. Data are annual, and cover the period 1959–1996.

8. Such a property is also known as the "Noah effect" [Mandelbrot (1969)].

9. See G. Steinmetz and C. Quintanilla, "Tough target: Whirpool expected easy going in Europe, and it got a big shock," *Wall Street Journal*, April 10, 1998; and I. Katz, "Whirlpool: In the wringer," *Business Week*, December 14, 1998.

10. Mirowski (2004) provides an illuminating account of how the economics community reacted to the pioneering work of Mandelbrot.

11. The parameterization of the characteristic function S_0 is particularly convenient because the density and the distribution functions are jointly continuous in all four parameters.

12. A comprehensive reference is Samorodnitsky and Taqqu (1994).

13. Furthermore, fat-tailed distributed productivity shocks can explain the equity premium puzzle, as shown in recent papers by Barro (2005) and Weitzman (2005).

14. Recall that when $\alpha < 1$, the population mean of the process does not exists. In this case, it is trivial to show that averaging would result in amplification of volatility.

15. For a discussion on this point, see Sornette (2000, pp. 90–93).

16. Available at http://www.nber.org/nberces/nbprod96.htm.

17. The empirical density plot has been obtained by smoothing the data with a Gaussian kernel with the width parameter equal to 2(inter-quartile range) $N^{-1/3}$. We also tried the Epanechnikov kernel smoothing, but the empirical density appears to be virtually insensitive to the particular kernel used.

18. In particular, we performed both a Bera-Jarque test and a Lilliefors version of the Kolmogorov-Smirnov test.

19. These three estimation methods are described and compared in, for example, Weron (2004).

20. Parameter estimation and computation of theoretical density functions were performed using the software package STABLE by John Nolan (1997).

21. Although it is generally accepted that inventions and discoveries may cause positive shifts in TFP of the order of magnitude at hand, the possibility of such extreme decreases in TFP due to technological regressions are generally viewed with suspicion, if not aversion [Zarnowitz (1992)]. In fact, a decline in TFP measures not only negative technology shocks but also unfavorable regulations and relative price movements. Another possible reason for large TFP declines has been recently advanced by Jovanovic (2006), according to whom the introduction of a new technology may imply unexpectedly large training costs, which force measured TFP to decline sensibly.

22. This transformation, originally suggested in Michael (1983), is aimed at detecting with higher precision a lack of fit near the extremes of the distribution, given that standard p-p plots emphasize behavior around the mode and squeeze the curve near the tails of the distribution.

23. The practice of reducing the number of points in q-q plots (hence, thinned q-q plots) is recommended by Nolan (1999) for large data sets.

REFERENCES

Axtell, Robert (2001) Zipf distribution in US firm sizes. *Science* 293, 181–182.

Barro, Robert (2005) Rare Disasters and Asset Markets in the Twentieth Century. Mimeo, Harvard University.

Bartelsman, Eric and Winn Grey (1996) The NBER Manufacturing Productivity Database. NBER Technical working paper no. 205.

Dupor, Bill (1999) Aggregation and irrelevance in multi-sector models. *Journal of Monetary Economics* 43, 391–409.

Gabaix, Xavier (2005) The Granular Origins of Aggregate Fluctuations. Mimeo, MIT.

Gnedenko, Boris and Andrey Kolmogorov (1954) *Limit Distributions for Sums of Independent Random Variables*. Reading, MA: Addison-Wesley.

Haltiwanger, John (1997) Measuring and analyzing aggregate fluctuations: the importance of building from microeconomic evidence. *Federal Reserve Bank of St. Louis Economic Review* 79, 55–77.

Hansen, Gary and Edward Prescott (1993) Did technology shocks cause the 1990–1991 recession? *American Economic Review* 83, 280–286.

Horvath, Michael (1998) Cyclicality and sectoral linkages: aggregate fluctuations from independent sectoral shocks. *Review of Economic Dynamics* 1, 781:808.

Horvath, Michael (2000), Sectoral shocks and aggregate fluctuations. *Journal of Monetary Economics* 45, 69–106.

Hulten, Christopher (1978) Growth accounting with intermediary inputs. *Review of Economic Studies* 45, 511–518.

Jovanovic, Boyan (1987) Micro shocks and aggregate risk. *Quarterly Journal of Economics* 102, 395–409.

Jovanovic, Boyan (2006) Asymmetric cycles. *Review of Economic Studies* 73, 145–162.

Kogon, Stephen and Douglas Williams (1998) Characteristic function based estimation of stable parameters. In Robert Adler, Raisa Feldman and Murad Taqqu (eds.), *A Practical Guide to Heavy Tails*, pp. 311–335. Berlin: Birkhäuser.

Lévy, Paul (1925) *Calcul de Probabilités*. Paris: Gauthier Villars.

Long, John and Charles Plosser (1983) Real business cycles. *Journal of Political Economy* 91, 39–69.

Lucas, Robert (1981) *Studies in Business Cycle Theory*. Cambridge, MA: MIT Press.

Mandelbrot, Benoit (1960) The Pareto-Lévy law and the distribution of income. *International Economic Review* 1, 79–106.

Mandelbrot, Benoit (1963) The variation of certain speculative prices. *Journal of Business* 36, 394–419.

Mandelbrot, Benoit (1969) Long-run linearity, locally Gaussian processes, H-spectra, and infinite variances. *International Economic Review* 10, 82–111.

Mantegna, Rosario and Eugene Stanley (1994) Stochastic process with ultraslow convergence to Gaussian: The truncated Lévy flight. *Physical Review Letters* 73, 2946–2949.

McCulloch, John (1986) Simple consistent estimators of stable distribution parameters. *Communications in Statistics: Simulations* 15, 1109–1136.

McCulloch, John (1998) Linear regression with stable disturbances. In Robert Adler, Raisa Feldman, and Murad Taqqu (eds.), *A Practical Guide to Heavy Tails*, pp. 489–499. Berlin: Birkhäuser.

Michael, John (1983) The stabilized probability plot. *Biometrika* 70, 11–17.

Mirowski, Philip (2004) Mandelbrot's economics after a quarter-century. In Philip Mirowsky (ed.), *The Effortless Economy of Science*?, pp. 251–270. Durham, NC: Duke University Press.

Murphy, Kevin, Andrei Shleifer, and Robert Vishny (1989) Building blocks of market clearing business cycle models. In Olivier Blanchard and Stanley Fischer (eds.), *NBER Macroeconomics Annual 1989*, pp. 247–301. Cambridge, MA: MIT Press.

Nolan, John (1997) Numerical computation of stable densities and distribution functions. *Communications in Statistics: Stochastic Models* 13, 759–774.

Nolan, John (1999) Fitting data and assessing goodness of fit with stable distributions. In *Proceedings of the Conference on Applications of Heavy Tailed Distributions in Economics, Engineering and Statistics*. *CD-Rom version, Tails42*. Washington, DC: American University.

Nolan, John (2007) *Stable Distributions*. Boston: Birkhäuser.

Pagan, Adrian (1996) The econometrics of financial markets. *Journal of Empirical Finance* 3, 15–102.

Prescott, Edward (1998) Needed: A theory of total factor productivity. *International Economic Review* 39, 525–552.

Samorodnitsky, Alex and Murad Taqqu (1994) *Stable Non-Gaussian Random Processes*. New York: Chapman & Hall.

Scheinkman, José and Michael Woodford (1994) Self-organized criticality and economic fluctuations. *American Economic Review* 84, 417–421.

Sornette, Didier (2000) *Critical Phenomena in Natural Sciences*. Berlin: Springer.

Weitzman, Martin (2005), Risk, Uncertainty, and Asset-Pricing "Puzzles." Mimeo, Harvard University.

Weron, Rafal (2004) Computationally intensive value at risk calculations. In James Gentle, Wolfgang Härdle, and Yuichi Mori (eds.), *Handbook of Computational Statistics*, pp. 911–950. Berlin: Springer.

Zarnowitz, Victor (1992) *Business Cycles: Theory, History, Indicators and Forecasting*. Chicago: University of Chicago Press.