INSTITUTE
OF ECONOMICS

Scuola Superiore
Sant'Anna

LEM | Laboratory of Economics and Management

Institute of Economics
Scuola Superiore Sant'Anna

Piazza Martiri della Libertà, 33 - 56127 Pisa, Italy
ph. +39 050 88.33.43
institute.economics@sssup.it

LEM

WORKING PAPER SERIES

Zipf Law and the Firm Size Distribution: a critical discussion of popular estimators

Giulio Bottazzi[§]
Davide Pirino[§°]
Federico Tamagni[§]

[§] Institute of Economics and LEM, Scuola Superiore Sant'Anna, Pisa, Italy
[°] Scuola Normale Superiore, Pisa, Italy

2013/17 September 2014
ISSN (online) 2284-0400

Zipf Law and the Firm Size Distribution: a critical discussion of popular estimators*

Giulio Bottazzi[†∘], Davide Pirino[‡], and Federico Tamagni[∘]

[∘]IE and LEM, Scuola Superiore Sant'Anna, Pisa, Italy

[‡]Scuola Normale Superiore, Pisa, Italy

September 15, 2014

Abstract

The upper tail of the firm size distribution is often assumed to follow a Power Law behavior. Several recent papers, using different estimators and different data sets, conclude that the Zipf Law, in particular, provides a good fit, implying that the fraction of firms whose size is above a given value is inversely proportional to the value itself. In this article we compare the asymptotic and small sample properties of different methods through which this conclusion has been reached. We find that the family of estimators most widely adopted, based on an OLS regression, is in fact unreliable and basically useless for appropriate inference. This finding rises some doubts about previously identified Zipf Laws. Based on extensive numerical analysis we recommend the adoption of the Hill estimator over any other method when individual observations are available.

JEL codes: L11; C15; C46; D20

Keywords: Firm size distribution; Hill estimator; Power-like distribution; Tail estimators; Zipf Law.

---

*This work was partially supported by the Italian Ministry of University and Research, grant PRIN 2009 "The growth of firms and countries: distributional properties and economic determinants", prot. 2009H8WPX5.

[†] *Corresponding Author*: Giulio Bottazzi, Scuola Superiore Sant'Anna, *E-mail*: bottazzi@sssup.it.

1 Introduction

While existing models of firm growth and firm dynamics predict a wide range of distribution functions to describe the size distribution of firms (see de Wit, 2005, for a review), recent empirical studies, following the influential contribution by Axtell (2001), come to the conclusion that the tail behavior of the firm size distribution is well approximated by a Zipf Law.[1] The result has rapidly become a stylized fact, which more recent models of firm-industry dynamics explicitly aim to agree with (see e.g. Luttmer, 2007).

The common starting point in the investigation of the Zipf Law hypothesis is to assume that the distribution of firm size S is well described by a Power Law or Pareto-type distribution, at least above a certain minimum size threshold S_0. This implies that the fraction of firms whose size is above S_0 is proportional to $S^{-\alpha}$, with α a positive constant. Within this "Power Law approach" to firm size distribution, the effort is devoted to obtain an estimate of α from real data, and compare it with the benchmark value $\alpha = 1$, corresponding to the Zipf Law. As we review in details in Section 3, different methods are used in the literature to infer about the value of α: the Ordinary Least Squares (OLS) log-log methods based on Axtell (2001), the closely related Rank$-1/2$ estimator introduced in Gabaix and Ibragimov (2011), and the classical maximum likelihood estimator for tail behavior based on Hill (1975).[2] The statistical significance of the agreement between estimated α and the Zipf Law is based on different strategies. Some studies merely present a crude graphical comparison of the empirical data with a theoretical Power Law with $\alpha = 1$. In other cases, studies resort to the joint consideration of point estimates and standard errors of the estimated Power Law coefficient, or to the related t-test for the null $H_o : \alpha = 1$.[3]

As we shall explore below in the paper, however, not all the estimators adopted in the literature have sound properties and inference based on them can be problematic, even when a more formal assessment via t-test is provided. This is particularly true when relatively small samples are considered. Indeed, we know that an unbiased version of the Hill estimator is available, and that the Rank$-1/2$ estimator is less biased than the other OLS estimators. Yet, the lack of knowledge about the small sample correction to point estimates and the

[1] This "Law" was originally proposed by Zipf to explain the frequencies of words in a given language, see Zipf (1932). The studies we refer to are discussed in Section 3.

[2] Of course many other estimators of Power Law behavior exist outside the literature on firm size, see Newman (2005) and Gabaix (2009) for reviews.

[3] Notice however that the common practice is to not explicitly report the t-test. An exception is in di Giovanni and Levchenko (2010).

non-normality of the estimators make usual tests unreliable and it is unclear a-priori which method could perform better in practice. Since it is not at all uncommon for applied research on firm size to have data only on a limited number of observations in the tail, a small sample characterization of the estimators certainly represents a relevant issue.

This paper re-examines and criticizes the established consensus on the Zipf Law behavior of firm size distributions by highlighting the pitfalls that arise when the wrong estimators are used as the basis for formal hypothesis testing. First, we perform a Monte Carlo study of the small sample bias and variance of the different estimators in samples generated under exact Zipf distribution. Second, since t-test-like procedures represent the common basis for the assessment of Zipf Law in the literature, we explore the behavior of the t-statistics associated to each estimator.[4] This step involves three sets of Monte Carlo exercises. First, we consider data generating processes with $\alpha \neq 1$, and compare the power to reject the null $\alpha = 1$ under the different estimators. Second, we explore the sensitivity of the point estimates and the ensuing rejection rates when sub-asymptotic corrections to the an exact Zipf behaviour are introduced in the data generating process. Third, and finally, we consider the case in which generated data are not independent and identically distributed, but rather exhibit dependence over time, a case which is likely to arise in practical applications when different years of data on firm sizes are pooled together. In agreement with the vast literature on Gibrat's law in firm size dynamics, we simulate panel samples with an autoregressive structure and Laplace growth shocks, and check the distortion this generates in the upper tail behaviour of the pooled distribution.

2 Popular Methods

The methods employed in the "Power Law approach" to firm size distribution are essentially of two types. The more natural estimator, introduced in Hill (1975), applies maximum likelihood to the estimation of extreme events, based on the theory of order statistics. Assume that the distribution of firm size S follows a Power Law

[4] In principle, there exist alternative ways to test the validity of the Zipf Law. For instance, one can compare goodness-of-fit measures in the upper tail of the estimated distribution or rely upon information criteria based on likelihood ratios for nested models.

$$F(x) = \text{Prob}\{S \leq x\} = 1 - \left(b\frac{x}{x_0}\right)^{-\frac{1}{\gamma}} , \tag{2.1}$$

where $\gamma > 0$ governs the tail behaviour, $b > 0$ is a scale parameter, and $x_0 > 0$ is the minimum threshold or, alternatively, the value above which $F(x)$ holds.[5] Let $s_{(1)} \geq s_{(2)} \geq \ldots \geq s_{(K)}$ denote the K-th largest observations of a sample of size N. The Hill point estimates based on these ordered statistics read

$$\hat{\gamma} = \frac{1}{K-1}\sum_{j=1}^{K} \ln s_{(j)} - \frac{K}{K-1} \ln s_{(K)} \tag{2.2}$$

and

$$\hat{b} = \left(\frac{N}{K}\right)^{\hat{\gamma}} \frac{1}{S_{(K)}} , \tag{2.3}$$

where the expression in (2.2) already includes a correction for small sample bias. The estimator $\hat{\gamma}$ enjoys the desirable properties of any ML estimator when the distribution to be estimated is smooth: it is asymptotically Normal and efficient.[6] Specifically, it holds that

$$E[\hat{\gamma}] = \gamma \quad \text{and} \quad V[\hat{\gamma}] = \frac{1}{K-1}\gamma^2 \quad . \tag{2.4}$$

The alternative approach is based on a class of estimators which apply OLS regressions to different log-log transformation of the data. These methods are, so to speak, more heuristic. Their popularity, even outside firm size studies, is due to simplicity of application, although they can be strongly biased in small sample.

Consider the survival function associated with (2.1) which, with a convenient re-parametrization, can be written as

[5] This parametrization is labeled as Pareto type-I in Kleiber and Kotz (2003) and goes back to the classical Pareto (1886) study of income inequality. See also Johnson et al. (1994) for a discussion.

[6] A huge literature studies the asymptotic and small sample behavior of the original Hill statistic under departures from the assumption of Power Law distributed data. The common approach is to focus on the case where the underlying distribution obeys conditions defining max-stable laws. Along these lines, weak consistency was proved in Mason (1982) under the condition that, as $N \to \infty$, $k \to \infty$ and $k/N \to 0$; Hall (1982) established asymptotic normality; bias and asymptotic variance are studied in Pictet et al. (1998); while Resnick and Starica (1998) provide an extension to dependent observations. Relatedly, another line of research compares the performance of the Hill statistic against other tail index estimators, when confronted with data artificially generated from a number of different distributions with different tail behaviors (see Pictet et al., 1998; De Haan and Peng, 1998; Weron, 2001, and the references therein).

$$R(x) = \text{Prob}\{S > x\} = Cx^{-\alpha} \quad , \tag{2.5}$$

where $\alpha = 1/\gamma$ and $C = b^{-1/\gamma}$. The empirical survival function is easily estimated with $\hat{R}(s_{(j)}) = j/N$, that is the rank of the j-th observation divided by the total number of observations. By taking a log-log transformation one obtains an estimate of the tail exponent via an OLS fit of the linear regression

$$\log \hat{R}(s_{(j)}) = \alpha \log s_{(j)} + c + \epsilon_j \ , j \in \{1, \ldots, K\} \ . \tag{2.6}$$

We refer to this estimator as the OLS-Rank (or simple Rank) estimator. A closely related estimator arises for the situation where the empirical data at hand are organized in classes of firm size, while the single observations on size of each firm are not available. In this case, the researcher has to resort to binned regression. The number of observations becomes the number of bins and one regresses

$$\log \hat{R}(b_r) = \alpha \log b_r + c + \epsilon_r \ , r \in \{1, \ldots, N_b\} \quad , \tag{2.7}$$

where N_b is the number of bins, b_r is the lower bound of the r-th bin and $\hat{R}(b_r)$ is the fraction of observations in the r-th and subsequent bins. We refer to this estimator as the OLS-CDF (or binned CDF) estimator.

Alternatively, one can start from the density

$$f(x) \sim C\alpha s^{-(\alpha+1)} \tag{2.8}$$

associated to (2.5), partition the sample in N_b bins, typically equispaced, and then compute the within-bin empirical probability density $Q_r = \text{Prob}\{b_r \leq S < b_{r+1}\}$, with b_r the lower bound of the r-th bin. This is just the fraction of observations laying in the semi-open interval $(b_r, b_{r-1}]$. Then one can perform the log-log OLS regression

$$\log Q_r = (\alpha + 1) \ (\log b_r + \log b_{r+1})/2 + c + \epsilon_r \quad , \tag{2.9}$$

to obtain what we refer to as the OLS-PDF (or binned PDF) estimator of the tail parameter α.

The use of bins in (2.7) strongly reduces the number of available observations and makes the estimate more noisy. The same problem arises with (2.9), but it is even worse in this case, since the regression is in fact only approximated. This approximation is afflicted by the typical problem of any density estimator (Silverman, 1986): a trade-off between an asymptotically finite bias, when the number of bins is low, and an exploding variance, when the number of bins is large. Moreover, all the OLS estimators of $\hat{\alpha}$ work under the usual assumption that the error terms ϵ are independent from the quantity on the right hand side of the equation, i.e. the survival function or the density. This condition does not generally holds, however. The issue is also present in the model in (2.6), which in fact provides biased estimates of the tail exponent. A solution to the bias is proposed in Gabaix and Ibragimov (2011) by estimating via OLS the alternative regression

$$\log\left(j - 1/2\right) = \alpha \ \log\left(s_{(j)}\right) + c + \epsilon_j \ , j \in \{1, \ldots, K\} \ . \tag{2.10}$$

This is equivalent to the Rank estimator in (2.6), apart from the $-1/2$ correction (whence the name Rank$-1/2$ estimator). It turns out that this modification corrects the leading order of the downward bias of the original Rank estimator. The corresponding corrected asymptotic standard error is $|\alpha| \, (2/K)^{1/2}$.

3 Previous findings

Studies within the Power Law approach to firm size distribution employ one or more of the above estimators. Table 1 summarizes the salient features and results of each study.

Okuyama et al. (1999) provide evidence from a sample of Japanese firms and from a cross-country dataset, also disaggregating the analysis by industrial sector of activity. The size proxy is firm income before taxes, and data are pooled over 4 years for Japanese firms and over 7 years for the other countries. The analysis is mainly graphical, but an OLS fit of the slope of the CDF plot is also reported, delivering the simple OLS-Rank estimates. Results show variation in the estimates of the Power Law exponent. At the aggregate economy level,

Table 1: Power Law approach to Firm Size Distribution - Literature Overview

Article	**Sample**	**Method**	**#Obs or #Bins**	**Findings**	**Assessment of Zipf**
Okuyama et al. (1999)	Japan firms, divided by sector and pooled over 4 years; Cross-country dataset, pooled over 7 years.	OLS-Rank	For Japan: thousands of firms in aggregate; from 700 to 11,000 at sectoral level. About 10,000 firms for U.S and 11,000 for other countries.	Size is income before taxes. $\alpha \simeq 1$ in Japan and Italy, $\alpha \simeq 1.44$ in US; $\alpha \in [0.7, 1.2]$ at sectoral level in Japan.	Un-conclusive on Zipf, via graphical analysis.
Axtell (2001)	U.S. Census Bureau in 1992 and 1997.	OLS-CDF, binned OLS-PDF, binned	10-15 bins	With employment, $\alpha \simeq 0.99$ in 1992 and $\alpha \simeq 1.06$ in 1997; with sales, $\alpha \simeq 0.99$ in 1997.	Zipf not rejected, based on S.E.
Fujiwara et al. (2003)	EU firms from AMADEUS in 1992-2002, estimates by year.	OLS-Rank	#Firms vary by country, year and size-proxy. From 8,000 to 15,000 firms approximately.	$\alpha \in [0.89, 0.99]$ for most years, size-proxies and countries.	Zipf not rejected, based on S.E.
Gabaix and Landier (2008)	U.S. COMPUSTAT over 1978-2004, estimates by year.	OLS Rank$-1/2$ Hill estimator	Top 500 firms	Size is market value. In 2004: Rank$-1/2$ gives $\alpha = 1.01$, Hill not reported. Averaging over time: Rank$-1/2$ gives $\alpha = 0.87$, Hill gives $\alpha = 1.09$.	Zipf not rejected, based on S.E.
Podobnik et al. (2010)	U.S. firms filing for bankruptcy in 1999-2009, estimates by year.	OLS-Rank	About 2800 firms in total, but top 100, 200, 300, 500 and 2500 enter the estimates.	With assets: $\alpha = 1.11$ pre-filing, $\alpha = 1.44$ post-filing. For the NASDAQ sub-sample: $\alpha = 1.1$ with market capitalization, and $\alpha = 1.02$ with equity.	Zipf not rejected, based on S.E.
di Giovanni and Levchenko (2010)	Cross-Country from ORBIS, yearly data in 2006-2008.	OLS-CDF, binned OLS-PDF, binned Rank$-1/2$	Several thousands or millions of firms in each country. Only countries with at least 1,000 firms; number of bins not reported.	Size is sales, estimates of α vary by country.	Zipf not rejected in most countries, via t-test for $\alpha = 1$. $\alpha < 1$ when rejected.
di Giovanni et al. (2011)	All French firms 2006, exporters vs. non-exporters, and by sector.	OLS-CDF, binned OLS-PDF, binned Rank$-1/2$	About 150,000 firms in total. From 50 to 20,000 firms at sectoral level. Number of bins not reported.	Aggregate manufacturing: with sales $\alpha \simeq 1.02$ for both CDF and PDF, $\alpha = 0.82$ for Rank$-1/2$; with employees $\alpha \simeq 1.08$ for both CDF and PDF, $\alpha = 0.79$ for Rank$-1/2$. Values smaller for exporters, and vary by sector.	Zipf not rejected in most cases, based on S.E. Though recognize the heuristic nature of CDF and PDF methods.

an estimate of α close to 1 is reported for Japanese and Italian firms, while $\alpha = 1.4$ best approximates U.S. data. Sectoral analysis for Japanese firms provides α in the range [0.7,1.2], while results for other countries are un-conclusive, due to limitation in the number of firms. In general, the authors tend to conclude that the Zipf Law is a reasonable description, although no formal test is provided.

The study by Axtell (2001) analyzes the size distribution of theoretically the entire population of US firms, through Census Bureau data. Size is measured in terms of both number of employees and revenues. Data on employment report the number of firms in successive classes, where the size of the classes increases in powers of three. The analysis in this case applies the binned OLS-PDF estimator, yielding $\alpha = 1.059$ (S.E. 0.054, R^2 0.994) in 1997, and $\alpha = 0.995$ (S.E. 0.043, R^2 0.994) in 1992. With revenues, firms are tabulated in classes whose width increases in powers of 10. Estimates are in this case obtained via the binned OLS-CDF estimator. The findings are reported for 1997 only, with $\alpha = 0.994$ (S.E. 0.064, R^2 0.976). Overall, the results are interpreted as supporting the Zipf Law.

Fujiwara et al. (2003) present evidence about a sample of European countries over the period 1992-2002. Reported results mainly focus on the UK, France, Italy and Spain. There is a minimum size thresholds which a firm must pass in order to be included in the dataset. The threshold is defined by the data collection process, and it varies according to the different proxies of firm size employed in the study: 150 units on employment, 150 ML Euros on operating revenues, 30 ML Euros on total assets. The statistical analysis is based on the simple OLS-Rank estimator, separately by year and country. For a collection of variables on France and UK data in 2001, inference on validity of $\alpha = 1$ is based on standard errors. Estimates for the other countries, years and proxies are discussed graphically. Overall, the authors conclude that the data support Zipf law.

Podobnik et al. (2010) apply the OLS-Rank estimator, too. The study covers a sample of about 2,800 U.S. firms filing for bankruptcy over the period 1999-2009. The focus is on measures of assets and market value. The authors run separate estimates by year, and compare different tail cut-offs (successively including the top-100, -200, -300, -500, -2500 firms). Reported results using book value of assets show that $\alpha = 1.11$ (S.E. 0.01) in the pre-filing period, and $\alpha = 1.44$ (S.E. 0.01) in the post-filing years. Further, for a sub-sample of NASDAQ firms, the authors report an estimated $\alpha = 1.10$ (S.E. 0.02) on market capitalization, and $\alpha = 1.02$ (S.E. 0.01) on firm equity. It is unclear, however, whether these estimates refer to the whole sample or to a specific tail cut-off. Based on reported standard errors, the estimate are interpreted as supporting the Zipf Law, although graphical analysis reveals deviations when the top-500 or more firms are included in the tail.

Gabaix and Landier (2008) is to our knowledge the only paper applying the Hill estimator to infer about the firm size distribution, within a study of CEO pays. The data cover US listed firms from the COMPUSTAT database, over the period 1978-2004. The size proxy is firm market value, measured as debt plus equity, with a minimum size cut-off implicitly defined by considering only the top-500 firms in the estimation. Averaging over the estimates obtained in each year, the Hill estimator gives $\alpha = 1.095$ (ST.DEV. 0.063), while the Rank$-1/2$ estimator yields $\alpha = 0.869$ (ST.DEV. 0.071). Rank$-1/2$ estimates are also provided for 2004, where it turns out that $\alpha = 1.01$ (ST.DEV. 0.063). This is suggested as further evidence in favor of Zipf law in firm size.

Finally, two recent works apply OLS log-log techniques to study the shape of the firm size distribution within the literature on firm heterogeneity in international trade. di Giovanni and Levchenko (2010) study the universe of French firms in 2006, while di Giovanni et al. (2011) analyze a large set of countries from the AMADEUS-ORBIS database.

The study on French firms uses both sales and employment as size proxy, both at the aggregate and sectoral level, and also comparing exporting vs. non-exporting firms. Following Gabaix (2009), a minimum size cut-off is identified graphically, corresponding to the value of sales above which the log-rank against log-size relation becomes approximately linear.[7] This yields a total of about 150,000 firms included in the analysis. The estimates are based on the binned OLS-CDF and the binned OLS-PDF estimators, and also checked against the Rank$-1/2$ estimator. The number of bins is not reported, however. At the aggregate level, results for sales give $\alpha \simeq 1.02$ (S.E. 0.030) from the OLS-CDF and OLS-PDF methods, respectively, and $\alpha = 0.825$ (S.E. 0.003) from the Rank$-1/2$ estimator. With employment, the estimates are $\alpha \simeq 1.08$ with OLS-CDF and OLS-PDF, $\alpha = 0.79$ for Rank$-1/2$. For both size proxies there is variation across exporters and non-exporters, with the latter group having greater α (still below 1 when considering the Rank$-1/2$ estimator). When breaking down the data by industrial sectors, there is considerable variation in the estimated α across sectors, although the ranking between exporters and non-exporters appears robust within sectors. It

[7] It is reported that this cut-off roughly corresponds to an institutional threshold on annual sales (750,000 Euro) that defines different accounting standards in place in France for firms above or below the threshold.

is also mentioned that consistent results are obtained on the full sample without size cut-off.

In the paper performing cross-country analysis from the ORBIS dataset, the size proxy is sales, and the minimum size cut-off (not reported) varies by country. It is mentioned that the data cover several thousands of firms per country (even millions, in some cases like for the U.S.), but reported estimates only concern countries with at least 1,000 firms available in the considered year. The analysis exploits the binned OLS-CDF and binned OLS-PDF estimators (number of bins not reported), and the two estimators deliver consistent results. There is considerable variation in the estimates of α across the different countries, with α in between 0.69 and 1.18. It is mentioned that Rank$-1/2$ estimates (not reported) are in agreement with these findings. The authors also provide the p-value of a t-test for the null of deviation from Zipf law ($\alpha \neq 1$). Despite the null cannot be rejected in a non-negligible number of countries, the results are interpreted, once again, as supporting Zipf Law as a good first order approximation.

4 Small sample properties of the estimators

Despite their widespread use in the study of firm size, we lack a systematic knowledge of the properties of the estimators in small sample. The Hill estimator in (2.2) is, at least theoretically, the mostly reliable estimator. The expression in (2.2) already contains a correction for bias, and its variance decreases with $1/N$. Gabaix and Ibragimov (2011) provide formal results for small sample properties of their Rank$-1/2$ estimator: the 1/2 correction reduces to a leading order the downward bias of the original OLS-Rank estimator, and the variance decreases with $2/N$. The binned OLS-CDF and the binned OLS-PDF estimators are known to be biased, but there is no systematic study of their behavior. Further, the binning procedure underlying both methods can severely impact on their performance, since in this case regressions are typically performed on a very small number of data points (10-15 bins in Axtell, 2001).

In this Section we present our main Monte Carlo exercises assessing the reliability of the different estimators in discriminating about the Zipf Law. The design of the simulations tries to keep comparability with the features of the above reviewed studies. In particular, we do not explore variation of estimators' properties with respect to endogeneous choice of the tail starting point. A relevant issue discussed in the statistical literature is indeed that the estimators' behavior can depend on the number of ordered observations belonging to the tail, and thus entering the estimation. This has spurred considerable work concerning methods to "optimally" select the tail starting point in practical applications, especially with reference to the Hill estimator.[8] This issue is however not considered in size distributions studies, where researchers either use all the data at hand or impose the cut-off on the basis of exogenously given criteria.

4.1 Bias and variance

We start by studying bias and variance of the estimators when simulated data obey exact Zipf Law. We generate $R = 10,000$ independent samples of size N, drawn from the Pareto type-I Power law in (2.1) with unitary tail exponent. On each sample we apply the OLS estimators (Rank, Rank$-1/2$, CDF and PDF) and the bias-corrected Hill statistic in (2.2), and then compute the average bias and the variance of the α estimated across the replications. In order to remain reasonably close to the sample size typically met in the empirical studies (recall Table 1), N varies from 50 to 2000.[9] Also, for the binned methods, we compare estimates obtained with 15 and 40 bins.

Figure 1 provides a graphical presentation of the findings. Plots in the left column report about the estimators exploiting individual data (Hill, Rank and Rank$-1/2$), while the binned CDF and PDF estimators are reported in the right column. The bias of the Hill estimator is practically zero, as expected (see top panel). Fluctuations here are artifacts of the numerical analysis. The bias of the Rank estimator is substantial and decreases proportionally to $1/N$. The Rank$-1/2$ estimator effectively improves upon the simple Rank estimator. Its bias, however, despite relatively small does not exactly converges to the Hill value even at N =2,000, due to the presence of sub-asymptotic corrections. Conversely, the binned CDF and binned PDF estimators are both afflicted by a severe bias, well above their standard deviation and clearly larger than the bias affecting the Rank estimator, even for large N. In particular, the PDF estimator turns out as the worst performing among all the methods considered. Its behavior is also sensibly affected by the binning, as the bias indeed increases if we increase from 15 to 40 bins. This is likely a result of the inconsistency of the binned estimate of

[8]Several formal algorithms have been developed. See, for instance, DuMouchel (1983), Hall and Welsh (1985), Beirlant et al. (1996), Beirlant et al. (1999), Resnick and Starica (1997), Danielsonn et al. (2001), Pictet et al. (1998).

[9]Convergence is already reached with this sample size, so results are informative also for even larger sample sizes encountered in the literature, such as in di Giovanni et al. (2011).

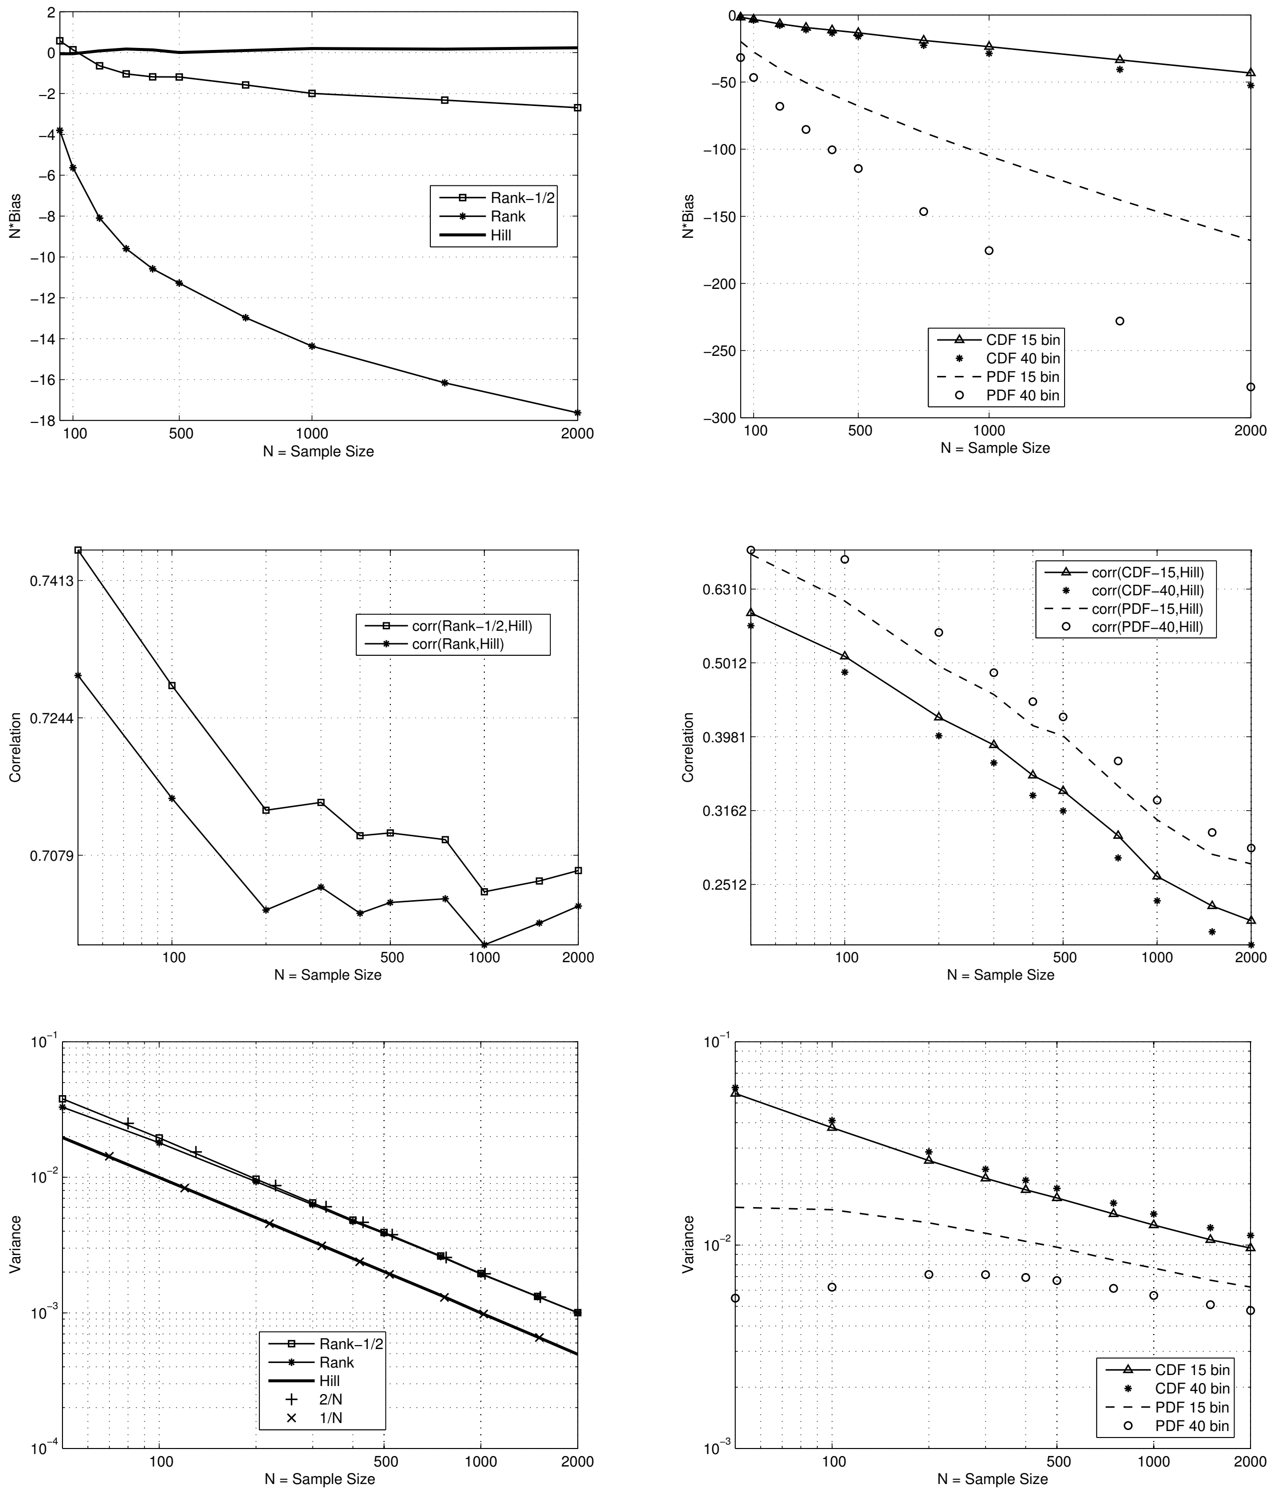

Figure 1: Properties of estimators: bias, plotted as $N * bias$ (top panel); correlation with the bias-corrected Hill estimator (middle panel); and variance (bottom panel). Figures computed over 10,000 simulated samples of different sample size N, drawn from the Power Law in (2.1) with unitary tail exponent.

the probability density. The effect does not disappear increasing the sample size. The CDF method is more stable in this respect.

We also explore (see middle panel) the correlation between the different point estimates and the benchmark provided by the biased-corrected Hill estimates. The Rank$-1/2$ and the Rank estimator display similar performance, with a correlation coefficient of about 0.7 for most values of N, and sensibly higher till $N = 200$. The CDF and PDF estimators perform worst relative to the benchmark, with very small correlation (about 0.2) for large N.

Concerning the variance (see bottom panel), we reproduce the theoretical results that the variance of the Hill estimator decreases with $1/N$ and that the variance of the Rank$-1/2$ estimator decreases with $2/N$. The simple Rank estimator has variance close to the Rank$-1/2$ variance for all N. The variance of the CDF estimator is larger, always above 0.01 (even for N =2,000) and different binning does not have any impact. The variance of the PDF estimator has a smoother behavior and it is sensibly smaller if we take 40 bins. The general lesson is therefore that the standard errors associated to all the OLS log-log estimators may be too large in many practical situations, preventing reliable inference about the actual value of the tail parameter α. In the next Section we present a more formal study of the properties of the t-tests associated with the different estimators.

4.2 Power of t-tests for the Zipf null

As known, a t-test for the null $H_o : \alpha = 1$ amounts to compare the estimated t-statistic under a given estimator of α, $\hat{t} = (\hat{\alpha} - 1)/\hat{\sigma}$, against tabulated values of the standardized Normal distribution, for given confidence level.

To assess the power of the t-tests based on the different estimators, we perform a Monte Carlo study of the probability to reject the null $\alpha = 1$ when the true data generating process (DGP) is a Power Law distribution with varying tail index. We fix a value of α and generate R artificial samples of size N drawn from the Pareto type-I distribution in (2.1) under that specific value of the tail index. On each sample, we estimate α and the corresponding t-statistics through all the estimators, and count the number of times that $\alpha = 1$ is rejected at 5% confidence level. Next, we repeat the procedure for another set of R samples generated under a different value of α. The ideal estimator should reject the null 5% of the times if the true DGP indeed has $\alpha = 1$, and always reject the null when the DGP has $\alpha \neq 1$, or at least

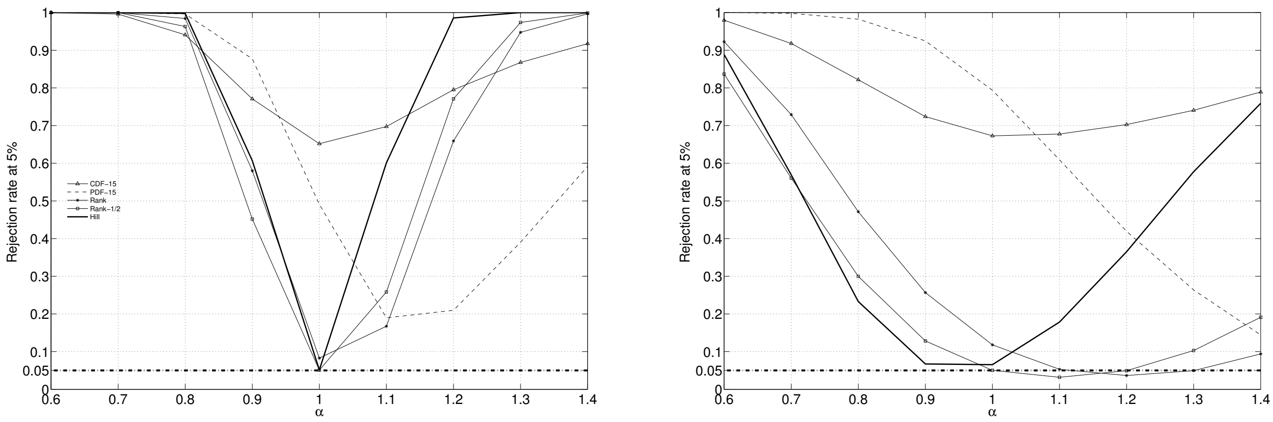

Figure 2: Comparing estimators' power to reject Zipf Law. Percentage of times (y-axis) that a t-test rejects $H_0 : \alpha = 1$, when the true DGP is generated as a Power Law with different values of α (x-axis). Results over 10,000 generated samples for each value of α, with sample size 500 (left plot) and 50 (right plot).

it should tend toward this optimal behavior as the number of replications increases.[10]

Figure 2 shows the results obtained over $R = 10,000$ samples. With sample size $N = 500$ (see left plot), we find that, in agreement the analysis of bias and variance, only the Hill and the Rank$-1/2$ estimators meet the theoretical 5% confidence level when the DGP has $\alpha = 1$. The Rank estimator has also some merit, although it slightly over-rejects the null (rejection rate is about 10%). The CDF and PDF estimators cannot be considered as reliable, with rejection rates of about $60 - 70\%$. Comparing Hill, Rank$-1/2$ and Rank estimators' power to correctly reject a simulated $\alpha \neq 1$, we tend to conclude that the Hill estimator turns out as the most solid basis for inference. Indeed, the associated rejection rates increase more rapidly toward 100%. The Rank estimator has similar performance for true $\alpha < 1$, while both the Rank and the Rank$-1/2$ estimators significantly under-reject for true $\alpha > 1$. For instance, for $\alpha = 1.1$, the rejection rate of the Hill estimator is about 70%, while the two competing estimators reject only in the $15 - 20\%$ of the cases. Results are identical with $N = 1000$. Thus, $N = 500$ represents a valid benchmark also for larger sample sizes encountered in the reference literature on firm size.

We repeat the same exercise with sample size $N = 50$, providing information for practical situations similar to the lower bound of the estimation sample sizes employed in the reference literature. As the results show (see right plot), the power of the tests reduces considerably for all the estimators. This fact is expected given the reduction in number of observations considered in the estimates. Despite the implied larger noise, however, we can broadly confirm

[10]We report CDF and PDF estimates with 15 bins, given the smaller bias found above as compared to the 40 bins case.

the patterns described above. First, the t-tests based on the CDF and the PDF estimators are completely unreliable, as indeed they cannot discriminate between different values of simulated α. Second, for true $\alpha = 1$, we still observe the comparatively superior performance of the Hill and Rank$-1/2$ estimators, as well as a satisfactory performance of the Rank estimator: as for $N = 500$, the associated rejection rates hit the theoretical 5% level (10% for the simple Rank). Third, the Hill estimator is definitely the only reliable estimator when there are reasons to believe that the true $\alpha > 1$, while Hill and Rank$-1/2$ have similar power for true $\alpha < 1$.

4.3 Power of t-tests under sub-asymptotic deviation from Zipf law

We next study the reliability of point estimates and t-tests when the DGP is of the Zipf form ($\alpha = 1$), but a sub-asymptotic perturbation is present. We consider the distribution function

$$P(X > x) = x^{-1}(1 + c(x^{-1} - 1)) \quad , x > 1,\, c \in [0, 1) \quad , \tag{4.1}$$

which gives exact Zipf law if $c = 0$, while for $c \neq 0$ the Zipf behavior is perturbed by a factor $\sim 1/x^2$. This correction is sub-asymptotic in the sense that it tends to zero for very large values of x, while it has more weight the wider the range of observations included in the tail.

Gabaix and Ibragimov (2011) exploit this distribution to study the sensitivity of their Rank$-1/2$ estimator to different values of c. We extend the same robustness analysis to the Hill, Rank, CDF and PDF estimators.[11] To keep comparability, the Monte Carlo design exactly follows the reference article. We generate R random samples of size N =2,000 extracted from the process in (4.1), for a given value of c. Next, on each sample we consider two different tail width for estimation, i.e. including either the top-50 or the top-500 observations. The procedure is then repeated for different c.

The results are presented in Table 2. As in Gabaix and Ibragimov (2011), we report the average point estimates across $R = 10,000$ runs together with asymptotic (theoretical) and sampled standard errors. In addition, we also explore the results of a t-test of the true null of unitary tail index. However, differently from Gabaix and Ibragimov (2011), who test whether the across-replication average point estimate deviates from 1, we compute the percentage of times that the null $\alpha = 1$ is rejected (at 5% confidence level) by a t-test performed at each

[11] As in previous section, CDF and PDF estimates are computed with 15 bins, given the superior performance as compared to the 40-bins version.

Table 2: Robustness to sub-asymptotic deviation from Zipf Law

	Top 50					Top 500				
c	Hill	Rank−1/2	Rank	CDF	PDF	Hill	Rank−1/2	Rank	CDF	PDF
0.0	1.022 (0.07)	1.010 (0.05)	0.923 (0.12)	0.964 (0.68)	0.605 (0.80)	1.002 (0.05)	0.998 (0.05)	0.978 (0.08)	0.973 (0.66)	0.864 (0.50)
	(0.147) (0.149)	(0.202) (0.195)	(0.185) (0.182)	(0.055) (0.238)	(0.138) (0.124)	(0.045) (0.045)	(0.063) (0.063)	(0.062) (0.063)	(0.033) (0.131)	(0.070) (0.099)
0.5	1.046 (0.09)	1.024 (0.05)	0.935 (0.11)	0.974 (0.68)	0.615 (0.77)	1.157 (0.90)	1.084 (0.22)	1.061 (0.15)	1.014 (0.66)	0.915 (0.26)
	(0.151) (0.154)	(0.205) (0.200)	(0.187) (0.187)	(0.056) (0.244)	(0.141) (0.128)	(0.052) (0.056)	(0.069) (0.076)	(0.067) (0.075)	(0.036) (0.148)	(0.082) (0.118)
0.8	1.182 (0.28)	1.108 (0.05)	1.010 (0.08)	1.036 (0.68)	0.673 (0.57)	1.475 (1.00)	1.313 (0.94)	1.284 (0.91)	1.141 (0.70)	1.053 (0.24)
	(0.171) (0.184)	(0.222) (0.235)	(0.202) (0.219)	(0.062) (0.281)	(0.158) (0.153)	(0.066) (0.074)	(0.083) (0.110)	(0.081) (0.109)	(0.049) (0.201)	(0.112) (0.168)

Note: Estimates of tail index from DGP following $P(X > x) = x^{-1}(1 + c(x^{-1} - 1))$, $x > 1$, $c \in [0, 1)$. Results over 10,000 Monte Carlo simulations with sample size $N = 2000$ and varying tail width (Top-50 vs. Top-500 observations), for different values of c. CDF and PDF estimates computed with 15 bins. For each c the first line reports point estimates of tail index averaged over the replications and, in parenthesis, the percentage of times the null of unitary tail index is rejected (at 5% confidence level); the second line shows, in parenthesis, the theoretical standard errors (usual OLS for Rank, CDF and PDF estimators; propagated via Taylor expansion of the asymptotic variance as in Gabaix and Ibragimov (2011) for the Rank−1/2 estimator; and given in 2.4 for the Hill estimator) together with sample standard errors.

run. So, for instance, if we take the top-50 firms and focus on the benchmark case of pure Zipf (i.e., $c = 0$) the average Rank$-1/2$ point estimate is 1.010 and the sample rejection rate of the Zipf null is 5% (reported in the first raw), while the theoretical and sample standard errors are 0.202 and 0.195, respectively (reported in the second raw).

If an estimator is well-behaved, then it is expected that the sample rejection rate is equal to the theoretical confidence level of 5%. Not surprisingly, when we switch the correction off (c=0), the Hill and Rank-1/2 estimators clearly outperform the other methods, with point estimates very close to 1 and properly sized rejection rates. Results change if we set $c \neq 0$. When we focus on the very extreme of the tail (top-50 observations), the Rank$-1/2$ estimator beats all the other estimators: the rejection rate indeed always equals the expected 5%. Though, in contrast with Gabaix and Ibragimov (2011), who report a substantial invariance of the Rank$-1/2$ point estimates with respect to c, we do find that point estimates vary with c. The simple Rank estimator also does a good job, in that the point estimates are not too far from 1, and the associated rejection rates remain low, close to 10%. A similar pattern is found for the Hill estimator, although this latter exhibits higher sensitivity to the sub-asymptotic correction and it rejects in 28% of the times for $c = 0.8$. The CDF estimator provides point estimates not that far from 1, but the rejection rates are clearly too high to provide a reliable basis for inference. The PDF estimator is the worst performing, with large bias in average point estimates and extremely high rejection rates.

Moving to a larger tail (top-500 observations), i.e. where the sub-asymptotic correction becomes theoretically more relevant, we confirm the superior performance of Hill and Rank-1/2 estimators for the case of no correction (c=0). When $c \neq 0$ instead, the general result is that point estimates of all estimators tend to increase as compared to the top-50 case, and rejection rates signal that all methods substantially over-reject the true null of unitary tail index.

Gabaix and Ibragimov (2011) use (4.1) to also investigate the behavior of their Rank$-1/2$ estimator under dependent data, by simulating an AR(1) or MA(1) process with innovations distributed according to (4.1). For completeness, a full extension of their analysis to all the other estimators is presented in Table 4 and Table 5 in Appendix. These forms of time dependence across observations are especially important for applications in finance (see Resnick and Starica, 1998), but they are not meaningful when studying the distribution of firm sizes over a yearly cross-section of firms. Conversely, controlling for an autoregressive structure can indeed be important when the behavior of the firm size distribution is inferred from pooled cross-sections (as in Okuyama et al., 1999), so that the same firm is observed repeatedly over time. We deal with this issue in the next section.

4.4 Pooling over time and Laplace growth shocks

A well established result in the empirics of firm growth is that the intertemporal evolution of the size of a firm i is well approximated by an autoregressive multiplicative model

$$S_{i,t} = S_{i,t-1}^{\beta}\, e^{\epsilon_{i,t}} \quad , \tag{4.2}$$

with β typically found to equal 1 and the growth shocks $\epsilon_{i,t}$ to follow a Laplace distribution.[12]

Accounting for this type of temporal dependence is relevant when one seeks to estimate the firm size distribution from data pooled over time. With an yearly cross-section of firms, indeed, independence across observations is the implicit assumption underlying the construction of estimators. With pooled data, repeated observations on the same firm, correlated over time, enter the sample, creating an obvious source of potential bias.

We present a Monte Carlo exercise that explores how the different estimators and related t-tests behave with respect to this issue. To simulate pooling over time, we generate a sample of 500 "initial" firm sizes drawn from the Pareto type-I Power law in (2.1), with a given value α of the tail exponent. Next, we evolve the firm size forward in time for 3 subsequent years, according to the process in (4.2) with fixed β and Laplace distributed shocks. Then we pool all the years together to obtain a sample of $2,000$ observations and take the top-500 observations of this pooled dataset as the tail on which we estimate the tail index and apply a t-test (at 5% confidence level) for the null of unitary tail index under the different estimators. Finally, the procedure is repeated for $R = 10,000$ independent samples, varying the tail index α and considering different values of the autoregressive parameter β.

Table 3 shows the results obtained with three different values of β, all close to 1 in accordance with the empirical evidence on Gibrat's Law. Growth shocks are drawn from a Laplace distribution with variance $\sigma = 0.01$.[13] We show point estimates and rejection rates of the

[12]Gibrat's Law of proportionate effects prescribes $\beta = 1$ and i.i.d. shocks. Deviations with $\beta < 1$ are usually observed form smaller firms, together with a negative relationship between variance of growth shocks and initial size, which also contradicts the Law. The Laplacian nature of the shocks has been found to be robust and invariant across countries and also across sectors, even at different level of sectoral aggregation. See Amaral et al. (1997); Bottazzi and Secchi (2006); Bottazzi et al. (2011, 2014).

[13]Results with different σ are presented in Appendix. The main findings do not change significantly. We

Table 3: Robustness to pooling over time with Laplace growth shocks

α	β	Hill	Rank−1/2	Rank	CDF	PDF
0.5	0.8	0.630 (1.00)	0.603 (1.00)	0.591 (1.00)	0.578 (1.00)	1.519 (1.00)
		(0.028) (0.047)	(0.038) (0.065)	(0.037) (0.064)	(0.017) (0.103)	(0.041) (0.080)
	0.9	0.576 (1.00)	0.571 (1.00)	0.560 (1.00)	0.569 (1.00)	1.509 (1.00)
		(0.026) (0.048)	(0.036) (0.068)	(0.035) (0.067)	(0.018) (0.111)	(0.032) (0.081)
	1.0	0.506 (1.00)	0.505 (1.00)	0.495 (1.00)	0.483 (1.00)	1.379 (1.00)
		(0.023) (0.045)	(0.032) (0.063)	(0.031) (0.062)	(0.022) (0.091)	(0.052) (0.061)
0.7	0.8	0.882 (0.70)	0.845 (0.70)	0.827 (0.76)	0.809 (0.88)	1.727 (0.92)
		(0.040) (0.066)	(0.053) (0.091)	(0.052) (0.090)	(0.024) (0.144)	(0.057) (0.113)
	0.9	0.806 (0.93)	0.800 (0.83)	0.784 (0.87)	0.797 (0.87)	1.713 (0.95)
		(0.036) (0.068)	(0.051) (0.096)	(0.050) (0.094)	(0.025) (0.156)	(0.046) (0.113)
	1.0	0.708 (1.00)	0.706 (0.98)	0.693 (0.98)	0.677 (0.97)	1.531 (1.00)
		(0.032) (0.064)	(0.045) (0.088)	(0.044) (0.087)	(0.030) (0.127)	(0.073) (0.085)
0.9	0.8	1.133 (0.69)	1.086 (0.35)	1.064 (0.31)	1.040 (0.77)	1.934 (0.37)
		(0.051) (0.085)	(0.069) (0.117)	(0.067) (0.115)	(0.031) (0.186)	(0.074) (0.145)
	0.9	1.037 (0.33)	1.028 (0.30)	1.008 (0.30)	1.025 (0.78)	1.917 (0.51)
		(0.046) (0.087)	(0.065) (0.123)	(0.064) (0.121)	(0.033) (0.200)	(0.059) (0.145)
	1.0	0.910 (0.56)	0.908 (0.48)	0.892 (0.53)	0.870 (0.74)	1.685 (0.86)
		(0.041) (0.082)	(0.057) (0.113)	(0.056) (0.112)	(0.039) (0.163)	(0.095) (0.110)
1.0	0.8	1.259 (0.98)	1.207 (0.69)	1.182 (0.62)	1.156 (0.83)	2.038 (0.34)
		(0.056) (0.095)	(0.076) (0.130)	(0.075) (0.128)	(0.035) (0.206)	(0.082) (0.161)
	0.9	1.152 (0.74)	1.143 (0.51)	1.120 (0.46)	1.139 (0.84)	2.019 (0.44)
		(0.052) (0.097)	(0.072) (0.137)	(0.071) (0.135)	(0.036) (0.223)	(0.065) (0.161)
	1.0	1.011 (0.33)	1.009 (0.32)	0.991 (0.33)	0.967 (0.67)	1.762 (0.61)
		(0.045) (0.091)	(0.064) (0.125)	(0.063) (0.124)	(0.044) (0.181)	(0.106) (0.123)
1.1	0.8	1.385 (1.00)	1.328 (0.91)	1.300 (0.88)	1.271 (0.87)	2.142 (0.48)
		(0.062) (0.104)	(0.084) (0.143)	(0.082) (0.141)	(0.038) (0.227)	(0.090) (0.177)
	0.9	1.267 (0.97)	1.257 (0.78)	1.232 (0.73)	1.253 (0.88)	2.120 (0.54)
		(0.057) (0.107)	(0.079) (0.150)	(0.078) (0.148)	(0.040) (0.245)	(0.072) (0.177)
	1.0	1.112 (0.59)	1.110 (0.44)	1.090 (0.40)	1.063 (0.68)	1.839 (0.35)
		(0.050) (0.100)	(0.070) (0.138)	(0.069) (0.136)	(0.048) (0.199)	(0.116) (0.136)
1.3	0.8	1.637 (1.00)	1.569 (1.00)	1.536 (0.99)	1.502 (0.94)	2.350 (0.75)
		(0.073) (0.123)	(0.099) (0.169)	(0.097) (0.167)	(0.045) (0.268)	(0.107) (0.209)
	0.9	1.497 (1.00)	1.486 (0.98)	1.456 (0.97)	1.481 (0.93)	2.324 (0.78)
		(0.067) (0.126)	(0.094) (0.178)	(0.092) (0.175)	(0.047) (0.289)	(0.085) (0.209)
	1.0	1.314 (0.98)	1.312 (0.86)	1.288 (0.82)	1.257 (0.80)	1.993 (0.11)
		(0.059) (0.118)	(0.083) (0.163)	(0.081) (0.161)	(0.057) (0.236)	(0.138) (0.162)
1.5	0.8	1.889 (1.00)	1.810 (1.00)	1.773 (1.00)	1.733 (0.97)	2.557 (0.88)
		(0.085) (0.142)	(0.114) (0.195)	(0.112) (0.193)	(0.052) (0.310)	(0.123) (0.241)
	0.9	1.728 (1.00)	1.714 (1.00)	1.680 (1.00)	1.709 (0.96)	2.528 (0.90)
		(0.077) (0.145)	(0.108) (0.205)	(0.106) (0.202)	(0.055) (0.334)	(0.098) (0.242)
	1.0	1.516 (1.00)	1.514 (0.98)	1.486 (0.98)	1.450 (0.89)	2.148 (0.20)
		(0.068) (0.136)	(0.096) (0.188)	(0.094) (0.186)	(0.065) (0.272)	(0.160) (0.187)

Note: Estimation samples obtained from initial samples of 500 observations with DGP following the Power Law in (2.1), then evolved forward in time for 3 time periods according to (4.2) with Laplace shocks of variance $\sigma = 0.01$, and finally taking the top-500 observations resulting from pooling the observations across time periods. Results over $10,000$ Monte Carlo simulations for varying tail exponent α and varying AR parameter β. CDF and PDF estimates computed with 15 bins. For each combination of α and β, the first line reports point estimates of tail index averaged over the replications and, in parenthesis, the percentage of times the null of unitary tail index is rejected (at 5% confidence level); the second line shows, in parenthesis, the theoretical standard errors (usual OLS for Rank, CDF and PDF estimators; propagated via Taylor expansion of the asymptotic variance as in Gabaix and Ibragimov (2011) for the Rank−1/2 estimator; and given in (2.4) for the Hill estimator) together with sample standard errors.

t-test for the null of unitary tail index (first row), and theoretical and sampled standard errors (in parenthesis, second row).[14]

Apart from the PDF estimator, which always delivers wrong point estimates, the point estimates obtained from the other methods do not differ much from the true value of α. They are quite precise, in particular, for $\beta = 1$, while becomes more biased for less persistent time series (smaller β). However, there is a clear tendency of associated t-tests to over-reject the null. This is obviously a desired property if the simulated α is indeed different from 1. When the true α is 1, and we would expect rejection rates close to the theoretical 5%, the best performing estimators are the Hill, Rank$-1/2$ and simple Rank estimators. Their rejection rate is about 30%, which is clearly too high.

5 Conclusion

In this paper we have reviewed the literature that has created a consensus about the Power Law behavior in firm size distribution, and examined the methods leading these studies to conclude that a Zipf Law represents an ubiquitous property of the data.

The CDF and PDF log-log estimators have poor properties. These estimators, despite their widespread use, do not provide a solid basis for tail inference. Indeed, the associated t-tests for the null of unitary tail index perform very poorly, both when the true data generating process is exactly Zipf and when we introduce sub-asymptotic deviations from the Zipf law. The PDF estimator gives particularly unreliable results with pooled data.

The Rank and Rank$-1/2$ share similar behavior, not surprisingly, but the latter is to be preferred both for the superior theoretical properties and for its better performance in most of the Monte Carlo analysis we have presented here.

As a result, the Hill and Rank$-1/2$ estimators stand out as the two more solid methods. In general, the Hill estimator is attractive because there is an explicit correction for its small sample bias and because of its maximum likelihood nature. Our exercises on reliability of associated t-tests show that the two estimators compete each other in small samples.

First, if the focus lies on the ability to reject the null of Zipf law when the true tail index indeed differs from 1, then the t-test associated to the Hill estimator is definitely more powerful for sample size around 500 observations. For very small sample size, i.e. of about

also experimented with Gaussian growth shocks, results are analogous and available upon request.

[14] Binned estimators employ 15 bins.

50 observations, the Hill estimator is still more valid if there is a-priori suspect (for instance, after a first step estimate) that the true α is above 1. Instead, the Hill and the Rank$-1/2$ estimators are equivalent and both well behaved for true α below 1.

Second, if the focus is on sensitivity to sub-asymptotic deviations from the standard i.i.d. Zipf process, the Rank$-1/2$ has some merit in this case. The associated t-test for rejecting a unitary tail index is more powerful when one considers a second order deviation from the Zipf law. Such superiority, however, vanishes as the tail width increases, i.e. in the range of observations where the correction is a-priori more relevant.

Third, pooling over time does not seem to make a big difference in terms of the ability of the estimators to deliver reasonable values of point estimates, but produces unreliable inference when there is reason to suspect that the data are exactly Zipf distributed.

Our results cast doubts on the validity of the conclusions reported in previous studies. They also provide guidance for future researchers who, subscribing to the Power Law approach, wish to estimate the tail behavior of the size distribution of firms. We remind, however, that the estimators studied here only provide alternative methods to discriminate which particular distribution within the Power Law family best approximates the data. They cannot answer the question of whether the Zipf law, in particular, or the Power Law, in general, does a better job than other distributions in describing the empirical firm size distribution.

References

AMARAL, L., S. BULDYREV, S. HAVLIN, P. MAASS, M. SALINGER, H. STANLEY, AND M. STANLEY (1997): "Scaling behavior in economics: The problem of quantifying company growth," *Physica A*, 244, 1–24.

AXTELL, R. L. (2001): "Zipf Distribution of U.S. Firm Sizes," *Science*, 293, 1818–1820.

BEIRLANT, J., G. DIERCKX, Y. GOEGEBEUR, AND G. MATTHYS (1999): "Tail Index Estimation and an Exponential Regression Model," *Extremes*, 2, 177–200.

BEIRLANT, J., P. VYNCKIER, AND J. TEUGELS (1996): "Tail Index Estimation, Pareto Quantile Plots, and Regression Diagnostics," *Journal of the American Statistical Association*, 91, 1659–1667.

BOTTAZZI, G., A. COAD, N. JACOBY, AND A. SECCHI (2011): "Corporate Growth and Industrial Dynamics: Evidence from French Manufacturing," *Applied Economics*, 43.

BOTTAZZI, G. AND A. SECCHI (2006): "Explaining the Distribution of Firms Growth Rates," *The RAND Journal of Economics*, 37, 235–256.

BOTTAZZI, G., A. SECCHI, AND F. TAMAGNI (2014): "Financial constraints and firm dynamics," *Small Business Economics*, 42, 99–116.

DANIELSONN, J., L. D. HAAN, L. PENG, AND C. G. D. VRIES (2001): "Using a Bootstrap Method to Choose the Sample Fraction in Tail Index Estimation," *Journal of Multivariate Analysis*, 76, 226–248.

DE HAAN, L. AND L. PENG (1998): "Comparison of tail index estimators," *Statistica Neerlandica*, 52, 60–70.

DE WIT, G. (2005): "Firm size distributions: An overview of steady-state distributions resulting from firm dynamics models," *International Journal of Industrial Organization*, 23, 423–450.

DI GIOVANNI, J. AND A. A. LEVCHENKO (2010): "Firm Entry, Trade, and Welfare in Zipf's World," NBER Working Papers 16313, National Bureau of Economic Research.

di Giovanni, J., A. A. Levchenko, and R. Ranciére (2011): "Power laws in firm size and openness to trade: Measurement and implications," *Journal of International Economics*, 85, 42–52.

DuMouchel, W. H. (1983): "Estimating the Stable Index α in Order to Measure Tail Thickness: A Critique," *The Annals of Statistics*, 11, 1019–1031.

Fujiwara, Y., C. D. Guilmi, H. Aoyama, M. Gallegati, and W. Souma (2003): "Do Pareto-Zipf and Gibrat laws hold true? An analysis with European Firms," Quantitative Finance Papers cond-mat/0310061, arXiv.org.

Gabaix, X. (2009): "Power Laws in Economics and Finance," *Annual Review of Economics*, 1, 255–293, also available as NBER Working Paper n. 14299.

Gabaix, X. and R. Ibragimov (2011): "Rank-1/2: A Simple Way to Improve the OLS Estimation of Tail Exponents," *Journal of Business and Economic Statistics*, 29, 24–39.

Gabaix, X. and A. Landier (2008): "Why Has CEO Pay Increased So Much?" *The Quarterly Journal of Economics*, 123, 49–100.

Hall, P. (1982): "On Some Simple Estimates of an Exponent of Regular Variation," *Journal ofvthe Royal Statistical Society. Series B (Methodological)*, 44, 37–42.

Hall, P. and A. Welsh (1985): "Adaptive Estimates of Regular Variation," *The Annals of Statistics*, 13, 331–341.

Hill, B. (1975): "A Simple General Approach to Inference About the Tail of a Distribution," *The Annals of Statistics*, 3, 1163–1174.

Johnson, N. L., S. Kotz, and N. Balakrishnan (1994): *Continuous univariate distributions*, Wiley and Sons: New York.

Kleiber, C. and S. Kotz (2003): *Statistical Size Distributions in Economics and Actuarial Sciences*, Wiley and Sons Statistics: New York.

Luttmer, E. G. J. (2007): "Selection, growth and the size distribution of firms," *The Quarterly Journal of Economics*, 122, 1103–1144.

Mason, D. M. (1982): "Laws of Large Numbers for Sums of Extreme Values," *The Annals of Probability*, 10, 754–764.

Newman, M. (2005): "Power Laws, Pareto Distributions and Zipf's Law," *Contemporary Physics*, 46, 323–351.

Okuyama, K., M. Takayasu, and H. Takayasu (1999): "Zipf's law in income distribution of companies," *Physica A*, 269, 125–131.

Pareto, V. (1886): "Sur la courbe de la répartition de la richesse, Université de Lousanne," English translation: *Rivista di Politica Economica*, 87 (1997), 645-700.

Pictet, O. V., M. M. Dacorogna, and U. A. Muller (1998): "Hill, Bootstrap and Jacknife Estimator for Heavy Tails," in *A Practical Guide to Heavy Tails*, ed. by A. R. J., R. E. Feldman, and M. S. Taqqu, Birkhauser: Boston, 283–310.

Podobnik, B., D. Horvatic, A. M. Petersen, B. Urosevic, and E. H. Stanley (2010): "Bankruptcy risk model and empirical tests," *Proocedings of the National Academy of Sciences of USA*, 107, 18325–18330.

Resnick, S. and C. Starica (1997): "Smoothing the Hill estimator," *Advances in Applied Probability*, 29, 273–293.

——— (1998): "Tail Index Estimation for Dependent Data," *The Annals of Applied Probability*, 8, 1156–1183.

Silverman, B. (1986): *Density estimation for statistics and data analysis*, Chapman and Hall, London.

Weron, R. (2001): "Levy-stable distributions revisited: tail index > 2 does not exclude the Levy-stable regime," *International Journal of Modern Physics C*, 12, 209–223.

Zipf, G. K. (1932): *Selected Studies of the Principle of Relative Frequency in Language*, Librairie du Recuil Sirey, Paris.

6 Appendix

To keep comparability with Gabaix and Ibragimov (2011), in Table 4 and Table 5 we extend their analysis of AR(1) and MA(1) data to all the estimators considered in this article, although these types of time dependence are more compelling for applications in finance. We design the Monte Carlo exactly as in Gabaix and Ibragimov (2011).

For the case of AR(1) DGP, we generate R =10,000 random samples of size N =2,000 extracted from the AR(1) process

$$Y_i = \rho Y_{i-1} + \epsilon_i \quad , i \geq 1, \, Y_0 = 0 \quad , \tag{6.1}$$

with ϵ extracted from (4.1), for a given combination of the values of c and ρ. On each sample we apply all the estimators for two different tail width, i.e. including either the top-50 or the top-500 observations in the tails. We then repeat the Monte Carlo for different values of the parameters.

In Table 4 we report the average of point estimates across the 10,000 runs, together with asymptotic (theoretical) and sampled standard errors, as well as rejection rates of a t-test (at 5% level) of the true null of unitary tail index performed at each run. First consider the sensitivity to the AR(1) structure, setting aside the impact of the sub-asymptotic correction (i.e., set $c = 0$ and vary ρ), and take the case when the top-50 observations are considered. Although all the rejection rates are above the theoretical 5%, the results provide a clear ranking. First, the CDF and PDF estimators both severely over-reject. Second, among the other three estimators, the Rank$-1/2$ is over-performing the others. However, if we take the top-500 observations in the tail, then the frequency at which the true null of unitary tail index is mistakenly rejected rapidly grows to above 50% for all the estimators. Similar conclusions emerge when we let both c and ρ vary at the same time.

Table 5 replicates the analysis to study the properties under the MA(1) process

$$Y_i = \epsilon_i + \theta \epsilon_{i-1} \quad , i \geq 1 \quad , \tag{6.2}$$

with $\epsilon \sim$ (4.1). As before, we simulate R =10,000 random samples of size N =2,000 with varying c and θ, and again compare the behavior of the estimators for different tail width (top-50 and top-500 observations). The findings for $\theta = 0$ obviously replicate the analysis on AR(1) with $\rho = 0$. Further, if we switch off the sub-asymptotic correction (i.e. set $c = 0$, and vary θ), we observe that, first, the CDF and PDF estimators are once again unreliable, with very high rejection rates. Second, although rejection rates are above the theoretical 5% for all the methods, the Rank and Rank$-1/2$ estimators perform better (smaller rejection rates) than the other methods. The Rank performs slightly better if the tail includes the top-50 observations, while the Rank$-1/2$ is slightly better for the top-500 observations. Third, the patterns are similar when we let c and θ vary together. If anything, we notice that the rejection rates associated to all the estimators rapidly increase to above 20% if we take the top-500 observations in the tail. Conversely, they are less dependent from the parameters in the top-50 exercise.

Finally, in Table 6 and Table 7 we replicate the pooling exercise, experimenting with growth shocks drawn from a Laplace distribution with variance 0.05 and 0.10, respectively. Results show that the conclusions presented in the main text are entirely robust to these different set-ups.

Table 4: AR(1) data with sub-asymptotic deviation from Zipf Law

		Top 50					**Top 500**				
c	ρ	Hill	Rank−1/2	Rank	CDF	PDF	Hill	Rank−1/2	Rank	CDF	PDF
0.0	0.0	1.022 (0.07)	1.010 (0.05)	0.923 (0.12)	0.964 (0.68)	0.605 (0.80)	1.002 (0.05)	0.998 (0.05)	0.978 (0.08)	0.973 (0.66)	0.864 (0.50)
		(0.147) (0.149)	(0.202) (0.195)	(0.185) (0.182)	(0.055) (0.238)	(0.138) (0.124)	(0.045) (0.045)	(0.063) (0.063)	(0.062) (0.063)	(0.033) (0.131)	(0.070) (0.099)
0.0	0.5	1.119 (0.31)	1.174 (0.15)	1.077 (0.15)	1.170 (0.82)	0.718 (0.53)	1.163 (0.78)	1.124 (0.51)	1.102 (0.46)	1.111 (0.84)	1.000 (0.44)
		(0.161) (0.253)	(0.235) (0.321)	(0.215) (0.297)	(0.051) (0.355)	(0.142) (0.186)	(0.052) (0.100)	(0.071) (0.146)	(0.070) (0.145)	(0.032) (0.223)	(0.065) (0.162)
0.0	0.8	1.315 (0.58)	1.483 (0.41)	1.369 (0.35)	1.485 (0.94)	0.891 (0.43)	1.306 (0.89)	1.261 (0.74)	1.237 (0.72)	1.260 (0.88)	1.120 (0.65)
		(0.190) (0.452)	(0.297) (0.559)	(0.274) (0.515)	(0.079) (0.602)	(0.149) (0.334)	(0.059) (0.199)	(0.080) (0.268)	(0.078) (0.264)	(0.040) (0.348)	(0.071) (0.273)
0.5	0.0	1.046 (0.09)	1.024 (0.05)	0.935 (0.11)	0.974 (0.68)	0.615 (0.77)	1.157 (0.90)	1.084 (0.22)	1.061 (0.15)	1.014 (0.66)	0.915 (0.26)
		(0.151) (0.154)	(0.205) (0.200)	(0.187) (0.187)	(0.056) (0.244)	(0.141) (0.128)	(0.052) (0.056)	(0.069) (0.076)	(0.067) (0.075)	(0.036) (0.148)	(0.082) (0.118)
0.5	0.5	1.142 (0.33)	1.189 (0.16)	1.091 (0.15)	1.184 (0.83)	0.729 (0.51)	1.303 (0.97)	1.199 (0.67)	1.175 (0.63)	1.154 (0.84)	1.050 (0.46)
		(0.165) (0.264)	(0.238) (0.330)	(0.218) (0.305)	(0.051) (0.364)	(0.145) (0.192)	(0.058) (0.122)	(0.076) (0.169)	(0.074) (0.167)	(0.034) (0.245)	(0.075) (0.185)
0.5	0.8	1.342 (0.59)	1.507 (0.42)	1.390 (0.37)	1.508 (0.94)	0.910 (0.43)	1.442 (0.95)	1.339 (0.80)	1.313 (0.78)	1.315 (0.89)	1.186 (0.67)
		(0.194) (0.471)	(0.301) (0.577)	(0.278) (0.532)	(0.079) (0.621)	(0.152) (0.347)	(0.065) (0.236)	(0.085) (0.303)	(0.083) (0.298)	(0.042) (0.385)	(0.079) (0.308)
0.8	0.0	1.182 (0.28)	1.108 (0.05)	1.010 (0.08)	1.036 (0.68)	0.673 (0.57)	1.475 (1.00)	1.313 (0.94)	1.284 (0.91)	1.141 (0.70)	1.053 (0.24)
		(0.171) (0.184)	(0.222) (0.235)	(0.202) (0.219)	(0.062) (0.281)	(0.158) (0.153)	(0.066) (0.074)	(0.083) (0.110)	(0.081) (0.109)	(0.049) (0.201)	(0.112) (0.168)
0.8	0.5	1.255 (0.47)	1.268 (0.21)	1.162 (0.18)	1.251 (0.85)	0.786 (0.42)	1.699 (1.00)	1.442 (0.91)	1.410 (0.89)	1.298 (0.85)	1.205 (0.56)
		(0.181) (0.313)	(0.254) (0.376)	(0.232) (0.347)	(0.054) (0.408)	(0.159) (0.223)	(0.076) (0.176)	(0.091) (0.241)	(0.089) (0.237)	(0.043) (0.317)	(0.106) (0.259)
0.8	0.8	1.466 (0.66)	1.611 (0.48)	1.485 (0.42)	1.609 (0.95)	0.988 (0.42)	1.896 (1.00)	1.613 (0.90)	1.580 (0.89)	1.513 (0.89)	1.404 (0.72)
		(0.212) (0.556)	(0.322) (0.659)	(0.297) (0.607)	(0.082) (0.703)	(0.166) (0.397)	(0.085) (0.354)	(0.102) (0.428)	(0.100) (0.420)	(0.053) (0.514)	(0.107) (0.429)

Note: Estimates of tail index for the AR(1) process $Y_i = \rho Y_{i-1} + \epsilon_i$, with innovations ϵ_i following $P(X > x) = x^{-1}(1 + c(x^{-1} - 1))$, $x > 1$, $c \in [0, 1)$. Results over 10,000 Monte Carlo simulations with sample size $N = 2000$ and varying tail width (Top-50 vs. Top-500 observations), for different values of c and ρ. CDF and PDF estimates computed with 15 bins. For each combination: the first line reports point estimates of tail index averaged over the replications and, in parenthesis, the percentage of times the null of unitary tail index is rejected (at 5% confidence level); the second line shows, in parenthesis, the theoretical standard errors (usual OLS for CDF and PDF estimators; propagated via Taylor expansion of the asymptotic variance as in Gabaix and Ibragimov (2011) for the Rank−1/2 estimator; given in 2.4 for the Hill estimator) together with the sample standard errors.

Table 5: MA(1) data with sub-asymptotic deviation from Zipf Law

		Top 50					Top 500				
c	θ	Hill	Rank−1/2	Rank	CDF	PDF	Hill	Rank−1/2	Rank	CDF	PDF
0.0	0.0	1.022 (0.07)	1.010 (0.05)	0.923 (0.12)	0.964 (0.68)	0.605 (0.80)	1.002 (0.05)	0.998 (0.05)	0.978 (0.08)	0.973 (0.66)	0.864 (0.50)
		(0.147) (0.149)	(0.202) (0.195)	(0.185) (0.182)	(0.055) (0.238)	(0.138) (0.124)	(0.045) (0.045)	(0.063) (0.063)	(0.062) (0.063)	(0.033) (0.131)	(0.070) (0.099)
0.0	0.5	1.065 (0.17)	1.077 (0.11)	0.987 (0.15)	1.048 (0.74)	0.650 (0.69)	1.073 (0.40)	1.052 (0.21)	1.031 (0.18)	1.021 (0.73)	0.920 (0.38)
		(0.154) (0.202)	(0.215) (0.277)	(0.197) (0.257)	(0.056) (0.328)	(0.139) (0.164)	(0.048) (0.066)	(0.067) (0.095)	(0.065) (0.094)	(0.035) (0.175)	(0.068) (0.129)
0.0	0.8	1.075 (0.20)	1.077 (0.12)	0.988 (0.17)	0.998 (0.73)	0.598 (0.67)	1.078 (0.43)	1.055 (0.22)	1.034 (0.19)	0.999 (0.69)	0.893 (0.30)
		(0.155) (0.218)	(0.215) (0.292)	(0.198) (0.272)	(0.067) (0.340)	(0.164) (0.197)	(0.048) (0.068)	(0.067) (0.098)	(0.065) (0.097)	(0.038) (0.167)	(0.090) (0.130)
0.5	0.0	1.046 (0.09)	1.024 (0.05)	0.935 (0.11)	0.974 (0.68)	0.615 (0.77)	1.157 (0.90)	1.084 (0.22)	1.061 (0.15)	1.014 (0.66)	0.915 (0.26)
		(0.151) (0.154)	(0.205) (0.201)	(0.187) (0.187)	(0.056) (0.244)	(0.141) (0.128)	(0.052) (0.055)	(0.069) (0.076)	(0.067) (0.075)	(0.036) (0.148)	(0.082) (0.118)
0.5	0.5	1.087 (0.20)	1.091 (0.11)	0.999 (0.15)	1.059 (0.75)	0.660 (0.66)	1.221 (0.96)	1.131 (0.47)	1.108 (0.39)	1.062 (0.73)	0.968 (0.32)
		(0.157) (0.209)	(0.218) (0.284)	(0.200) (0.264)	(0.057) (0.336)	(0.141) (0.168)	(0.055) (0.080)	(0.072) (0.112)	(0.070) (0.111)	(0.038) (0.197)	(0.078) (0.149)
0.5	0.8	1.097 (0.22)	1.090 (0.13)	1.000 (0.17)	1.009 (0.72)	0.608 (0.64)	1.226 (0.96)	1.133 (0.48)	1.110 (0.40)	1.037 (0.68)	0.944 (0.21)
		(0.158) (0.225)	(0.218) (0.300)	(0.200) (0.279)	(0.068) (0.348)	(0.167) (0.202)	(0.055) (0.083)	(0.072) (0.115)	(0.070) (0.114)	(0.041) (0.188)	(0.102) (0.151)
0.8	0.0	1.182 (0.28)	1.108 (0.05)	1.010 (0.08)	1.036 (0.68)	0.673 (0.57)	1.475 (1.00)	1.313 (0.94)	1.284 (0.91)	1.141 (0.70)	1.053 (0.24)
		(0.171) (0.184)	(0.222) (0.235)	(0.202) (0.219)	(0.062) (0.281)	(0.158) (0.153)	(0.066) (0.074)	(0.083) (0.110)	(0.081) (0.109)	(0.049) (0.201)	(0.112) (0.168)
0.8	0.5	1.208 (0.36)	1.169 (0.14)	1.070 (0.15)	1.125 (0.77)	0.717 (0.52)	1.588 (1.00)	1.368 (0.92)	1.338 (0.89)	1.198 (0.75)	1.114 (0.40)
		(0.174) (0.246)	(0.234) (0.325)	(0.214) (0.301)	(0.060) (0.378)	(0.155) (0.196)	(0.071) (0.112)	(0.087) (0.164)	(0.085) (0.162)	(0.049) (0.265)	(0.109) (0.215)
0.8	0.8	1.216 (0.37)	1.167 (0.15)	1.070 (0.16)	1.075 (0.73)	0.669 (0.51)	1.599 (1.00)	1.371 (0.92)	1.341 (0.89)	1.169 (0.70)	1.097 (0.28)
		(0.175) (0.263)	(0.233) (0.343)	(0.214) (0.318)	(0.073) (0.394)	(0.181) (0.232)	(0.072) (0.117)	(0.087) (0.169)	(0.085) (0.167)	(0.056) (0.257)	(0.138) (0.218)

Note: Estimates of tail index for the MA(1) process $Y_i = \epsilon_i + \theta\epsilon_{i-1}$, with innovations ϵ_i following $P(X > x) = x^{-1}(1 + c(x^{-1} - 1))$, $x > 1$, $c \in [0, 1)$. Results over 10,000 Monte Carlo simulations with sample size $N = 2000$ and varying tail width (Top-50 vs. Top-500 observations), for different values of c and θ. CDF and PDF estimates computed with 15 bins. For each combination: the first line reports point estimates of tail index averaged over the replications and, in parenthesis, the percentage of times the null of unitary tail index is rejected (at 5% confidence level); the second line shows, in parenthesis, the theoretical standard errors (usual OLS for CDF and PDF estimators; propagated via Taylor expansion of the asymptotic variance as in Gabaix and Ibragimov (2011) for the Rank−1/2 estimator; given in 2.4 for the Hill estimator) together with sample standard errors.

Table 6: Robustness to pooling over time with Laplace growth shocks

α	β	Hill	Rank−1/2	Rank	CDF	PDF
0.5	0.8	0.630 (1.00)	0.603 (1.00)	0.591 (1.00)	0.578 (1.00)	1.519 (1.00)
		(0.028) (0.047)	(0.038) (0.065)	(0.037) (0.064)	(0.017) (0.103)	(0.041) (0.080)
	0.9	0.576 (1.00)	0.571 (1.00)	0.560 (1.00)	0.570 (1.00)	1.509 (1.00)
		(0.026) (0.048)	(0.036) (0.068)	(0.035) (0.067)	(0.018) (0.112)	(0.032) (0.080)
	1.0	0.505 (1.00)	0.505 (1.00)	0.495 (1.00)	0.483 (1.00)	1.386 (1.00)
		(0.023) (0.045)	(0.032) (0.063)	(0.031) (0.062)	(0.022) (0.091)	(0.054) (0.064)
0.7	0.8	0.882 (0.69)	0.845 (0.70)	0.827 (0.76)	0.809 (0.89)	1.727 (0.92)
		(0.040) (0.066)	(0.053) (0.091)	(0.052) (0.090)	(0.024) (0.145)	(0.057) (0.112)
	0.9	0.806 (0.93)	0.800 (0.83)	0.784 (0.87)	0.799 (0.86)	1.713 (0.95)
		(0.036) (0.068)	(0.051) (0.096)	(0.050) (0.094)	(0.026) (0.157)	(0.046) (0.113)
	1.0	0.708 (1.00)	0.707 (0.98)	0.694 (0.98)	0.676 (0.97)	1.545 (0.99)
		(0.032) (0.063)	(0.045) (0.088)	(0.044) (0.087)	(0.030) (0.127)	(0.077) (0.091)
0.9	0.8	1.134 (0.70)	1.087 (0.35)	1.064 (0.31)	1.040 (0.77)	1.935 (0.37)
		(0.051) (0.085)	(0.069) (0.117)	(0.067) (0.116)	(0.031) (0.186)	(0.074) (0.145)
	0.9	1.037 (0.33)	1.029 (0.30)	1.008 (0.30)	1.028 (0.78)	1.917 (0.51)
		(0.046) (0.087)	(0.065) (0.123)	(0.064) (0.121)	(0.033) (0.202)	(0.059) (0.145)
	1.0	0.910 (0.56)	0.909 (0.48)	0.892 (0.53)	0.869 (0.75)	1.706 (0.77)
		(0.041) (0.081)	(0.057) (0.113)	(0.056) (0.112)	(0.039) (0.163)	(0.100) (0.119)
1.0	0.8	1.260 (0.98)	1.207 (0.69)	1.182 (0.62)	1.156 (0.83)	2.039 (0.34)
		(0.056) (0.095)	(0.076) (0.130)	(0.075) (0.128)	(0.035) (0.207)	(0.082) (0.161)
	0.9	1.152 (0.74)	1.143 (0.51)	1.120 (0.46)	1.142 (0.84)	2.018 (0.44)
		(0.052) (0.097)	(0.072) (0.137)	(0.071) (0.135)	(0.037) (0.225)	(0.065) (0.161)
	1.0	1.011 (0.32)	1.010 (0.32)	0.991 (0.33)	0.966 (0.67)	1.787 (0.52)
		(0.045) (0.090)	(0.064) (0.125)	(0.063) (0.124)	(0.043) (0.182)	(0.111) (0.133)
1.1	0.8	1.386 (1.00)	1.328 (0.91)	1.301 (0.88)	1.272 (0.87)	2.143 (0.48)
		(0.062) (0.104)	(0.084) (0.143)	(0.082) (0.141)	(0.038) (0.228)	(0.091) (0.177)
	0.9	1.267 (0.97)	1.257 (0.78)	1.233 (0.73)	1.257 (0.88)	2.120 (0.54)
		(0.057) (0.107)	(0.080) (0.150)	(0.078) (0.148)	(0.041) (0.247)	(0.072) (0.177)
	1.0	1.112 (0.59)	1.111 (0.44)	1.090 (0.40)	1.062 (0.68)	1.868 (0.28)
		(0.050) (0.099)	(0.070) (0.138)	(0.069) (0.136)	(0.047) (0.200)	(0.123) (0.147)
1.3	0.8	1.639 (1.00)	1.570 (1.00)	1.537 (0.99)	1.503 (0.94)	2.351 (0.75)
		(0.073) (0.123)	(0.099) (0.169)	(0.097) (0.167)	(0.045) (0.269)	(0.107) (0.209)
	0.9	1.498 (1.00)	1.486 (0.98)	1.457 (0.97)	1.487 (0.93)	2.324 (0.79)
		(0.067) (0.126)	(0.094) (0.178)	(0.092) (0.175)	(0.048) (0.293)	(0.085) (0.209)
	1.0	1.314 (0.98)	1.313 (0.86)	1.289 (0.82)	1.255 (0.80)	2.033 (0.12)
		(0.059) (0.117)	(0.083) (0.163)	(0.081) (0.161)	(0.055) (0.237)	(0.145) (0.175)
1.5	0.8	1.892 (1.00)	1.812 (1.00)	1.774 (1.00)	1.735 (0.97)	2.559 (0.88)
		(0.085) (0.142)	(0.115) (0.195)	(0.112) (0.193)	(0.052) (0.311)	(0.123) (0.242)
	0.9	1.729 (1.00)	1.715 (1.00)	1.681 (1.00)	1.716 (0.96)	2.528 (0.90)
		(0.077) (0.145)	(0.108) (0.205)	(0.106) (0.202)	(0.056) (0.338)	(0.098) (0.242)
	1.0	1.516 (1.00)	1.515 (0.98)	1.487 (0.98)	1.449 (0.89)	2.198 (0.27)
		(0.068) (0.134)	(0.096) (0.188)	(0.094) (0.186)	(0.064) (0.274)	(0.167) (0.204)

Note: Estimates on samples obtained from a DGP following the Power Law in (2.1), then brought forward in time for 4 time periods according to (4.2) with Laplace shocks with variance $\sigma = 0.05$, and finally pooling the Top-500 observations in each time period. Results over 10,000 Monte Carlo simulations with sample size $N = 2000$, for varying tail exponent α and varying AR parameter β. CDF and PDF estimates computed with 15 bins. For each combination of α and β, the first line reports point estimates of tail index averaged over the replications and, in parenthesis, the percentage of times the null of unitary tail index is rejected (at 5% confidence level); the second line shows, in parenthesis, the theoretical standard errors (usual OLS for Rank, CDF and PDF estimators; propagated via Taylor expansion of the asymptotic variance as in Gabaix and Ibragimov (2011) for the Rank−1/2 estimator; and given in (2.4) for the Hill estimator) together with sample standard errors.

Table 7: Robustness to pooling over time with Laplace growth shocks

α	β	Hill	Rank−1/2	Rank	CDF	PDF
0.5	0.8	0.630 (1.00)	0.604 (1.00)	0.591 (1.00)	0.578 (1.00)	1.519 (1.00)
		(0.028) (0.047)	(0.038) (0.065)	(0.037) (0.064)	(0.017) (0.103)	(0.041) (0.080)
	0.9	0.576 (1.00)	0.571 (1.00)	0.560 (1.00)	0.571 (1.00)	1.509 (1.00)
		(0.026) (0.048)	(0.036) (0.068)	(0.035) (0.067)	(0.019) (0.112)	(0.033) (0.080)
	1.0	0.505 (1.00)	0.505 (1.00)	0.496 (1.00)	0.483 (1.00)	1.393 (1.00)
		(0.023) (0.045)	(0.032) (0.063)	(0.031) (0.062)	(0.021) (0.091)	(0.056) (0.067)
0.7	0.8	0.883 (0.69)	0.845 (0.70)	0.828 (0.76)	0.809 (0.88)	1.727 (0.92)
		(0.040) (0.066)	(0.053) (0.091)	(0.052) (0.090)	(0.024) (0.145)	(0.058) (0.113)
	0.9	0.807 (0.93)	0.800 (0.83)	0.784 (0.86)	0.801 (0.86)	1.713 (0.95)
		(0.036) (0.068)	(0.051) (0.096)	(0.050) (0.094)	(0.026) (0.158)	(0.046) (0.113)
	1.0	0.708 (1.00)	0.707 (0.98)	0.694 (0.98)	0.676 (0.97)	1.558 (0.99)
		(0.032) (0.063)	(0.045) (0.088)	(0.044) (0.087)	(0.030) (0.128)	(0.078) (0.095)
0.9	0.8	1.136 (0.70)	1.088 (0.35)	1.065 (0.31)	1.041 (0.77)	1.936 (0.37)
		(0.051) (0.086)	(0.069) (0.117)	(0.067) (0.116)	(0.031) (0.187)	(0.074) (0.145)
	0.9	1.037 (0.34)	1.029 (0.30)	1.009 (0.30)	1.030 (0.78)	1.917 (0.50)
		(0.046) (0.087)	(0.065) (0.123)	(0.064) (0.121)	(0.034) (0.203)	(0.059) (0.145)
	1.0	0.910 (0.56)	0.909 (0.47)	0.892 (0.53)	0.870 (0.75)	1.725 (0.72)
		(0.041) (0.080)	(0.057) (0.113)	(0.056) (0.111)	(0.038) (0.165)	(0.100) (0.125)
1.0	0.8	1.262 (0.98)	1.209 (0.69)	1.184 (0.63)	1.157 (0.83)	2.040 (0.34)
		(0.057) (0.095)	(0.076) (0.130)	(0.075) (0.129)	(0.035) (0.208)	(0.082) (0.162)
	0.9	1.153 (0.74)	1.143 (0.51)	1.121 (0.46)	1.146 (0.84)	2.019 (0.44)
		(0.052) (0.097)	(0.072) (0.137)	(0.071) (0.135)	(0.038) (0.226)	(0.066) (0.161)
	1.0	1.011 (0.32)	1.010 (0.32)	0.991 (0.33)	0.967 (0.68)	1.810 (0.45)
		(0.045) (0.089)	(0.064) (0.125)	(0.063) (0.124)	(0.043) (0.184)	(0.111) (0.140)
1.1	0.8	1.389 (1.00)	1.330 (0.92)	1.303 (0.88)	1.273 (0.87)	2.144 (0.48)
		(0.062) (0.105)	(0.084) (0.143)	(0.082) (0.142)	(0.038) (0.229)	(0.091) (0.178)
	0.9	1.268 (0.97)	1.258 (0.78)	1.233 (0.73)	1.261 (0.88)	2.120 (0.53)
		(0.057) (0.106)	(0.080) (0.150)	(0.078) (0.148)	(0.042) (0.248)	(0.073) (0.177)
	1.0	1.112 (0.58)	1.111 (0.44)	1.091 (0.40)	1.064 (0.69)	1.896 (0.24)
		(0.050) (0.098)	(0.070) (0.138)	(0.069) (0.136)	(0.047) (0.203)	(0.122) (0.155)
1.3	0.8	1.644 (1.00)	1.574 (1.00)	1.541 (0.99)	1.506 (0.94)	2.353 (0.75)
		(0.074) (0.124)	(0.100) (0.170)	(0.097) (0.168)	(0.045) (0.271)	(0.107) (0.211)
	0.9	1.499 (1.00)	1.487 (0.98)	1.458 (0.97)	1.491 (0.93)	2.324 (0.78)
		(0.067) (0.126)	(0.094) (0.178)	(0.092) (0.175)	(0.049) (0.294)	(0.087) (0.209)
	1.0	1.314 (0.99)	1.314 (0.86)	1.289 (0.82)	1.259 (0.81)	2.068 (0.17)
		(0.059) (0.115)	(0.083) (0.163)	(0.082) (0.161)	(0.055) (0.243)	(0.142) (0.185)
1.5	0.8	1.899 (1.00)	1.817 (1.00)	1.780 (1.00)	1.739 (0.97)	2.563 (0.89)
		(0.085) (0.144)	(0.115) (0.196)	(0.113) (0.194)	(0.052) (0.313)	(0.124) (0.243)
	0.9	1.731 (1.00)	1.717 (1.00)	1.683 (1.00)	1.721 (0.96)	2.528 (0.90)
		(0.078) (0.145)	(0.109) (0.205)	(0.106) (0.202)	(0.057) (0.339)	(0.102) (0.241)
	1.0	1.516 (1.00)	1.516 (0.99)	1.488 (0.98)	1.456 (0.89)	2.243 (0.39)
		(0.068) (0.132)	(0.096) (0.188)	(0.094) (0.185)	(0.064) (0.283)	(0.161) (0.216)

Note: Estimates on samples obtained from a DGP following the Power Law in (2.1), then brought forward in time for 4 time periods according to (4.2) with Laplace shocks with variance $\sigma = 0.1$, and finally pooling the Top-500 observations in each time period. Results over 10,000 Monte Carlo simulations with sample size $N = 2000$, for varying tail exponent α and varying AR parameter β. CDF and PDF estimates computed with 15 bins. For each combination of α and β, the first line reports point estimates of tail index averaged over the replications and, in parenthesis, the percentage of times the null of unitary tail index is rejected (at 5% confidence level); the second line shows, in parenthesis, the theoretical standard errors (usual OLS for Rank, CDF and PDF estimators; propagated via Taylor expansion of the asymptotic variance as in Gabaix and Ibragimov (2011) for the Rank−1/2 estimator; and given in (2.4) for the Hill estimator) together with sample standard errors.