

Scale-free networks well done

Ivan Voitalov,^{1,2,*} Pim van der Hoorn,^{1,2,3,†} Remco van der Hofstad,^{3,‡} and Dmitri Krioukov^{1,2,4,5,§}

¹Department of Physics, Northeastern University, Boston, Massachusetts 02115, USA

²Network Science Institute, Northeastern University, Boston, Massachusetts 02115, USA

³Department of Mathematics and Computer Science, Eindhoven University of Technology, Postbus 513, 5600 MB Eindhoven, Netherlands

⁴Department of Mathematics, Northeastern University, Boston, Massachusetts 02115, USA

⁵Department of Electrical and Computer Engineering, Northeastern University, Boston, Massachusetts 02115, USA



(Received 7 November 2018; revised manuscript received 26 May 2019; published 18 October 2019)

We bring rigor to the vibrant activity of detecting power laws in empirical degree distributions in real-world networks. We first provide a rigorous definition of power-law distributions, equivalent to the definition of regularly varying distributions that are widely used in statistics and other fields. This definition allows the distribution to deviate from a pure power law arbitrarily but without affecting the power-law tail exponent. We then identify three estimators of these exponents that are proven to be statistically consistent—that is, converging to the true value of the exponent for any regularly varying distribution—and that satisfy some additional niceness requirements. In contrast to estimators that are currently popular in network science, the estimators considered here are based on fundamental results in extreme value theory, and so are the proofs of their consistency. Finally, we apply these estimators to a representative collection of synthetic and real-world data. According to their estimates, real-world scale-free networks are definitely not as rare as one would conclude based on the popular but unrealistic assumption that real-world data come from power laws of pristine purity, void of noise, and deviations.

DOI: 10.1103/PhysRevResearch.1.033034

I. INTRODUCTION

Scale-free and *power-law* are sacral words in network science, a mature field that studies complex systems in nature and society by representing these systems as networks of interacting elements [1–4]. The most basic property of any network, second only to the network size and average degree, is the degree distribution, and the early days of network science were filled with the surprising and exciting news that degree distributions in many real-world networks of completely different origins are scale-free, i.e., “close to power laws.” This property means that the node degrees in a network are highly variable and lack a characteristic scale, with a multitude of profound and far-reaching implications for a wide spectrum of structural and dynamical properties of networks [1–8]. These implications are the reason why these scale-free findings were extremely impactful, and why they steered the whole field of network science in the direction it has followed for nearly two decades. They impacted essentially all the key aspects of

network science, from the basic tasks of network modeling, all the way down to concrete applications, such as prediction and control of the dynamics of real-world complex systems, or identifying their vulnerabilities [1–4].

Yet there is one glaring problem behind all these exciting developments. The problem is that scale-free networks do not have any widely agreed-upon rigorous definitions. Specifically, it is quite unclear what it really means for a degree sequence in a given real-world network to be power-law or “close” to a power law. This lack of rigor has led and still leads to confused controversy and never-ending heated debates [9–27]. This controversy has culminated in the recent work [20] that concluded that “scale-free networks are rare.” Here we arrive at quite a different conclusion based on a state-of-the-art statistical analysis and a more general definition of power laws.

Faced with the question whether a given real-world network is scale-free or not, one first has to decide how much the data can be trusted—how well does the measured degree sequence reflect the actual degree sequence in the network? We do not address this question here, and assume that we can trust the data. Under this assumption, the next questions are:

(1) What exactly does it mean that a distribution is approximately a power law?

(2) What are the correct, i.e., statistically consistent, methods to estimate the tail exponent of this power law from the measured degree sequence?

(3) How likely is it that the measured sequence comes from a power law with the estimated exponent?

Here we address all these three questions.

*i.voitalov@northeastern.edu

†w.l.f.v.d.hoorn@tue.nl

‡r.w.v.d.hofstad@tue.nl

§dima@northeastern.edu

One of the most frequently seen formula in the early days of network science was

$$P(k) \sim k^{-\gamma}. \quad (1)$$

It intended to say that the fraction $P(k)$ of nodes of degree k in a network under consideration decays with k approximately as a power law with exponent γ . The symbol “ \sim ” could mean anything, but usually its intended meaning was something like “roughly proportional.” The literature was also abundant with plots of empirical probability mass/density functions (PMFs/PDFs) $P(k)$ and complementary cumulative distribution functions (CCDFs) $\bar{F}(k)$ of degrees k drawn on the loglog scale to illustrate that these functions are “roughly straight lines,” so that the network is power-law, thus deserving a publication.

The first attempt to introduce some rigor into this vibrant activity, which became overwhelmingly popular in network science, came in [19], when network science was about a decade old. In Ref. [19], Eq. (1) was taken literally to mean that $P(k)$ for $k \geq k_{\min}$ is *exactly* proportional to $k^{-\gamma}$, i.e.,

$$P(k) = c k^{-\gamma}, \quad (2)$$

where c is the normalization constant.

But complexly mixed stochastic processes driving evolution of many different real-world networks are of different origins and nature. Worse, they all are prone to different types and magnitudes of noise and fluctuations. Therefore basic common sense suggests that these processes can hardly produce beautifully clean power-law dependencies void of any deviations from (2). This is similar to how one cannot expect Newton’s laws on Earth with friction to yield results as beautiful as Newton’s laws in the empty space without friction. That is why it is not surprising that if one looks for such idealized power-law dependencies in real-world networks, one is doomed to find them quite rare [20]. And as far as power-law network *models* are concerned, even the most basic such model, preferential attachment, is known to have a degree distribution with a power-law tail, but the exact expression for the degree distribution in preferential attachment networks is not a pure power law (2), as shown in Refs. [28–30]. In fact, power-law network models with the pristine purity of (2) are an exception rather than a rule.

For all these reasons, in statistics one considers the class of *regularly varying distributions* [31–34] instead of the pure power laws (2). Compared to the rather restrictive distribution class (2), the class of regularly varying distributions is much larger. In particular, it contains all the distributions whose PDFs are given by

$$P(k) = \ell(k)k^{-\gamma}, \quad (3)$$

thus allowing for deviations from pure power laws by means of a slowly varying function $\ell(k)$, i.e., a function that varies slowly at infinity, classic examples including functions converging to constants or $\ln^a k$ for any constant a . The exact definition of regularly varying distributions requires their CCDFs to be of the form

$$\bar{F}(k) := \sum_{k' > k} P(k') = \ell'(k)k^{-\alpha}, \quad (4)$$

where $\alpha = \gamma - 1$, and $\ell'(k)$ is also a slowly varying function. The class of distributions that satisfy (4) is even more general than (3): if (3) holds for a distribution, then so does (4), but not necessarily the other way around.

Compared to (2), any distribution in the class (4) has the same power-law tail exponent γ , but it can have drastically different shapes for finite degrees. The exact shape of $\ell(k)$ is of much less significance than the value of the tail exponent γ , because it is γ , and not $\ell(k)$, that is solely definitive for a number of important structural and dynamical properties of networks in the limit of large network size [5–8,35–41]. As the simplest example, the value of γ determines how many moments of the degree distribution remain bounded in the large-graph limit, affecting many important network properties. Yet we also note that some properties of finite-size networks may and usually do depend on a specific form of $\ell(k)$.

For all these reasons, and following the well-established tradition in statistics, in Sec. II we *define* a distribution to be *power-law* if it is *regularly varying*, i.e., if its CCDF satisfies (4).

The next question, that we address in Sec. III, is how to properly estimate the value of γ under the assumption that a given degree sequence comes from a regularly varying distribution. This question has attracted extensive research attention in probability, statistics, physics, engineering, and finance [31–34,42–52], where a variety of estimators have been developed for this task, all based on extreme value theory. We identify the maximal subset of such estimators that, to the best of our knowledge, are the only currently existing estimators that (1) are applicable to any regularly varying distribution; (2) are statistically consistent, i.e., have been proven to converge to the true γ , if applied to increasing-length sequences sampled from any regularly varying distribution; and (3) can be fully automated by the means of the double bootstrap method that has been proven to yield the optimal estimation of γ for any finite sequence of numbers sampled from any regularly varying distribution.

It is important to stress here that (2) is just one representative of the extremely wide class of regularly varying distributions (4). Therefore, as opposed to the methods in Refs. [19,20] that are consistent only under the assumption that a given degree sequence comes from a pure power law (2) above a certain minimal degree threshold, the estimators that we discuss in Sec. III are proven to be consistent under the much more general assumption that the sequence comes from any impure power law, including any distribution that satisfies (3) or even (4) with any nontrivial slowly varying functions $\ell(k)$ and $\ell'(k)$.

In Sec. IV, we evaluate these estimators by applying them to a wide range of synthetic sequences sampled from a variety of regularly varying distributions, as well as to degree sequences in paradigmatic network models—the configuration model, preferential attachment, and random hyperbolic graphs. In all the considered cases, all the considered estimators converge as expected. We also compare their performance to that of the PLFIT algorithm from Refs. [19,20], which is believed to represent the state of the art in network science. We find that PLFIT tends to show much worse performance when applied to distributions with nontrivial slowly varying

functions. Remarkably, one example of such nontrivial distributions is the degree distribution in the “harmonic oscillator” of power laws—the preferential attachment model.

The key strength behind the estimators considered in this paper is that most of them have been proven to be consistent not only under the assumption that the sampling distribution $P(k)$ is regularly varying, but also under the even more general assumption that it is *any* distribution belonging to the maximum domain of attraction of *any* extreme value distribution with *any* index ξ , which is the main parameter of an extreme value distribution. The extreme value distributions are the $n \rightarrow \infty$ limit distributions of rescaled maximum values among n samples from any given distribution $P(k)$. If $P(k)$ is regularly varying, then ξ is strictly positive, and the tail exponent γ and extreme value index ξ are related by

$$\xi = \frac{1}{\gamma - 1}. \quad (5)$$

If $P(k)$ is not regularly varying, then ξ is either negative or zero, in which case the tail exponent γ is undefined. None of the considered estimators estimates γ directly. They all are based on extreme value theory, and estimate the index ξ instead.

The last question from the list of the three questions above is about hypothesis testing. Given any degree sequence, one can always apply to it any ξ estimator that will always return some ξ estimate $\hat{\xi}$. How likely is it that this sequence comes from a regularly varying distribution with exponent $\gamma = 1 + 1/\hat{\xi}$? Clearly, if $\hat{\xi}$ is either negative or zero, then this question is ill-posed since one cannot even tell what the γ is. However, what if $\hat{\xi}$ is positive?

Section V is dedicated to the explanation that even in this case one cannot devise any hypothesis test to answer the above question. The popular p -value approach used often in hypothesis testing is deeply problematic and should be avoided, as has been long known and recently well documented in a statement article by the American Statistical Association [53], followed by a special issue of *The American Statistician* [54]. However, it is not that p -values are bad, and there is a better way. *Hypothesis testing is simply impossible with regularly varying distributions.* Intuitively, the main reason for this impossibility is the infinite number of “degrees of freedom” contained in the space of slowly varying functions $\ell(k)$ that make the space of regularly varying distributions nonparametric. In particular, there is an infinite number of regularly varying distributions such that for any *finite* sequence length, degree sequences of this length sampled from these distributions do not appear to be regularly varying, or the other way around, there is an infinite number of distributions that are not regularly varying, but such that random sequences of any *finite* length sampled from these distributions appear as regularly varying.

In view of this extremely important but badly misunderstood observation, which is one of the key points in this paper, the best strategy one can follow is to consult as many consistent γ estimators as possible to see whether they agree on the ranges of their γ estimates on a given sequence [31], and this is indeed the strategy we follow in Sec. V to *define* what it means for a given degree sequence to be power-law. If at least one of the considered estimators returns a negative

or zero value of $\hat{\xi}$, then we call the degree sequence *not power-law*, but if all the estimators agree that $\hat{\xi} > 1/4$, then we say that the sequence is *power-law*. If neither of these conditions are satisfied, then we call the degree sequence *hardly power-law*. The threshold $\hat{\xi} = 1/4$ between the power-law and hardly power-law ranges is completely arbitrary, and one is free to choose any nonnegative value of ξ for this threshold, determining the value of γ above which one can hardly call a network power-law. We chose this value to be $\gamma = 1 + 1/\hat{\xi} = 5$ for the reasons discussed in Sec. V.

Finally, in Sec. VI, we implement all the considered estimators in a software package [55] available to the public, and apply them to the degree sequences of 115 real-world networks with more than 1000 nodes collected from the KONECT database [56]. The collection contains many paradigmatic networks from different domains. Some of them were found to be power-law in the past (the Internet, for instance), while others were documented not to be power-law (road networks are a classic example). We find that the considered consistent estimators mostly agree with this classification, while overall, according to the definitions above, these estimators report that 49% of the considered undirected networks have degree sequences that are power-law. Among the considered directed networks, 24% have both in- and out-degree sequences that are power-law, while 82% have either in- or out-degree sequence that is power-law. The bipartite networks exhibit a similar picture according to the estimators: 35% of them have power-law degree sequences for both types of nodes, while in 74% of them at least one type of nodes has a power-law degree sequence.

In summary, if we relax the unrealistic requirement that degree distributions in real-world networks must be pure power laws, and allow for real-world impurity via regularly varying distributions, then upon the application of the state-of-the-art methods in statistics to detect such distributions in empirical data, we find that one can definitely not call scale-free networks “rare.”

II. POWER-LAW DISTRIBUTIONS

We *define* a distribution to be *power-law* if it is *regularly varying*. A distribution with PDF $P(k)$ is called *regularly varying* [32,33] if its CCDF

$$\bar{F}(k) := 1 - F(k) = \sum_{k' > k} P(k') \quad (6)$$

satisfies

$$\bar{F}(k) = \ell(k)k^{-\alpha}, \quad (7)$$

where $\alpha > 0$, and $\ell(k)$ is a slowly varying function. A function $\ell(x)$ is called *slowly varying* if

$$\lim_{x \rightarrow \infty} \frac{\ell(tx)}{\ell(x)} = 1 \quad (8)$$

for any $t > 0$. If the PDF of a distribution satisfies (3) with some slowly varying function, then the distribution is regularly varying, i.e., its CCDF satisfies (7) with some other slowly varying function. The converse may or may not be true, as discussed in Appendix A.

If a distribution is regularly varying, but its slowly varying function $\ell(k)$ in (7) does not vary at all, i.e., if it is constant, then we call such a distribution a *pure power law*. If k is integer-valued, $k = k_{\min}, k_{\min} + 1, \dots$, where k_{\min} is a natural number, then this pure power law is known as the generalized zeta distribution with PDF

$$P(k) = \frac{k^{-\gamma}}{\zeta(\gamma, k_{\min})}, \quad (9)$$

where γ is the PDF tail exponent, and $\zeta(\gamma, k_{\min})$ is the Hurwitz zeta function. If $k = x$ is real and $x \geq x_{\min} > 0$, then this pure power law is known as the Pareto distribution whose PDF is

$$P(x) = \alpha x_{\min}^{\alpha} x^{-\gamma}, \quad (10)$$

where $\alpha = \gamma - 1$. In both cases, the constant slowly varying functions are simply the normalization constants. Clearly, pure power laws form a small subset of general power laws, i.e., regularly varying distributions.

The definition of power-law distributions as regularly varying distributions formalize the point that the distribution exhibits a power-law tail at high degrees, but has an arbitrary shape at small degrees. They follow the well-established convention in probability, statistics, physics, engineering, and finance [31–34,42–52], where regularly varying distributions are the best studied subclass of much larger classes of distributions, such as *heavy-tailed* and others, see Appendix A.

We also note that the rigorous definition of regularly varying distributions in (7) perfectly formalizes the common traditional intuition behind the “ \sim ” sign in the nonrigorous “scale-free formula” (1). Indeed, if the regularly varying functions $\ln(ck)k^{-\alpha}$ and $Ck^{-\alpha}$, for example, are drawn on the loglog scale, one would see nothing but straight lines at large k in both cases, even though the first case is not a pure power law. This observation is formalized by Potter’s theorem [32, Theorem 1.5.6], stating that $\lim_{k \rightarrow \infty} \ell(k)k^{-\delta} = 0$ for any slowly varying function $\ell(k)$ and any $\delta > 0$. Therefore, in both cases, one would be tempted to write $\bar{F}(k) \sim k^{-\alpha}$, so that the power-law definition (7) is indeed a perfect way to hide any distributional peculiarities that do not asymptotically influence the power-law shape of the distribution tail.

We emphasize here that due to the nature of slowly varying functions, definition (7) is intrinsically asymptotic, dealing with the $k \rightarrow \infty$ limit. In particular, this implies that a distribution satisfying (7) can take any form for all degrees $k < K$ below an arbitrarily large but fixed threshold $K > 0$. This observation, and more generally, the asymptotic nature of power laws is the key factor responsible for the impossibility of hypothesis testing with regularly varying distributions, Sec. V.

The simplest and most frequently seen examples of regularly varying distributions can be found in Appendix A.

III. CONSISTENT ESTIMATORS OF THE TAIL EXPONENT

We now turn to the question of how to estimate the tail exponent of a regularly varying distribution given a finite collection of samples (e.g., node degrees) from it. We employ three estimators—Hill [57], moments [58], and kernel

[59]—that have been long proven to be statistically consistent at this task. *Consistency* means that as the number of samples increases, the estimated values of the exponent $\hat{\xi}$ are guaranteed to converge to the true exponent value ξ *regardless* of the slowly-varying function $\ell(k)$.

We note that all the considered estimators are consistent under the assumption that the data that they are applied to is a collection of i.i.d. (independent, identically distributed) samples from a regularly varying distribution. There is no, and cannot be any, hypothesis testing procedure that will tell whether a given sequence [of degrees in a (real-world) network] is an i.i.d. sequence from a regularly varying distribution, as we explain in detail in Sec. V. Therefore the application of these estimators to degree sequences of real-world networks can be justified only indirectly. In particular, their consistency has been recently proven for a wide range of preferential-attachment models, in which degree sequences are not exactly i.i.d. [60]. In case of the configuration model [61–63], it is known that a degree sequence sampled i.i.d.’ly from a distribution with finite variance is graphical with positive probability [5, Theorem 7.21]. This probability is very close to 1/2 for any distribution with a finite mean that takes odd values with positive probability, a surprising fact proven in [64]. This means that random graphs with a power-law degree distribution can be sampled by first sampling i.i.d.’ly a degree sequence from the distribution, and then constructing a graph with this degree sequence using known techniques [65]. Such a graph exists with nonzero probability because the degree sequence is graphical with this probability. Applied to graphs constructed this way, the estimators are consistent because the degree sequences in these graphs are i.i.d. Yet proving the consistency of these estimators applied to other network models is an open research area, which is only tangentially related to justifying their application to real-world networks, since there cannot be any “ultimately best” model for any real-world network. We also note that these estimators are actively employed in practice, in particular in financial mathematics [43,45,66–68], where regularly varying distributions are abundant, where the estimation of rare events is of key importance (e.g., for portfolio or fund management), and where the i.i.d. assumption cannot be checked to hold in real-world data either.

All the considered estimators do not estimate either the PDF or CCDF tail exponents γ or $\alpha = \gamma - 1$ directly. They are all based on extreme value theory, so that instead of estimating γ or α , they estimate the extreme value index ξ of the distribution. Given a sequence of n i.i.d. samples x_1, \dots, x_n from a distribution, extreme value theory is concerned with the behavior of the maximum value $m_n = \max_{i=1}^n x_i$ in this sample. In particular, one is typically interested in finding n -dependent constants c_n and d_n such that the distribution of $(m_n - d_n)/c_n$ has a nondegenerate limit. This limit distribution, if it exists, is called an *extreme value distribution (EVD)*, and the distribution of x ’s is then said to belong to the *maximum domain of attraction (MDA)* of this EVD. One of the key results in extreme value theory [69] is that there are only three families of EVDs. They are parameterized by a real number ξ , called the *extreme value index*. The three families are Weibull with $\xi < 0$, Gumbel with $\xi = 0$, and Fréchet with $\xi > 0$.

The reason why extreme value estimators are the standard tool in statistics to infer the tail exponent of regularly varying distribution, is the fundamental fact proven in Ref. [70]. It states that the class of all distributions that belong to the Fréchet MDA with $\xi > 0$ is exactly the class of all regularly varying distributions, i.e., those distributions whose CCDFs satisfy (7). Moreover, the PDF and CCDF tail exponents γ and α are related to the extreme value index ξ in this case by

$$\xi = \frac{1}{\gamma - 1} = \frac{1}{\alpha}. \quad (11)$$

It is this intimate relation between regularly varying distributions and extreme value theory that provides a rigorous and well-explored framework to analyze regularly varying distributions and make inferences concerning them.

We note that while the Hill estimator is consistent under the assumption that a given sequence is sampled only from a regularly varying distribution, i.e., that it is in the Fréchet MDA, the other considered estimators—that is, the moments and kernel estimators—are consistent for degree sequences sampled from *any* distribution belonging to the MDA of *any* extreme value distribution. This means that if these estimators are applied to increasing-length sequences sampled from distributions belonging to the Fréchet, Gumbel, or Weibull MDAs, then in all these three cases the estimates of these estimators are guaranteed to converge to the true values of ξ that are positive, zero, and negative, respectively. As a side note, while the Fréchet MDA is exactly all the regularly varying distributions, the Weibull MDA consists of distributions with upper-bounded supports, while the Gumbel MDA contains all other distributions that can be either light-tailed or heavy-tailed, but not regularly varying. Appendix B contains all the relevant details.

The key point here, which we rely upon in the next section, is that if the estimators, applied to a particular degree sequence, return either negative values of ξ , or values of ξ close to zero, then this sequence is quite unlikely to come from the Fréchet MDA, i.e., from a regularly varying distribution. Yet again, there is no way to quantify this unlikeliness rigorously, as explained in Sec. V.

Applied to n data samples x_1, x_2, \dots, x_n , the estimators operate by first sorting the data in nonincreasing order, $x_{(1)} \geq x_{(2)} \geq \dots \geq x_{(n)}$, and then limiting their consideration only to the κ largest data samples $x_{(1)}, x_{(2)}, \dots, x_{(\kappa)}$, where κ is a free parameter. Since the κ -th order statistic is a random variable representing the κ^{th} largest element among n i.i.d. samples from a distribution, the κ parameter is known as the *number of order statistics*. The estimators thus operate only on the κ tail of the empirical distribution represented by the κ order statistics. Given this tail, different estimators provide different expressions, documented in Appendix B, for the estimated value $\hat{\xi}_{\kappa,n}$ of ξ , which depends on κ . These expressions rely on different aspects of the order statistics contained in the tail, but all these expressions are consistent, meaning that

$$\hat{\xi}_{\kappa,n} \rightarrow \xi \quad \text{as } \kappa, n \rightarrow \infty, \kappa/n \rightarrow 0, \quad (12)$$

for all the estimators. The convergence above is usually in probability, although in some cases some stronger results, such as almost sure convergence or asymptotic normality, are available under additional assumptions on the data.

It is important to note here that in proving this convergence, the number of order statistics κ cannot be fixed, it must diverge with the number of samples n to incorporate more and more data in the tail, so that the estimated value of ξ is less and less affected by the fluctuations in the tail. Yet κ cannot be equal to n either, since in this case the estimated ξ would be affected by the slowly varying function $\ell(k)$. This implies that if applied to finite-size data samples, these estimators will not give a good estimate of ξ for either small or large values of κ . One option to deal with this problem in practice is to investigate the plot of $\hat{\xi}_{\kappa,n}$ as a function of κ in order to find the value of κ where this function is “most flat/constant.” This subjective approach can clearly not be rigorous. Worse, on real-world data, these functions can behave violently, see, for instance, the figures in Chap. 4 of Ref. [31] or in Ref. [71], so that finding such a flat region of $\hat{\xi}_{\kappa,n}$ may be quite problematic.

Fortunately, for the three estimators that we consider, the double bootstrap method documented in Appendix C is proven to find the optimal value κ^* of κ . *Optimality* means here that the error between the estimated and true values of ξ is minimized, Appendix C. The double bootstrap method is also proven not to break consistency, meaning that as a function of n , the value of κ_n^* diverges sublinearly, so that in view of (12), the estimated value of ξ , $\hat{\xi}_{\kappa_n^*,n}$, converges to the true ξ :

$$\hat{\xi}_{\kappa_n^*,n} \rightarrow \xi \quad \text{as } n \rightarrow \infty. \quad (13)$$

In addition to the Hill, moments, and kernel estimators, the Pickands estimator [72] and its generalized versions [73] are also often considered. However, only for one of these generalizations has the double bootstrap method been proven to be consistent, Appendix B. Worse, in application to real-world data, the Pickands estimator has been shown to be unstable and volatile [73,74] and to have poor efficiency [59,75]. Many other ξ estimators exist, see Ref. [76] for a review, but the proofs of consistency of the double bootstrap method are available only for the Hill, moments, kernel, and Pickands estimator.

Therefore, to the best of our knowledge, the Hill, moments, and kernel estimators are the maximal subset of consistent, stable, and efficient estimators, for which the double bootstrap method that automatically determines the optimal value of κ , is proven to be both optimal and consistent. The reason we consider not one but all such estimators is mentioned above: since as we explain in Sec. V there can be no hypothesis test to tell whether a given degree sequence is an i.i.d. sequence sampled from a regularly varying distribution, the best one can do is to consider as many consistent estimators as possible, testing as many different aspects of the degree sequence as possible, and see whether they agree in their estimations [31].

IV. EVALUATION OF ESTIMATOR PERFORMANCE

In Appendix D, we perform an in-depth evaluation of all the three estimators based on extreme value (EV) theory from the previous section. We apply them to a collection of random sequences sampled from various distributions, as well as to degree sequences in three popular network models—the configuration model, preferential attachment, and random hyperbolic graphs. We also juxtapose these validation

results against the performance of the PLFIT algorithm from Refs. [19,20].

The results of these experiments are as expected: all the EV estimators converge to the true value of ξ if the distribution is regularly varying, and they do not converge if it is not. They also converge even in the case where we sample not from a fixed regularly varying distribution, but from a sequence of distributions that are not regularly varying but that converge to a regularly varying distribution—the case with a Pareto distribution with the diverging natural exponential cutoff. On degree sequences in network models where individual degrees are not i.i.d. samples from a fixed degree distribution, the estimators converge as well, even though the i.i.d. assumption no longer holds.

For PLFIT, we find in Appendix D that if the sample distribution is sufficiently “nice,” then the estimation accuracy and convergence rates of the PLFIT are comparable to those of the EV estimators. However, in cases where the distribution is not so nice and is further from a pure power law, the EV estimators perform significantly better than the PLFIT. This is the case, for example, with distributions that can be fitted by power laws with wrong exponents in the region of small degrees. Remarkably, one example of such a distribution is the degree distribution in the preferential attachment model, a “harmonic oscillator” of power laws in network science [28–30]. For these and a number of other lower-level technical reasons, all documented in Appendix D and fully supported in a more recent and detailed focused study [24], we exclude the PLFIT from the subsequent considerations here.

V. POWER-LAW DEGREE SEQUENCES

There is no way to test the hypothesis that any given number was sampled from any given distribution that contains the number in its support. Yet if one has a long sequence of numbers, then there is a multitude of hypothesis testing procedures to measure how likely it is that this sequence was sampled from the distribution. The longer the sequence, the more reliable such procedures are, and any good procedure will give a definitive answer as the sequence length approaches infinity. This statistical methodology is widely known to work not only for a fixed distribution, but also for many parametric families of distributions. In the latter case, the testing involves one additional step: the parameters of the distribution are first to be estimated from the sequence using a consistent estimator.

A variation of this standard approach is at the core of Refs. [19,20], where the parametric family of distributions consists of pure power laws—the zeta or Pareto distributions. Their parameters, the tail exponents, are estimated using a combination of the likelihood maximization and Kolmogorov-Smirnov (KS) distance minimization techniques documented in Appendix D. Finally, the hypothesis testing procedure is the KS test, yielding a popular p -value number reflecting roughly how likely a given sequence comes from the pure power law with the estimated exponent.

We now come to the key point that this or any other hypothesis testing approach is not, and cannot be, applicable to regularly varying distributions, simply because these distributions do *not* form a parametric family of distributions. Instead, they are a nonparametric class of distributions of

an asymptotic nature with an uncountably infinite number of “degrees of freedom” contained in the slowly varying functions $\ell(k)$ (Appendix A). Testing whether a given *finite* collection of numbers was sampled from such an infinite-dimensional family of distributions is akin to testing whether a given number was sampled from a given distribution, which clearly is impossible as mentioned above.

Situations of this type are quite familiar for a physicist or network scientist. Phase transitions are a classic example: true phase transitions occur only in the thermodynamic limit, while for any finite system, we can only observe their signs. The simplest example in network science is graph sparsity. The definition of sparse graphs applies only to family of graphs whose size tends to infinity, and one cannot say anything at all about how sparse any given finite-size graph is, even if this is an empty graph of $n = 10^{10}$ nodes, simply because this graph can be considered as a typical Erdős-Rényi graph with the connection probability $p = 10^{-10}$, which is dense.

Yet, for a variety of reasons, these matters, including the impossibility of hypothesis testing with nonparametric families of distributions, as well as various consequences of this impossibility, are routinely overlooked and misunderstood. For these reasons, we first discuss the general picture behind this impossibility, and then illustrate it with a collection of examples.

First, the general picture is as follows. Recall that the consistency of an exponent estimator means that if we sample i.i.d. increasingly larger numbers n of random numbers k_i , $i = 1, \dots, n$, from a fixed regularly varying distribution with exponent γ and any slowly varying function $\ell(k)$, then the estimates $\hat{\gamma}_n$ that the estimator returns are guaranteed to converge to γ . Observe that while γ is a fixed number, $\hat{\gamma}_n$ is a random number, i.e., a random variable, because the k_i ’s are random. That is why one has to be careful with statements concerning in what particular sense the random variable $\hat{\gamma}_n$ converges to number γ . As stated in Sec. III, the convergence is usually in probability, but in some cases one can prove that $\hat{\gamma}_n$ converges to a normally distributed random variable with mean γ and some vanishing variance. For different definitions of convergence of random variables, we refer to any textbook on probability, such as Ref. [77].

It is crucially important to recognize that the convergence in probability does not mean that for any *finite* n there are any guarantees on how close the estimate $\hat{\gamma}_n$ will be to the true γ . To see why, observe that the slowly varying function $\ell(k)$ can be arbitrarily bad, breaking pure power laws for any arbitrarily large number of degree samples or range of degrees, while the true tail of the distribution can be inferred only in the limit of infinitely long sampled sequences, which one never has in practice.

We thus see that this general picture is very different from the one with hypothesis testing with a parametric family of distributions, such as the normal distributions or pure power laws. If we employ MLE, for instance, to estimate the parameters of such distributions, we usually know all we need to tell how close our estimates are expected to be to the true values for any given sample of size n . We often even know the full distribution of these estimates as random variables, and we then have the luxury to employ any reasonable hypothesis

test of our choice, or to compute p values to quantify chances if we wish. With regularly varying distributions, the situation is very different because if we do not know $\ell(k)$, we simply do not know how large n must be so that our estimators and hypothesis tests start showing any signs of convergence, simply because $\ell(k)$ can be arbitrarily bad.

To illustrate this extremely important point, we consider several examples next. The first two are of artificial/adversarial nature, while the last one is a well-studied network model.

The first example is a regularly varying distribution with support on $[c_1, \infty)$ and PDF with $\gamma > 1$ and constants $c_2 > c_1 \geq 0$ and $f \in [0, 1]$:

$$P(x) = \ell(x)x^{-\gamma}, \text{ where}$$

$$\ell(x) = \begin{cases} \frac{1-f}{c_2-c_1}x^\gamma, & \text{if } x \in [c_1, c_2], \\ (\gamma-1)fc_2^{\gamma-1}, & \text{if } x \in (c_2, \infty). \end{cases} \quad (14)$$

In words, this distribution is uniform on the interval $[c_1, c_2]$, and a pure power law (Pareto) for $x > c_2$. The parameter f is the fraction of the distribution mass falling within the Pareto region. This distribution is regularly varying for any constants $c_2 > c_1 \geq 0$, $f \in (0, 1]$, and $\gamma > 1$ because for $x > c_2$ the slowly varying function of its CCDF is constant, or in simpler terms, because it has a pure power law tail. However, if we sample $n < 1/f$ random numbers from this distribution, then there is no way to infer from these samples that the distribution is regularly varying with exponent γ because the expected number of samples in the Pareto region is below 1, so that all samples are expected to be from the uniform part of the distribution. Only if the number of samples n is sufficiently larger than $1/f$, can we expect to start seeing signs of the presence of a power-law tail. Figure 1 confirms that this is indeed the case. Clearly, one can replace the uniform part of the distribution with an arbitrary function, thus reflecting the reality of degree sequences observed in many real-world networks much more closely.

As another example, consider the Pareto distribution supported on $[1, \infty)$ with a fixed exponential cutoff at $c > 1$:

$$P(x) = \frac{c^{\gamma-1}}{\Gamma(1-\gamma, 1/c)}x^{-\gamma}e^{-x/c}, \quad (15)$$

where Γ denotes the upper incomplete gamma function. This distribution is not regularly varying, yet if our sample size n satisfies $n < c^{\gamma-1}$, then we will be tempted to conclude that the distribution is regularly varying, and that the exponent is γ , because almost all samples will be from the Pareto part of the distribution. Only if the number of samples n is sufficiently larger than $c^{\gamma-1}$, will we see signs of that this distribution does not really have any power-law tail, as confirmed in Fig. 1.

To see that such deceiving situations can occur in quite reasonable network models we refer to superlinear preferential attachment. In this model of growing networks, new nodes join a network one at a time, and connect to existing nodes of degree k with probability proportional to k^δ , where $\delta > 1$. For any such δ , the limit degree distribution is not regularly varying: the number of nodes with degrees exceeding a certain fixed threshold is finite [78]. Yet this threshold becomes larger if δ approaches 1. The threshold is also a growing function of the average degree \bar{k} , i.e., the number of links that new nodes

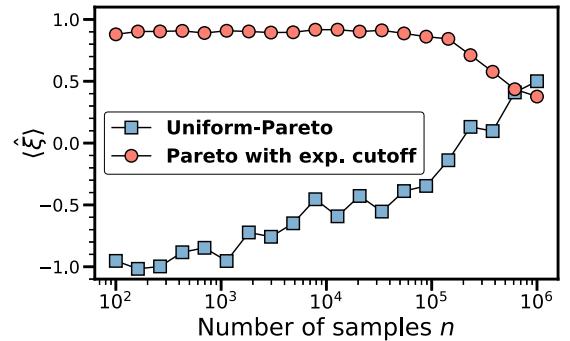


FIG. 1. The extreme value index estimates for the two adversarial examples in Sec. V. The sequences of varying length n are sampled from the distributions defined by Eq. (14) with $c_1 = 1$, $c_2 = 2$, $f = 2 \times 10^{-4}$, and $\gamma = 2$ (blue squares), and by Eq. (15) with $c = 500$ and $\gamma = 2$ (red circles). The data shown are the estimates of the moments estimator with the double bootstrap procedure applied to these samples. The results are averaged over 100 random sequences for each data point. In the case of blue squares, the distribution is regularly varying with $\gamma = 2$, so that $\xi = 1$. However, if n is not sufficiently larger than $1/f$, the sequences sampled from this distribution appear as if sampled from the uniform distribution that belongs to the Weibull MDA with $\xi = -1$. In the case of red circles, the distribution is not regularly varying. It belongs to the Gumbel MDA with $\xi = 0$. However, if n is not sufficiently larger than $c^{\gamma-1}$, the sequences sampled from this distribution appear as if sampled from a regularly varying distribution with $\gamma = 2$ and $\xi = 1$. The examples illustrate that for any finite n there is absolutely no way to tell what distribution class (regularly varying or not) the samples are coming from.

establish. More importantly, the larger this threshold, the more slowly the degree distribution approaches its limit, appearing as a reasonably “clean” power law in its vast preasymptotic regime. For example, for $\delta = 1.15$ and $\bar{k} = 4$, there are no noticeable deviations from this seemingly pure power-law behavior until the network size reaches about 10^{17} [78].

All these examples illustrate the point that based on any given *finite* degree sequence (of a real-world network), there is absolutely no way to tell how likely the hypothesis is that this sequence was sampled from a regularly varying distribution. In view of this impossibility, the best strategy one can follow is to simply rely on the estimates of the consistent estimators discussed in the previous section [31]. If the estimates of ξ that these estimators report on a given sequence are all positive, then it might be the case that this sequence comes from a regularly varying distribution. Yet if these estimates are negative or close to zero, then the chances of that are slim. However, *there is absolutely no, and cannot be any, rigorous way to quantify these chances, using p values or any other methods*, for the reasons above. This is the key point in our paper.

In view of these considerations, we take a conservative approach, and propose the following definition of a power-law degree sequences, based on the values of ξ that the three estimators from the previous section return upon their application to the sequence:

- (1) *A degree sequence is not power-law (NPL) if at least one estimator returns a negative or zero value of ξ ;*

(2) *A degree sequence is hardly power-law (HPL)* if all the estimators return positive values of ξ , and if at least one estimator returns a value of $\xi \leq 1/4$;

(3) *A degree sequence is power-law (PL)* if all the estimators return values of $\xi > 1/4$.

In purely intuitive and nonrigorous terms, the larger the ξ , the more likely it is that the degree sequence comes from a distribution with a power-law tail. These chances are the smaller, the closer the positive ξ is to zero, and we take a conservative approach to doubt that the degree sequence is power-law if $\xi \leq 1/4$. If $\xi \leq 0$, these chances are really slim. Unfortunately, as discussed above, it is principally impossible to attach any rigorous quantifiers to this intuition.

Yet we note that one important advantage of this classification scheme is that it tries to make a decision based on information from several estimators that are known to be consistent, instead of just one of unknown consistency. It is also possible to include other consistent estimators to collect more information about a degree sequence. We reiterate that we employ the Hill, moments, and kernel estimators here because they are the only three consistent estimators that are known to be stable on real-world data, and for which the double bootstrap procedure has been proven to be consistent.

We also note that the choice of the hardly power-law $\xi = 1/4$ threshold is completely arbitrary, and in view of the considerations above we should not have defined any *hardly* power-law regime, and call a degree sequence power-law if all ξ 's are positive. Yet if $\xi = 0.01$, for instance, then $\gamma = 101$. To call a degree sequence with such γ a power law is an unsatisfactory stretch of terminology. Another reason to define this threshold is that it is very difficult to tell whether a very small value $\hat{\xi} > 0$ that an estimator returns is an estimation of $\xi = 0$ or of a very small $\xi > 0$. In the latter case, the sequence was sampled from a regularly varying distribution, while in the latter case it was sampled from a distribution in the Gumbel MDA. This MDA consists of all kinds of distributions, including both light-tailed and heavy-tailed, but not regularly varying. The lognormal distribution, for example, is not regularly varying, but it is heavy-tailed and belongs to the Gumbel MDA, see Appendix B. Yet if the task is to tell whether a sequence was sampled from a regularly varying distribution or not, then classifying the sequence as regularly varying based on a small value $\hat{\xi}$ increases the chances of false positives because this small $\hat{\xi}$ may be an estimate of $\xi = 0$, in which case the source distribution is not regularly varying. To minimize the chances of such false positives, we do define the hardly power-law threshold $\xi = 1/4$, so that if at least one estimator thinks that $\xi \leq 1/4$, we doubt that the sequence is coming from a regularly varying distribution. We set this threshold to $\xi = 1/4$ here by selecting the largest value of γ that is known to us to still matter. That is, we are unaware of any value of γ that would correspond to any critical point, and that is larger than $\gamma = 1 + 1/\xi = 5$ in the Ising model on random graphs in Ref. [79].

Power-law degree sequences whose distributions have divergent second moments, meaning $\gamma < 3$ and $\xi > 1/2$, are of particular interest to network science for a variety of reasons. For example, networks with such degree sequences are particularly robust thanks to the absence of the percolation threshold [6], they are ultrasmall worlds versus small worlds [80], the

degree correlations in them are unavoidable due to structural constraints [81], etc. Therefore we also define a subclass of power-law degree sequences with divergent second moments of their distributions:

A power-law degree sequence has a divergent second moment (DSM) if all the estimators return values of $\xi > 1/2$.

We note that we do not put any restriction on how close to each other the estimated values reported by the different estimators must be in the definitions above. The main reason for that is that the speeds of convergence of these estimators are not known. They may converge to the true ξ at different rates. However, as discussed above, if the data size is relatively large, and all the estimators report values $\hat{\xi} > 1/4$, the chances that the degree sequence does not come from a regularly varying distribution ought to be slim. The power-law sequence definitions above represent one of many possible classification schemes. However, if a degree sequence is classified as PL or DSM according to this scheme, and if all the estimators report values that are close to each other, then one can be confident about the true values of ξ and γ . Unfortunately, there is no way to quantify this confidence. Since the convergence speeds are unknown, one cannot attach any rigorous bounds on how close the estimated values must be to yield any given accuracy in the estimation of the true γ . It is also important to recognize that these considerations apply not only to the estimators considered here, but also to any other estimator, including the PLFIT [19], whose convergence speed on regularly varying distributions is not known.

Finally, if a network is simple unweighted undirected unipartite single-layer and static, then it has only one degree sequence associated with it, so that it is straightforward to call such a network power law if its degree sequence is power law. However, in more complicated situations, such as directed, multipartite, multilayer, multiplex, and/or temporal networks, there are not one but many degree sequences associated with the network. To call such a network power law based only on one, or all, or some percentage (as in Ref. [20]) of the total number of its degree sequences, is purely a matter of taste. What usually does matter is a specific question, e.g., the spread of a disease, posed for the network, and different degree sequences, e.g., in versus out degrees, are of interest for different questions. Therefore we do not propose to classify such networks as power-law or not and, instead, report the data for each degree sequence separately in the next section.

VI. REAL-WORLD NETWORKS

Here we apply the Hill, moments, and kernel estimators to a collection of degree sequences in real-world networks from the KONECT database [56]. The database is a curated collection of real-world networks categorized by several network attributes such as size, (un)directedness, (un)weightedness, etc. The database uses a unified edge list format to store the data, which simplifies the automation of data processing. Better yet, the database allows one to sort networks by their properties, and to filter out networks with possibly incomplete information. This is in contrast to other large network

collections, such as ICON [82], that link their entries to third-party databases of various formats and origins, which makes it quite difficult to collect and process the data in an automated manner.

To streamline data processing, we do not consider networks in the database that are not downloadable in the KONECT edge list format. We also ignore temporal networks to avoid arbitrariness in selecting the temporal scale for data

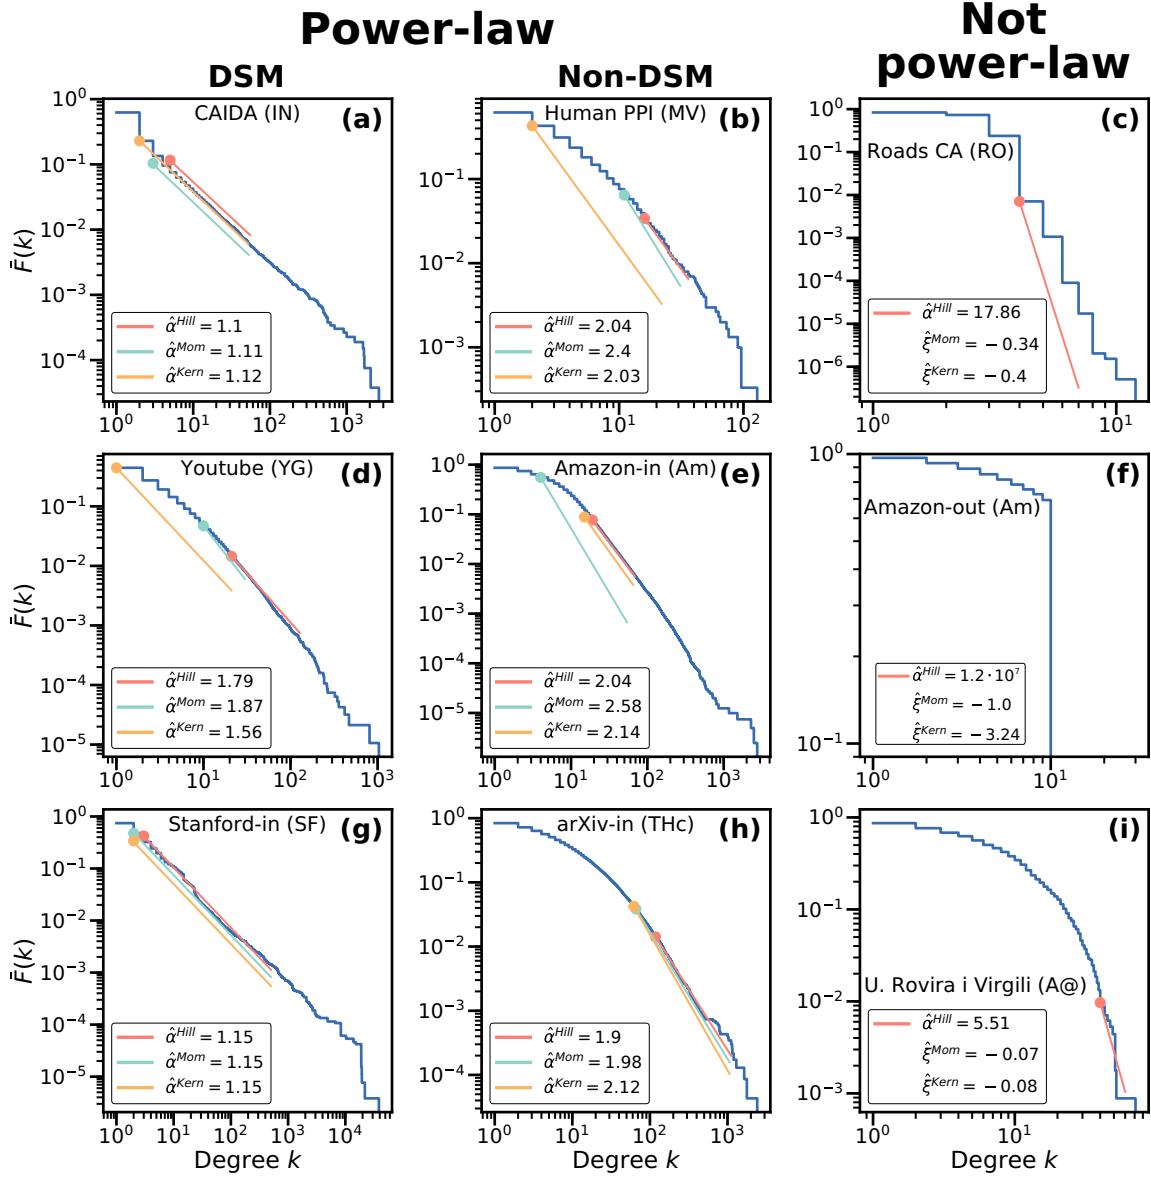


FIG. 2. The results of the Hill, moments, and kernel estimators applied to the degree sequences of nine real-world networks. The degree sequences belong to different classes defined in Sec. V. [(a), (d), and (g)] Power-law DSM: (a) CAIDA, the undirected network of the Internet at the autonomous system level; (d) Youtube, the user degree distribution of the bipartite network of Youtube users and their group memberships; (g) Stanford, the in-degree distribution of the directed network of hyperlinks between the WWW pages at the Stanford University website. [(b), (e), and (h)] Power-law non-DSM: (b) human PPI, the undirected network of human protein-protein interactions; (e) Amazon, the in-degree distribution of the directed network of product recommendations at Amazon; (h) arXiv, the in-degree distribution of the directed citation network of publications on high-energy physics theory at arXiv. [(c), (f), and (i)] Not power-law: (c) roads CA, the undirected network of road intersections in the state of California; (f) Amazon, the out-degree distribution of the same network as in (e); (i) U. Rovira i Virgili, the undirected email communication network at the University of Rovira i Virgili. The shown network names are their codenames used in the KONECT database [56], and they also appear in Tables I–III. Each panel shows the empirical complementary cumulative distribution functions (CCDFs) $\bar{F}(k)$ of the degree sequences on log-log scale. The straight lines visualize the estimated values of the CCDF exponents $\alpha = 1/\xi$. The filled circles are the optimal values of the number of order statistics κ^* found by the double bootstrap method. The estimators operate only on degrees larger than κ^* . The estimated values of α are $\hat{\alpha} = 1/\hat{\xi}(\kappa^*)$, where $\hat{\xi}(\kappa)$ is the estimated value of the tail index ξ as a function of κ . For nonpositive values of $\hat{\xi}(\kappa^*)$, the $\hat{\alpha}$ is undefined, so that the legends in (c), (f), and (i) show $\hat{\xi} = \hat{\xi}(\kappa^*)$ instead. Hardly power-law examples are not shown as they are not particularly interesting, lying somewhere in between power-law and not power-law examples.

TABLE I. The tail exponent estimation results for the 35 real-world undirected networks collected from the KONECT database [56]. Each network name is followed by its KONECT code in braces. The estimators return estimates $\hat{\xi}$ of ξ that are translated to $\hat{\gamma} = 1 + 1/\hat{\xi}$. If $\hat{\xi} \leq 0$, then $\hat{\gamma}$ is set to ∞ . The estimates are colored according to the definitions in Sec. V: (1) not power-law networks, at least one estimate is nonpositive $\hat{\xi} \leq 0$ (red); (2) hardly power-law networks, all the estimates are positive $\hat{\xi} > 0$ and at least one estimate is $\hat{\xi} \leq 1/4$ (yellow); (3) power-law networks with a divergent second moment, all the estimates are $\hat{\xi} > 1/2$ (green); and (4) other power-law networks, the rest of the cases (blue).

Network name	n	\bar{k}	$\hat{\gamma}_{\text{Hill}}$	$\hat{\gamma}_{\text{mom}}$	$\hat{\gamma}_{\text{kern}}$
CAIDA (IN)	26 475	4.03	2.1	2.11	2.11
Skitter (SK)	1 696 415	13.08	2.38	2.36	2.43
Actor collaborations (CL)	382 219	173.28	3.71	6.7×10^3	2.36
Amazon (CA)	334 863	5.53	3.99	3.48	3.44
arXiv (AP)	18 771	21.1	4.41	5.78	7.29
Bible names (MN)	1773	10.3	3.09	3.36	2.88
Brightkite (BK)	58 228	7.35	3.51	3.8	2.96
Catster (Sc)	149 684	72.8	2.09	2.06	1.98
Catster/Dogster (Scd)	623 748	50.33	2.1	2.11	2.04
Chicago roads (CR)	1467	1.77	77.92	∞	∞
DBLP (CD)	317 080	6.62	6.59	13.99	3.06
Dogster (Sd)	426 816	40.03	2.15	2.15	2.12
Douban (DB)	154 908	4.22	4.42	6.88	1.86
U. Rovira I Virgili (A@)	1133	9.62	6.49	∞	∞
Euro roads (ET)	1174	2.41	4.73	44.48	29.57
Flickr (LF)	1 715 254	18.13	3.94	4.29	5.02
Flickr (FI)	105 938	43.74	6.18	1.79	1.65
Flixster (FX)	2 523 386	6.28	53.63	1.93	1.95
Gowalla (GW)	196 591	9.67	2.8	2.8	2.86
Hamsterster (Shf)	1858	13.49	4.45	8.09	3.51
Hamsterster (Sh)	2426	13.71	4.57	25.39	6.32
Hyves (HY)	1 402 673	3.96	2.98	2.23	1.99
LiveJournal (Lj)	5 203 763	18.72	3.86	4.04	3.15
Livemocha (LM)	104 103	42.13	9.13	∞	2.39
Orkut (OR)	3 072 441	76.28	3.58	2.65	3.35
Power grid (UG)	4941	2.67	6.62	7.76	9.2
Proteins (Mp)	1846	2.39	3.09	3.31	3.87
Reactome (RC)	6229	46.93	4.86	34.33	∞
Roads CA (RO)	1 965 206	2.82	18.86	∞	∞
Roads PA (RD)	1 088 092	2.83	18.24	∞	∞
Roads TX (R1)	1 379 917	2.79	21.83	∞	∞
Route views (AS)	6474	3.88	2.13	2.16	2.14
WordNet (WO)	146 005	9.00	2.86	2.68	2.61
Youtube (CY)	1 134 890	5.27	2.48	2.58	2.17
Human PPI (MV)	3023	4.07	3.04	3.4	3.03

aggregation. Among database entries that possibly represent the same real-world network [for example, Wikipedia (EN) hyperlinks and Wikipedia (EN) links, both representing the English Wikipedia], we select only one entry. We also ignore networks that are marked as *incomplete* in the database. Finally, since the estimation of ξ cannot be reliable for networks of a small size, we only consider networks consisting of at least $n = 1000$ nodes.

The KONECT database contains not only undirected networks, but also directed and bipartite. For the latter two classes, we obtain not one, but two degree sequences for each network: the in- and out-degree sequences for directed networks, and one degree sequence for each of the two types of nodes in bipartite networks. We also remove all self-loops and multiedges from each collected network. After these filtering steps, we are left with 115 networks of three different types:

undirected (35), directed (49), and bipartite (31). The degree sequences of these networks are available at the software package repository [55].

We then feed the obtained degree sequences to the three estimators. Figure 2 shows the exponent estimation results that the estimators produce on some paradigmatic real-world networks in different domains, while Tables I, II, and III contain the full lists of these estimations for the undirected, directed, and bipartite networks, respectively. We see that many networks that were found to be power-law in the past have degree sequences that are classified as such by these estimators as well. These include the Internet, WWW, human protein interactions, social group memberships, citations, and product recommendation networks. The other way around, degree sequences of networks that are known not to be power-law are classified as not power-law—the California

TABLE II. The tail exponent estimation results for the 49 real-world directed networks collected from the KONECT database [56]. The style and notations are the same as in Table I. The estimators and coloring are applied to the in- and out-degree sequences separately. Since the operation of the estimators is undefined on zeros, the zero entries in the in- and out-degree sequences are removed, explaining the different lengths of the in- and out-degree sequences n_{in} and n_{out} for the same network, as well as the different values of the in and out average degrees \bar{k}_{in} and \bar{k}_{out} .

Network name	n_{in}	n_{out}	\bar{k}_{in}	\bar{k}_{out}	$\hat{\gamma}_{\text{in}}^{\text{Hill}}$	$\hat{\gamma}_{\text{in}}^{\text{mom}}$	$\hat{\gamma}_{\text{in}}^{\text{kern}}$	$\hat{\gamma}_{\text{out}}^{\text{Hill}}$	$\hat{\gamma}_{\text{out}}^{\text{mom}}$	$\hat{\gamma}_{\text{out}}^{\text{kern}}$
Adolescent health (ME)	2441	2313	5.31	5.61	6.68	∞	∞	501.0	∞	∞
Advogato (AD)	4422	4009	10.66	11.76	3.16	3.35	3.0	3.29	3.43	4.37
Air traffic control (TC)	1210	1083	2.16	2.41	6.0	∞	7.58	4.28	23.22	8.3
Amazon (Am)	403 312	402 439	8.4	8.42	3.04	3.58	3.14	1.2×10^7	∞	∞
arXiv (PHc)	28 226	32 153	14.93	13.11	3.51	4.16	4.18	3.79	4.89	6.46
arXiv (THc)	23 176	25 055	15.22	14.08	2.9	2.98	3.11	4.09	9.13	3.84
Baidu (BAi)	1 241 374	1 654 404	14.33	10.75	2.47	2.51	2.57	4.58	5.22	5.57
Baidu (BAr)	277 991	394 482	11.81	8.33	2.72	2.79	3.26	19.52	2.73	2.59
Berkley Stanford (BS)	617 094	680 486	12.32	11.17	2.09	2.08	2.06	334.33	3.25	3.36
Blogs (Mg)	990	1064	19.21	17.88	3.27	4.11	6.1	5.15	14.7	7.21
CiteSeer (CS)	194 959	336 024	8.95	5.19	3.46	3.84	3.31	4.08	5.55	4.23
Cora (CC)	13 879	21 201	6.59	4.32	3.03	3.2	2.95	5.1	4.94	6.68
DBLP (Pi)	11 564	3158	4.3	15.75	3.1	3.4	2.98	3.42	2.81	2.84
Edinburgh thesaurus (EA)	22 675	8210	13.75	37.98	6.13	3.94	4.97	30.41	∞	∞
Ego Google Plus (GP)	23 591	131	1.66	199.56	5.83	4.21	4.83	3.91	11.75	∞
Ego Twitter (TL)	22 964	938	1.44	35.29	4.46	3.4	3.79	7.45	∞	48.62
Epinions (ES)	51 957	60 341	9.79	8.43	3.43	4.38	3.62	3.54	3.77	3.98
FOLDOC (FO)	13 309	13 356	9.03	9.0	3.16	2.67	2.5	5.42	7.9	10.01
Gnutella (GN)	62 283	16 387	2.37	9.02	5.85	∞	3.67	8.14	3.15	2.3
Google (GO)	714 545	739 454	7.14	6.9	2.5	2.55	2.68	5.1	3.34	2.97
Google (GC)	15 762	12 447	10.81	13.68	2.16	2.17	2.04	3.12	2.71	2.23
Hudong (HUi)	798 202	1 725 741	18.39	8.51	3.23	3.65	2.34	3.72	2.53	2.63
Hudong (HUr)	991 745	2 232 238	19.01	8.45	3.29	3.46	3.07	126.0	3.62	2.64
JDK dependencies (DJ)	2375	6369	63.57	23.71	2.07	2.07	1.71	10.62	2.36	2.52
JUNG/JAVAX dependencies (Dj)	2208	6055	62.82	22.91	2.0	2.07	1.71	10.62	2.4	2.54
Libimseti (LI)	168 791	135 359	102.85	128.25	4.21	2.57	2.67	2.54	2.49	2.56
Linux (LX)	12 013	25 619	17.77	8.33	1.97	1.96	2.03	4.42	11.75	5.39
Notre Dame (ND)	325 729	136 934	4.51	10.73	2.05	2.57	1.99	2.55	2.73	2.26
Open flights (OF)	3418	3409	19.8	19.85	6.38	3.63	1.84	6.21	4.72	1.84
Pokec (PL)	1 519 452	1 432 693	20.15	21.37	4.97	6.95	7.76	4.1	4.94	6.52
Slashdot Zoo (SZ)	65 220	45 598	7.9	11.3	2.99	3.34	4.02	29.57	∞	∞
Standford (SF)	261 588	281 731	8.84	8.21	2.15	2.15	2.15	63.5	5.55	3.09
TREC WT10g (WT)	1 295 841	1 532 051	6.22	5.26	2.98	3.67	2.53	2.19	2.17	2.27
Twitter ICWSM (Ws)	465 016	2502	1.8	333.65	2.54	2.57	2.64	1.6×10^4	∞	∞
Twitter MPI (TF)	49 395 940	43 983 853	39.75	44.64	2.02	2.31	2.01	1.91	1.99	2.34
Twitter WWW (TW)	35 689 148	40 103 281	41.14	36.61	1.93	2.01	2.43	2.04	2.03	1.98
US airports (AF)	1504	1478	18.77	19.1	13.2	∞	2.46	14.89	∞	2.26
US patents (PC)	3 258 983	2 089 345	5.07	7.91	4.28	4.6	4.52	3.43	3.6	3.82
Wikipedia links DE (Wde)	2 262 745	3 221 527	36.06	25.33	2.42	2.4	2.15	5.17	4.25	3.22
Wikipedia links EN (Wen)	7 549 312	12 114 964	50.08	31.21	2.78	2.79	3.03	4.94	8.58	5.9
Wikipedia links FR (Wfr)	2 127 693	2 993 436	48.11	34.2	2.4	2.37	2.54	4.0	5.24	3.61
Wikipedia links IT (Wit)	1 488 860	1 855 986	61.47	49.31	2.88	2.76	3.17	6.81	5.46	2.93
Wikipedia links JA (Wja)	1 253 659	1 609 718	56.67	44.14	2.48	2.47	2.68	5.37	3.73	4.27
Wikipedia links PL (Wpl)	1 196 546	1 528 795	48.04	37.6	2.59	2.59	2.95	3.82	4.46	4.12
Wikipedia links PT (Wpt)	1 137 929	1 591 426	43.07	30.8	2.58	2.51	2.72	14.7	4.03	4.5
Wikipedia links RU (Wru)	1 834 424	2 852 544	44.72	28.76	2.48	2.46	2.62	7.8	251.0	4.89
Yahoo ads (YD)	194 317	653 260	15.09	4.49	2.2	2.18	2.23	4.02	3.98	∞
Human PPI (MF)	2033	338	3.17	19.09	27.32	∞	∞	2.4	2.86	2.04
Human PPI (MS)	1698	1597	3.63	3.86	4.62	2.44	2.57	6.75	2.42	2.63

road network or the out-degree distribution in the directed network of Amazon product recommendations, for instance.

We emphasize again the importance of using as many consistent estimators as possible: on any finite degree

sequence, different estimators are not guaranteed to return the same ξ estimation, as they may explore different parts of the distribution, especially if the slowly varying function $\ell(k)$ is nontrivial, Appendix C. That is why we use the maximal

TABLE III. The tail exponent estimation results for the 31 real-world bipartite networks collected from the KONECT database [56]. The style and notations are the same as in Table II. The estimators and coloring are applied to the degree sequences of nodes of types 1 and 2 (domains 1 and 2) separately.

Network name	n_{d_1}	n_{d_2}	\bar{k}_{d_1}	\bar{k}_{d_2}	$\hat{\gamma}_{d_1}^{\text{Hill}}$	$\hat{\gamma}_{d_1}^{\text{mom}}$	$\hat{\gamma}_{d_1}^{\text{kern}}$	$\hat{\gamma}_{d_2}^{\text{Hill}}$	$\hat{\gamma}_{d_2}^{\text{mom}}$	$\hat{\gamma}_{d_2}^{\text{kern}}$
Movies/Actors (AM)	383 640	127 823	3.83	11.5	4.92	7.67	5.24	6.13	11.99	6.13
arXiv (AC)	22 015	16 726	2.66	3.5	9.33	12.9	13.99	5.33	6.05	6.75
Book Crossing implicit (BX)	340 523	105 278	3.38	10.92	3.54	2.26	2.31	3.08	3.37	1.88
Book Crossing ratings (Bx)	185 955	77 802	2.33	5.57	2.35	2.39	2.43	2.83	2.82	2.86
Countries DBPedia (CN)	2302	590 112	276.77	1.08	1.43	1.44	1.93	9.77	3.03	∞
DBLP (PA)	4 000 150	1 425 813	6.07	2.16	5.22	4.3	6.75	2.22	2.23	2.14
Discogs labels/artists (DI)	270 771	1 754 823	53.24	8.21	2.08	2.15	1.93	3.12	3.11	3.29
Discogs genres/artists (Da)	15	1 754 823	1.3×10^6	10.85	7.67	2.72	1.0	3.12	3.07	3.21
Discogs genres/labels (Dr)	15	270 771	2.8×10^5	15.32	1.9	2.56	7.76	2.2	2.2	2.17
Flickr (FG)	103 631	395 979	82.46	21.58	2.66	3.7	2.85	5.22	6.52	6.29
Genres DBPedia (GE)	7783	258 769	59.55	1.79	2.16	2.57	∞	5.61	12.49	∞
Github (GH)	120 867	56 519	3.64	7.79	2.08	2.07	2.07	2.77	4.07	3.31
Jester100 (J1)	100	73 421	4.1×10^4	56.34	4.3×10^4	∞	∞	5.1×10^5	∞	∞
Jester150 (J2)	140	50 692	1.2×10^4	34.1	3.26	3.32	5.46	251.0	3.29	2.73
LiveJournal (LG)	7 489 073	3 201 203	15.0	35.08	1.78	1.8	1.8	1.0×10^3	∞	4.32
Locations DBPedia (LO)	53 407	172 079	5.5	1.71	2.02	2.02	2.05	6.52	3.9	38.04
Movies DBPedia (ST)	81 085	76 098	3.47	3.7	3.94	4.85	5.98	5.52	13.35	8.69
Occupations DBPedia (OC)	101 730	127 571	2.47	1.97	1.73	1.74	1.73	6.56	19.18	1.0×10^3
Orkut (OG)	8 730 857	2 783 196	37.46	117.5	1.88	1.89	1.89	53.63	2.4	2.63
Producers DBPedia (PR)	138 839	48 833	1.49	4.24	3.69	3.89	4.13	4.09	4.58	2.12
Labels DBPedia (RL)	18 421	168 268	12.66	1.39	2.05	2.08	2.09	5.95	5.18	6.05
Reuters (RE)	283 911	781 265	213.34	77.53	6.38	1.55	1.56	6.46	∞	∞
TREC (TR)	1 173 225	551 787	71.28	151.56	1.61	1.63	1.6	3.93	8.87	2.94
TV tropes DBPedia (DBT)	87 678	64 415	36.86	50.18	3.18	3.73	3.37	3.34	3.07	2.97
Teams DBPedia (TM)	34 461	901 130	39.65	1.52	8.46	12.76	2.2	21.83	∞	∞
vi.sualize.us images/tags (Vti)	495 402	82 035	4.64	28.02	3.15	3.34	2.97	1.81	1.85	1.86
vi.sualize.us tags/users (Vut)	82 035	17 122	28.02	134.26	1.81	1.85	1.85	2.54	2.57	2.13
vi.sualize.us images/users (Vui)	495 402	17 122	4.64	134.26	3.15	3.31	2.98	2.49	2.58	2.51
Web trackers (WT)	12 756 244	27 665 730	11.02	5.08	2.03	2.03	2.04	2.62	2.1	2.05
Writers DBPedia (WR)	46 213	89 355	3.12	1.62	3.77	4.1	4.06	5.88	7.62	15.08
Youtube (YG)	30 087	94 238	9.75	3.11	2.31	2.36	2.45	2.79	2.87	2.56

subset of stable and efficient estimators for which the double bootstrap method to determine the optimal number of order statistics κ^* is proven to be consistent.

Finally, Fig. 3 summarizes the estimation results for $\hat{\gamma} = 1 + 1/\hat{\xi}$ in Tables I–III by classifying the degree sequences of all the considered networks into the not power-law (NPL), hardly power-law (HPL), and power-law (PL) classes, the latter containing the subclass of power-law networks with divergent second moments (DSM), defined in the previous section. We see that the percentages of power-law and DSM degree sequences in undirected networks are 49% and 29%, respectively. Among the considered directed networks, 24% and 6% have both in- and out-degree sequences that are power-law and DSM, while 82% and 45% of these networks have either in- or out-degree sequence which is power-law and DSM, with a majority of those being in-degree sequences. The bipartite networks exhibit a similar picture: 35% and 13% of them are power-law and DSM according to both types of nodes, while 74% and 55% are power-law and DSM according to at least one type of nodes.

While one cannot directly compare these results to the ones in Ref. [20], they present quite a different picture than painted there.

VII. CONCLUSION AND DISCUSSION

In summary, we call a distribution power-law if it is regularly varying. The pure power laws—the Pareto and zeta distributions—are a small subset of this more general, realistic, and well-studied class of distributions. This class constitutes the most inclusive theoretical framework capable of formalizing all the aspects of the “*straight line on log-log scale*” intuition behind power-law observations in real-world networks. Utilizing the connection between this class of distributions and the maximum domain of attraction of the Fréchet distribution in extreme value theory, we identify state-of-the-art statistical tools to estimate the tail exponent γ in a given degree sequence. These are then deployed to design a classification scheme for degree sequences. The application of this scheme to a representative collection of degree sequences in real-world networks reveals that significant fractions of these networks have power-law degree sequences.

We note that the problem of classifying a given degree sequence as power-law or not has nothing to do with possible mechanisms that may lead to the emergence of power-law distributions in real data, and that are of great interest to network science in general. The reason why such mechanisms

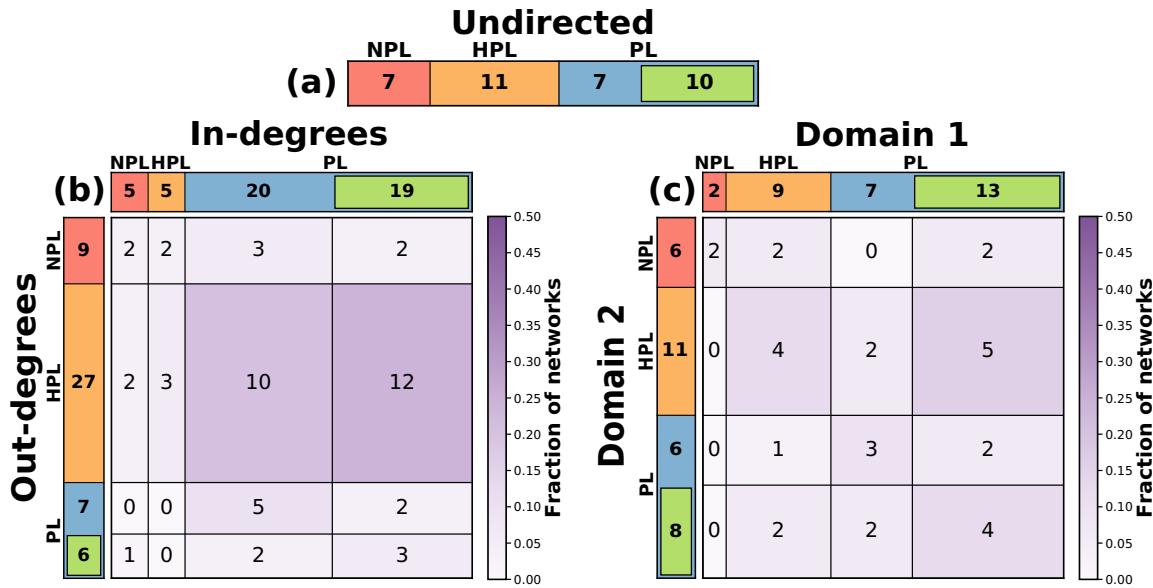


FIG. 3. The breakdown of the degree sequences of the considered real-world networks into the three classes defined in Sec. V: not power-law (NPL), hardly power-law (HPL), and power-law (PL), the latter containing the divergent second moment (DSM) subclass, shown in green. The full data appear in Tables I–III. The data are shown for (a) undirected, (b) directed, and (c) bipartite networks. The numbers in the boxes are the numbers of networks falling within the corresponding (sub)class. In directed networks, the in- and out-degree sequences are classified separately. Similarly, in bipartite networks, the degree sequences of nodes of types 1 and 2 (domains 1 and 2) are classified separately as well. The numbers of networks with all possible combinations of the two classifications are then shown in the squares, along with the marginals outside of the squares. The color shading in the squares shows the fraction of directed and bipartite networks in each category.

are a completely different subject altogether is simple. We can think of different mechanisms as different network models approximating stochastic processes that drive the evolution of real-world networks, and it is quite well known that completely different network models and thus completely different network formation mechanisms may lead to networks that have exactly the same degree distribution. That is, these networks may certainly be very different in all respects other than the degree distribution [83]. Therefore the question of what mechanism causes this or that degree distribution is completely irrelevant and ill-posed, as it is impossible in principle to infer it based only on the degree distribution.

The impossibility of hypothesis testing for regularly varying distributions is the reason why one cannot attach any statistical weight, such as a p value, to the statement that a given finite sequence is regularly varying or not. Yet many other aspects of the current state of affairs in statistics related to detecting power laws in empirical network data do allow for improvement, so we comment on some of them here.

Fundamental limitations of estimators based on extreme value theory. The existing consistent estimators of tail exponents are based on extreme value (EV) theory. These estimators cannot generally differentiate between heavy-tailed and light-tailed distributions, simply because the maximum domain of attraction of the Gumbel EV distribution contains distributions of both types—the light-tailed normal and heavy-tailed lognormal distributions, for example, Appendix B. Since for many applications in network science an important question is whether a degree distribution is heavy- or light-tailed, versus regularly varying or not, it is of

particular interest to devise other estimators, not based on EV theory, that would be capable of differentiating between these two types of distributions. Some initial steps in this direction have recently been made [84]. Even more generally, it is often of interest whether a given degree sequence comes from a distribution with an infinite or finite second moment, versus power-law or not, so that it would be desirable to develop statistically consistent methods to test the infiniteness of the second moment. Such tests cannot be based on EV theory either.

Yet even for EV-based estimators there are many paths to improve their applicability and rigorous guarantees, which we discuss next.

The i.i.d. assumption. First, it would be nice to relax the i.i.d. assumption for these estimators, and to prove their consistency in application to network models. The first step in this direction was made in Ref. [60]. We saw in our experiments in Appendix D that all the considered estimators converge in all the considered network models, but there are no proofs for this convergence for any network model other than preferential attachment, to the best of our knowledge.

Convergence speed. Another important open problem is the convergence speed. All we currently know is that the considered estimators converge to the true value of the power-law exponent γ on sequences of random numbers of increasing length n sampled i.i.d. from any regularly varying distribution with this γ , but we do not know how quickly this convergence occurs, so that, for instance, there is no way to tell how close the estimates of different estimators on the same finite- n sequence are supposed to be, even if this sequence

is sampled i.i.d.'ly from a regularly varying distribution. The speed of this convergence depends not only on γ but also on the slowly varying function $\ell(k)$. Thus the problem is to obtain bounds, as functions of n , on the error of estimation of γ for a given γ and $\ell(k)$. Can such bounds be obtained for certain classes of $\ell(k)$'s?

Not one sequence but many sequences. More pertinent to networks, and also closely related to the convergence speed, is the question of dealing with not one sequence but with sequences of sequences. For some real-world networks, there exist data not only on one snapshot of the network but also on a historical series of such snapshots. In this case, we have not one degree sequence but a series of degree sequences. One can then apply the estimators to these series, obtaining a series of estimates. Given such an estimate series and the length of the sequence attached to each element of the series, i.e., the network size, can one extract any additional information about the convergence of the series, and possibly devise some tests of the hypothesis that the series comes from a regularly varying generative process? To the best of our knowledge, these questions are wide open.

Integer-valued sequences. Another network-specific issue is that degree sequences are integer-valued, while the considered EV estimators were designed with real-valued data in mind. As a consequence, these estimators are known to be unstable and to converge quite slowly in the case of integer-valued regularly varying distributions, Appendix C. We circumvent this issue in our experiments by adding symmetric uniform noise, but it would be nice to design estimators that work reliably on integer-valued data directly.

The second-order condition. Another down-to-earth issue is the second-order condition needed to prove the consistency of the double bootstrap method, Appendix C. This condition is violated by pure power-law distributions, the Pareto and zeta distributions. We saw in our experiments in Appendix D that the estimators equipped with the double bootstrap method converge in these cases as well, but there are no proofs of the consistency of the double bootstrap method in these cases.

Cutoffs. Finally, we comment on the important issue of cutoffs that often causes much confusion. Here we have to differentiate between many possibilities of what a *cutoff* might mean. Two classes of such possibilities are finite-size effects and true cutoffs. In the first case, a cutoff is just an illusion due to a finite sample size. If one samples an insufficiently large number of i.i.d. samples from a regularly varying distribution, the empirical distribution of these samples may appear to have a cutoff, even though the distribution we are sampling from does not have any cutoffs by definition of it being regularly varying. In simple terms, the tail of the empirical distribution may bend downwards, but this effect is simply due to the insufficient number of samples. In such cases, if one explores the empirical distribution tail, one finds only a few data points there. We note that EV theory gives not only the expected value of the maximum among these samples, but also the exact distribution of this properly rescaled maximum in the limit, Appendix B.

In networks, however, this maximum can simply not be greater than the network size n which is equal to the degree sequence length, and there are other kinds of degree correlations

and degree sequence constraints that are forced by the network structure, many documented in Ref. [81], for instance. These constraints can be such that the degree distribution does have true cutoffs. More generally, it may very well happen that the process driving the evolution of a given network is such that its degree distribution does converge to a distribution with true cutoffs, sharp or soft. Examples are the preferential attachment model with a preference kernel which is constant above a certain degree threshold [85, Sec. 4], or the causal set of the universe [86].

In these cases, one has to further differentiate between the following two possibilities. First, the cutoff can be constant, that is, independent of the network size/degree sequence length. In this case, the distribution is not regularly varying by definition, so that one cannot call it power-law. If one still wishes to estimate γ in samples from, for example, the distribution class $P(k) = \ell(k)k^{-\gamma}e^{-k/c}$ where $\ell(k)$ is a slowly varying function, and $c > 0$ a constant, then it is yet another open problem since EV-based estimators can clearly not be employed for this estimation, simply because the distributions in this class are not regularly varying. Neither are we aware of any consistent estimators that can do this estimation. In fact, such estimators are quite unlikely to exist, simply because this task appears to be ill-defined. Indeed, $\ell(k)$ can be arbitrarily bad for any finite k . All we know about this function is that it varies slowly at infinity. However, we also know from the shape of the distribution that it is exponential at infinity.

The other possibility is that the cutoff diverges with the network size. In this case we have a scenario that can possibly be modeled by random sequences of varying length n sampled from a sequence of distributions parameterized by n . If their cutoff diverges with n , then the latter sequence may or may not converge to a regularly varying distribution in the $n \rightarrow \infty$ limit. In Appendix D, we considered an example of this sort, diverging natural exponential cutoffs, where the n -dependent distributions $P_n(k) = C_n k^{-\gamma} e^{-k/n^{\xi}}$ do converge to the regularly varying Pareto distribution. We saw there that even in this case, the considered estimators converge to the true values of γ , even though the key assumptions behind the proofs of their convergence are violated. Proving the consistency of these and other estimators for sequences of random numbers sampled from sequences of distributions converging to regularly varying distributions, is thus yet another open problem.

Notwithstanding these open problems, the consistent estimators considered in this paper represent the current state of the art in the rigorous detection of power laws in empirical data. Their implementation is available in Ref. [55], and their application to a representative collection of degree sequences in real-world networks confirms that scale-free networks are not rare.

ACKNOWLEDGMENTS

We thank S. Resnick, S. Foss, A. Vespignani, A. Broido, and J. Kelley for useful discussions and suggestions. This work was supported by NSF Grant No. IIS-1741355, ARO Grant Nos. W911NF-16-1-0391 and W911NF-17-1-0491, NWO VICI Grant No. 639.033.806, and the Gravitation NETWORKS Grant No. 024.002.003.

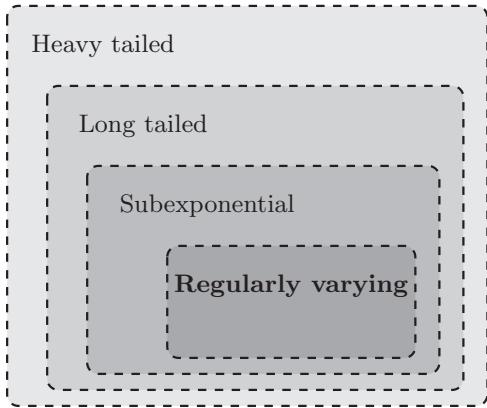


FIG. 4. Schematic overview of the landscape of heavy-tailed distributions, containing regularly varying distributions.

APPENDIX A: CLASSES OF DISTRIBUTIONS WITH HEAVY TAILS

Here we briefly review the taxonomy of distributions with heavy tails and provide the definition of the simplest and most frequently seen regularly varying distributions. All the distribution classes mentioned here are characterized by the key property that their tails decay more slowly than exponentially. The most general class is that of the heavy-tailed distributions. We note that “fat-tailed” distributions are also mentioned sometimes in the literature, but do not appear to have any rigorous definition. We focus on distributions with support on \mathbb{R}_+ . Chapters 2 and 3 in Ref. [33] contain further details.

1. Heavy-tailed distributions

A distribution with cumulative distribution function (CDF) $F(x)$ is said to be *heavy-tailed* [33, Theorem 2.6] if its complementary CDF (CCDF) $\bar{F}(x)$ satisfies, for any $t > 0$,

$$\limsup_{x \rightarrow \infty} e^{tx} \bar{F}(x) = \infty.$$

In words, this definition literally says that the tail of the distribution $\bar{F}(x)$ decays more slowly than exponentially.

The class of heavy-tailed distributions is quite vast and general which makes it rather difficult to work with them in their full generality. Therefore many different narrower and more tractable subclasses of heavy-tailed distributions have been defined and studied, see Fig. 4 for an overview of the landscape of heavy-tailed distributions. For completeness, we briefly discuss two important subclasses that encapsulate regularly varying distributions, which are our main interest.

a. Long-tailed distributions. A distribution with CDF $F(x)$ is called long-tailed [33, Definition 2.21] if its CCDF satisfies, for any fixed $y > 0$,

$$\lim_{x \rightarrow \infty} \frac{\bar{F}(x+y)}{\bar{F}(x)} = 1, \quad (\text{A1})$$

meaning that any finite shift does not asymptotically affect the tail of the distribution. This property is nice and useful as, for instance, if X is a random variable which has a long-tailed distribution, and Y a random variable that only takes values

on a finite set, then the tail of the distribution of $X + Y$ is asymptotically equivalent to that of X [33, Corollary 2.32].

Long-tailed distributions are heavy-tailed [33, Lemma 2.17], but not all heavy-tailed distributions are long-tailed. A simple example of a heavy-tailed function which is not long-tailed is

$$f(x) = \sum_{k=1}^{\infty} 2^{-k} \mathbb{1}\{2^{(k-1)} \leq x < 2^k\},$$

where $\mathbb{1}$ is the indicator function. Indeed, for any $t > 0$,

$$\limsup_{x \rightarrow \infty} e^{tx} f(x) \geq \limsup_{k \rightarrow \infty} e^{t2^k} 2^{-k} = \infty,$$

so that f is heavy-tailed, but

$$\liminf_{x \rightarrow \infty} \frac{f(x+1)}{f(x)} \leq \liminf_{k \rightarrow \infty} \frac{f(2^k+1)}{f(2^k)} = \frac{1}{2} \neq 0,$$

so that f is not long-tailed.

b. Subexponential distributions. Let $(F * F)(x)$ be the convolution of CDF $F(x)$ with itself. That is, $F * F$ is the CDF of $X + X'$, where X and X' are independent random variables with CDF F . A distribution with CDF $F(x)$ is said to be subexponential [33, Definition 3.1] if

$$\lim_{x \rightarrow \infty} \frac{\overline{(F * F)}(x)}{\overline{F}(x)} = 2. \quad (\text{A2})$$

This definition means that if X and X' are independent samples from a subexponential distribution, then the CCDF of $X + X'$ is asymptotically twice as large as the CCDF of the original distribution. This property implies, for instance, that if the sum $\sum_{i=1}^n X_i$ of n independent samples from a subexponential distribution exceeds some large threshold, then it is because just one X_i has exceeded this threshold. This is in contrast to independent samples from a Poisson distribution, for instance, as their sums exceeding a large threshold do not contain, with high probability, any terms exceeding this threshold.

The class of subexponential distributions is contained in that of long-tailed distribution [33, Lemma 3.2], hence they are heavy-tailed. In fact, it is strictly contained. However, unlike the case for heavy-tailed versus long-tailed distributions, examples of long-tailed distributions that are not subexponential are more involved, see Sec. 3.7 in Ref. [33].

Our main interest is in regularly varying distributions, which form a subclass of subexponential distributions [33, Theorem 3.29]. This hierarchy endows regularly varying distributions with all the nice theoretical properties of the subexponential and long-tailed ones, but in contrast to these more general classes, regularly varying distributions are equipped with a concise and tractable representation that makes them very convenient to work with in statistical inference settings.

2. Regularly varying distributions

A function $f(x)$ is said to be *regularly varying at infinity with index α* [32,33] if there exists a slowly varying function $\ell(x)$, such that

$$f(x) = \ell(x)x^{-\alpha}, \quad (\text{A3})$$

where a slowly varying function $\ell(x)$ is defined to be a function satisfying, for any $t > 0$,

$$\lim_{x \rightarrow \infty} \frac{\ell(tx)}{\ell(x)} = 1.$$

The simplest examples of slowly varying functions are functions converging to constants or $\ln^a(bx)$ for any $a \in \mathbb{R}$ and $b > 0$.

The full class of slowly varying functions is of course much richer, and it is fully characterized by Karamata's representation theorem [31, Corollary 2.1] stating that

$$\ell(x) = c(x) \exp \left\{ \int_1^x t^{-1} \varepsilon(t) dt \right\},$$

for some functions $c, \varepsilon : \mathbb{R}_+ \mapsto \mathbb{R}_+$ satisfying

$$\lim_{x \rightarrow \infty} c(x) = c \in (0, \infty), \quad \lim_{x \rightarrow \infty} \varepsilon(x) = 0.$$

The theory of regular variations is a rich and well-developed one, and for further details we refer to Ref. [32].

A distribution is *defined* to be regularly varying if its CCDF $\bar{F}(x)$ is a regularly varying function. In Sec. II, we also define a *power-law distribution* to be a regularly varying distribution.

We note that if the PDF of a distribution is regularly varying, then so is its CCDF with another slowly varying function $\ell'(x)$,

$$P(x) = \ell(x)x^{-\gamma} \Rightarrow \bar{F}(x) = \ell'(x)x^{-\alpha}, \text{ where } \alpha = \gamma - 1,$$

according to Karamata's theorem [31, Theorem 2.1] in the case of continuous distributions, and to [87, Lemma 9.1] in the case of discrete ones. The converse is not generally true, and depends on the exact form of the slowly varying function $\ell'(x)$. A simple example of a distribution whose PDF is not regularly varying but whose CCDF is, is given by the PDF $P(x) = c \sin(x)^2 x^{-3}$ with support $x \geq 1$, and c the normalization constant. This PDF is not regularly varying since $\ell(x) = c \sin^2 x$ is not slowly varying. However, the CCDF of this distribution $\bar{F}(x) = \ell'(x)x^{-2}$, where $\ell'(x) = (c/2)(\sin^2 x + x \sin(2x) - 2x^2 \text{Ci}(2x))$ and $\text{Ci}(x) = -\int_x^\infty \cos(t)/t dt$ is the cosine integral, is regularly varying because $\ell'(x)$ is slowly varying: it converges to the constant $c/4$ at $x \rightarrow \infty$.

Another important property of regularly varying distributions is that if the sum $X + Y$ of two random variables is regularly varying, and $\lim_{z \rightarrow \infty} \bar{F}_Y(z)/\bar{F}_{X+Y}(z) = 0$, then X is regularly varying as well, and the tail exponents γ of $X + Y$ and X are the same [88, Lemma 3.12]. In application to directed networks, this means for instance that if the total degree distribution is regularly varying and either the in-degree or out-degree distribution is not heavy-tailed, then the other distribution must be regularly varying with the same exponent as the total degree distribution.

As a subclass of heavy-tailed distributions, regularly varying distributions can model data with high variability, yet here we stress again that they are far from being as general as heavy-tailed distributions, which means, in particular, that if a given data fails to be regularly varying, it does not necessarily mean that it is not heavy-tailed or even subexponential. The simplest example of a subexponential distribution which is not regularly varying is the lognormal distribution. Yet on the other hand, regularly varying distributions are a vast

generalization of pure power laws exclusively considered in Refs. [19,20], i.e., of the Pareto distribution (10) if x is continuous, or of the generalized zeta distribution (9) if x is integer-valued.

3. Simplest examples of regularly varying distributions

To make the definition (A3) more concrete, here we give the simplest examples of regularly varying distributions, both continuous and integer-valued ones.

The simplest example is the continuous Pareto distribution with scale x^* and shape α , or exponent $\gamma = \alpha + 1$:

$$P_{\text{Pareto}}(x) = \begin{cases} \alpha(x^*)^\alpha x^{-\gamma} & \text{if } x \geq x^*, \\ 0 & \text{else.} \end{cases}$$

Here the slowly varying function is simply the constant $\ell_{\text{Pareto}}(x) = \alpha(x^*)^\alpha$, which does not vary at all.

There are two simple ways to turn a continuous regularly varying distribution into an integer-valued one, both of which again belong to the class of regularly varying distributions. In the first example, we simply take the integer k to be the floor of the continuous value x : $k = \lfloor x \rfloor$. If x is Pareto-distributed, then since $P(k) = \bar{F}_{\text{Pareto}}(k-1) - \bar{F}_{\text{Pareto}}(k)$, it follows that for all $k \geq \lfloor x^* \rfloor$

$$P_{\text{floorP}}(k) = \left(\frac{k-1}{x^*} \right)^{-\alpha} - \left(\frac{k}{x^*} \right)^{-\alpha} = \ell_{\text{floorP}}(k)k^{-\gamma},$$

where $\ell_{\text{floorP}}(k)$ converges to $\alpha(x^*)^\alpha$ as $k \rightarrow \infty$. We note that in this example the slowly varying function is not a constant. Yet it approaches a constant asymptotically.

The second example is a mixed Poisson distribution [89] with Pareto mixing. The easiest way to define a mixed Poisson distribution is via the procedure to sample from it: as its name suggests, first sample x from the Pareto distribution, and then sample k from the Poisson distribution with mean x . The resulting PDF of k is thus

$$P_{\text{mPois}}(k) = \int_{x^*}^{\infty} \frac{x^k e^{-x}}{k!} P_{\text{Pareto}}(x) dx = \alpha(x^*)^\alpha \frac{\Gamma(k+1-\gamma, x^*)}{\Gamma(k+1)} = \ell_{\text{mPois}}(k)k^{-\gamma}, \quad (\text{A4})$$

$$\ell_{\text{mPois}}(k) = \alpha(x^*)^\alpha k^\gamma \frac{\Gamma(k+1-\gamma, x^*)}{\Gamma(k+1)},$$

where $\Gamma(k, x)$ is the upper incomplete Gamma function. The function $\ell_{\text{mPois}}(k)$ is slowly varying, and its $k \rightarrow \infty$ limit is, as in the previous example, $\alpha(x^*)^\alpha$.

Mixed Poisson distributions appear often as exact degree distributions in network models with hidden variables [90], also known in mathematics as inhomogeneous random graphs [91], or more generally, graphon-based W -random graphs [92]. Both the expected value and the tail exponent of mixed Poisson k are equal to those of Pareto x , versus floored Pareto in which the expected value of k is $\langle k \rangle = \zeta(\langle x \rangle)$, where ζ is the Riemann zeta function.

APPENDIX B: CONSISTENT ESTIMATORS FOR TAIL EXPONENTS OF REGULARLY VARYING DISTRIBUTIONS

Here we give the definitions of the three consistent estimators of the tail of a regularly varying distribution that we use to infer the tail exponents in synthetic and real-world degree sequences. The two other consistent estimators that are also included in our software package [55] are defined here as well.

Although we work only with regularly varying distributions, the used estimators are actually designed to estimate the *index* of an *extreme value distribution*. In fact, the consistency results are proven under the assumption that the distribution belongs to the *maximum domain of attraction* of an extreme value distribution. It turns out that any regularly varying distribution satisfies this assumption. Therefore we start with a brief review of extreme value distributions and their maximum domains of attraction, and then explain how these concepts are employed by the consistent estimators of tail exponents.

1. Extreme value distributions and their maximum domains of attraction

Let x_1, \dots, x_n be an i.i.d. sequence sampled from some distribution $P(x)$, and denote by $m_n = \max_{1 \leq i \leq n} x_i$ the largest value in the sequence. Extreme value theory is concerned with the properties of the distribution of m_n , whose CDF is given by the order statistics $F^n(x)$. The typical question is whether there is a nondegenerate limit law, i.e., a distribution which is not a delta function, for $\mu_n = (m_n - d_n)/c_n$ for some appropriately chosen n -dependent constants $c_n > 0$ and $d_n \in \mathbb{R}$. A degenerate limit for μ_n exists for any distribution as one can always select $d_n = 0$ and any c_n growing with n faster than the expected value of m_n , in which case the distribution of μ_n would approach the delta-function distribution centered at zero. However, a nondegenerate limit exists [43, Theorem 3.1.3] if and only if the CDF $F(x)$ of the distribution satisfies

$$\lim_{x \rightarrow X_F} \frac{1 - F(x)}{F(x) - F(x-)} = 1 \quad \text{and} \quad F(X_F-) = 1, \quad (\text{B1})$$

where $X_F = \sup\{x ; F(x) < 1\}$ is the right endpoint of the distribution, which can be infinite, and $F(x-) = \lim_{t \rightarrow \infty} F(x - 1/t)$ is the left limit of F at x . In words, this requirement states that $F(x)$ must be sufficiently flat at its right end and must not jump there. Many distributions frequently appearing in practice do satisfy this requirement, but not all. Notable examples of distributions that do not satisfy it, are the Poisson [43, Example 3.1.4] and geometric [43, Example 3.1.5] distributions. Indeed, for a distribution with support on non-negative integers, the limit in (B1) is equivalent to $\lim_{k \rightarrow \infty} \bar{F}(k)/\bar{F}(k-1) = 1$. For the Poisson distribution with mean λ , we have

$$\frac{\bar{F}(k)}{\bar{F}(k-1)} \leq 1 - \left(1 + \frac{\lambda}{k-\lambda}\right)^{-1},$$

which tends to 0 as $k \rightarrow \infty$, while for the geometric distribution with success probability p , $\bar{F}(k)/\bar{F}(k-1) \rightarrow 1 - p$, violating (B1) in both cases.

If a distribution $P(x)$ does satisfy (B1), so that a nondegenerate limit distribution of $\mu_n = (m_n - d_n)/c_n$ does exist, this latter distribution $\mathcal{P}(\mu)$ is called an *extreme value distribution* [93, 94]. An important result [69] (see also Ref. [93,

Proposition 0.3]) states that extreme value distributions are parameterized by an *index* parameter $\xi \in \mathbb{R}$, and that the class of extreme value distributions consists of just three subclasses—Fréchet, Gumbel, and Weibull distributions—corresponding, respectively, to $\xi > 0$, $\xi = 0$, and $\xi < 0$. The CDFs $\mathcal{F}(\mu)$ of these three distributions can be grouped into the CDF of the generalized extreme value distribution

$$\mathcal{F}(\lambda) = e^{-\lambda}, \quad \text{where}$$

$$\lambda = \begin{cases} (1 + \xi v)^{-1/\xi}, & \text{if } \xi \neq 0, \\ e^{-v}, & \text{otherwise,} \end{cases} \quad \text{where} \quad (\text{B2})$$

$$v = \frac{\mu - l}{s},$$

where $l \in \mathbb{R}$ and $s > 0$ are known as, respectively, the location and scale parameters. The supports of the distributions are $v \geq -1/\xi$ for $\xi > 0$, $v \leq -1/\xi$ for $\xi < 0$, and $v \in \mathbb{R}$ for $\xi = 0$.

A distribution $P(x)$ is said to belong to the *maximum domain of attraction* (MDA) of an extreme value distribution $\mathcal{P}(\mu)$ if there exist n sequences of constants $c_n > 0$ and $d_n \in \mathbb{R}$ such that the distribution of $\mu_n = (m_n - d_n)/c_n$ converges to $\mathcal{P}(\mu)$. The crucially important fact, originally proven in [70], is that the regularly varying distributions are exactly all the distributions comprising the MDA of the Fréchet distribution, see also Refs. [93, Proposition 1.11] and [94, Theorem 1.4.20], so that any regularly varying distribution with PDF and CCDF tail exponents γ and α belongs to the MDA of a Fréchet distribution with index

$$\xi = \frac{1}{\alpha} = \frac{1}{\gamma - 1}. \quad (\text{B3})$$

The sequences d_n and c_n in this regularly varying/Fréchet case are [93, Proposition 1.11]

$$d_n = 0, \quad c_n = F^{-1}\left(1 - \frac{1}{n}\right),$$

where F^{-1} is the inverse CDF of the distribution $P(x)$, while the location and scale parameters of the Fréchet distribution in (B2) are

$$l = 1, \quad s = \xi,$$

so that the distribution of the largest values m_n among n i.i.d. samples from any regularly varying distribution has the following limit upon rescaling by c_n :

$$\lim_{n \rightarrow \infty} F^n(c_n \mu) = \mathcal{F}(\mu) = e^{-\mu^{-1/\xi}} \quad (\text{B4})$$

with support $\mu \geq 0$.

If the distribution $P(x)$ is Pareto, for example, then

$$c_n = F^{-1}\left(1 - \frac{1}{n}\right) = x^* n^\xi, \quad (\text{B5})$$

related to the known observations that the expected maximum degree among n samples from a power-law networks with exponent γ is proportional to $n^{1/(\gamma-1)}$ [81].

We note that the expressions above specify not only the expected values but also the full limit distributions of such maxima. We also note that the mean of the limit Fréchet distribution $\mathcal{F}(\mu) = e^{-\mu^{-1/\xi}}$ is $\langle \mu \rangle = \Gamma(1 - \xi)$ if $\xi < 1$ ($\gamma > 2$),

and that this mean is infinite if $\xi \geq 1$ ($\gamma \leq 2$), so that if $\gamma > 2$ and n is large, one can approximate the expected value of m_n in Pareto by

$$\langle m_n \rangle \approx \langle \mu \rangle c_n = \Gamma\left(\frac{\gamma-2}{\gamma-1}\right) x^* n^{1/(\gamma-1)}. \quad (\text{B6})$$

To complete the picture of the classification of distributions based on their MDAs, the MDA of the Weibull distribution consists of all distributions with an upper-bounded support, $X_F < \infty$, and CCDFs satisfying $\bar{F}(X_F - 1/t) = \ell(t)t^{1/\xi}$ for $t \rightarrow \infty$, some slowly varying function $\ell(t)$, and $\xi < 0$, which is the same ξ as in (B2) [43, Theorem 3.3.12]. This requirement says that the CCDF approaches its right end as a regularly varying function. Examples are the uniform or beta distributions on $[0, 1]$.

By exclusion, all other distributions satisfying (B1) are in the MDA of the Gumbel distribution. However, there are more insightful characterizations of the Gumbel MDA (roughly, it consists of all von Mises functions and tail-equivalent distributions) [43, Theorems 3.3.26-3.3.27]. Examples are the normal [43, Example 3.3.29] and exponential [43, Example 3.3.19] distributions, which are not heavy-tailed, but also heavy-tailed distributions that are not regularly varying—the subexponential lognormal distribution, for example [43, Example 3.3.31].

The key point, however, is that if a distribution is regularly varying with tail exponent γ , then it is in the MDA of the Fréchet distribution with index $\xi = 1/(\gamma - 1)$ which all the following estimators actually estimate.

2. Hill's estimator

Hill's estimator [57] was introduced to analyze the tail behavior of a distribution without any assumptions about its shape, other than that it belongs to the Fréchet MDA. Given an i.i.d. sample x_i , $i = 1, \dots, n$, and its order statistics $x_{(1)} \geq x_{(2)} \geq \dots \geq x_{(n)}$, the estimator is defined by

$$\hat{\xi}_{\kappa,n}^{\text{Hill}} = \frac{1}{\kappa} \sum_{i=1}^{\kappa} \ln\left(\frac{x_{(i)}}{x_{(\kappa+1)}}\right) \quad (\text{B7})$$

Theorems 4.1 and 4.2 in Ref. [31] prove that if $\kappa/n \rightarrow 0$ and $\kappa \rightarrow \infty$ as $n \rightarrow \infty$, then this estimator is statistically consistent, i.e., satisfies (12), for any distribution in the MDA of the Fréchet distribution. In other words, the estimator is statistically consistent for any regularly varying distribution with any tail exponent $\gamma > 1$, or equivalently any index $\xi > 0$.

3. Moments estimator

The moments estimator [58] is a modification of Hill's estimator that is statistically consistent not only for distributions from the MDA of the Fréchet distribution, but also for distributions from the MDAs of the Gumbel or Weibull distributions, i.e., for any $\xi \in \mathbb{R}$. To define it, denote

$$\hat{\xi}_{\kappa,n}^{\text{Hill},2} = \frac{1}{\kappa} \sum_{i=1}^{\kappa} \left(\ln \frac{x_{(i)}}{x_{(\kappa+1)}} \right)^2.$$

With this notation, the moments estimator is

$$\hat{\xi}_{\kappa,n}^{\text{Moment}} = \hat{\xi}_{\kappa,n}^{\text{Hill}} + 1 - \frac{1}{2} \left(1 - \frac{(\hat{\xi}_{\kappa,n}^{\text{Hill}})^2}{\hat{\xi}_{\kappa,n}^{\text{Hill},2}} \right)^{-1}. \quad (\text{B8})$$

Consistency of $\hat{\xi}_{\kappa,n}^{\text{Moment}}$ is proven in Ref. [58, Theorem 2.1]. It converges almost surely if $\kappa/n \rightarrow 0$ and $\kappa \rightarrow \infty$ as $n \rightarrow \infty$, and there exists a constant $\delta > 0$ such that $\ln(n)^\delta/\kappa \rightarrow 0$.

4. Kernel estimator

Similar to the moments estimator, the Kernel estimator [59] is consistently applicable to distributions with any $\xi \in \mathbb{R}$. As its name suggests, the Kernel estimator uses a *kernel*, which is a function $\phi : [0, 1] \rightarrow [0, \infty)$ that can be chosen by the user, and that must satisfy a set of conditions for the estimator to be consistent [59]. The estimator also employs a parameter $\lambda > 1/2$ to get rid of possible singularities. Finally, instead of using an integer-valued κ to determine the range of the order statistics to consider for ξ estimation, the estimator relies on a continuous bandwidth parameter $h > 0$ for that purpose. Thanks to this modification, as a function of h , the estimator tends to be smoother compared to the other estimators.

Given the chosen kernel ϕ , denote $\phi_h(u) := \phi(u/h)/h$. With this notation, the Kernel estimator is

$$\begin{aligned} \hat{\xi}_{h,n}^{\text{Kernel}} &= \hat{\xi}_{h,n}^{\text{pos}} - 1 + \frac{\hat{q}_{h,n}^{(1)}}{\hat{q}_{h,n}^{(2)}}, \text{ where} \\ \hat{\xi}_{h,n}^{\text{pos}} &= \sum_{i=1}^{n-1} \frac{i}{n} \phi_h\left(\frac{i}{n}\right) \ln\left(\frac{x_{(i)}}{x_{(i+1)}}\right), \\ \hat{q}_{h,n}^{(1)} &= \sum_{i=1}^{n-1} \left(\frac{i}{n}\right)^\lambda \phi_h\left(\frac{i}{n}\right) \ln\left(\frac{x_{(i)}}{x_{(i+1)}}\right), \\ \hat{q}_{h,n}^{(2)} &= \sum_{i=1}^{n-1} \frac{\partial}{\partial u} [u^{\lambda+1} \phi_h(u)]_{u=i/n} \ln\left(\frac{x_{(i)}}{x_{(i+1)}}\right). \end{aligned}$$

The consistency of this estimator for $n \rightarrow \infty$, $h \rightarrow 0$, and $hn \rightarrow \infty$ is proven in Ref. [59].

For the experiments in this paper, which are also the default settings in Ref. [55], we prepare a list of fractions of order statistics h_1, \dots, h_s . The estimator is then evaluated at each h value h_i , $i = 1, \dots, s$. These fractions h_i are logarithmically spaced in the interval $[1/n, 1]$, where n is the sequence length. The number of different h values is set to $s = [0.3n]$. The logarithmic binning is chosen to scan the tail of the degree sequence more densely, while the choice of s guarantees that the sample sizes used in the double bootstrap procedure described in Appendix C1 exceed s , so that kernel smoothing is applied to both bootstrap samples as well. For λ , we use the setting in Ref. [59] where $\lambda = 0.6$. In our software package [55], the values of λ and s can be changed to any other values $\lambda > 1/2$ and $s > 0$. We note that the estimator is proven to be consistent for any choice of s , h_i , and λ satisfying the requirements above. Package [55] also implements the bi- and triweight kernels from [59]

$$\begin{aligned} \phi^{(1)}(u) &= \frac{15}{8}(1-u^2)^2, \\ \phi^{(2)}(u) &= \frac{35}{16}(1-u^2)^3, \end{aligned}$$

where $\phi^{(1)}$ is used for the tail estimation, and the combination of $\phi^{(1)}$ and $\phi^{(2)}$ is used to find the optimal h^* as described in Appendix C 1. Once such an h^* is found, the value of κ^* is set to $\lfloor nh^* \rfloor$ in Ref. [55].

5. Smooth Hill estimator

Although Hill's estimator is consistent, it, as a function of the number of order statistics κ , can be highly irregular for finite-size data samples. The plots of such functions are even known as Hill horror plots, Sec. 4.4.2 in Ref. [31]. These horrors make it essentially impossible to examine these plots in search of the stable regime of $\hat{\xi}_{\kappa,n}^{\text{Hill}}$, i.e., the region of κ 's where $\hat{\xi}_{\kappa,n}^{\text{Hill}}$ is approximately constant. The value that the estimator yields in this constant regime is then one's best estimate of ξ , but if these plots are highly irregular, then this estimation procedure is unavoidably subjective. Even though no results presented in this paper rely on such subjective manipulations—instead we rely on the statistically consistent double bootstrap method to find κ^* , Appendix C 1—in practice one may wish to investigate such plots to get deeper insight into the data at hand. To this end, one usually uses either the smoothed version of Hill's estimator or Pickands estimator, which are both included in Ref. [55].

The smooth Hill estimator [95] is defined for any integer $r \geq 2$, which is a parameter, by

$$\hat{\xi}_{\kappa,n}^{\text{smooH}} = \frac{1}{(r-1)\kappa} \sum_{j=\kappa+1}^{r\kappa} \hat{\xi}_{j,n}^{\text{Hill}}, \quad (\text{B9})$$

which is just an average of Hill's estimators over the range $[\kappa+1, r\kappa]$. This estimator is also statistically consistent, for any r , as proven in Ref. [95]. The practical advantage of this smooth estimator compared to the original Hill's estimator is that by averaging the latter, the former suppresses its erratic behavior, making it easier to identify its stable region.

6. Pickands estimator

The Pickands estimator [72] is defined by

$$\hat{\xi}_{\kappa,n}^{\text{Pickands}} = \frac{1}{\ln 2} \ln \left(\frac{x_{(\kappa)} - x_{(2\kappa)}}{x_{(2\kappa)} - x_{(4\kappa)}} \right). \quad (\text{B10})$$

Consistency of the estimator is proven also in Ref. [72]. Similarly to the moments and kernel estimators, it is consistently applicable to distributions in the MDAs of extreme value distributions with any $\xi \in \mathbb{R}$. In practice this estimator provides a simple way to check whether the assumption that the data come from a regularly varying distribution makes sense. Specifically, if the function $\hat{\xi}_{\kappa,n}^{\text{Pickands}}$ of κ is all negative, then this assumption can hardly be true.

In contrast to the other estimators, the Pickands estimator has an issue dealing with integer-valued data containing ties. For instance, if there are many data points with the same value (many nodes with the same degree), it can happen that for some κ , $x_{(2\kappa)} = x_{(4\kappa)}$, in which case $\hat{\xi}_{\kappa,n}^{\text{Pickands}}$ is undefined. This drawback can however be remedied, in a provably consistent manner, by adding uniform noise to the integer-valued data, as explained in Appendix C 2.

It is known that in practice the Pickands estimator is quite volatile as a function of the number of order statistics κ , and that it has large asymptotic variance [73] and poor efficiency [59]. Attempts to cure this poor behavior resulted in a number of different versions of generalized Pickands estimators [73, 96–100], all of them using linear combinations of logarithmic spacings of the order statistics $\kappa_{(1)}, \dots, \kappa_{(n)}$. Yet, to the best of our knowledge, the consistency of the double bootstrap method has been proven [101] for only one version, the one defined in Ref. [96], so that we implemented only this version in Ref. [55].

APPENDIX C: ESTIMATING THE TAIL EXPONENT OF AN EMPIRICAL DEGREE SEQUENCE

Here we discuss the technical details concerning the application of the consistent estimators discussed in the previous section to empirical degree sequences coming from either synthetic or real-world networks.

1. Finding the optimal number of order statistics

All the estimators in Appendix B depend on the number of order statistics κ . That is, all the estimators operate only on the κ largest-value data samples (degrees). The consistency of all the estimators is proven only in the limit of both κ and n , the number of samples (nodes), tending to infinity. Therefore, when applied to a finite empirical degree sequence, these estimators have the value of κ as a free parameter. The main focus of this section is the double bootstrap method that algorithmically identifies an optimal κ value κ^* in a statistically consistent manner, meaning that the value of $\hat{\xi}_{\kappa^*,n}$ estimated by these estimators with $\kappa = \kappa^*$ provably converges to the true value of ξ as $n \rightarrow \infty$.

The identification of an optimal value κ^* has been an active research topic in extreme value statistics for several decades [101, 102]. The existing methods for choosing κ^* can be roughly split into two classes: (1) heuristic approaches that propose to study tail index estimates plotted as functions of κ and (2) theoretical approaches based on the minimization of the asymptotic mean squared error (AMSE) of the estimator.

The heuristic methods mainly consider various ways of identifying regions of κ where estimators show *stable behavior*, i.e., where the estimator plot is relatively flat as a function of κ . Examples of such approaches are the automated eyeball method [95], or picking a fixed small percentage (typically 5% or 10%) of the largest-value data samples. Such methods, involving (semi-)subjective *ad hoc* choices, may not be robust.

The main idea behind the theoretical methods is as follows. Suppose x_1, \dots, x_n is an i.i.d. sequence sampled from a distribution that belongs to the domain of attraction of the generalized extreme value distribution (B2) with a given ξ . Denote by $\hat{\xi}_{\kappa,n}$ the estimated value of ξ returned by a given estimator applied to the κ largest elements in this sequence. Observe that since the sequence is random, $\hat{\xi}_{\kappa,n}$ is a random variable. Define the *asymptotic mean squared error* between the true and estimated ξ 's as

$$\text{AMSE}(n, \kappa) = \mathbb{E}[(\hat{\xi}_{\kappa,n} - \xi)^2]. \quad (\text{C1})$$

The main goal is to find the optimal κ value κ^* that minimizes this error:

$$\kappa^* = \arg \min_{\kappa} \text{AMSE}(n, \kappa). \quad (\text{C2})$$

To estimate κ^* in this paper, we use the AMSE-based double bootstrap method developed in Refs. [59,101–103] because of its proved consistency, stability, and applicability to the considered estimators. The method finds a consistent optimal value κ^* for a given consistent estimator by employing not only this estimator, but also another consistent estimator. The two estimators are applied to two collections of bootstrap samples from the original data, estimating ξ at all possible values of κ in these collections, and the value κ^* is then determined as the value of κ at which the two estimators agree most in their estimation of ξ according to the empirical AMSE evaluated on the bootstrap collections.

Specifically, the double bootstrap method operates using the following steps with two parameters: r denotes the number of bootstrap samples, and $t \in (0, 1)$ defines the first and second bootstrap sample sizes as $n_1 = \sqrt{tn}$ and $n_2 = tn$, where n is the original sequence length. In all the experiments in this paper, and in the software package [55], these parameters are set to $r = 500$ and $t = 1/2$ by default, so that the size of the second bootstrap sample is $n_2 = n/2$. (1) Sample $r > 1$ bootstrap samples of size $n_1 = [\sqrt{tn}]$ from the original data. (2) Using the two consistent estimators, estimate $\xi_{\kappa_1, j}^{(1)}$ and $\xi_{\kappa_1, j}^{(2)}$ for each value of $\kappa_1 = 1, \dots, n_1$ in each bootstrap sample $j = 1, \dots, r$. (3) Find κ_1^* that minimizes the empirical AMSE between the two estimates with respect to the r bootstrap samples, i.e.,

$$\kappa_1^* = \arg \min_{\kappa_1} \frac{1}{r} \sum_{j=1}^r (\xi_{\kappa_1, j}^{(1)} - \xi_{\kappa_1, j}^{(2)})^2.$$

(4) Repeat the same procedure for a smaller bootstrap sample size $n_2 = [tn]$ and find κ_2^* in the same manner. (5) The optimal value of κ for the original data is given by

$$\kappa^* = A(\kappa_1^*, n_1, n) \frac{(\kappa_1^*)^2}{\kappa_2^*}, \quad (\text{C3})$$

where $A(\kappa_1^*, n_1, n)$ is a prefactor that depends on κ_1^* , n_1 , n , and whose exact form depends on the two estimators used.

Following the derivations in Refs. [59,101–103], we use the following combinations of consistent estimators for the double bootstrap procedure applied to the Hill, kernel, and moments estimators: (1) the first (Hill) and second moment estimators for the Hill double bootstrap; (2) the second and third moment estimators for the moments double bootstrap; and (3) the biweight and triweight kernel estimators for the kernel double bootstrap. We note that in principle any combination of consistent estimators can be used in the double bootstrap method, but proofs of the consistency of such combinations must be carried out for each combination, so that we use the combinations that are already proven to be consistent and optimal.

We also note that these proofs are based on an additional assumption that the regularly varying distribution of the samples satisfies the *second-order condition* [104], [105, Definition 2.3.1]. This condition is often invoked to prove

asymptotic normality of estimators [76,106], but it is known to be either difficult or impossible to check in real-world data [107]. To define it for a given distribution with CDF $F(x)$, let $U(x) = F^{-1}(1 - 1/x)$ be the inverse of the CCDF $1 - F(x)$. If the distribution is in an MDA of some extreme value distribution, then it is known [105, Theorem 1.1.6] that there exists a positive function $a(x)$ such that, for any $t > 0$,

$$\lim_{x \rightarrow \infty} \frac{U(tx) - U(x)}{a(x)} = b_{\xi}(t) := \begin{cases} \frac{t^{\xi} - 1}{\xi}, & \text{if } \xi \neq 0, \\ \ln t, & \text{otherwise.} \end{cases}$$

The second-order condition concerns the scaling of $(U(tx) - U(x))/a(x) - b_{\xi}(t)$ as $x \rightarrow \infty$. The distribution is said to satisfy the second-order condition if there exist functions $A(x)$ with $\lim_{x \rightarrow \infty} A(x) = 0$ and a nondegenerate $H(t) \neq cb_{\xi}(t)$ for any $c \neq 0$, such that for any $t > 0$

$$\lim_{x \rightarrow \infty} \left(\frac{U(tx) - U(x)}{a(x)} - \frac{t^{\xi} - 1}{\xi} \right) / A(x) = H(t). \quad (\text{C4})$$

A simple example of a regularly varying CDF that satisfies the second-order condition is

$$F(x) = 1 - x^{-\alpha} - dx^{-\delta},$$

where $d > 0$, $\delta > \alpha > 0$, and $x \geq x^*$, where x^* is the root of $F(x) = 0$. The simplest example of a distribution that does not satisfy the second-order condition is a Pareto distribution. To see this, note that in case of Pareto $U(x) = x^* x^{1/\alpha}$, so that $a(x) = \alpha^{-1} x^* x^{1/\alpha}$, and

$$\frac{U(tx) - U(x)}{a(x)} = \alpha(t^{1/\alpha} - 1) = \frac{t^{\xi} - 1}{\xi}.$$

Hence the left-hand side in (C4) is always zero, meaning that no nondegenerate function $H(t)$ exists.

Since the proofs of consistency of the double bootstrap method rely on the second-order condition, nothing can be said regarding the convergence and consistency of the considered estimators equipped with the double bootstrap method, if they are applied to sequences sampled from distributions that do not satisfy the second-order condition. However, in our experiments, we find that even in these cases the double bootstrap procedure performs well, and the values of $\hat{\xi}_{\kappa^*, n}$ quickly converge to the true ξ 's as $n \rightarrow \infty$ in most such cases, Appendix D 1.

Further technical details on the double bootstrap procedure for the Hill estimator can be found in Refs. [102,103], for the moments estimator in Ref. [101], and for the kernel estimator in Ref. [59], where the consistency of double bootstrapping applied to these estimators is also proven.

2. Working with integer data

A common issue with all the known consistent estimators is their instability, i.e., erratic behavior of $\hat{\xi}_{\kappa, n}$'s as functions of sampled sequences and the number of order statistics κ , on integer-valued sequences [71,108], which is the case with degree sequences. For instance, just rounding samples in sequences sampled from a continuous regularly varying distribution makes the estimators unstable [71], even though such rounded sequences are still regularly varying with the same exponent, Sec. A 3. In other words, the estimators remain

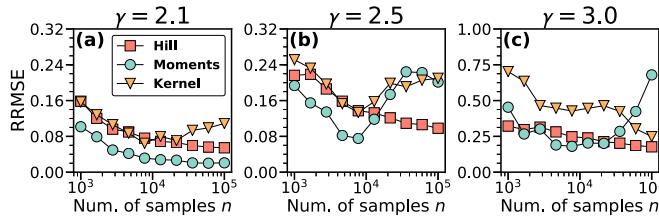


FIG. 5. The relative root mean squared error (RRMSE) (D7) of the three estimators for the i.i.d. sequences of random integers k of varying length n sampled from the zeta distributions with different values of exponent γ , Appendix D1. The integers are fed to the estimators as is, without adding the uniform symmetric noise. The RRMSE is larger than with noise, cf. Appendix D1.

consistent on integer-valued regularly varying distributions, but they tend to be unstable and exhibit slow convergence in such cases.

To resolve this issue we add uniform symmetric noise to the integer-valued sequences x_i , $i = 1, \dots, n$, that is, to all sequences considered in this paper. Specifically, instead of applying the estimators to x_i , we apply them to $y_i = x_i + u_i$, where u_i 's are i.i.d. samples from the uniform distribution on $[-1/2, 1/2]$. This does not affect the tail exponent: if x is a regularly varying random variable with tail index $\xi > 0$, and u is a uniform random variable on $[-1/2 \times 10^{-p}, 1/2 \times 10^{-p}]$, where $p \geq 0$, then $x + u$ is also regularly varying with the same exponent [108, Theorem 5.3.1]. Adding such noise greatly improves the stability and convergence of the estimators, see Fig. 5 and compare it with Appendix D1.

3. Example of the estimator operation using the double bootstrap method

To emphasize the importance of using as many consistent estimators as possible in application to degree sequences in real-world networks, here we consider an example of how the estimators work in conjunction with the double bootstrap method, showing that different estimators may explore

different parts of the empirical degree distribution for any finite sequence, thus explaining why they may return different estimations on such sequences, especially if the slowly varying function $\ell(k)$ is not trivial.

Figure 6 shows that the Hill estimator yields a higher estimation of $\alpha = 1/\xi$ than the other two estimators applied to the in-degree sequence of the *Libimseti* network. This happens because the value of the optimal number of order statistics κ^* returned by Hill's double bootstrap is substantially lower than for the other two estimators, so that the Hill estimator considers a smaller part of the distribution tail. The value of Hill's κ^* is smaller because it is based on finding the minimum of the AMSE as a function of the number of order statistics κ , and as we can see in the figure these minima occur at quite different values of κ for the Hill versus the two other estimators.

This effect is actually expected in small-sized sequences sampled from regularly varying distributions with nontrivial slowly varying functions $\ell(k)$. Figure 7 shows the details behind estimator convergence in two different cases, with a “nice” and “not so nice” slowly varying function $\ell(k)$. The figure illustrates the point that the farther the $\ell(k)$ is from a constant, the larger the network size must be for all the estimators to yield similar values of κ^* and $\hat{\xi}$. The estimators are guaranteed to converge to the true value of ξ for any $\ell(k)$, but only in the infinite sample limit $n \rightarrow \infty$, and, to the best of our knowledge, there are no results (for the bounds) on the speed of this convergence, partly because this speed may depend in an unknown way on some properties of $\ell(k)$. That is why using as many consistent estimators as possible in application to real-world data is the best strategy one can follow.

APPENDIX D: EVALUATION ON SYNTHETIC SEQUENCES AND NETWORK MODELS

Here we show that the estimators based on extreme value (EV) theory—the Hill, moments, and kernel estimators

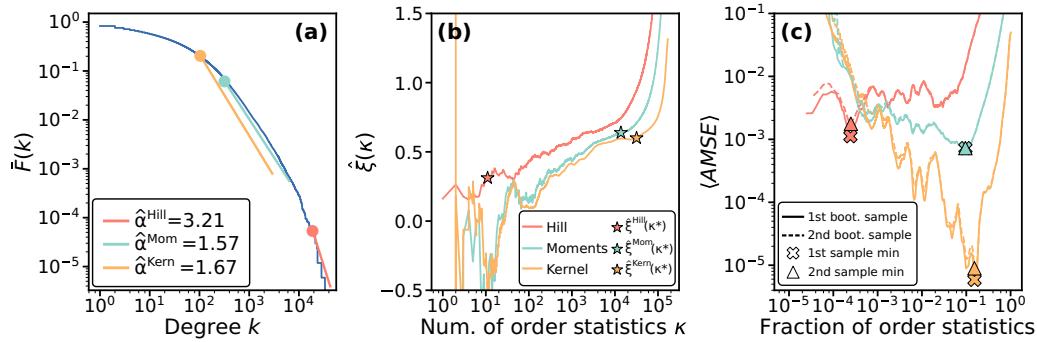


FIG. 6. An example of the estimator operation on the in-degree sequence of the *Libimseti* network, an online dating social network. (a) follows the same style and notations as in Fig. 2. (b) shows the estimated values $\hat{\xi}(\kappa)$ of the extreme value index ξ as a function of the number of order statistics κ . The filled symbols correspond to $\hat{\xi}(\kappa^*)$, where κ^* is the optimal value of κ found by the double bootstrap algorithm. (c) shows the averaged asymptotic mean squared error (AMSE), defined in (C1), as a function of the fraction $f = \kappa/n$ of the number of order statistics used in the two bootstrap samples of sizes $[n/\sqrt{2}]$ and $[n/2]$, corresponding to the “1st boot.” and “2nd boot.” curves in the figure. Their minima, given by (C2), are shown by the triangle and cross markers. The values of the number of order statistics κ corresponding to these minima— κ_1^* and κ_2^* for the first and second bootstrap samples, respectively—are then used to identify the optimal value κ^* via (C3).

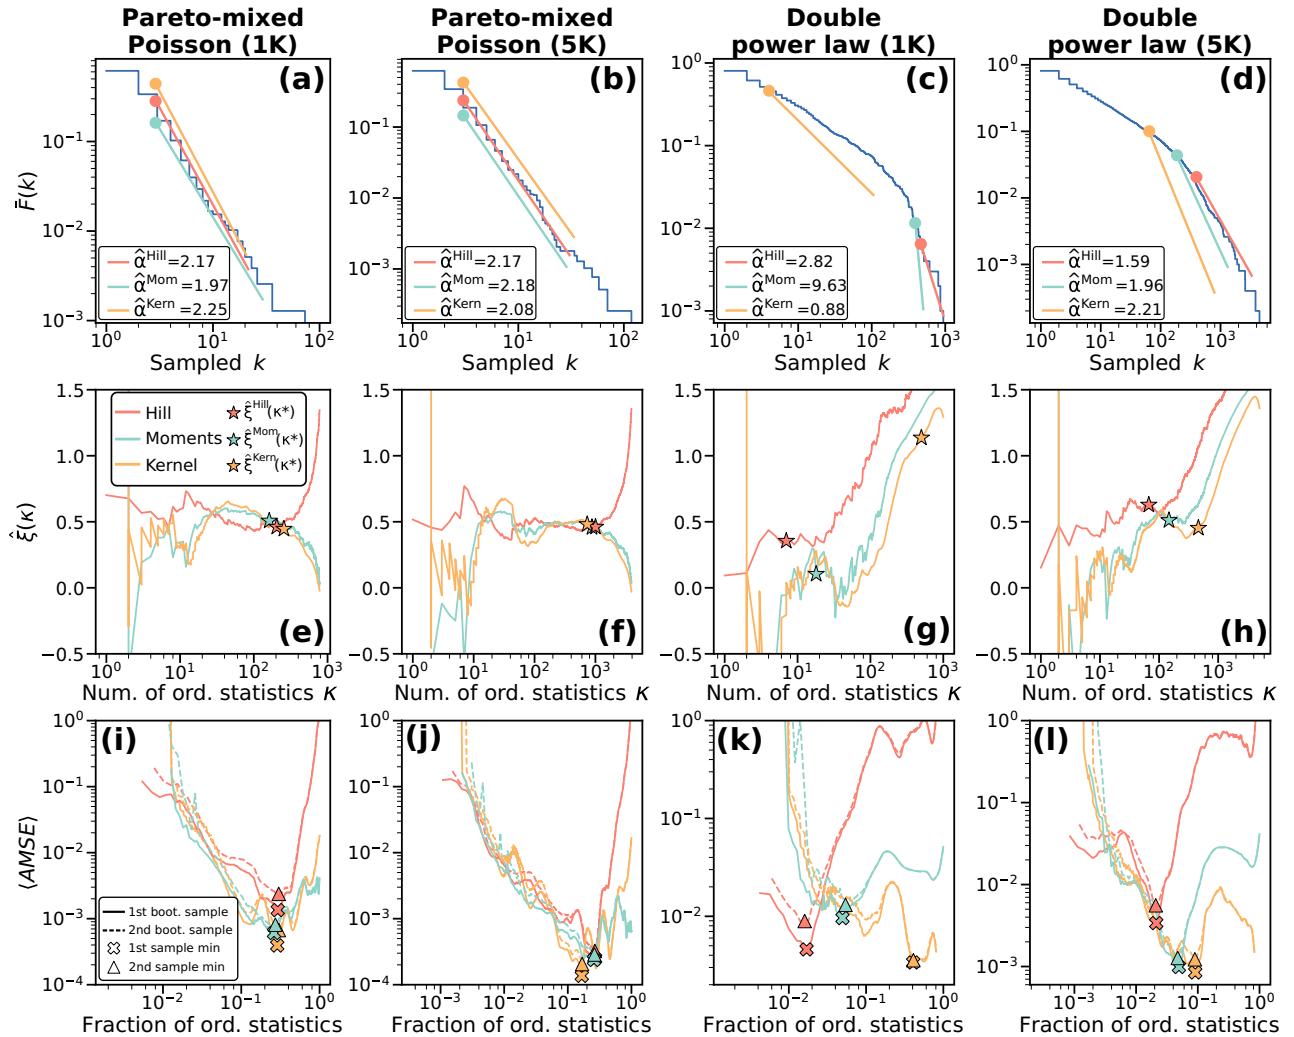


FIG. 7. An example of convergence of the estimators on the i.i.d. sequences of two different sizes, “small” $n_s = 1000$ and “large” $n_l = 5,000$, sampled from two regularly varying distributions with the same tail exponent $\gamma = 3$, but with different slowly varying functions $\ell(k)$: the Pareto-mixed Poisson distribution and double power law, see Sec. D 1 for details. The parameters for the Pareto-mixed Poisson are $\gamma = 3$ and $x_0 = 1$, while for the double power law, they are: $\gamma = 3$, $\gamma_0 = 1.5$, $c = 500$, $r = 0.1$, and $x_0 = 1$. The first, second, and third rows follow the same style and notations as panels (a,b,c) in Fig. 6, respectively. One can see that while all the estimators yield similar estimates for the sequences sampled from the Pareto-mixed Poisson distribution for both small and large sequence sizes, the estimates of different estimators for the small sequence size n_s are far apart in the case of the double power law distribution with “uglier” $\ell(k)$. In the latter case, the estimators start to agree on their estimates only for the large sequence size n_l . One can also see that the main reason for this effect is that the AMSE minima occur at far-apart locations if the sample size is small and $\ell(k)$ is not so “nice.”

equipped with the double bootstrap procedure, the code in Ref. [55]—yield the expected results when applied to synthetic degree sequences and to network models. We also compare the estimations that these estimators produce with the ones by the PLFIT [19], which is based on a combination of techniques inspired by maximum-likelihood estimation (MLE) and Kolmogorov-Smirnov (KS) distance minimization. We use the `plfit.m` code version 1.0.11 by Aaron Clauset, which is widely used and publicly available at [109]. As stated in the code comments of the `fit.py` script in Ref. [110]—the PYTHON implementation of the PLFIT used in Ref. [20]—this implementation is based on the original

MATLAB code [109], so that the results obtained using any of these two implementations [109,110] should be identical.

1. Synthetic sequences

Here we sample different numbers n of positive integers $k \in \mathbb{N}_+$ from the distributions listed below, so that the sampled sequence length is always n . The set of chosen distributions is intended to be diverse and representative of distributions claimed to be observed in real-world networks. In cases where the distribution has support on non-negative integers $k \in \mathbb{N}$, we discard all the zero entries from the sequence since

they would correspond to nodes of degree $k = 0$ in networks. The parameter γ in all the distributions below can be any real number greater than 1.

Zeta distribution. The distribution PDF (or PMF, to be precise) is

$$P(k) = \frac{k^{-\gamma}}{\zeta(\gamma)}, \quad k \in \mathbb{N}_+, \quad (\text{D1})$$

where $\zeta(\gamma)$ is the Riemann zeta function. This is the “clean” integer-valued power-law distribution with constant slowly varying function $\ell(k) = 1/\zeta(\gamma)$.

Pareto-mixed Poisson distribution. For each sample, we first sample a real number x from the Pareto distribution, and then sample an integer k from the Poisson distribution with mean x :

$$P(k|x) = \frac{x^k e^{-x}}{k!}, \quad k \in \mathbb{N}, \quad (\text{D2})$$

$$P(x) = \alpha x_0^\alpha x^{-\gamma}, \quad x \geq x_0 > 0, \quad (\text{D3})$$

where $\alpha = \gamma - 1$, and we set $x_0 = 1$ in the experiments. The Pareto-mixed Poisson distribution is ubiquitous in network models with hidden variables [90], also known in mathematics as inhomogeneous random graphs [91], and more generally, as graphon-based W -random graphs [92]. This is one of the simplest regularly varying distribution with nonconstant $\ell(k)$ that converges to a constant, $\ell(k) \rightarrow \alpha x_0^\alpha$, Appendix A 3.

Pareto distribution with natural exponential cutoff. We sample a random number x from the Pareto distribution with the exponential cutoff at n^ξ ,

$$P_n(x) = \frac{x_0^\alpha}{E_\gamma(x_0/n^\xi)} x^{-\gamma} e^{-x/n^\xi}, \quad x \geq x_0 > 0, \quad (\text{D4})$$

where E_γ is the exponential integral function, $\xi = 1/\alpha$, and $\alpha = \gamma - 1$. We then round x to the closest integer $k = [x]$. We set $x_0 = 1$. The value n^ξ of where the exponential decay becomes prominent corresponds to the natural cutoff [81], which is proportional to the exact expected maximum value (B6) among n i.i.d. samples from the Pareto distribution with exponent γ . This is an example of not a fixed distribution, but of an n -dependent family of distributions. For any fixed n , the distribution is not regularly varying since it has an exponential tail instead of a power-law tail. Yet as n increases, the location n^ξ of the “soft beginning” of the exponential tail diverges, so that in the $n \rightarrow \infty$ limit the distributions in this family converge to the pure Pareto distribution with exponent γ , which is regularly varying.

Pareto distribution with a constant exponential cutoff. The sampling is the same as in the previous example, except that the location of the exponential cutoff does not depend on n , and is fixed to be 10 instead of n^ξ . This is an example of a distribution which is not regularly varying.

Double power law. We sample a random number x from the double power-law distribution with the PDF

$$P(x) = \beta x^{-\gamma_0} \left(1 + \left(\frac{x}{c}\right)^{\alpha_0/r}\right)^{-r}, \quad x \geq x_0 > 0, \quad (\text{D5})$$

$$\alpha_0 = \gamma - \gamma_0, \quad \gamma \geq \gamma_0 > 1, \quad (\text{D6})$$

where c is the location of the switch between the two power laws with exponents γ_0 for $x \ll c$ and γ for $x \gg c$, r is the

parameter that controls how smooth this switch is, and β is the normalizing constant given by

$$\beta = \frac{\alpha x_0^\alpha}{c^{\alpha_0} {}_2F_1(r, r\alpha/\alpha_0, 1 + r\alpha/\alpha_0, -(c/x_0)^{\alpha_0/r})},$$

where $\alpha = \gamma - 1$ and ${}_2F_1$ is the Gauss hypergeometric function. As in all other examples, given this random x , we round it to integer $k = [x]$. In our experiments, we set $\gamma_0 = 1.5$, $c = 500$, $r = 0.1$, and $x_0 = 1$. This distribution is regularly varying with exponent γ , which we vary in the experiments. Yet, as discussed in Sec. V, it may be difficult for the estimators to see that it is indeed γ and not γ_0 if n is small. Distributions of this form characterize the degree distribution in the causal set of the universe [86], and they also frequently appear in astrophysics [111].

To assess the accuracy of the estimators, we sample $s = 100$ random sequences for each combination of the distributions listed above, the values of γ , and the numbers of samples n in a sequence. On each sampled sequence j , $j = 1, \dots, s$, each estimator returns an estimated value $\hat{\xi}_j$ of ξ . Given a collection of these $\hat{\xi}_j$ ’s, we compute the relative root mean squared error (RRMSE), a standard measure used to assess the accuracy of the tail index estimation [112,113], defined as

$$\text{RRMSE} = \frac{\sqrt{\frac{1}{s} \sum_{j=1}^s (\hat{\xi}_j - \xi)^2}}{\xi}, \quad (\text{D7})$$

and show the results in Fig. 8 both for the extreme value (EV) estimators (Hill, kernel, moments), and for the MLE-based PLFIT [19].

We observe that all the results are as expected. On sequences sampled from distributions that are regularly varying, all the EV estimators converge. They also converge in the case where the distributions are not regularly varying for any finite sample size n , but where they converge to a regularly varying distribution at $n \rightarrow \infty$ —the Pareto distribution with the diverging natural cutoff. No estimator converges in the case of a fixed distribution which is not regularly varying—the Pareto distribution with a fixed exponential cutoff.

Also as expected, the PLFIT yields a lower estimation error in case of the zeta distribution. This is because the zeta distribution satisfies PLFIT’s main assumption of a clean power law with constant $\ell(k)$, but does not satisfy the second-order condition, thus affecting the optimality of the double bootstrap, Appendix C 1. In other cases with reasonably “nice” regularly varying distributions with $\ell(k)$ quickly converging to a constant, the accuracy and convergence rates of the EV and PLFIT estimators are comparable. However, as soon as the regularly varying distribution is not really nice—the double power law case with nonconstant $\ell(k)$ over a wide range of degrees k —the PLFIT estimations are completely off, as opposed to the EV estimators. This is also expected for the reasons discussed in Appendix D 3.

2. Network models

The main motivation to test the performance of the EV estimators not only on synthetic sequences of numbers sampled from various distributions, but also on degree sequences in network models, is to see whether and how their performance

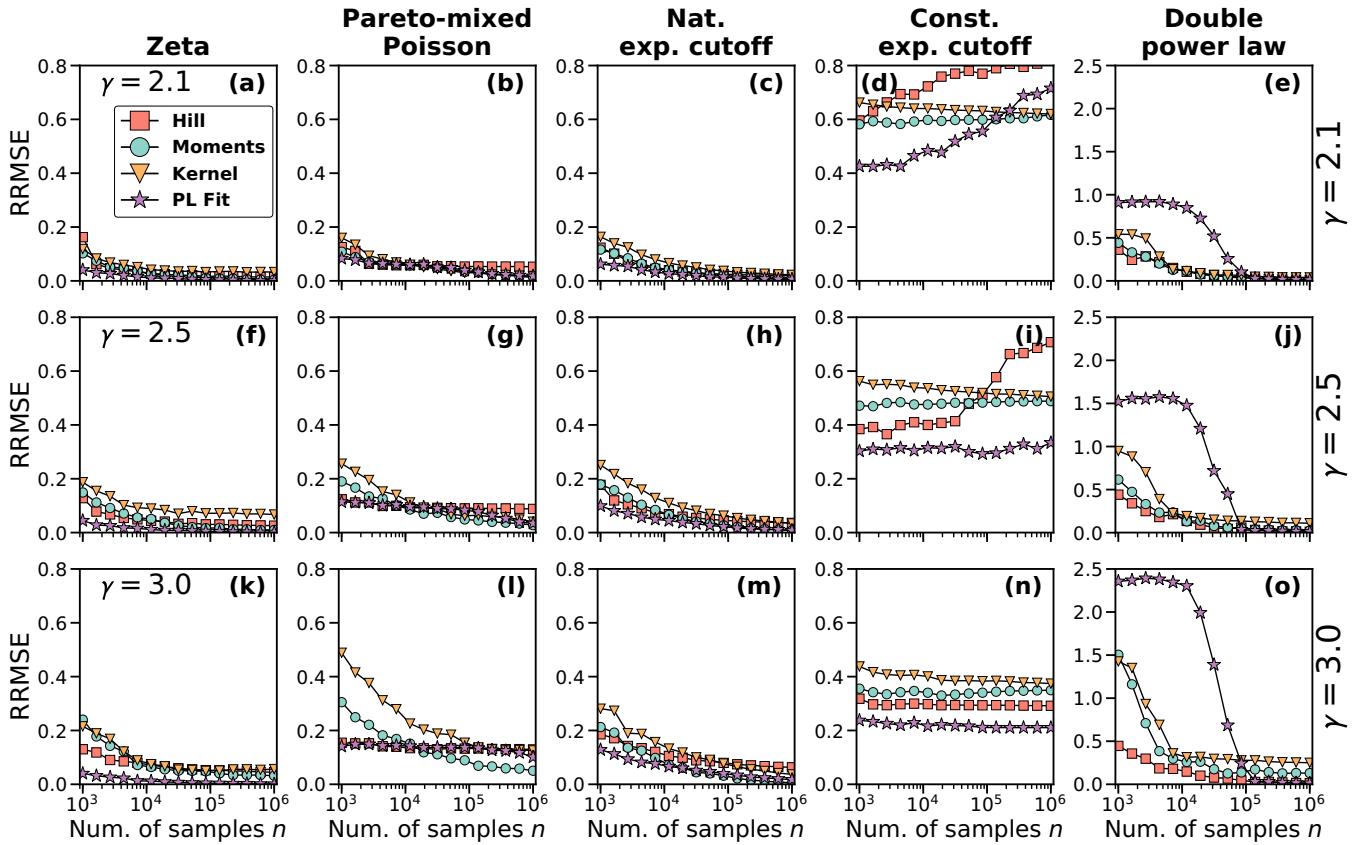


FIG. 8. The relative root mean squared error (RRMSE) (D7) of the three EV-based estimators and the MLE-based PLFIT algorithm [19] applied to i.i.d. sequences of random numbers sampled from the distributions described in Sec. D 1. The columns correspond to these distributions, while the rows are for the three values of γ —2.1, 2.5, and 3.0—used in these distributions to sample the number sequences. The sequence length varies in the range $n \in [10^3, 10^6]$ everywhere.

is affected by possible non-i.i.d.-ness of the latter sequences. To this end, we consider three paradigmatic network models in which the degree distributions have been proven to converge to a regularly varying distribution, and in which the degree sequences are not i.i.d.: (1) the erased configuration model (ECM) [114], (2) preferential attachment (PA) [115], and (3) hyperbolic random graphs (HRG) [116].

Erased configuration model. We sample varying-length i.i.d. sequences of random integers from the zeta distributions with different values of the exponent, and then either accept or reject the sequence based on whether the sum of its elements is even or odd. Each number in the sequence is the number of stubs attached to a node in a network to be formed. We match pairs of stubs uniformly at random, and then delete loops and multiedges. For any finite sample size, the degree sequence in the resulting network is neither zeta-distributed nor i.i.d., but it converges to the original zeta distribution as the sample size tends to infinity [114, Theorem 2.1].

Preferential attachment. We use the redirection implementation in Ref. [117]: starting with the first node of degree 0, nodes arrive one by one, and each new node picks an already existing node uniformly at random, and then connects either to it with probability $1 - r$, or to its random neighbor with probability r . The only exception is the second node that connects to the first node with probability 1. We use this redirection probability to control the exponent of the

power-law tail of the degree distribution, because this distribution converges to the following regularly varying distribution with exponent $\gamma = 1 + 1/r$ [117]:

$$P(k) = (\gamma - 1) \frac{\Gamma(2\gamma - 3)}{\Gamma(\gamma - 2)} \frac{\Gamma(k + \gamma - 3)}{\Gamma(k + 2\gamma - 3)}. \quad (\text{D8})$$

Hyperbolic random graphs. The degree distribution in random geometric graphs in hyperbolic spaces converges to regularly varying Pareto-mixed Poisson distributions (A4), and, as opposed to the previous two models, these graphs also have nonvanishing average local clustering coefficients [116]. We use the software package developed in [118] and available at [119] to generate these graphs. We fix the average degree parameter to $\bar{k} = 10$, the temperature parameter to $T = 0$ corresponding to strongest clustering, and vary the γ parameter.

For each model, we vary the γ over the three values $\gamma = 2.1, 2.5$, and 3.0 , and vary the network size n from 10^3 to 10^6 . For each combination of the model, γ , and n , we generate 100 random networks, read off their degree sequences, and feed them to all the considered estimators. We then compute the RRMSE (D7), and show the results in Fig. 9.

We observe that in all the considered cases, all the EV estimators converge, even though the degree sequences are not i.i.d. The slow convergence in some cases is explained by the slow convergence of the degree distributions in these

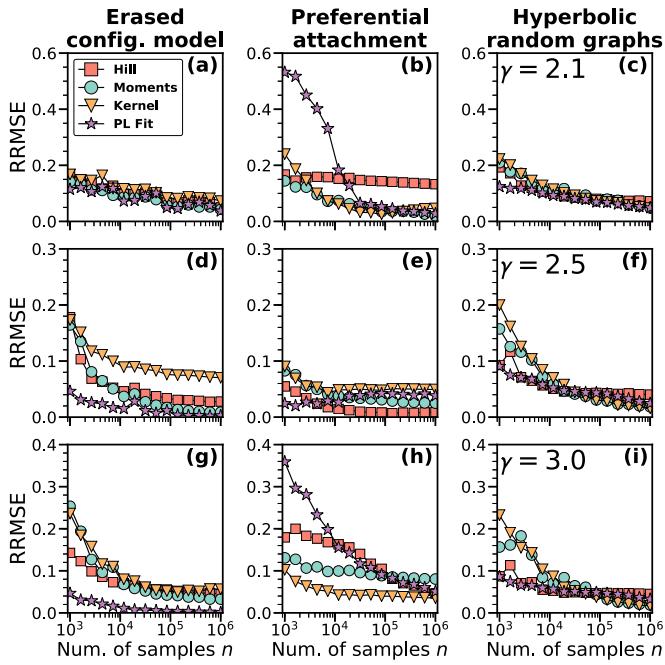


FIG. 9. The relative root mean squared error (RRMSE) (D7) of the three EV-based estimators and the MLE-based PLFIT algorithm [19] applied to the network models described in Appendix D 2. The first, second, and third columns show the results for the erased configuration model, preferential attachment, and hyperbolic random graphs, respectively. The first, second, and third rows show the results for $\gamma = 2.1$, 2.5 , and 3.0 , respectively. The network size varies in the range $n \in [10^3, 10^6]$ everywhere.

finite-sized networks to their limits. This is the case, for example, in the HRGs with $\gamma = 2.1$: the degree distribution in HRGs converges to its Pareto-mixed Poisson limit the more slowly, the closer the γ is to 2 [116].

The most notable results are for PA. Here the EV estimators clearly outperform the PLFIT if $\gamma = 2.1$ or 3, while all the estimators are on par if $\gamma = 2.5$ for the reasons that we discuss in the next section.

3. Anatomy of the PLFIT

To better understand the slow convergence of the PLFIT in the double power law and preferential attachment cases in the previous two sections, it is instructive to recall first how exactly the PLFIT algorithm works. The algorithm is a variation of estimators in [120,121] based on maximum-likelihood estimation (MLE). The starting point of the PLFIT operations is a sequence of possible γ values γ_s to experiment with. By default, this sequence is linearly spaced in the region $[1.5, 3.5]$ with step size 0.01 in the code [109] released with Ref. [19]. These default settings have to be manually changed for the code to be applicable to degree sequences coming from distributions with $\gamma > 3.5$. The default values of γ_s in the code [110] used in Ref. [20] are linearly spaced in $[1.01, 6.50]$ with step size 0.01.

Given the sequence γ_s and a degree sequence of length n supplied as input data, the PLFIT algorithm first finds the sequence of unique degree values k_t appearing in the degree

sequence. For each value k_t , the algorithm computes the vector of logarithmic-likelihood values

$$\mathcal{L}_{ts} = -n_t \ln \zeta(\gamma_s, k_t) - \gamma_s \sum_{i=1}^n \mathbb{1}\{k_i \geq k_t\} \ln k_i, \quad (\text{D9})$$

where n_t is the number of nodes with degrees $k_i \geq k_t$, $\mathbb{1}$ the indicator function, and

$$\zeta(\gamma_s, k_t) = \sum_{k=k_t}^{\infty} k^{-\gamma_s} \quad (\text{D10})$$

is the Hurwitz zeta function. This likelihood is based on the assumption that the degrees that are greater than or equal to k_t form a sequence of i.i.d. samples from a pure power law with exponent γ_s , i.e., from the generalized zeta distribution (9) with parameters γ_s , $k_{\min} = k_t$, and the normalization constant $c = 1/\zeta(\gamma_s, k_t)$. Among all the considered values γ_s , the algorithm then identifies the one, γ_t^* , that corresponds to the maximum value of \mathcal{L}_{ts} for the given k_t . This γ_t^* serves as an approximation of the MLE of γ for the degrees that are greater than or equal to k_t . For the same k_t , the algorithm then computes the Kolmogorov-Smirnov (KS) distance D_t^{KS} between the generalized zeta distribution with parameters γ_t^* and $k_{\min} = k_t$, and the empirical CDF of degrees $k_i \geq k_t$. This procedure is then repeated for each k_t observed in the sequence, and the estimates $\hat{\gamma}$ and \hat{k}_{\min} that the algorithm eventually returns are those that correspond to the minimum $D_{t^*}^{\text{KS}}$ of D_t^{KS} across all possible values of k_t , i.e., $\hat{k}_{\min} = k_{t^*}$ and $\hat{\gamma} = \gamma_{t^*}$.

The algorithm is thus a mixture of two optimization strategies: one is based on likelihood maximization, while the other one deals with the KS distance minimization. We note that since the algorithm does not implement MLE exactly, it trivially cannot be consistent if the true value of γ does not belong to the finite set of γ_s values, because it can never report any γ estimate $\hat{\gamma}$ that does not belong to the finite set of γ_s 's. More importantly, even though the correct implementation of MLE with a fixed and known k_{\min} had long been proven to be consistent [19], the consistency of MLE in combination with KS-distance minimization has been proven only very recently in Ref. [24], and only for pure power laws, i.e., for the Pareto or generalized zeta distributions. If the distribution is not a pure power law but a general regularly varying distribution, the consistency of the algorithm is a question that has not been rigorously explored at all, except the conjectures in [24] that this MLE-KS combination appears to be consistent for regularly varying distributions satisfying the second-order condition, and for regularly varying distributions whose slowly varying functions $\ell(k)$ converge to constants, and that the algorithm is likely not to be consistent for all other classes of regularly varying distributions. That is, in all these other cases the algorithm may be consistent, or it may not be.

The problem is that there is the following logical inconsistency in the algorithm: it operates under the assumption that above a certain k_{\min} , the distribution of degrees k is a pure power law, but then it recognizes that the distribution may be not a pure power law, and tries to account for that by finding a reasonable value of k_{\min} such that above this value the assumption would hold “approximately.” If the distribution

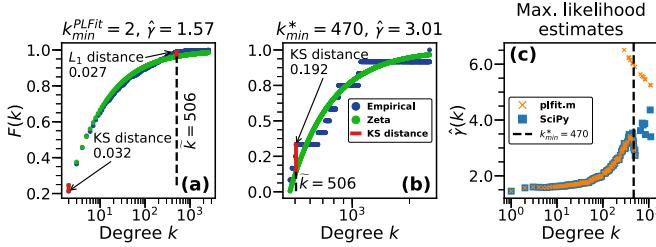


FIG. 10. Anatomy of Kolmogorov-Smirnov (KS) distance minimization on a sequence of random numbers sampled from the double power-law distribution with the parameters as in Appendix D 1, tail exponent $\gamma = 3$ and size $n = 1000$. Applied to this degree sequence, the PLFIT algorithm estimates $k_{\min}^{\text{PLFIT}} = 2$ and $\hat{\gamma}(k_{\min}^{\text{PLFIT}}) = 1.57$, the MLE estimate of γ with $k_{\min} = 2$. (a) shows the empirical and theoretical (generalized zeta) cumulative distribution functions (CDFs) with these parameters. The KS distance (0.032) is marked. (b) shows the same CDFs but for different $k_{\min}^* = 470$ and $\hat{\gamma}(k_{\min}^*) = 3.01$, the MLE estimate of γ with $k_{\min} = 470$. This k_{\min} is optimal in this sequence, in the sense that the MLE value of γ with this k_{\min} is closest to the true γ across all other degree values present in the sequence to which k_{\min} can be set. Yet the marked KS distance (0.192), achieved at $\tilde{k} = 506$, is greater than the KS distance achieved at a different location in (a), where the L^1 distance (0.027) between the two CDFs at \tilde{k} is also shown. (c) shows a collection of numerical errors produced by the PLFIT if modified to compute the MLE values of γ for large k_{\min} 's. The errors are due to the numerically incorrect computations of the Hurwitz zeta function in the `plfit.m` code [109].

was a pure power law, then the search for this k_{\min} would not be necessary, since the value of k_{\min} would be equal with high probability to the smallest value observed in the sequence. However, if the distribution is not a pure power law, then such a value of k_{\min} simply does not exist, since for any k_{\min} the distribution of $k > k_{\min}$ is *not* a pure power law. Yet the distribution of such k 's may converge to a pure power law, but only if the k_{\min} value that the algorithm finds diverges with the sample size n , and only if the slowly varying function $\ell(k)$ converges to a constant. Therefore, for this subclass of regularly varying distributions with $\ell(k)$'s converging to constants, the algorithm is likely to be consistent, yet the full proof is currently lacking [24]. If $\ell(k)$ does not converge to a constant, then the consistency of the algorithm is quite unclear at present.

In all the regularly varying distributions considered in the previous two sections, the function $\ell(k)$ does converge to a constant, and indeed in all these cases the PLFIT appears to eventually converge. Yet in two of these cases, namely, double power laws and preferential attachment, its convergence is worse than that of any of the considered EV estimators. To see why, we analyze the two components of the PLFIT, KS distance minimization and likelihood maximization, separately—in Figs. 10 and 11, respectively.

Figure 10 illustrates that the KS distance minimization component of the PLFIT drives the values of k_{\min} that the PLFIT attempts to estimate to erroneously low values, in full agreement with more recent and in-depth investigations in Ref. [24]. This happens because the smaller the k_{\min} , the smaller the deviations of the empirical CDF at degrees k

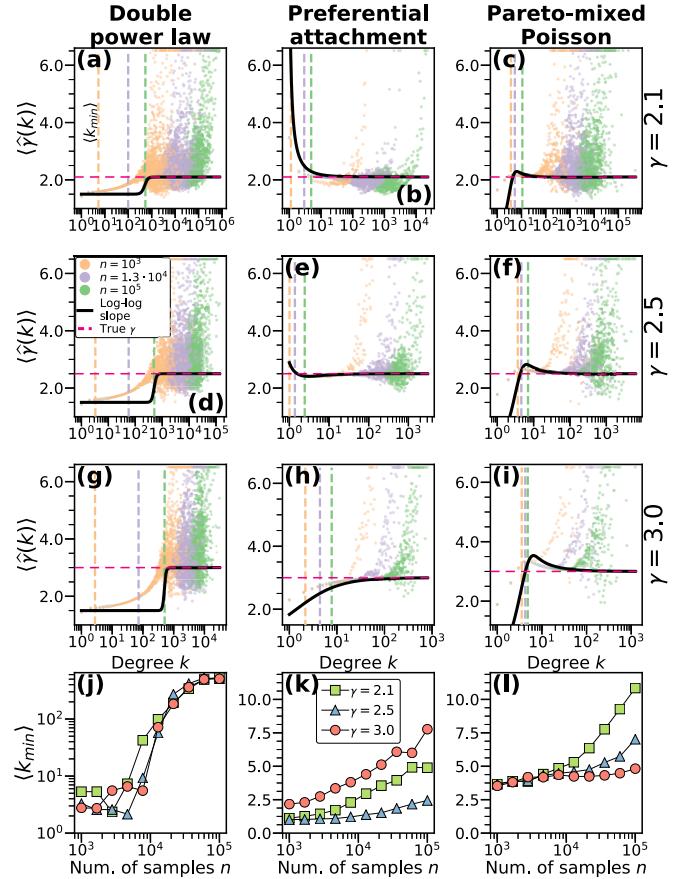


FIG. 11. Anatomy of maximum-likelihood estimation (MLE) of γ on different sequences of different sizes. The columns, left to right, correspond to (1) number sequences sampled from the double power-law distribution (D5) with the same parameters as in Appendix D 1; (2) degree sequences of random graphs in the preferential attachment model with the degree distribution given by (D8); and (3) number sequences sampled from the Pareto-mixed Poisson distribution (A4). The first three rows show the results for $\gamma = 2.1, 2.5$, and 3.0 , while the last row show the average values of k_{\min} found by the PLFIT algorithm for different γ 's as functions of sequence length n . The first three rows show the results only for three different sequence sizes n , marked by different colors. The colored dots are the average values, computed using the SCIPY package [122], of the MLE $\hat{\gamma}$ estimates obtained by the minimization of (D9) over γ_s linearly spaced in $[1.01, 6.50]$ with step 0.01, shown as functions of $k_{\min} = k_t$, where k_t is a degree value appearing at least once in the collection of sequences. For any combination of the model, γ , and n , the results are averaged over 100 random sequences. The solid black curves in the first three rows are the loglog slopes of the corresponding theoretical PDFs considered as functions of continuous k . The horizontal dashed lines are the true γ 's, while the vertical dashed lines color-correspond to the n 's, showing the average k_{\min} 's for these n 's from the bottom row.

right above k_{\min} from the theoretical CDF, because if the distribution is regularly varying, there are more nodes with smaller degrees. The larger the k_{\min} , the larger are these deviations caused by “sparser statistics” in the distribution tail, and as a consequence the KS distance grows larger. If k_{\min} is set to a small value, the deviations in the tail are

suppressed as they are getting “squished” in the high-degree region of the CDF close to 1, cf. panels (a) and (b) in the figure.

Panel (c) in Fig. 10 shows that the `plfit.m` code [109], both originally released with [19] as well as its current version, cannot be used to compute the MLE values of γ for large k_{\min} ’s, because it contains errors in computing the Hurwitz zeta function with the required accuracy [121], leading to numerical errors. Therefore we use a SciPy [122] implementation instead in Fig. 11.

Figure 11 shows that if k_{\min} is small—and it is, thanks to the KS distance minimization part of the PLFIT—then the MLE component of the PLFIT does very little other than trying to fit the log-log slope of the PDF evaluated at this k_{\min} . The reasons behind this problem are the same as those behind the KS distance minimization problems discussed above: since there is a lot of data with degrees right above a small value of k_{\min} , and since the MLE is primarily concerned with fitting as much data as possible, it tries to fit the part of the distribution with degrees k right above k_{\min} , versus the true tail of the distribution with large k ’s, thus getting bad estimates of the tail exponent.

In other words, the errors in PLFIT’s estimates are due to the combination of the following two factors related, ironically, to the two key ideas behind the PLFIT: (1) the small values of k_{\min} returned by the KS distance minimization part of the algorithm, and (2) the MLE part of the PLFIT that estimates not the tail exponent but, roughly, the loglog slope of the PDF evaluated close to this k_{\min} . If this slope is different from the slope at large k ’s, i.e., the true tail exponent, then the PLFIT does not really fit any power-law tail. However, if the distribution is such that at least one of these conditions is not satisfied, then the PLFIT estimates are more accurate, cf. the Pareto-mixed Poisson or the $\gamma = 2.5$ preferential attachment cases in Fig. 11.

If these two conditions are satisfied, which is the case with the double power law and preferential attachment with $\gamma = 2.1$ and 3.0 in Fig. 11, then the PLFIT estimates of the tail exponent are quite off. However, if they are off, and if one

then performs KS hypothesis testing using these inaccurate estimates, then the hypothesis that the degree sequence comes from a pure power law with the estimated exponent will be rejected with high probability, simply because the true tail exponent is different. For these reasons, if one applies the PLFIT to preferential attachment networks with these exponents and then deploys the KS hypothesis tests, one will likely find that these networks are not power-law [20].

Finally, Fig. 11 also shows that the whole idea of using MLE to estimate tail exponents of regularly varying distributions is quite problematic to begin with, explaining why it has not been seriously explored in statistics. Indeed, for such an estimation procedure to be accurate, the values of k_{\min} must be large and diverging in the $n \rightarrow \infty$ limit for the reasons discussed above. However, the larger the k_{\min} , the smaller the second term in (D9). On the other hand, as a function of γ_s , the first term in (D9) grows monotonically at a much higher rate than the linear rate of growth of the second term. Therefore, if k_{\min} is above a certain threshold, then the MLE will do nothing but select the largest possible value of γ_s to maximize the likelihood via the first term. This is exactly what we see in Fig. 11, where for many instances of large k_{\min} , the MLE-selected values of γ are the largest possible values within the range that we offer the MLE to operate with. Therefore the correctness of MLE depends on whether there exists a “sweet-spot” range of values of k_{\min} that are not too large and not too small. It might be the case that such a range simply does not exist for some regularly varying distributions. Worse, even if it can be proven to always exist, which is unclear at present, we have seen above that the KS distance minimization procedure is quite unlikely to be a correct, statistically consistent, procedure to identify this range. At least, the KS distance minimization has not been proven to be such a procedure for general regularly varying distributions. Whether a required procedure exists at all, is also unclear. After all, for the reasons mentioned above, even the required sweet-spot range of k_{\min} ’s is quite unlikely to exist for regularly varying distributions whose slowly varying functions do not converge to constants.

- [1] A.-L. Barabási, *Network Science* (Cambridge University Press, Cambridge, 2016).
- [2] M. E. J. Newman, *Networks*, 2nd ed. (Oxford University Press, Oxford, 2018).
- [3] A. Barrat, M. Barthélémy, and A. Vespignani, *Dynamical Processes on Complex Networks* (Cambridge University Press, Cambridge, 2008).
- [4] *Handbook of Graph and Networks: From the Genome to the Internet*, edited by S. Bornholdt and H. G. Schuster (Wiley-VCH, Berlin, 2002).
- [5] R. van der Hofstad, *Random Graphs and Complex Networks* (Cambridge University Press, Cambridge, 2016), Vol. 1.
- [6] S. N. Dorogovtsev, A. V. Goltsev, and J. F. F. Mendes, Critical phenomena in complex networks, *Rev. Mod. Phys.* **80**, 1275 (2008).
- [7] A. Arenas, A. Díaz-Guilera, J. Kurths, Y. Moreno, and C. Zhou, Synchronization in complex networks, *Phys. Rep.* **469**, 93 (2008).
- [8] L. Dall’Asta, Dynamical phenomena on complex networks, Ph.D. thesis, Université Paris Sud, 2006.
- [9] W. Willinger, R. Govindan, S. Jamin, V. Paxson, and S. Shenker, Scaling phenomena in the internet: Critically examining criticality, *Proc. Natl. Acad. Sci. USA* **99**, 2573 (2002).
- [10] L. Li, D. Alderson, J. C. Doyle, and W. Willinger, Towards a theory of scale-free graphs: Definition, properties, and implications, *Internet Math.* **2**, 431 (2005).
- [11] *Network Science*, edited by C. B. Duke (The National Academies Press, Washington, 2006).
- [12] D. Krioukov, K. Claffy, M. Fomenkov, F. Chung, A. Vespignani, and W. Willinger, The workshop on internet topology (WIT) report, *ACM SIGCOMM Comput. Commun. Rev.* **37**, 69 (2007).
- [13] W. Willinger, D. Alderson, and J. C. Doyle, Mathematics and the internet: A source of enormous confusion and great potential, *Not. Am. Math. Soc.* **56**, 586 (2009).

[14] M. Mitzenmacher, A brief history of generative models for power law and lognormal distributions, *Internet Math.* **1**, 226 (2004).

[15] R. Khanin and E. Wit, How scale-free are biological networks, *J. Comput. Biol.* **13**, 810 (2006).

[16] M. P. H. Stumpf and M. A. Porter, Critical truths about power laws, *Science* **335**, 665 (2012).

[17] Á. Corral, F. Font, and J. Camacho, Noncharacteristic half-lives in radioactive decay, *Phys. Rev. E* **83**, 066103 (2011).

[18] Á. Corral and Á. González, Power-law distributions in geo-science revisited, *Earth Space Sci.* **6**, 673 (2019).

[19] A. Clauset, C. R. Shalizi, and M. E. Newman, Power-law distributions in empirical data, *SIAM Rev.* **51**, 661 (2009).

[20] A. D. Broido and A. Clauset, Scale-free networks are rare, *Nat. Commun.* **10**, 1017 (2019).

[21] E. Klarreich, Scant Evidence of Power Laws Found in Real-World Networks (2018).

[22] P. Holme, Rare and everywhere: Perspectives on scale-free networks, *Nat. Commun.* **10**, 1016 (2019).

[23] J. Lee, L. F. James, S. Choi, and F. Caron, A Bayesian model for sparse graphs with flexible degree distribution and overlapping community structure, [arXiv:1810.01778](https://arxiv.org/abs/1810.01778).

[24] H. Drees, A. Janßen, S. I. Resnick, and T. Wang, On a minimum distance procedure for threshold selection in tail analysis, [arXiv:1811.06433](https://arxiv.org/abs/1811.06433).

[25] M. Gerlach and E. G. Altmann, Testing Statistical Laws in Complex Systems, *Phys. Rev. Lett.* **122**, 168301 (2019).

[26] M. Serafino, G. Cimini, A. Maritan, S. Suweis, J. R. Banavar, and G. Caldarelli, Scale-free networks revealed from finite-size scaling, [arXiv:1905.09512](https://arxiv.org/abs/1905.09512).

[27] A. Charpentier and E. Flachaire, Extended Scale-Free Networks, [arXiv:1905.10267](https://arxiv.org/abs/1905.10267).

[28] S. N. Dorogovtsev, J. F. F. Mendes, and A. N. Samukhin, Structure of Growing Networks with Preferential Linking, *Phys. Rev. Lett.* **85**, 4633 (2000).

[29] P. L. Krapivsky, S. Redner, and F. Leyvraz, Connectivity of Growing Random Networks, *Phys. Rev. Lett.* **85**, 4629 (2000).

[30] B. Bollobás, O. Riordan, J. Spencer, and G. Tusnády, The degree sequence of a scale-free random graph process, *Random Struct. Alg.* **18**, 279 (2001).

[31] S. I. Resnick, *Heavy-Tail Phenomena: Probabilistic and Statistical Modeling* (Springer, New York, 2007).

[32] N. H. Bingham, C. M. Goldie, and J. L. Teugels, *Regular Variation* (Cambridge University Press, Cambridge, 1989).

[33] S. Foss, D. Korshunov, and S. Zachary, *An Introduction to Heavy-Tailed and Subexponential Distributions* (Springer, New York, 2011), Vol. 6.

[34] J. Beirlant, Y. Goegebeur, J. Segers, and J. L. Teugels, *Statistics of Extremes: Theory and Applications* (Wiley, Chichester, 2006).

[35] C. Stegehuis, R. van der Hofstad, A. J. E. M. Janssen, and J. S. H. van Leeuwaarden, Clustering spectrum of scale-free networks, *Phys. Rev. E* **96**, 042309 (2017).

[36] R. van der Hofstad, G. Hooghiemstra, and D. Znamenski, Distances in random graphs with finite mean and infinite variance degrees, *Electron. J. Probab.* **12**, 703 (2007).

[37] P. van der Hoorn and M. Olvera-Cravioto, Typical distances in the directed configuration model, *Ann. Appl. Probab.* **28**, 1739 (2018).

[38] M. Boguñá, D. Krioukov, and K. C. Claffy, Navigability of complex networks, *Nat. Phys.* **5**, 74 (2009).

[39] S. A. Delre, W. Jager, T. H. Bijmolt, and M. A. Janssen, Will it spread or not? the effects of social influences and network topology on innovation diffusion, *J. Prod. Innov. Manag.* **27**, 267 (2010).

[40] B. Doerr, M. Fouz, and T. Friedrich, Why rumors spread so quickly in social networks, *Commun. ACM* **55**, 70 (2012).

[41] K. Bringmann, R. Keusch, J. Lengler, Y. Maus, and A. Molla, in *Proceedings of the ACM Symposium on Principles of Distributed Computing - PODC '17* (ACM, New York, NY, USA, 2017), pp. 371–380.

[42] A. Ameraoui, K. Boukhetala, and J.-F. Dupuy, Bayesian estimation of the tail index of a heavy tailed distribution under random censoring, *Comput. Stat. Data Anal.* **104**, 148 (2016).

[43] P. Embrechts, C. Klüppelberg, and T. Mikosch, *Modelling Extremal Events: For Insurance and Finance* (Springer, Berlin, 2013), Vol. 33.

[44] S. Boucheron and M. Thomas, Tail index estimation, concentration and adaptivity, *Electron. J. Stat.* **9**, 2751 (2015).

[45] A. J. McNeil and R. Frey, Estimation of tail-related risk measures for heteroscedastic financial time series: An extreme value approach, *J. Empir. Finance.* **7**, 271 (2000).

[46] D. W. Jansen and C. G. De Vries, On the frequency of large stock returns: Putting booms and busts into perspective, *Rev. Econ. Stat.* **73**, 18 (1991).

[47] J. H. McCulloch, 13 financial applications of stable distributions, *Handbook of statistics* **14**, 393 (1996).

[48] M. Kotulski, Asymptotic distributions of continuous-time random walks: a probabilistic approach, *J. Stat. Phys.* **81**, 777 (1995).

[49] R. Metzler and J. Klafter, The random walk's guide to anomalous diffusion: A fractional dynamics approach, *Phys. Rep.* **339**, 1 (2000).

[50] S. Lu and F. J. Molz, How well are hydraulic conductivity variations approximated by additive stable processes? *Adv. Environ. Res.* **5**, 39 (2001).

[51] S. I. Resnick, Heavy tail modeling and teletraffic data: Special invited paper, *Ann. Stat.* **25**, 1805 (1997).

[52] C. L. Nikias and M. Shao, *Signal Processing with Alpha-Stable Distributions and Applications* (Wiley-Interscience, New York, 1995).

[53] R. L. Wasserstein and N. A. Lazar, The ASA's statement on p-values: Context, process, and purpose, *Am. Stat.* **70**, 129 (2016).

[54] R. L. Wasserstein, A. L. Schirm, and N. A. Lazar, Moving to a world beyond “ $p < 0.05$ ”, *Am. Stat.* **73**, 1 (2019).

[55] I. Voitalov, Tail Index Estimation for Degree Sequences of Complex Networks, <https://github.com/ivanvoitalov/tail-estimation> (2018).

[56] J. Kunegis, in *Proceedings of the 22nd International Conference on World Wide Web* (ACM, New York, NY, USA, 2013), pp. 1343–1350.

[57] B. M. Hill, A simple general approach to inference about the tail of a distribution, *Ann. Stat.* **3**, 1163 (1975).

[58] A. L. Dekkers, J. H. Einmahl, and L. De Haan, A moment estimator for the index of an extreme-value distribution, *Ann. Stat.* **17**, 1833 (1989).

[59] P. Groeneboom, H. Lopuhaä, and P. De Wolf, Kernel-type estimators for the extreme value index, *Ann. Stat.* **31**, 1956 (2003).

[60] T. Wang and S. I. Resnick, Consistency of Hill estimators in a linear preferential attachment model, *Extremes* **22**, 1 (2019).

[61] E. A. Bender and E. R. Canfield, The asymptotic number of labeled graphs with given degree sequences, *J. Comb. Theory, Ser. A* **24**, 296 (1978).

[62] B. Bollobás, A probabilistic proof of an asymptotic formula for the number of labelled regular graphs, *Eur. J. Comb.* **1**, 311 (1980).

[63] N. C. Wormald, Some problems in the enumeration of labelled graphs, *B. Aust. Math. Soc.* **21**, 159 (1980).

[64] R. Arratia and T. M. Liggett, How likely is an i.i.d. degree sequence to be graphical? *Ann. Appl. Probab.* **15**, 652 (2005).

[65] C. I. Del Genio, H. Kim, Z. Toroczkai, and K. E. Bassler, Efficient and exact sampling of simple graphs with given arbitrary degree sequence, *PLoS One* **5**, e10012 (2010).

[66] K. F. Chan and P. Gray, Using extreme value theory to measure value-at-risk for daily electricity spot prices, *Int. J. Forecast.* **22**, 283 (2006).

[67] J. Danielsson and C. G. De Vries, Value-at-risk and extreme returns, *Ann. Econ. Statist.*, 239 (2000).

[68] M. Gilli and E. Kellezi, An application of extreme value theory for measuring financial risk, *Comput. Econ.* **27**, 207 (2006).

[69] R. A. Fisher and L. H. C. Tippett, Limiting forms of the frequency distribution of the largest or smallest member of a sample, *Math. Proc. Cambridge Philos. Soc.* **24**, 180 (1928).

[70] B. Gnedenko, Sur la distribution limite du terme maximum d'une série aléatoire, *Ann. Math.* **44**, 423 (1943).

[71] M. Matsui, T. Mikosch, and L. Tafakori, Estimation of the tail index for lattice-valued sequences, *Extremes* **16**, 429 (2013).

[72] J. Pickands III, Statistical inference using extreme order statistics, *Ann. Stat.* **3**, 119 (1975).

[73] J. Segers, Generalized pickands estimators for the extreme value index, *J. Stat. Plan. Infer.* **128**, 381 (2005).

[74] W. L. Shinyie, N. Ismail, and A. A. Jemain, Semi-parametric estimation for selecting optimal threshold of extreme rainfall events, *Water Resour. Manage.* **27**, 2325 (2013).

[75] S. Müller and K. Rufibach, Smooth tail-index estimation, *J. Stat. Comput. Sim.* **79**, 1155 (2009).

[76] M. I. Gomes and A. Guillou, Extreme value theory and statistics of univariate extremes: A review, *Int. Stat. Rev.* **83**, 263 (2015).

[77] P. Billingsley, *Convergence of Probability Measures* (Wiley, New York, 2013).

[78] P. Krapivsky and D. Krioukov, Scale-free networks as preasymptotic regimes of superlinear preferential attachment, *Phys. Rev. E* **78**, 026114 (2008).

[79] S. Dommers, C. Giardinà, and R. van der Hofstad, Ising critical exponents on random trees and graphs, *Commun. Math. Phys.* **328**, 355 (2014).

[80] R. Cohen and S. Havlin, Scale-Free Networks are Ultrasmall, *Phys. Rev. Lett.* **90**, 058701 (2003).

[81] M. Boguñá, R. Pastor-Satorras, and A. Vespignani, Cut-offs and finite size effects in scale-free networks, *Eur. Phys. J. B* **38**, 205 (2004).

[82] A. Clauset, E. Tucker, and M. Sainz, The Colorado Index of Complex Networks, <https://icon.colorado.edu/> (2016).

[83] C. Orsini, M. M. Dankulov, P. Colomer-de Simón, A. Jamakovic, P. Mahadevan, A. Vahdat, K. E. Bassler, Z. Toroczkai, M. Boguñá, G. Caldarelli, S. Fortunato, and D. Krioukov, Quantifying randomness in real networks, *Nat. Commun.* **6**, 8627 (2015).

[84] S. A. Hill, A measure for characterizing heavy-tailed networks, [arXiv:1907.04808](https://arxiv.org/abs/1907.04808).

[85] N. Berger, C. Borgs, J. T. Chayes, R. M. D'Souza, and R. D. Kleinberg, Degree distribution of competition-induced preferential attachment graphs, *Comb. Probab. Comput.* **14**, 697 (2005).

[86] D. Krioukov, M. Kitsak, R. S. Sinkovits, D. Rideout, D. Meyer, and M. Boguñá, Network cosmology, *Sci. Rep.* **2**, 793 (2012).

[87] D. Yao, P. van der Hoorn, and N. Litvak, Average nearest neighbor degrees in scale-free networks, *Internet Math.*, 1 (2018).

[88] H. Jessen and T. Mikosch, Regularly varying functions, *Publ. l'Institut. Math.* **80**, 171 (2006).

[89] G. E. Willmot and X. S. Lin, in *Lundberg Approximations for Compound Distributions with Insurance Applications* (Springer, New York, 2001), pp. 37–49.

[90] M. Boguñá and R. Pastor-Satorras, Class of correlated random networks with hidden variables, *Phys. Rev. E* **68**, 036112 (2003).

[91] B. Bollobás, S. Janson, and O. Riordan, The phase transition in inhomogeneous random graphs, *Random Struct. Alg.* **31**, 3 (2007).

[92] L. Lovász and B. Szegedy, Limits of dense graph sequences, *J. Comb. Theory, Ser. B* **96**, 933 (2006).

[93] S. I. Resnick, *Extreme Values, Regular Variation and Point Processes* (Springer, New York, 2013).

[94] T. Mikosch, *Regular Variation, Subexponentiality and Their Applications in Probability Theory*, Report Eurandom Vol. 99013 (Eurandom, Eindhoven, 1999).

[95] S. Resnick and C. Stărică, Smoothing the Hill estimator, *Adv. Appl. Probab.* **29**, 271 (1997).

[96] T. T. Pereira, in *Extreme Value Theory and Applications III, Proc. Gaithersburg Conference (NIST Special Publication)* (NIST, Gaithersburg, 1993), p. 165.

[97] M. Falk, Efficiency of convex combinations of pickands estimator of the extreme value index, *J. Nonparametr. Stat.* **4**, 133 (1994).

[98] M. F. Alves, Estimation of the tail parameter in the domain of attraction of an extremal distribution, *J. Stat. Plan. Infer.* **45**, 143 (1995).

[99] H. Drees, Refined pickands estimators of the extreme value index, *Ann. Stat.* **23**, 2059 (1995).

[100] S. Yun, On a generalized pickands estimator of the extreme value index, *J. Stat. Plan. Infer.* **102**, 389 (2002).

[101] G. Draisma, L. de Haan, L. Peng, and T. T. Pereira, A bootstrap-based method to achieve optimality in estimating the extreme-value index, *Extremes* **2**, 367 (1999).

[102] Y. Qi, Bootstrap and empirical likelihood methods in extremes, *Extremes* **11**, 81 (2008).

[103] J. Danielsson, L. de Haan, L. Peng, and C. G. de Vries, Using a bootstrap method to choose the sample fraction in tail index estimation, *J. Multivar. Anal.* **76**, 226 (2001).

[104] L. de Haan and U. Stadtmüller, Generalized regular variation of second order, *J. Aust. Math. Soc.* **61**, 381 (1996).

[105] L. de Haan and A. Ferreira, *Extreme Value Theory: An Introduction* (Springer, New York, 2007).

[106] L. de Haan and S. Resnick, Second-order regular variation and rates of convergence in extreme-value theory, *Ann. Probab.* **24**, 97 (1996).

[107] P. Jordanova and M. Stehlík, Flexible extreme value inference and Hill plots for a small, mid and large samples, [arXiv:1509.06718](https://arxiv.org/abs/1509.06718).

[108] J. J. Velthoen, Estimation of the extreme value index for imprecise data, B.S. thesis, Delft University of Technology, 2014.

[109] A. Clauset, PLFIT Matlab Script, <http://tuvalu.santafe.edu/~aaronc/powerlaws/plfit.m> (2012).

[110] A. Broido, Scale-free network analysis, <https://github.com/adbroido/SFAnalysis> (2019).

[111] D. Huppenkothen, A. L. Watts, P. Uttley, A. J. Van der Horst, M. Van der Klis, C. Kouveliotou, E. Göögüs, J. Granot, S. Vaughan, and M. H. Finger, Quasi-periodic oscillations and broadband variability in short magnetar bursts, *Astrophys. J.* **768**, 87 (2013).

[112] M. Brzezinski, Robust Estimation of the pareto tail index: A Monte Carlo analysis, *Empir. Econ.* **51**, 1 (2016).

[113] P. Jordanova, Z. Fabián, P. Hermann, L. Střelec, A. Rivera, S. Girard, S. Torres, and M. Stehlík, Weak properties and robustness of t-hill estimators, *Extremes* **19**, 591 (2016).

[114] T. Britton, M. Deijfen, and A. Martin-Löf, Generating simple random graphs with prescribed degree distribution, *J. Stat. Phys.* **124**, 1377 (2006).

[115] A.-L. Barabási and R. Albert, Emergence of scaling in random networks, *Science* **286**, 509 (1999).

[116] D. Krioukov, F. Papadopoulos, M. Kitsak, A. Vahdat, and M. Boguná, Hyperbolic geometry of complex networks, *Phys. Rev. E* **82**, 036106 (2010).

[117] P. L. Krapivsky and S. Redner, Organization of growing random networks, *Phys. Rev. E* **63**, 066123 (2001).

[118] R. Aldecoa, C. Orsini, and D. Krioukov, Hyperbolic graph generator, *Comput. Phys. Commun.* **196**, 492 (2015).

[119] C. Orsini and R. Aldecoa, Hyperbolic Graph Generator, <https://github.com/named-data/Hyperbolic-Graph-Generator/releases> (2016).

[120] M. L. Goldstein, S. A. Morris, and G. G. Yen, Problems with fitting to the power-law distribution, *Eur. Phys. J. B* **41**, 255 (2004).

[121] H. Bauke, Parameter estimation for power-law distributions by maximum likelihood methods, *Eur. Phys. J. B* **58**, 167 (2007).

[122] E. Jones, T. Oliphant, and P. Peterson, SCIPY: Open Source Scientific Tools for PYTHON (2001-), <http://www.scipy.org/>.