



CENTRE DE RECERCA MATEMÀTICA

This is a preprint of: *Dependence of exponents on text length versus
finite-size scaling for word-frequency distributions*
Journal Information: *Physical Review E*,
Author(s): A. Corral and F. Font-Clos.
Volume, pages: 96 1-9, DOI:[10.1103/PhysRevE.96.022318]

Dependence of exponents on text length versus finite-size scaling for word-frequency distributions

Álvaro Corral^{1,2,3,4} and Francesc Font-Clos⁵

¹*Centre de Recerca Matemàtica, Edifici C, Campus Bellaterra, E-08193 Barcelona, Spain*

²*Departament de Matemàtiques, Facultat de Ciències, Universitat Autònoma de Barcelona, E-08193 Barcelona, Spain*

³*Barcelona Graduate School of Mathematics, Edifici C, Campus Bellaterra, E-08193 Barcelona, Spain*

⁴*Complexity Science Hub Vienna, Josefstädter Straße 39, 1080 Vienna, Austria*

⁵*ISI Foundation, Via Chisola 5, 10126 Torino, Italy*

(Received 4 May 2017; published 22 August 2017)

Some authors have recently argued that a finite-size scaling law for the text-length dependence of word-frequency distributions cannot be conceptually valid. Here we give solid quantitative evidence for the validity of this scaling law, using both careful statistical tests and analytical arguments based on the generalized central-limit theorem applied to the moments of the distribution (and obtaining a novel derivation of Heaps' law as a by-product). We also find that the picture of word-frequency distributions with power-law exponents that decrease with text length [X. Yan and P. Minnhagen, *Physica A* **444**, 828 (2016)] does not stand with rigorous statistical analysis. Instead, we show that the distributions are perfectly described by power-law tails with stable exponents, whose values are close to 2, in agreement with the classical Zipf's law. Some misconceptions about scaling are also clarified.

DOI: [10.1103/PhysRevE.96.022318](https://doi.org/10.1103/PhysRevE.96.022318)

I. INTRODUCTION

Many complex processes in biology, social science, economy, Internet science, and cognitive science are mimicked by the occurrence of words in texts. Indeed, the statistics of insects in plants [1], molecules in cells [2], inhabitants in cities [3], followers of religions [4], telephone calls to people [4], employees in companies [5], links to sites on the World Wide Web [6], chords in musical pieces [7], etc., share with word-frequency distributions the property of being broadly distributed, or “heavy tailed.” And most of these phenomena are described, at least asymptotically, by power-law distributions with exponents close to 2; in such cases one can talk about the fulfillment of Zipf's law [8,9].

A fundamental problem is how these systems evolve, in particular, how they grow to reach a state for which a power law, or even Zipf's law, holds [10–12]. In Ref. [13], Bernhardsson, da Rocha, and Minnhagen challenge the “Zipf's view,” proposing that the distribution of word frequencies in a text or collection of texts (by the same author) changes with the text length as

$$D_L(k) = A \frac{e^{-k/(c_0 L)}}{k^{\gamma(L)}}, \quad (1)$$

where k is the absolute frequency (number of tokens) of the different words (word types), L is the text length in number of tokens, $D_L(k)$ is the probability mass function of k (i.e., the distribution of word frequencies), $\gamma(L)$ is a power-law exponent, c_0 is a constant parameter (independent of L), and A is a normalizing constant. Bernhardsson *et al.*'s equation [Eq. (1)] should apply to individual texts or collections when one considers parts of length L of the whole. The key ingredient of that approach to model the change of $D_L(k)$ with L is the explicit dependence of the exponent γ on the text length L , decreasing with increasing L . Note also that Eq. (1) implies that, for the largest k , the word-frequency distribution decays exponentially (in contrast to the algebraic decay proposed in Zipf's law).

Alternatively, in Ref. [14], we (together with Boleda) argue that the variability of the statistics of words in a text with its length is more simply explained by a scaling law,

$$D_L(k) = \frac{1}{LV_L} g(k/L), \quad (2)$$

where V_L is the size of the vocabulary (number of different words, i.e., word types) for a fraction of text of length L , and $g(z)$ is an undefined scaling function, the same for any value of L (but not necessarily the same for different authors).

Note that the scaling-law paradigm, Ref. [14] and Eq. (2), does not assume any particular, parametric shape of $D_L(k)$ [in contrast to Ref. [13] and Eq. (1), which give a truncated gamma distribution]; the scaling-law paradigm only states that, for a fixed text, all the $D_L(k)$'s have the same shape no matter the value of L , but at different characteristic scales given by L . In other words, the shape parameters of the distributions do not change with L , whatever the form of this distribution is (we do not enter here into such debate [9,15–17]), and it is only a scale parameter that changes, proportionally to L . Both exponential tails and power-law tails are allowed by the scaling function $g(z)$; what is “forbidden” are text-length-dependent exponents $\gamma(L)$. Moreover, the scaling paradigm represented by Eq. (2) does not involve any free parameter, as there is no restriction on the scaling function g , and L and V_L are given directly by the text. In fact, the scaling law is just a finite-size scaling law [18–21] (which is not explicitly mentioned in Ref. [14]).

Subsequently, Yan and Minnhagen claimed that this scaling law is “fundamentally impossible” [22], “cannot be conceptually valid,” and “is not borne out by the data” [23]. These statements constitute good examples of the sometimes counterintuitive nature of scaling laws. Let us summarize the points of these authors to make it clear that their critique is not relevant.

i. First, in their Fig. 1 they find that the scaling law does not hold for $k = 1$.

ii. Second, in their Fig. 2 they show that the scaling law does not work well for, let us say, $k \leq 10$.




FIG. 1. Total number of words N_L with a relative frequency greater than or equal to k/L , for varying $L = L_{\text{tot}}/n$, with L_{tot} the length of the complete text. We have taken the same books as in Ref. [23], (a) *Moby-Dick* and (b) *Harry Potter*, exactly reproducing Figs. 2(a) and 2(b) in Ref. [23] but also including some additional values of n . Deviations from the scaling law are always in the regime of very low frequencies, as expected due to discreteness effects (which are due to the fact that word tokens are discrete).

iii. Third, it is argued that a “randomness view,” based in the concepts of “random group formation,” “random book transformation,” and Metabook, predicts the right form of $D_L(k)$, which is that in Refs. [13,24], i.e., Eq. (1) above.

It is quite clear that the first and second criticisms of Yan and Minnhagen [23] are not fundamental, as they simply imply that the scaling law can only be valid beyond the low-frequency limit, so the scaling law can be rewritten as

$$D_L(k) = \frac{1}{LV_L} g(k/L) \quad \text{for } k > 10.$$

This is not surprising at all, as it is well known in statistical physics that scaling laws are usually observed only asymptotically (see Appendix A). It is remarkable then that, for texts, scaling is attained just after the first decade in frequencies (i.e., for words that appear more than about 10 times). It is also remarkable that, despite the fact that Yan and Minnhagen [23] stretch the scaling hypothesis up to very small fragments of texts (215 525 tokens divided into 500 fragments, yielding

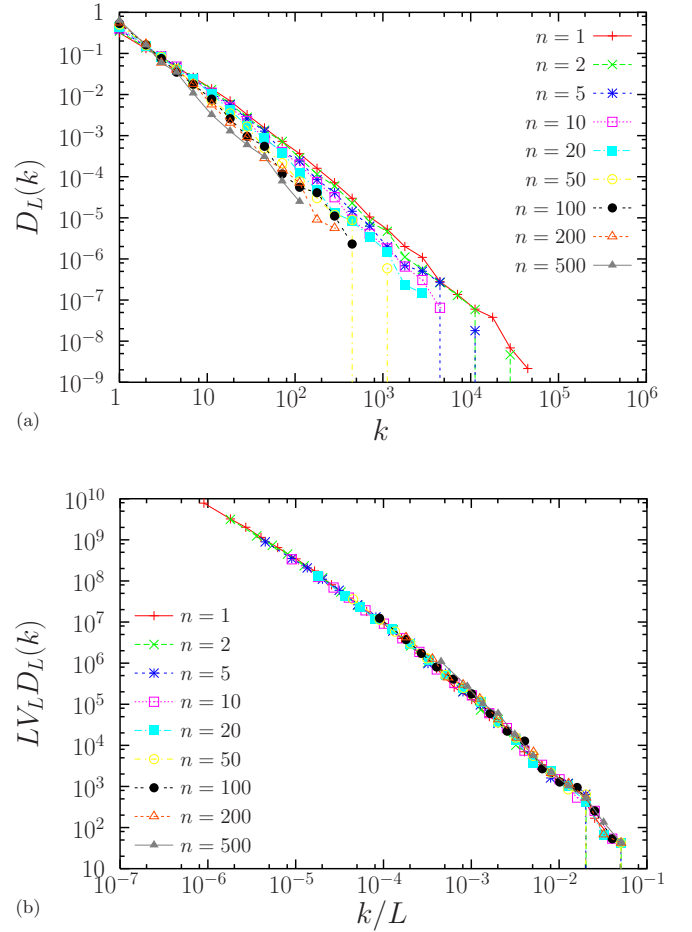


FIG. 2. (a) Probability mass function $D_L(k)$ of the absolute frequency k , for varying subsets of length $L = L_{\text{tot}}/n$ of *Harry Potter*, displaying a seeming change of shape. (b) The same curves, but plotting $D_L(k)LV_L$ versus k/L , as proposed in Ref. [14] and stated in Eq. (2). All curves collapse into a single, length-independent scaling function $g(k/L)$, in agreement with Eq. (2). Note the excellent data collapse: even the deviations for small k are negligible in this case. This is at odds with Eq. (1): a length-dependent exponent in $D_L(k)$, as proposed by Yan and Minnhagen, is not compatible with the data collapse shown here.

about 400 tokens in each one, for the case of *Moby-Dick*), the scaling law still is fulfilled reasonably well beyond the first decade in k , as we detail below. Naturally, the appropriate way to further test the validity of the scaling law is the opposite way, analyzing larger and larger texts. The third point of those authors [23] is also not justified, as the authors do not provide any statistical evidence supporting the claim that Eq. (1) fits the empirical data to an acceptable confidence level.

In this paper we revise the evidence for the finite-size scaling law for word-frequency distributions, Ref. [14] and Eq. (2), comparing this approach with the one in Ref. [13] and Eq. (1). In Sec. II we summarize the main claims in Ref. [13] and how they relate to the validity of the scaling law; in Sec. III we compare the performance of different fits related to the two approaches; in Sec. IV we use a direct method to test the validity of the scaling hypothesis applied to word-frequency distributions; and, finally, in Sec. V we compare our approach

with other ways of assessing errors in scaling. Section VI presents a novel theoretical calculation of the scaling of the moments of the distribution using the generalized central-limit theorem, connecting them with the scaling law (and yielding a derivation of Heaps' law as a by-product). The conclusions are given in Sec. VII, and Appendixes A and B follow. As the empirical evidence in favor of the scaling law used in Ref. [14] was essentially "visual" (collapse of rescaled plots in log-log) and the theoretical arguments were reduced to a heuristic derivation, the present paper provides a substantial improvement in support of the validity of a finite-size scaling law in word-frequency distributions.

II. VALIDITY OF THE FINITE-SIZE SCALING LAW FOR WORD-FREQUENCY DISTRIBUTIONS

Let us explain Yan and Minnhagen's points [23] in more detail. They base their analysis on the empirical value of the number of types with a frequency equal to or greater than k , defined as

$$N_L(\geq k) = V_L S_L(k) = V_L \sum_{k'=k}^{\infty} D_L(k'),$$

where $S_L(k)$ is the empirical complementary cumulative distribution of the frequency and $N_L(\geq k)$ turns out to be nothing but the empirical rank associated with frequency k . In terms of $S_L(k)$ the scaling law, (2), transforms into [14]

$$S_L(k) = \frac{1}{V_L} G(k/L) \quad \text{for } k > 10,$$

under a continuous approximation, with $G(z)$ a new scaling function directly related to $g(z)$ [14]. So, for $N_L(\geq k)$ one has that the scaling law, (2), can be written as

$$N_L(\geq k) = V_L S_L(k) = G(k/L) \quad \text{for } k > 10. \quad (3)$$

Then, in their Figs. 1(a) and 1(b) Yan and Minnhagen [23] compare this cumulative number evaluated for the complete text, $N_{L_{\text{tot}}}(\geq k)$, with the vocabulary size V_L for variable L , which verifies, by definition, $V_L = N_L(\geq 1)$ (note that L_{tot} is the length of the complete text). The disagreement between $N_{L_{\text{tot}}}(\geq k)$ and $N_L(\geq 1)$ for the same values of the ratio between frequency and length ($k/L_{\text{tot}} = 1/L$) makes it clear that the scaling law, for any $L \neq L_{\text{tot}}$ and under the form given by Eq. (3), does not work for the corresponding $k = 1$ (the *hapax legomena*; these are types that appear just once in a text sample of length L). Nevertheless, this disagreement does not invalidate the scaling law, (2), for $k > 1$.

Subsequently, in their Fig. 2, the same authors [23] plot $N_L(\geq k)$ versus k/L for all k and different L and indeed find deviations with respect to the scaling law, but let us note that these deviations are restricted to $k \leq 10$. Obviously, the reason for these deviations is just that the scaling law is expected to hold only asymptotically, which in practice means that $k > 10$ or so (see Appendixes A and B).

To make our thesis totally clear, in Fig. 1 we present $N_L(\geq k)$ for the same data as in Fig. 2 of Ref. [23] (i.e., *Moby-Dick*, by Herman Melville, and *Harry Potter*, books 1 to 7, by J. K. Rowling), but adding symbols (instead of only lines, as in Ref. [23]). It is apparent that even in the extreme case of 500 equal fragments of the full text, the scaling law

only fails for very low frequencies. We have verified that the scaling law holds for many other texts [14], even for *Finnegans Wake*, by James Joyce, which constitutes an extreme case of experimental literary creation, yielding an unusual, somewhat concave relation between $N_L(\geq k)$ and k (in log-log) [25]; in any case, the scaling function $g(z)$ is not affected by concavity or convexity.

Additionally, in Fig. 2 we perform the data collapse associated with the scaling law in terms of the probability mass function $D_L(k)$ for *Harry Potter*, presented in Ref. [23] as a counterexample to the scaling law. As shown, the collapse is excellent: after proper rescaling [Fig. 2(b)], all curves collapse into a single, length-independent function, even for very low frequencies (lower than 10). So we can write, for this text,

$$D_L(k) \simeq \frac{1}{LV_L} g(k/L) \quad \text{for all } k.$$

Note that this is the original form in which the scaling law for word-frequency distributions was presented [14], and not the one in Eq. (3). In other words, deviations from the scaling law play an even smaller role in the representation in terms of $D_L(k)$, in comparison to the representation in terms of $N_L(\geq k)$. Finding a functional form for $g(z)$ [whose particular shape in the case of *Harry Potter* is unveiled in our Fig. 2(b)] is a delicate issue, and it is not our interest here (in Ref. [14], when considering lemmatized texts, a double power law was proposed for the sake of illustration); nevertheless, in the next section we show that the empirical data, for each text and different values of L , are compatible with a unique $g(z)$ characterized by a power-law tail, with an exponent close to 2.

III. PROPER FITTING OF THE POWER-LAW TAIL

Looking at our Fig. 2(a), where $D_L(k)$ is shown with no rescaling, one could conclude, as did Yan and Minnhagen [23], that for different text lengths one gets different shapes for $D_k(L)$. Indeed, a visual inspection of the plot seems to show different slopes for different L values, corresponding to different exponents $\gamma(L)$. However, the data collapse after rescaling in Fig. 2(b) demonstrates that all distributions have the same shape, given by $g(z)$, but at different scales, given by L (remember that in log-log a rescaling is seen just as a shift). And the deviations for $k \leq 10$ do not play any relevant role. The apparent different slopes of the different curves in Fig. 2(a) are a visual artifact caused by the convex (but close to linear) log-log shape of the curves. In other words, the larger L , the larger portion of $g(z)$ one sees beyond the tail. As the body of the distribution does not decay as rapidly as the tail, the more of the body one sees, the smaller the apparent exponent, which is a sort of average between the body and the tail. This illustrates how a simple replotting under a rescaled form can unveil a common pattern in distributions given at different scales. And let us repeat that we are not interested here in providing a parametric model for this convex shape; for that, just see Ref. [9].

In order to support their point, Yan and Minnhagen [23] fit power laws [Eq. (1) with $c_0 = \infty$] to the word-frequency distributions for different L 's, finding a drift from $\gamma \simeq 3$ for the smallest fragments of text to $\gamma \simeq 1.6$ for the largest length (Zipf's law would correspond to $\gamma = 2$, strictly speaking). The

TABLE I. Results of a goodness-of-fit test [27] of a discrete power-law distribution $D_L(k) \propto 1/k^{\gamma_2}$ in the range $k \geq k_{\text{cut}}$ for *Harry Potter* (for which $L_{\text{tot}} = 1\,108\,955$). Note the stability of the γ_2 exponent for $k_{\text{cut}} \propto L$. Note also that for $L = L_{\text{tot}}/500$ the difference between Yan and Minnhagen's result ($\gamma \simeq 2.4$) and our result here with $k_{\text{cut}} = 2$ (yielding $\gamma \simeq 2.16$) is rather large. It has been pointed out that inappropriate fitting methods can lead to biased results [4,26].

L	γ_2	k_{cut}	V_L	p value
L_{tot}	2.06 ± 0.09	1024	22276	0.677
$L_{\text{tot}}/2$	2.05 ± 0.09	513	16361	0.504
$L_{\text{tot}}/5$	2.09 ± 0.09	205	10658	0.862
$L_{\text{tot}}/10$	2.09 ± 0.09	103	7431	0.263
$L_{\text{tot}}/20$	2.18 ± 0.09	52	5186	0.819
$L_{\text{tot}}/50$	2.14 ± 0.08	20	3240	0.554
$L_{\text{tot}}/100$	2.17 ± 0.09	11	2079	0.708
$L_{\text{tot}}/200$	2.14 ± 0.08	5	1353	0.903
$L_{\text{tot}}/500$	2.16 ± 0.07	2	774	0.683

authors do not mention which fitting method they use, or the fitting range, but we can demonstrate that power-law fits for the whole range of k in *Moby-Dick* and *Harry Potter* are in general rejected after rigorous goodness-of-fit tests, no matter whether one fits continuous [26] or discrete power laws [9]. Taking *Harry Potter* and applying the maximum-likelihood fitting plus the Kolmogorov-Smirnov test detailed in Ref. [9], we only get one case of a nonrejectable discrete power law defined in the range $k \geq 1$, which corresponds to $L = L_{\text{tot}}/500$, with an exponent $\gamma = 2.11 \pm 0.04$ and $p = 0.49$. For all other lengths of fragments, power laws defined for $k \geq 1$ are rejected at the 0.05 significance level.

On the contrary, if we fit a power-law tail, which is a power law

$$D_L(k) \propto \frac{1}{k^{\gamma_2}} \quad (4)$$

defined only for $k \geq k_{\text{cut}}$, with k_{cut} a cutoff value of k verifying $k_{\text{cut}} > 1$, we find nonrejectable power laws for all L with stable exponents when k_{cut} is large enough. For the case of *Harry Potter* we find that for $k_{\text{cut}} \simeq 0.0009L$, the exponents γ_2 turn out to be stabilized with values very close to 2. So, $g(z)$ has a power-law tail, valid for $z = k/L > 0.0009$ and with a stable exponent, at odds with Ref. [13]'s claims. Details are listed in Table I. Figure 3 provides more examples of this behavior, for different texts. Although these results are in agreement with Zipf's law [and in disagreement with the exponential tail represented by Eq. (1)], we are not interested here in the parametric form of the distribution and only report the stability of the exponents with L as a signature of the existence of a well-defined, L -independent scaling function $g(z)$.

IV. TESTING OF THE SCALING HYPOTHESIS

A more direct and nonparametric way to test the existence of scaling is to use the two-sample Kolmogorov-Smirnov test [28], which compares two data sets under the null hypothesis that both of them come from the same population and, therefore, have the same underlying theoretical distribution (which is unknown and remains unknown after the test). But




FIG. 3. Exponents of the power-law tails as a function of the text length, for seven texts in English, Spanish, and French. Fits are performed as in Refs. [8,27], with k_{cut} restricted to the last two decades of the distribution, and accepted for the k_{cut} that yields a p value larger than 0.20 (calculated from 1000 Monte Carlo simulations). Error bars correspond to one standard deviation. Types are defined at the word-lemma-tag level; more details on text processing are given in Ref. [8]. Note how for $L > 3000$ all exponents are below 2.2, and they become quite stable in the range $6000 < L < 2 \times 10^6$.

in the case of scaling we are dealing not with the same distribution, but with distributions which have the same shape at different scales, i.e., distributions that are the same except for a scale parameter; then, rescaling the distributions by their scale parameter would lead to the same distributions (under the null hypothesis that scaling holds). This procedure to test the fulfillment of scaling has been used before for continuous distributions [29,30].

The Kolmogorov-Smirnov test is probably the best accepted test for comparing continuous distributions, but word-frequency data are discrete and, after rescaling, become discretized over different sets (as the scaling factors L of the two distributions can be very different, in general). So, our first step, in order to avoid this problem, is to approximate the discrete empirical distributions with continuous ones, by adding to each frequency a random term; in this way $k \rightarrow k + u$ (where u is a uniform random number between -0.5 and 0.5). Although there are more sophisticated ways to continuize the distributions, this one uses no information from the data (except that the k 's are natural numbers). The second step is to remove low frequencies (remember that scaling is a large-scale property); in our case we remove values of k below 4. Next, the third step is to perform the rescaling

$$k \rightarrow \frac{\langle k \rangle k}{\langle k^2 \rangle},$$

where the moments of k are the original empirical ones (calculated for the discrete distribution). In a simple case (with no power laws involved [29,30]) we would have rescaled just by the mean $\langle k \rangle$; in this case the rescaling is a bit more involved [31,32]. Note that this rescaling is totally equivalent to dividing k by L , as shown below; nevertheless, our choice is more




FIG. 4. P values of two-sample Kolmogorov-Smirnov tests for word-frequency distributions rescaled as explained in the text. One of the samples has text length L and the other L' , ranging from $L_{\text{tot}}/50$ to L_{tot} . One of the data sets shown corresponds to testing the first fragment of length L with the same fragment of length L' ; for the other data set other fragments are chosen. The texts are the same seven used in the previous figure. P values below 0.05 should lead to the rejection of the scaling hypothesis with a 0.95 confidence level.

general and makes the rescaling applicable when the data do not come from a text.

Once these three steps have been done, the two-sample Kolmogorov-Smirnov test [28] is performed for all pairs of samples given by different L 's, restricting the samples to a common support, i.e., a common minimum value is taken as the minimum value of the sample with the smallest L (which has the largest minimum when the frequencies are rescaled). Figure 4 shows the P value of this test for several texts and different divisions of the texts, up to $L_{\text{tot}}/50$. The fact that the P value appears as uniformly distributed between zero and one is an indication that the scaling null hypothesis holds.

V. RELATIVE ERRORS OF THE SCALING LAW

Although the proper way to compare statistical distributions is by means of statistical tests (as done in the previous section), Yan and Minnhagen [23] instead use relative errors. They show numbers for the relative error provided by the scaling law and compare it with the error of the so-called random-group-formation hypothesis. We explain why their comparison is not appropriate. First, for the scaling law, the empirical values of $N_L(\geq k)$ are compared for a fixed ratio k/L with $N_{L_{\text{tot}}}(\geq k)$ and the errors are claimed to be large. Second, for the random group formation the error is claimed to be much smaller, but in this case the empirical data $N_L(\geq k)$ are compared with random samples of the same length L , and not with a distribution of a different length. It is obvious that this procedure has to yield better results, and this constitutes a totally biased comparison.

But further, the errors provided by Yan and Minnhagen [23] for the scaling law are inflated. Figure 5 shows the relative difference or error between the true value $N_L(\geq k)$ and the value approximated by the scaling law, $N_{L_{\text{tot}}}(\geq k')$, with




FIG. 5. Relative error $[N_{L_{\text{tot}}}(\geq k') - N_L(\geq k)]/N_L(\geq k)$ of the approximation to $N_L(\geq k)$ given by the scaling of the distribution of the whole text $N_{L_{\text{tot}}}(\geq k')$, with $k' = kL_{\text{tot}}/L$, for *Harry Potter*. Symbols correspond to the same values of n as in the previous figures.

$k' = L_{\text{tot}}k/L$, which is

$$\varepsilon_L(k) = \frac{N_{L_{\text{tot}}}(\geq k') - N_L(\geq k)}{N_L(\geq k)}.$$

Note that, in general, replacing $N_L(\geq k)$ with $N_{L_{\text{tot}}}(\geq k')$ in the denominator inflates the reported error, as, when there are deviations, this number is systematically below $N_L(\geq k)$. The results of our analysis (Fig. 5) show that the errors are not as large as reported by Yan and Minnhagen [23]. Dividing *Harry Potter* into up to 20 parts, the relative error provided by the scaling law is almost always below 0.2, with the remarkable exception of the case $k = 1$ for $L = L_{\text{tot}}/20$. Dividing the text into smaller parts yields that the relative error is always below 0.3 for $k > 10$. But the error for small k is further reduced if one uses for comparison the probability mass function $D_L(k)$ instead of $N_L(\geq k)$. Remember that the original form of the scaling law was reported for $D_L(k)$ and not for $N_L(\geq k)$. Our Fig. 2(b) speaks for itself.

VI. SCALING OF MOMENTS FROM THE GENERALIZED CENTRAL-LIMIT THEOREM, HEAPS' LAW, AND RELATION WITH THE SCALING LAW

We start this section dealing with a distribution $D_L(k)$ that has a power-law tail with an exponent in the range $1 < \gamma_1 < 2$. We consider the moments $\langle k \rangle$ and $\langle k^2 \rangle$ not as the moments of the theoretical distribution (which would be equal to ∞) but as the moments of a finite sample, whose size is just the size of the vocabulary V_L (by definition); that is,

$$\langle k \rangle = \frac{1}{V_L} \sum_{i=1}^{V_L} k_i \quad \text{and} \quad \langle k^2 \rangle = \frac{1}{V_L} \sum_{i=1}^{V_L} k_i^2.$$

Due to the power-law behavior for large k , the generalized central-limit theorem [32,33] allows one to obtain the scaling properties of these sums, assuming that the individual frequencies are independent (or weakly dependent). Indeed, $\sum_{i=1}^{V_L} k_i$ does not scale linearly with V_L but superlinearly, as $\sum_{i=1}^{V_L} k_i \propto V_L^{1/(\gamma_1-1)}$. Moreover, if k has a power-law-tailed

distribution with exponent γ_1 , so does k^2 , but with exponent γ'_1 fulfilling $\gamma'_1 - 1 = (\gamma_1 - 1)/2$ (and in the range $1 < \gamma'_1 < 2$), and then the generalized central-limit theorem also applies to k^2 , to give $\sum_{i=1}^{V_L} k_i^2 \propto V_L^{2/(\gamma_1-1)}$.

On the other hand, we can also use the exact result $\sum_{i=1}^{V_L} k_i = L$ (the definition of text length), from which we obtain the classical Heaps' law (also called Herdan's law in llinguistics) [14,34–37],

$$V_L \propto L^{\gamma_1-1}, \quad (5)$$

and therefore the moments fulfill

$$\langle k \rangle = L^{2-\gamma_1} \quad \text{and} \quad \langle k^2 \rangle = L^{3-\gamma_1} \quad (6)$$

(and, in general, $\langle k^m \rangle = L^{m+1-\gamma_1}$).

This result is compatible with a scaling law of the form

$$D_L(k) = \frac{1}{L^{\gamma_1}} g\left(\frac{k}{L}\right). \quad (7)$$

The case considered in the literature [32,38] assumes that $g(z)$ has an intermediate power-law decay with exponent γ_1 followed by a much faster decay (exponential or so) for the largest k 's. The pure power-law tail considered above is included in this framework when $g(z)$ goes to 0 abruptly, transforming the pure power law into a truncated power law. Indeed, if the power law is truncated at k_{\max} , using a continuous approximation we get

$$\begin{aligned} \langle k^m \rangle &= \int_1^\infty k^m D_L(k) dk \simeq \int_1^{k_{\max}} k^m D_L(k) dk \\ &= L^{m+1-\gamma_1} \int_{1/L}^{k_{\max}/L} z^m g(z) dz \sim L^{m+1-\gamma_1}, \end{aligned} \quad (8)$$

because (for $m > \gamma_1 - 1$) the integral tends to a constant when L is large, taking into account that k_{\max} is the maximum of k from a sample of size V_L and scales in the same way as $\sum k_i$, i.e., as L . To see this one can just calculate, for any P , the percentiles k_P of the distribution of the maximum of V_L frequencies, which verify $[1 - S_L(k_P)]^{V_L} = P$. Substituting a power law for $S_L(k)$ (with exponent $\gamma_1 - 1$) we get

$$k_P \propto \frac{1}{(1 - P^{1/V_L})^{1/(\gamma_1-1)}} \simeq \left(\frac{V_L}{-\ln P} \right)^{1/(\gamma_1-1)} \propto L.$$

So, as all the percentiles of the maximum scale with L , the distribution of the maximum scales with L too [we arrive at the last result using that $P^{1/V_L} = e^{(\ln P)/V_L} \simeq 1 + (\ln P)/V_L$, valid for large V_L , and also Heaps' law (5)].

Therefore, a power-law tail for the distribution of frequencies, with exponent γ_1 , is somehow equivalent to an upper truncated power-law tail, with an effective cutoff $k_{\max} \propto V_L^{\gamma_1-1} \propto L$, which means that the features of the distribution at the largest k have to scale with L . The scaling law, (2), follows directly from Eq. (7) using Heaps' law, (5), although note that the version of the law given by Eq. (2) is nonparametric, in the sense that the value of the exponent γ_1 does not appear in the law (which is good if the determination of the exponent contains errors).




FIG. 6. Linear relation between the second moment of the distribution of word frequencies and the ratio of the squared text length to the vocabulary size, for diverse texts used in the previous figures [8]. A line with a linear coefficient 0.01 is shown for comparison.

Moreover, as a by-product we obtain another form for the scaling law,

$$D_L(k) = \frac{\langle k \rangle^3}{\langle k^2 \rangle^2} g\left(\frac{k \langle k \rangle}{\langle k^2 \rangle}\right), \quad (9)$$

using the scaling, (6), of the moments with L and the scaling law, with the scaling function g being the same as before, except for proportionality factors. This scaling has been used previously for self-organized critical phenomena, but under different conditions [31]. Remember that here the moments are not those of the theoretical distribution but the ones corresponding to a sample of size V_L . The equivalence of both scaling laws, Eqs. (2) and (9), is empirically shown in Fig. 6 by means of the proportionality between $\langle k^2 \rangle$ and L^2/V_L .

However, real distributions of frequencies are not well described at the highest frequencies by scaling functions $g(z)$ that decay either exponentially or abruptly (i.e., are sharply truncated), as shown in Table I and Fig. 3. Instead, we expect that the tail of $g(z)$ [and therefore the tail of $D_L(k)$] is another power law, with an exponent $\gamma_2 > \gamma_1$. Remarkably, this framework is also described by the scaling law, (7), and the scaling of moments, (6), the key point being that one can change the upper limit of the integral from infinite to k_{\max} , and k_{\max} still scales linearly with L , so the derivation is the same as in Eq. (8).

In order to support empirically the fulfillment of a scaling law of the form given by Eq. (7) we follow the approach presented in Ref. [39]. If such a scaling law holds, the distance between the different rescaled distributions in log-scale

$$(\ln k_i - \delta \ln L, \ln D_L(k_i) + \gamma_1 \ln L)$$

should be minimum when the right values of the exponents are substituted. Note that we have introduced an extra exponent δ , which we expect becomes equal to 1. We proceed by minimizing such distances as a function of the exponents δ and γ_1 , resulting in values of δ very close to 1 indeed and

TABLE II. Exponents δ and γ_1 obtained after the minimization of the distance between the rescaled word-frequency distributions, optimizing the data collapse, for seven texts. The text length L varies from L_{tot} to $L_{\text{tot}}/10$. Different seeds are used in the algorithm in order to avoid local minima.

Text and author	Language	δ	γ_1
<i>Clarissa</i> by S. Richardson	English	0.96	1.50
<i>Moby-Dick</i> by H. Melville	English	1.05	1.87
<i>Ulysses</i> by J. Joyce	English	1.04	1.94
<i>El Quijote</i> by M. de Cervantes	Spanish	0.96	1.64
<i>La Regenta</i> by L. A. Clarín	Spanish	0.86	1.52
<i>Artamène</i> by Scudéry siblings	French	1.04	1.63
<i>Le Vicomte de Bragelonne</i> by A. Dumas (father)	French	0.96	1.58

values of γ_1 in the range 1.5 to 1.95, when different texts are used (see Table II).

VII. DISCUSSION AND CONCLUSIONS

As an important remark, we want to clarify that we are not against the so-called random-group-formation hypothesis [24], as in some previous research we have made use of randomness to explain real texts [40]. Our conclusion was that real texts are not random, but the first appearance of a word is close to random, so the word-frequency distribution (related to Zipf's law) and the type-token growth curve (related to Heaps' law) remain the same for real texts and for random versions of them. The reason is that the word-frequency distribution is independent of the word order, and the type-token growth curve only depends on the first appearance of a word. Other properties of real texts are different from those of random texts, such as interappearance distances [40–42].

Summarizing, the empirical facts are clear: a finite-size scaling law gives a very good approximation of the distribution of word frequencies for different fragments of text of length L . The shape of $D_k(L)$ is the same for all L , and it is only a scaling factor proportional to L that makes the difference for different L 's. It is the parametric proposal in Ref. [13] which is not well supported by solid statistical testing. In any case, if the theory held by Yan and Minnhagen [23] is valid, then it must contain the scaling law in some limit. If not, the theory is irrelevant for real texts.

In conclusion, we show how the sort of scaling arguments usual in statistical physics, in particular, finite-size scaling, (2), can describe complex processes much better than parametric formulas, (1). Finally, in order to avoid misunderstandings, let us state that although curve fitting is a very honorable approach in science (when done correctly [9,26,27]), a scaling approach has nothing to do with that [23].

ACKNOWLEDGMENTS

We are grateful to R. Ferrer-i-Cancho for drawing our attention to Ref. [13] and to G. Boleda for facilitating the beginning of this research. Yan and Minnhagen's criticisms have allowed us to explain in much more detail the validity of our scaling law and to arrive at the new results presented here.

Research projects in which this work is included are FIS2012-31324 and FIS2015-71851-P, from Spanish MINECO, and 2014SGR-1307, from AGAUR.

APPENDIX A

We explain here the difference between a power law and a scaling law, and how scaling laws in statistical physics usually only hold asymptotically. Let us start with a scale transformation. This is an operation that stretches and/or contracts a function, i.e.,

$$T[f(x, y)] = cf(x/a, y/b),$$

where $f(x, y)$ is a (in this example bivariate) real function, a , b , and c are constant and positive scale factors, and T is the scale transformation. If we ask the question which functions are invariant under scale transformations (i.e., which functions do not change when they are stretched and/or compressed), we find that a solution is

$$f(x, y) = \frac{1}{x^\beta} g\left(\frac{y}{x^\alpha}\right), \quad (\text{A1})$$

with $\alpha = \ln b / \ln a$, $\beta = -\ln c / \ln a$, and g an arbitrary function, called the scaling function (we also consider $x > 0$). Moreover, if we look for a solution valid for any real value of $a > 0$, the previous solution, (A1), turns out to be the only solution [38]. One refers to Eq. (A1) as a scaling form or a scaling law.

Note that a power law is a special case of a scaling law, just taking the arbitrary scaling function g to be a constant C . In fact, in one dimension [i.e., for univariate functions $f(x)$] the only scaling laws (the only scale-invariant functions for any value of the scale factor a) are the power laws [10,43], so $f(x) = C/x^\beta$. Although the terms “scaling law” and “power law” are sometimes taken as synonyms, it is clear that they are only equivalent for univariate functions. For bivariate (and multivariate) functions one needs to be more careful in distinguishing both concepts (as we do here).

Let us stress that scaling is a fundamental pillar of 20th-century statistical physics [44]. In our case, we propose that a (bivariate) scaling law holds for $D_L(k)$, so we identify $k = y$, $L = x$, and $D_L(k) = f(x, y)$ and assume Heaps' law for the usual scaling law to hold (more details in Ref. [14]). Alternatively, we may identify $N_L(\geq k)$ not with $D_L(k)$ but with $f(x, y)$, with no necessity of using Heaps' law. In any case this does not necessarily imply that the scaling function g has a power-law shape. We do not care here about the functional form of g ; this is just the shape shown in Fig. 2(b) (for the particular book under consideration there).

We provide in Fig. 7 a practical example of how scaling laws in statistical physics hold usually only for large x and y (i.e., large L and k). We display the rescaled size distribution $D_L(k)$ of a critical Galton-Watson branching process [45] with its number of generations bounded by a finite L and with its offspring distribution given by a binomial distribution with two trials. This process is totally equivalent to percolation in the Bethe lattice [46]. The figure shows the deviation from the scaling law for $k \leq 10$, but nevertheless, it has been proven analytically that finite-size scaling holds in this system [46]. Ironically, in this case the scaling function is




FIG. 7. Rescaled size distribution $D_L(k)$ as a function of the rescaled size k for a critical Galton-Watson branching process with a binomial offspring distribution (with a maximum of two offspring) and with different values of the maximum number of generations L . A finite-size scaling law, of the form $D_L(k) = g(k/L^2)/L^3$, holds, but only for $k > 10$, roughly. Curiously, when the offspring distribution is geometric (for which the size distribution coincides with the escape time of a random walker moving between an absorbing and a reflecting boundary [21]), the deviation takes the opposite sign (see Ref. [32]). But beyond the smallest values of k , the rescaled distributions coincide and universality shows up [44].

well approximated by the function proposed by Bernhardtsson *et al.* [13] [our Eq. (1)], but with a constant exponent $\gamma = \beta/\alpha = 3/2$.

APPENDIX B

Let us see how discreteness effects alter scaling. Naturally, the discrete nature of word-frequency distributions comes from the fact that the fundamental unit is the count of word tokens. If we assume that scaling holds for all k , even for $k = 1$ and $k = 2$, the finite-size scaling law, under the form given by Eq. (3), implies that

$$N_L(\geq 2) = N_{L/2}(\geq 1),$$

and we can relate this to the size of the vocabulary for each text length, so

$$V_L - n_L^s(1) = V_{L/2},$$

where $n_L(k)$ counts the number of types with frequency (exactly equal to) k , and the superscript s denotes that we are under the scaling hypothesis.

On the other hand, for a random text, $V_{L/2}$ can be calculated from V_L as

$$V_{L/2} = V_L - n_{L/2}(0) = V_L - \sum_{k \geq 1} h_{0,k} n_L(k),$$

where $h_{0,k}$ gives the probability of getting 0 tokens of a certain type when a fragment of text of length $L/2$ is taken from a text of length L in which the same type has frequency k (see Eq. (5) in Ref. [14]). Comparing both equations for $V_{L/2}$ we get $n_L^s(1) = \sum_{k \geq 1} h_{0,k} n_L(k)$. But using that $h_{0,k} < 1/2^k$ (see below) and $n_L(k) < n_L(1)$ for $k > 1$ [from empirical evidence; see Fig. 2(a), for instance], we arrive at

$$n_L^s(1) = n_{L/2}(0) = \sum_{k \geq 1} h_{0,k} n_L(k) < n_L(1) \sum_{k \geq 1} \frac{1}{2^k} = n_L(1)$$

(extending the sum to infinite). Thus, the scaling hypothesis yields, for a random text and for $k = 1$, fewer types than it should. This can be seen by looking carefully at some of the plots in Fig. 2 of Ref. [14], but not in Fig. 3 there or in Fig. 2(b) here, as the deviations are rather small [just note that the empirical $D_L(k)$ is proportional to $n_L(k)$].

The fact that $h_{0,k} < 1/2^k$ comes from the fact that $h_{0,k}$ is given by the hypergeometric distribution (as we assume that we take tokens from the larger text with no replacement), and then

$$h_{0,k} = \binom{k}{0} \binom{L-k}{L/2-0} / \binom{L}{L/2} \quad (\text{B1})$$

$$= \left[\frac{(L-k)!}{(L-k-L/2)!} \right] / \left[\frac{L!}{(L/2)!} \right] \quad (\text{B2})$$

$$= \left[\frac{(L/2)!}{(L/2-k)!} \right] / \left[\frac{L!}{(L-k)!} \right] \quad (\text{B3})$$

$$= \frac{(L/2)(L/2-1) \dots (L/2-k+1)}{L(L-1) \dots (L-k+1)} \quad (\text{B4})$$

$$= \prod_{j=0}^{k-1} \left(\frac{L/2-j}{L-j} \right), \quad (\text{B5})$$

where all factors are smaller than $1/2$, except the one for $j = 0$. This yields $h_{0,k} < 1/2^k$. In this way we show how discrete effects break scaling for the lowest frequencies, but as shown in the plots, this effect is very small.

- [1] S. Pueyo and R. Jovani, Comment on “A keystone mutualism drives pattern in a power function”, *Science* **313**, 1739c (2006).
- [2] C. Furusawa and K. Kaneko, Zipf’s Law in Gene Expression, *Phys. Rev. Lett.* **90**, 088102 (2003).
- [3] Y. Malevergne, V. Pisarenko, and D. Sornette, Testing the Pareto against the lognormal distributions with the uniformly most powerful unbiased test applied to the distribution of cities, *Phys. Rev. E* **83**, 036111 (2011).
- [4] A. Clauset, C. R. Shalizi, and M. E. J. Newman, Power-law distributions in empirical data, *SIAM Rev.* **51**, 661 (2009).

- [5] R. L. Axtell, Zipf distribution of U.S. firm sizes, *Science* **293**, 1818 (2001).
- [6] L. A. Adamic and B. A. Huberman, Zipf’s law and the Internet, *Glottometrics* **3**, 143 (2002).
- [7] J. Serrà, A. Corral, M. Boguñá, M. Haro, and J. Ll. Arcos, Measuring the evolution of contemporary western popular music, *Sci. Rep.* **2**, 521 (2012).
- [8] A. Corral, G. Boleda, and R. Ferrer-i-Cancho, Zipf’s law for word frequencies: Word forms versus lemmas in long texts, *PLoS ONE* **10**, e0129031 (2015).

- [9] I. Moreno-Sánchez, F. Font-Clos, and A. Corral, Large-scale analysis of Zipf's law in English texts, *PLoS ONE* **11**, e0147073 (2016).
- [10] M. E. J. Newman, Power laws, Pareto distributions and Zipf's law, *Cont. Phys.* **46**, 323 (2005).
- [11] B. Corominas-Murtra, R. Hanel, and S. Thurner, Understanding scaling through history-dependent processes with collapsing sample space, *Proc. Natl. Acad. Sci. USA* **112**, 5348 (2015).
- [12] V. Loreto, V. D. P. Servedio, S. H. Strogatz, and F. Tria, Dynamics on expanding spaces: Modeling the emergence of novelties, in *Creativity and Universality in Language*, edited by M. Degli Esposti *et al.* (Springer, Berlin, 2016), pp. 59–83.
- [13] S. Bernhardsson, L. E. Correa da Rocha, and P. Minnhagen, The meta book and size-dependent properties of written language, *New J. Phys.* **11**, 123015 (2009).
- [14] F. Font-Clos, G. Boleda, and A. Corral, A scaling law beyond Zipf's law and its relation to Heaps' law, *New J. Phys.* **15**, 093033 (2013).
- [15] R. Ferrer i Cancho and R. V. Solé, Two regimes in the frequency of words and the origin of complex lexicons: Zipf's law revisited, *J. Quant. Linguist.* **8**, 165 (2001).
- [16] M. A. Montemurro, Beyond the Zipf-Mandelbrot law in quantitative linguistics, *Physica A (Amsterdam)* **300**, 567 (2001).
- [17] W. Li, P. Miramontes, and G. Cocho, Fitting ranked linguistic data with two-parameter functions, *Entropy* **12**, 1743 (2010).
- [18] E. Brézin, An investigation of finite size scaling, *J. Phys.* **43**, 15 (1982).
- [19] M. N. Barber, Finite-size scaling, in *Phase Transitions and Critical Phenomena*, edited by C. Domb and J. L. Lebowitz (Academic Press, London, 1983), Vol. 8, pp. 145–266.
- [20] V. Privman, Finite-size scaling theory, in *Finite Size Scaling and Numerical Simulation of Statistical Systems*, edited by V. Privman (World Scientific, Singapore, 1990), pp. 1–98.
- [21] A. Corral, R. Garcia-Millan, and F. Font-Clos, Exact derivation of a finite-size scaling law and corrections to scaling in the geometric Galton-Watson process, *PLoS ONE* **11**, e0161586 (2016).
- [22] X.-Y. Yan and P. Minnhagen, Comment on 'A scaling law beyond Zipf's law and its relation to Heaps' law', *arXiv:1404.1461v1* (2014).
- [23] X. Yan and P. Minnhagen, Randomness versus specifics for word-frequency distributions, *Physica A* **444**, 828 (2016).
- [24] S. K. Baek, S. Bernhardsson, and P. Minnhagen, Zipf's law unzipped, *New J. Phys.* **13**, 043004 (2011).
- [25] J. Kwapień and S. Drozd, Physical approach to complex systems, *Phys. Rep.* **515**, 115 (2012).
- [26] A. Deluca and A. Corral, Fitting and goodness-of-fit test of non-truncated and truncated power-law distributions, *Acta Geophys.* **61**, 1351 (2013).
- [27] A. Corral, A. Deluca, and R. Ferrer-i-Cancho, A practical recipe to fit discrete power-law distributions, *arXiv:1209.1270* (2012).
- [28] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes in FORTRAN*, 2nd ed. (Cambridge University Press, Cambridge, UK, 1992).
- [29] A. Corral, Statistical tests for scaling in the inter-event times of earthquakes in California, *Int. J. Mod. Phys. B* **23**, 5570 (2009).
- [30] E. Lippiello, A. Corral, M. Bottiglieri, C. Godano, and L. de Arcangelis, Scaling behavior of the earthquake intertime distribution: Influence of large shocks and time scales in the Omori law, *Phys. Rev. E* **86**, 066119 (2012).
- [31] O. Peters, A. Deluca, A. Corral, J. D. Neelin, and C. E. Holloway, Universality of rain event size distributions, *J. Stat. Mech.* (2010) P11030.
- [32] A. Corral, Scaling in the timing of extreme events, *Chaos. Solit. Fract.* **74**, 99 (2015).
- [33] J.-P. Bouchaud and A. Georges, Anomalous diffusion in disordered media: Statistical mechanisms, models and physical applications, *Phys. Rep.* **195**, 127 (1990).
- [34] R. Baeza-Yates and G. Navarro, Block addressing indices for approximate text retrieval, *J. Am. Soc. Inform. Sci.* **51**, 69 (2000).
- [35] A. Kornai, How many words are there? *Glottometrics* **2**, 61 (2002).
- [36] L. Lü, Z.-K. Zhang, and T. Zhou, Zipf's law leads to Heaps' law: Analyzing their relation in finite-size systems, *PLoS ONE* **5**, e14139 (2010).
- [37] M. A. Serrano, A. Flammini, and F. Menczer, Modeling statistical properties of written text, *PLoS ONE* **4**, e5372 (2009).
- [38] K. Christensen and N. R. Moloney, *Complexity and Criticality* (Imperial College Press, London, 2005).
- [39] A. Deluca and A. Corral, Scale invariant events and dry spells for medium-resolution local rain data, *Nonlin. Proc. Geophys.* **21**, 555 (2014).
- [40] F. Font-Clos and A. Corral, Log-Log Convexity of Type-Token Growth in Zipf's Systems, *Phys. Rev. Lett.* **114**, 238701 (2015).
- [41] E. G. Altmann, J. B. Pierrehumbert, and A. E. Motter, Beyond word frequency: Bursts, lulls, and scaling in the temporal distributions of words, *PLoS ONE* **4**, e7678 (2009).
- [42] A. Corral, R. Ferrer-i-Cancho, and A. Díaz-Guilera, Universal complex structures in written language, *arXiv:0901.2924*.
- [43] A. Corral, Scaling and universality in the dynamics of seismic occurrence and beyond, in *Acoustic Emission and Critical Phenomena*, edited by A. Carpinteri and G. Lacidogna (Taylor and Francis, London, 2008), pp. 225–244.
- [44] H. E. Stanley, Scaling, universality, and renormalization: Three pillars of modern critical phenomena, *Rev. Mod. Phys.* **71**, S358 (1999).
- [45] A. Corral and F. Font-Clos, Criticality and self-organization in branching processes: Application to natural hazards, in *Self-Organized Criticality Systems*, edited by M. Aschwanden (Open Academic Press, Berlin, 2013), pp. 183–228.
- [46] R. Garcia-Millan, F. Font-Clos, and A. Corral, Finite-size scaling of survival probability in branching processes, *Phys. Rev. E* **91**, 042122 (2015).