

Scaling laws and fluctuations in the statistics of word frequencies

This content has been downloaded from IOPscience. Please scroll down to see the full text.

2014 New J. Phys. 16 113010

(<http://iopscience.iop.org/1367-2630/16/11/113010>)

View [the table of contents for this issue](#), or go to the [journal homepage](#) for more

Download details:

This content was downloaded by: gerlach_m

IP Address: 193.174.246.169

This content was downloaded on 04/11/2014 at 12:38

Please note that [terms and conditions apply](#).

Scaling laws and fluctuations in the statistics of word frequencies

Martin Gerlach and Eduardo G Altmann

Max Planck Institute for the Physics of Complex Systems, D-01187 Dresden, Germany
E-mail: gerlach@pks.mpg.de and edugalt@pks.mpg.de

Received 11 July 2014, revised 10 September 2014

Accepted for publication 19 September 2014

Published 4 November 2014

New Journal of Physics **16** (2014) 113010

doi:[10.1088/1367-2630/16/11/113010](https://doi.org/10.1088/1367-2630/16/11/113010)

Abstract

In this paper, we combine statistical analysis of written texts and simple stochastic models to explain the appearance of scaling laws in the statistics of word frequencies. The average vocabulary of an ensemble of fixed-length texts is known to scale sublinearly with the total number of words (Heaps' law). Analyzing the fluctuations around this average in three large databases (Google-ngram, English Wikipedia, and a collection of scientific articles), we find that the standard deviation scales linearly with the average (Taylor's law), in contrast to the prediction of decaying fluctuations obtained using simple sampling arguments. We explain both scaling laws (Heaps' and Taylor) by modeling the usage of words using a Poisson process with a fat-tailed distribution of word frequencies (Zipf's law) and topic-dependent frequencies of individual words (as in topic models). Considering topical variations lead to quenched averages, turn the vocabulary size a non-self-averaging quantity, and explain the empirical observations. For the numerous practical applications relying on estimations of vocabulary size, our results show that uncertainties remain large even for long texts. We show how to account for these uncertainties in measurements of lexical richness of texts with different lengths.

Keywords: scaling laws, stochastic processes, statistical fluctuations, natural language



Content from this work may be used under the terms of the [Creative Commons Attribution 3.0 licence](https://creativecommons.org/licenses/by/3.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

1. Introduction

Fat-tailed distributions [1–3], allometric scaling [4, 5], and fluctuation scaling [6–8] are the most prominent examples of scaling laws appearing in complex systems. Statistics of words in written texts provide some of the best studied examples: the fat-tailed distribution of word frequencies (Zipf’s law) [9] and the sublinear growth (as in allometric scalings) of the number of distinct words as a function of database size (Heaps’ law) [10, 11]. The connection between these two scalings has been known at least since Mandelbrot [12] and has been further investigated in recent years [13–15], especially for large databases [16], finite text sizes [17, 18], and more general distributions [19, 20]. In this paper, we report the existence of a third type of scaling in the statistics of words: fluctuation scaling. This scaling appears when investigating the fluctuations around the Heaps’ law, i.e., the variance of the vocabulary over different texts of the same size scales with the average. We show that this scaling results from topical aspects of written text that are ignored in the usual connection between Zipf’s and Heaps’ law.

The importance of looking at the fluctuations around Heaps’ law is that this law is used in different applications [21], e.g., (i) to optimize the memory allocation in inverse indexing algorithms [22]; (ii) to estimate the vocabulary of a language [23, 24]; and (iii) to compare the vocabulary richness of documents with different lengths [25–27]. Beyond linguistic applications, scalings of the number of unique items as a function of database size similar to Heaps’ law have been observed in other domains, e.g. the species-area relationship in ecology [28, 29], collaborative tagging [30], network growth [31], and in the statistics of chess moves [32]. These scaling laws have been analyzed from the general viewpoint of innovation dynamics [33] and sampling problems [34]. Our results allow for the quantification of uncertainties in the estimation of these scaling laws and lead to a rethinking of the statistical significance of previous findings.

We use as databases three different collections of texts: (i) all articles of the English Wikipedia [35], (ii) all articles published in the journal PlosOne [36], and (iii) the Google-ngram database [23], a collection of books published in 1520–2008 (each year is treated as a separate document). See appendix A for details on the data.

The manuscript is divided as follows. Section 2 reports our empirical findings with focus on the deviations from a Poisson null model. Section 3 shows how these deviations can be explained by including topicality, which plays the role of a quenched disorder and leads to a non-self averaging process. The consequences of our findings to applications, e.g. vocabulary richness, are discussed in section 4. Finally, section 5 summarizes our main results.

2. Empirical scaling laws

The most-prominent scaling in language is Zipf’s law [9], which states that the frequency, F , of the r th most frequent word (i.e., the fraction of times it occurs in the database) scales as

$$F_r \propto r^{-\alpha} \text{ for } r \gg 1. \quad (1)$$

Another well-studied scaling in language concerns the vocabulary growth and is known as Heaps’ law [10, 11]. It states that the number of different words, N , scales sublinearly with the total number of words, M , i.e.

$$N(M) \propto M^\lambda \text{ for } M \gg 1, \quad (2)$$

with $0 < \lambda < 1$. As a third case, we consider here the problem of the vocabulary growth for an ensemble of texts, and study the scaling of fluctuations by looking at the relation between the standard deviation, $\sigma(M) = \sqrt{\mathbb{V}[N(M)]}$, and the mean value, $\mu(M) = \mathbb{E}[N(M)]$, computed over the ensemble of texts with the same textlength M . In other systems, Taylor's law [6]

$$\sigma(M) \propto \mu(M)^\beta \text{ for } \mu(M) \gg 1 \quad (3)$$

with $1/2 \leq \beta \leq 1$ is typically observed [8].

The connection between scalings (1) and (2) (Zipf's and Heaps' law) can be revealed assuming the usage of each word r is governed by an independent Poisson process with a given frequency F_r . In this description, the number of different words, N , becomes a stochastic variable for which we can calculate the expectation value $\mathbb{E}[N(M)]$ and the variance $\mathbb{V}[N(M)]$ over the realizations of the Poisson process (see appendix B for details)

$$\mathbb{E}[N(M)] \equiv \mu(M) = \sum_r 1 - e^{-MF_r}, \quad (4)$$

$$\mathbb{V}[N(M)] \equiv \sigma(M)^2 \equiv \mathbb{E}[N(M)^2] - \mathbb{E}[N(M)]^2 = \sum_r e^{-MF_r} - e^{-2MF_r}. \quad (5)$$

Assuming Zipf's law (1), for $M \gg 1$ we recover Heaps' law (2), i.e., $\mathbb{E}[N(M)] \propto M^\lambda$, with a simple relation between the scaling exponents $\alpha = \lambda^{-1}$ [37] and Taylor's law (3) with $\beta = 1/2$.

In figure 1, we show empirical data of real texts for the scaling relations (1)–(3) and compare them with predictions from the Poisson null model in equations (4), (5). The Poisson null model correctly elucidates the connection between the scaling exponents in Zipf's and Heaps' law, but it suffers from two severe drawbacks. First, it is of limited use for a quantitative prediction of the vocabulary size for individual articles as it systematically overestimates its magnitude, see figures 1(b), (e) and (h). Second, it dramatically underestimates the expected fluctuations of the vocabulary size, yielding a qualitatively different behavior in the fluctuation scaling: whereas the Poisson null model yields an exponent $\beta \approx 1/2$ expected from central-limit-theorem-like convergence [8], the three empirical data (figures 1(c), (f) and (i)) exhibit a scaling with $\beta \approx 1$. This implies that relative fluctuations of N around its mean value μ for fixed M do not decrease with larger text size (the vocabulary growth, $N(M)$, is a non-self-averaging quantity) and remain of the order of the expected value. Indeed, we find that in all three databases

$$\sigma(M) \approx 0.1\mu(M). \quad (6)$$

Instead of looking at a single value (N, M) for each document, as described previously, an alternative approach is to count the number of different words, N , in the first M words of the document. This leads to a curve $N(M)$ for $M = 1, 2, \dots, M_{\max}$, where M_{\max} is the length of the document. This alternative approach was employed in figures 1(e) and (f) and leads to results equivalent to the ones obtained using single values (N, M), i.e., the $\mu(M)$ and $\sigma(M)$ obtained over different texts lead to identical Heaps' and Taylor's laws. In figure 1(f), we show that anomalous fluctuation scaling in the vocabulary growth is preserved if shuffling the word order of individual texts. This illustrates that in contrast to usual explanations of fluctuation scaling in terms of long-range correlations in time-series [8], here, the observed deviations from the Poisson null model are mainly due to fluctuations across different texts.

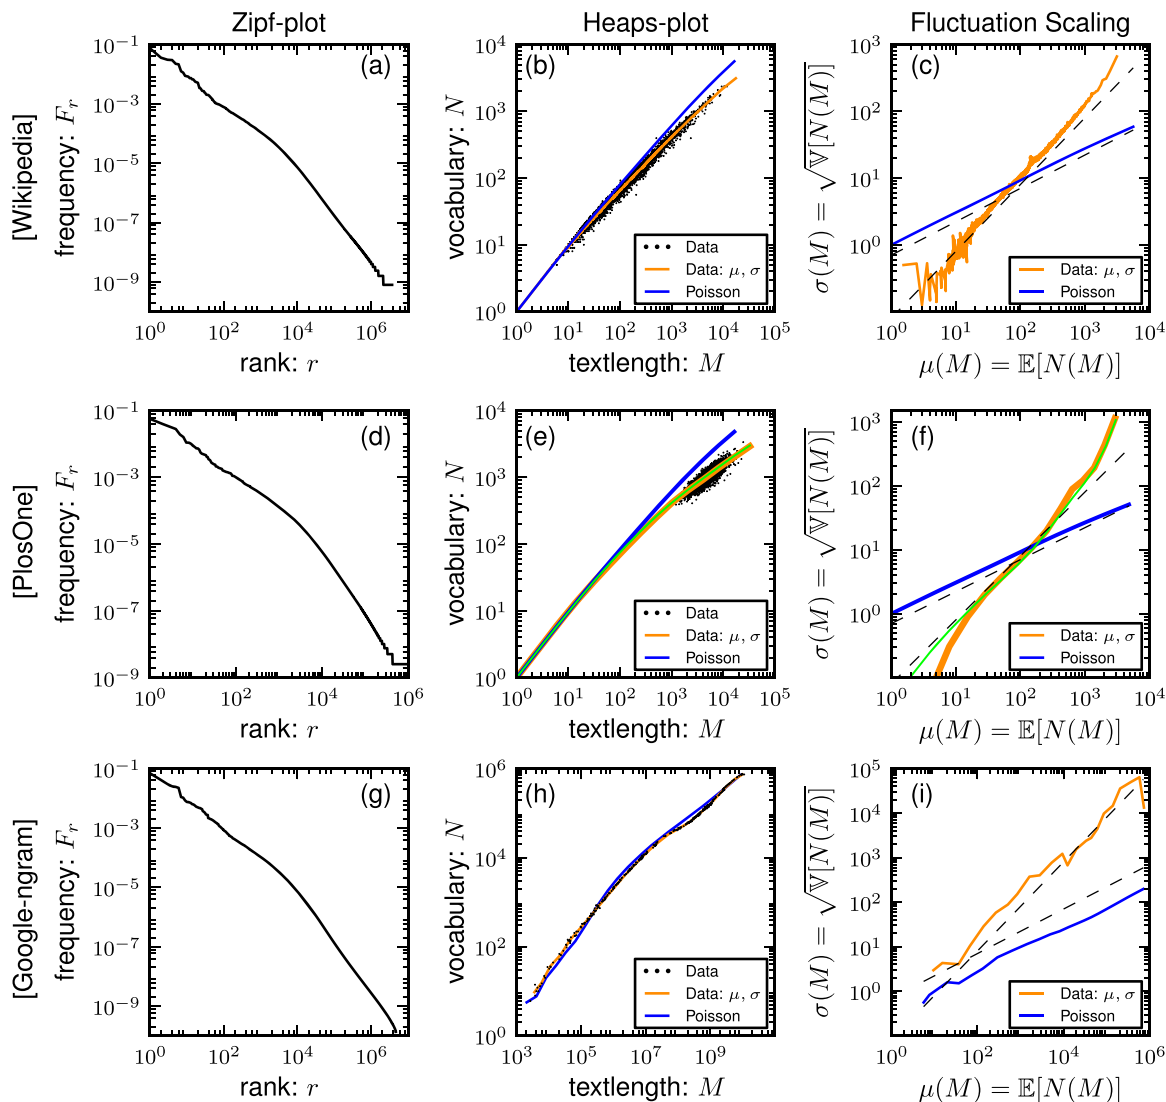


Figure 1. Scaling of Zipf's law (1), Heaps' law (2), and fluctuation scaling (3). Each row corresponds to one of the three databases used in our work. (a,d,g) Zipf's law: rank-frequency distribution F_r considering the full database (the double power-law nature of the curves is apparent [19]). (b,e,h) Heaps' law: the number of different words, N , as a function of textlength, M , for each individual article in the corresponding database (black dots). (c,f,i) Fluctuation scaling: standard deviation, $\sigma(M)$, as a function of the mean, $\mu(M)$, for the vocabulary $N(M)$ conditioned on the textlength M . Poisson (blue-solid) shows the expectation from the Poisson null model, equations (4) and (5), assuming the empirical rank-frequency distribution from (a,d,g), respectively. (Data: μ, σ) (yellow-solid) shows the mean, $\mu(M)$, and standard deviation, $\sigma(M)$, of the data $N(M)$ within a running window in M (see appendix A for the details on the procedure). Additionally, (e,f) show the results (Data: μ, σ) obtained after shuffling the word order for each individual article (thin green-solid). The fact that this curve is indistinguishable from the original curve shows that the results are not due to temporal correlations within the text. For comparison, we show in (c,f,i) the scalings $\sigma(M) \propto \mu(M)^{1/2}$ and $\sigma(M) \propto \mu(M)$ (dashed).

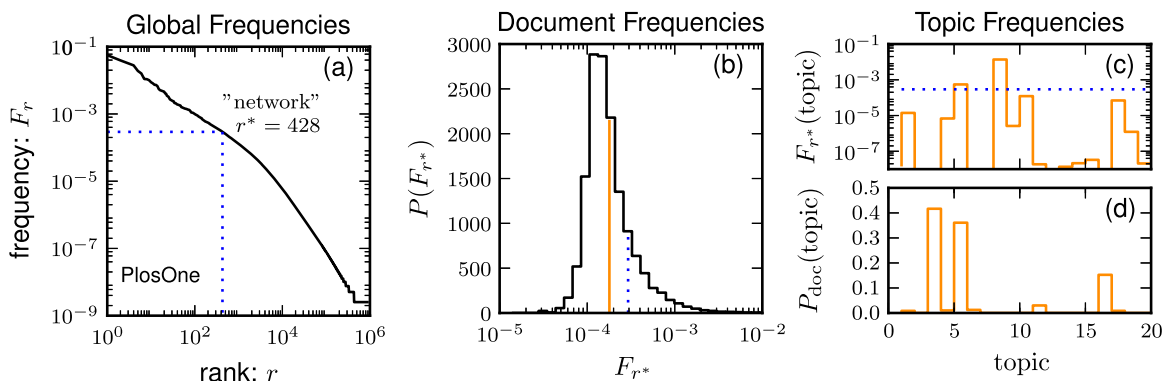


Figure 2. Variation of frequencies due to topicality in the PlosOne database. (a) Rank-frequency distribution considering the complete database. The word ‘network’ (dotted line) has $F_{r^*=428} \approx 2.9 \times 10^{-4}$. (b) Distribution $P(F_{r^*})$ of the local frequency F_{r^*} obtained from each article separately for the word ‘network’ with the global frequency from (a) (dotted). (c) Topic-dependent frequencies $F_{r^*}(\text{topic})$ inferred from LDA with $T = 20$ topics for the word ‘network’ with global frequency from (a) as comparison (dotted). (d) One realization for the topic composition of a single document, $P_{\text{doc}}(\text{topics})$, drawn from a Dirichlet distribution. For this realization, the effective frequency is $F_{r,\text{doc}} = \sum_{t=1}^T P_{\text{doc}}(t) F_r(t) \approx 2.0 \times 10^{-4}$ and is shown in (b) (solid).

In the following, we argue that these observations can be accounted for by considering topical aspects of written language, i.e., instead of treating word frequencies as fixed, we will consider them to be topic dependent ($F_r \mapsto F_r(\text{topic})$).

3. Topicality in vocabulary growth

3.1. Topicality

The frequency of an individual word varies significantly across different texts, meaning that its usage cannot be described alone by a single global frequency [38–40]. For example, consider the usage of the (topical) word ‘network’ in all articles published in the journal PlosOne. It has an overall rank $r^* = 428$ and a global frequency, $F_{r^*=428} \approx 2.9 \times 10^{-4}$, see figure 2(a). The local frequency obtained from each article separately varies over more than one decade, see figure 2(b). Note that, although in this case the local rank-ordering differs from document to document, the index r still refers to the globally determined rank and is used as a unique label for each word.

One popular approach to account for the heterogeneity in the usage of single words is topic models [41]. The basic idea is that the variability across different documents can be explained by the existence of (a smaller number of) topics. In the framework of a generative model, it assumes (i) that individual documents are composed of a mixture of topics (indexed by $t = 1, \dots, T$), with each topic represented in an individual document by the probabilities $P_{\text{doc}}(\text{topic} = t)$ and (ii) that the frequency of each word is topic dependent, i.e., $F_r(\text{topic} = t)$, which leads to a different effective frequency in each document, $F_{r,\text{doc}} = \sum_{t=1}^T P_{\text{doc}}(t) F_r(t)$. One particularly popular variant of topic models is Latent Dirichlet Allocation (LDA) [42], which assumes that the topic composition $P_{\text{doc}}(\text{topic})$ of each document is drawn from a

Dirichlet distribution, P_{Dir} , such that only a few topics contribute substantially to each document. Given a database of documents, LDA infers the topic-dependent frequencies, $F_r(\text{topic})$, from numerical maximization of the posterior likelihood of the generative model [43]. As an illustration, in figure 2(c), we show $F_{r,*}(\text{topic})$ obtained using LDA for the word ‘network’ in the PlosOne database. As expected from a meaningful topic model, we see that the conditional frequencies vary over many orders of magnitude, and that the global frequency $F_{r,*}$ is governed by few topics. The advantage of LDA is that, instead of measuring the distribution of frequencies of each individual word (or two-point distributions for assessing correlations) over different documents, it estimates the frequency of individual words for a finite (and small) number of topics. In combination with the generative model (e.g., drawing $P_{\text{doc}}(\text{topic})$ from a Dirichlet distribution), this not only yields a more compact description of topicality by dramatically reducing the number of parameters, but also allows for an easy extrapolation to unseen texts from a small training sample [42].

3.2. General treatment

In this section, we show how topicality can be included in the analysis of the vocabulary growth. The simplest approach is to consider again that the usage of each word is governed by Poisson processes, but this time to consider that frequencies are not fixed but are themselves random variables that vary across texts.

In this setting, the random variable representing the vocabulary size, N , for a text of length M can be written as

$$N(M) = \sum_r I[n_r(M, F_r)], \quad (7)$$

in which n_r is the integer number of times the word r occurs in a Poisson process of length M with frequency F_r and $I[x]$ is an indicator-type function, i.e., $I[x = 0] = 0$ and $I[x \geq 1] = 1$. The calculation of the expectation value now consists of two parts: (i) the average over realizations i of the Poisson processes $n_r^{(i)}(M, F_r^{(j)})$ for a given realization j of the set of frequencies $F_r^{(j)}$ and (ii) the average overall possible realizations j of the sets of frequencies $F_r^{(j)}$ (which vary due to topicality). In this framework, expectation values correspond to quenched averages (denoted by subscript q)

$$\mathbb{E}_q[N(M)] = \left\langle N(M)^{(i,j)} \right\rangle_{i,j} = \sum_r \left\langle I[n_r^{(i)}(M, F_r^{(j)})] \right\rangle_{i,j} = \sum_r 1 - \left\langle e^{-MF_r^{(j)}} \right\rangle_j, \quad (8)$$

where we used

$$\left\langle I[n_r^{(i)}(M, F_r^{(j)})] \right\rangle_i = 1 - P(n_r = 0; M, F_r^{(j)}) = 1 - e^{-MF_r^{(j)}}. \quad (9)$$

The last equation corresponds to the probability of word r not occurring for a Poisson process of duration M with frequency $F_r^{(j)}$, as in equation (4). For simplicity, hereafter $\langle \dots \rangle \equiv \langle \dots \rangle_j$ (the average over realizations of sets of frequencies $F_r^{(j)}$).

Using the inequality between arithmetic and geometric mean

$$e^{\langle \ln x \rangle} = \langle x \rangle_{\text{geometric}} \leq \langle x \rangle_{\text{arithmetic}} = \langle e^{\ln x} \rangle, \quad (10)$$

we obtain that

$$\mathbb{E}_q[N(M)] = \sum_r 1 - \langle e^{-MF_r} \rangle \leq \sum_r 1 - e^{-M \langle F_r \rangle} \equiv \mathbb{E}_a[N(M)]. \quad (11)$$

The right-hand side corresponds to the result of the Poisson null model (with fixed $F_r = \langle F_r \rangle$), see equation (4), and can be interpreted as an annealed average (denoted by subscript a). This implies that the heterogeneous dissemination of words across different texts leads to a reduction of the expected size of the vocabulary, in agreement with the first deviation of the Poisson null model reported in figures 1(b), (e) and (h).

For the quenched variance, we obtain (see appendix C)

$$\mathbb{V}_q[N(M)] \equiv \mathbb{E}_q[N(M)^2] - \mathbb{E}_q[N(M)]^2 \quad (12)$$

$$= \sum_r \langle e^{-MF_r} \rangle - \langle e^{-MF_r} \rangle^2 + \sum_r \sum_{r' \neq r} \text{Cov}[e^{-MF_r}, e^{-MF_{r'}}] \quad (13)$$

where $\text{Cov}[e^{-MF_r}, e^{-MF_{r'}}] \equiv \langle e^{-MF_r} e^{-MF_{r'}} \rangle - \langle e^{-MF_r} \rangle \langle e^{-MF_{r'}} \rangle$. Comparing to the Poisson case in equation (5), we see that the quenched average yields an additional term containing the correlations of different words. In general, this term does not vanish and is responsible for the anomalous fluctuation scaling with $\beta = 1$ observed in real text, explaining the second deviation from the Poisson null model reported in figures 1(c), (f) and (i).

3.3. Specific ensembles

In this section, we compute the general results from equation (8), (13) for particular ensembles of frequencies $F_r^{(j)}$ and compare them to the empirical results. In the absence of a generally accepted parametric formulation of such an ensemble, we propose two nonparametric approaches explained in the following.

In the first approach, we construct the ensemble $F_r^{(j)}$ directly from the collection of documents, i.e., the frequency $F_r^{(j)}$ corresponds to the frequency of word r in document j , such that

$$\langle e^{-MF_r} \rangle = \frac{1}{D} \sum_{j=1}^D e^{-MF_r^{(j)}}, \quad (14)$$

where D is the number of documents in the data, see figure 2(b).

In the second approach, we construct the ensemble from the LDA topic model [42], in which $F_r^{(j)} = F_r(\text{topic} = j)$ corresponds to the frequency of word r conditional on the topic $j = 1, \dots, T$, see figures 2(c) and (d). In this particular formulation, each document is assumed to consist of a composition of topics, $P_{\text{doc}}(\text{topic})$, which is drawn from a Dirichlet distribution, such that we get for the quenched average

$$\langle e^{-MF_r} \rangle = \int d\theta P_{\text{Dir}}(\theta|\alpha) e^{-MF_r(\theta)}, \quad (15)$$

in which $\theta = (\theta_1, \dots, \theta_T)$ are the probabilities of each topic, $F_r(\theta) = \sum_{j=1}^T \theta_j F_r(\text{topic} = j)$, and the integral is over a T -dimensional Dirichlet-distribution $P_{\text{Dir}}(\theta|\alpha)$ with concentration parameter α . We infer the $F_r(\text{topic})$ using Gensim [43] for LDA with $T = 100$ topics.

The results from both approaches are compared to the PlosOne database in figure 3. Figure 3(a) shows that both methods lead to a reduction in the mean number of different words.




Figure 3. Vocabulary growth for specific topic models. (a) Average vocabulary growth and (b) fluctuation scaling in the PlosOne database (Data) and in the calculations from equations (8), (13) for the two topic models based on the measured frequencies in individual articles (Real Freq) and on LDA (LDA Freq), compare equations (14), (15). For comparison, we show the results from the Poisson null model (Poisson), equations (4), (5), which do not consider topicality. The inset in (b) (same scale as main figure) shows the individual contributions to the fluctuations in equation (13): $\sum_r \langle e^{-MF_r} \rangle - \langle e^{-2MF_r} \rangle$ (dotted) and $\sum_r \sum_{r' \neq r} \text{Cov}[e^{-MF_r}, e^{-MF_{r'}}]$ (solid), illustrating that correlations between different words lead to anomalous fluctuation scaling. The solid lines for LDA-Freq and Real Freq in (b) show the calculations of the corresponding topic models replacing the Poisson by multinomial usage in the derivation of equations (8), (13) to avoid finite-size effects for $\mu(M) < 100$.

Whereas the direct ensemble, equation (14), almost perfectly matches the curve of the data, the LDA-ensemble, equation (15), still overestimates the mean number of different words in the data. This is not surprising because, due to the fewer number of topics (when compared to the number of documents), it constitutes a much more coarse-grained description than the direct ensemble. Additionally, the LDA-ensemble relies on a number of ad-hoc assumptions, e.g., the Dirichlet-distribution in equation (15) or the particular choice of parameters in the inference algorithm, which were not optimized here. More importantly, both methods correctly account for the anomalous fluctuation scaling with $\beta = 1$ observed in the real data, see figure 3(b), and even yield a similar proportionality factor in the quantitative agreement with the data. The comparison of the individual contributions to the fluctuations, equation (13), shown in the inset of figure 3(b) shows that the anomalous fluctuation scaling is due to correlations in the co-occurrence of different words (contained in the term $\text{Cov}[e^{-MF_r}, e^{-MF_{r'}}]$).

4. Applications

4.1. Adding texts

In thermodynamic terms, Heaps' law (as other allometric scalings) implies that the vocabulary size is neither extensive nor intensive ($N(M) < N(2M) < 2N(M)$, also for $M \rightarrow \infty$). Although this can be seen as a direct consequence of Zipf's law, our results show that Heaps' law depends also sensitively on the fluctuations of the frequency of specific words across different documents. To illustrate this, consider the problem of doubling the size of a text of size M . This can be done either by simply extending the size of the same text up to size $2M$

(denoted by $M' = 2 \cdot M$) or by concatenating another text of size M (denoted by $M' = 2 \times M$). The Poisson model (fixed frequency or annealed average) predicts the same expected vocabulary for both procedures

$$\mathbb{E}_a[N(2 \cdot M)] = \mathbb{E}_a[N(2 \times M)] = \sum_r 1 - e^{-2M\langle F_r \rangle}. \quad (16)$$

Taking fluctuations of individual frequencies across documents (quenched average) into account yields (see appendix D for details)

$$\mathbb{E}_q[N(2 \cdot M)] = \sum_r 1 - \langle e^{-2MF_r} \rangle \quad \text{and} \quad \mathbb{E}_q[N(2 \times M)] = \sum_r 1 - \langle e^{-MF_r} \rangle^2. \quad (17)$$

Using equation (10) and the fact that $\langle x^2 \rangle \geq \langle x \rangle^2$, we obtain the following general result

$$\mathbb{E}_q[N(2 \cdot M)] \leq \mathbb{E}_q[N(2 \times M)] \leq \mathbb{E}_a[N(2 \times M)] = \mathbb{E}_a[N(2 \cdot M)]. \quad (18)$$

This is consistent with the intuition that the concatenation of different texts (e.g., on different topics) leads to larger vocabulary than a single longer text. The preceding calculations remain true if the text is extended by a factor k (instead of 2), even for $k \rightarrow \infty$.

The fluctuations around the mean show a more interesting behavior, as revealed by repeating the preceding calculations for the variance. We consider the case of k texts each of length M , such that $M' = k \times M$, and focus on the terms containing correlations between different words shown to be responsible for the anomalous fluctuation scaling (see appendix D for details):

$$\mathbb{V}_q[N(k \times M)] \sim \sum_{r,r'} \langle e^{-MF_r} e^{-MF_{r'}} \rangle^k - \langle e^{-MF_r} \rangle^k \langle e^{-MF_{r'}} \rangle^k. \quad (19)$$

The individual terms can be written as

$$\sum_{r,r'} \langle e^{-MF_r} e^{-MF_{r'}} \rangle^k = \left\langle \left[\sum_r e^{-Mk\bar{F}_r^{(k)}} \right]^2 \right\rangle_{j_1, \dots, j_k}, \quad (20)$$

$$\sum_{r,r'} \langle e^{-MF_r} \rangle^k \langle e^{-MF_{r'}} \rangle^k = \left[\left\langle \sum_r e^{-Mk\bar{F}_r^{(k)}} \right\rangle_{j_1, \dots, j_k} \right]^2, \quad (21)$$

in which $\langle \cdot \rangle_{j_1, \dots, j_k}$ denotes the averaging over the realizations (j_1, \dots, j_k) of frequencies $F_r^{(j_i)}$ in each single text $i = 1, \dots, k$ and $\bar{F}_r^{(k)} = \frac{1}{k} \sum_{i=1}^k F_r^{(j_i)}$ is the k -sample average frequency based on the realizations (j_1, \dots, j_k) . In the limit $k \rightarrow \infty$: $\bar{F}_r^{(k)} \rightarrow \langle F_r \rangle$ such that

$$\sum_{r,r'} \langle e^{-MF_r} e^{-MF_{r'}} \rangle^k - \langle e^{-MF_r} \rangle^k \langle e^{-MF_{r'}} \rangle^k \rightarrow 0 \quad (22)$$

for $k \rightarrow \infty$. This implies that, for $k \gg 1$, (adding many different texts) the fluctuations in the vocabulary across documents (and therefore the correlations between different words) vanish and normal fluctuation scaling ($\beta = 1/2$) is recovered. This prediction can be tested in data. Starting from a collection of documents, we create a new collection by concatenating k randomly selected documents (each document is used once). We then compute for each

concatenated document the number of distinct words N up to size M for increasing M , $\mathbb{E}[N(M)]$, and $\mathbb{V}[N(M)]$. We observe a transition of the exponent β in the fluctuation scaling, equation (3), from $\beta \approx 1 \rightarrow \beta \approx 1/2$.

4.2. Vocabulary richness

When measuring vocabulary richness, we want a measure that is robust to different text sizes. The traditional approach is to use Herdan's C , i.e., $C = \log N / \log M$ [25–27]. Although quite effective for rough estimations, this approach has several problems. An obvious problem is that it does not incorporate any deviations from the original Heaps' law (e.g., the double scaling regime [19]). More seriously, it does not provide any estimation of the statistical significance or expected fluctuations of the measure. For instance, if two values are measured for different texts, one cannot determine whether one is significantly larger than the other. Our approach is to compare observations with the fluctuations expected from models in the spirit of section 3.2.

The computation of statistical significance requires an estimation of the probability of finding N different words in a text of length M , $P(N|M)$, which can be obtained from a given generative model (e.g., as presented in section 3). For a text with (N^*, M^*) , we compute the percentile $P(N > N^*|M^*)$, which allows for a ranking of texts with different sizes such that the smaller the percentile, the richer the vocabulary. An estimation of the significance of the difference in the vocabulary can then be obtained by comparison of the different percentile.

For the sake of simplicity, we illustrate this general approach by approximating $P(N|M)$ using a Gaussian distribution. In this case, the percentile are determined by the mean, $\mu(M) = \mathbb{E}[N(M)]$, and the variance, $\sigma(M) = \sqrt{\mathbb{V}[N(M)]}$, in terms of the z -score

$$z_{(N,M)} = \frac{N - \mu(M)}{\sigma(M)}, \quad (23)$$

which shows how much the measured value (N, M) deviates from the expected value $\mu(M)$ in units of standard deviations ($z_{(N,M)}$ follows a standard normal distribution: $z \stackrel{d}{\sim} \mathcal{N}(0, 1)$). If we consider our quantitative result on fluctuation scaling in the vocabulary in equation (6), i.e., $\sigma(M) \approx 0.1\mu(M)$, we can calculate the z -score of the observation (N, M) as

$$z_{(N,M)} \approx \frac{N - \mu(M)}{0.1\mu(M)} = 10 \left(\frac{N}{\mu(M)} - 1 \right), \quad (24)$$

in which we need to include the expected vocabulary growth, $\mu(M)$, from a given generative model (e.g., Heaps' law with two scalings [19]). We can now: (i) for a single text (N, M) , assign a value of lexical richness, the z -score $z_{(N,M)}$, considering deviations from the pure Heaps' law that should be included in $\mu(M)$; (ii) given two texts (N_1, M_1) and (N_2, M_2) , compare directly the respective z -scores $z_{(N_1, M_1)}$ and $z_{(N_2, M_2)}$ to assess which text has a higher lexical richness independent of the difference in the textlengths; and (iii) estimate the statistical significance of the difference in vocabulary by considering $\Delta z := z_{(N_1, M_1)} - z_{(N_2, M_2)}$, which is distributed according to $\Delta z \stackrel{d}{\sim} \mathcal{N}(0, 2)$ because $z \stackrel{d}{\sim} \mathcal{N}(0, 1)$. Point (iii) implies that the difference in the vocabulary richness of two texts is statistically significant on a 95%-confidence level if $|\Delta z| > 2.77$, i.e., in this case there is at most a 5% chance that the observed difference originates from topic fluctuations. As a general rule, for two texts of approximately the same length ($N(M) \approx \mu(M)$), the relative difference in the vocabulary must be larger than 27.7% to be sure on a 95%-confidence level that the difference is not due to expected topic fluctuations.




Figure 4. Measures of vocabulary richness. For 5000 randomly selected articles from the Wikipedia database (black dots), we compute the ratio between the number of different words $N(M)$ and the average number of different words $N_{\text{avg}}(M)$ (empirically determined from all articles with the same textlength M). We compare the predictions of different measures of vocabulary richness (solid lines): (a) Herdan's C and (b+c) z -score, equation (24), in which we calculate the expected null model, $\mu(M)$, according to equation (E.5) with parameters $\gamma = 1.77$, $\tilde{r} = 7830$ [19], and $a \rightarrow \infty$ (in b) or $a = 0.08$ (in c). The solid lines are contours corresponding to values of $N(M)$ that yield the same measure of vocabulary richness varying from rich (red: $C = 0.98$ and $z = 4$) to poor (purple: $C = 0.8$ and $z = -4$) vocabulary. The article with the richest vocabulary according to each measure is marked by \times (red).

We illustrate this approach for the vocabulary richness of Wikipedia articles. As a proxy for the true vocabulary richness, we measure how much the vocabulary of each article, $N(M)$, exceeds the average vocabulary $N_{\text{avg}}(M)$ with the same textlength M empirically determined from all articles in the Wikipedia. In practice however, when assessing the vocabulary richness of a single article, information of $N_{\text{avg}}(M)$ from an ensemble of texts is usually not available and measures such as the ones described previously are needed. In figure 4, we compare the accuracy of measures of vocabulary richness according to Herdan's C , figure 4(a), and the z -score, figures 4(b) and (c). For the latter, we use equation (24) and calculate $\mu(M)$ from Poisson word usage by fixing Zipf's law and assuming Gamma-distributed word frequencies across documents, see appendix E for details. We see in figure 4(a) that Herdan's C shows a strong bias towards assigning high values of C to shorter texts: following a line with constant C , we observe for $M \gtrsim 10$ articles with a vocabulary below average, whereas for $M > 1000$ articles with a vocabulary above average. A similar (weaker) bias is observed in figure 4(b) for the calculation of the z -score for the case in which we consider deviations from the pure Heaps' law but treat frequencies of individual words as fixed, i.e., ignoring topicality. The z -score calculations including topicality in figure 4(c) show that we obtain a measure of vocabulary richness which is approximately unbiased with respect to the textlength M (contour lines are roughly horizontal). Furthermore, in contrast to the two other measures, we correctly assign the highest z -score to the article with the highest ratio $N(M)/N_{\text{avg}}(M)$. Altogether, this implies that it is not only important to consider deviations from the pure Heaps' law, but that it is crucial to consider topicality in the form of a quenched average.

5. Discussion

In summary, we used large text databases to investigate the scaling between vocabulary size N (number of different text words) and database size M . Besides the usual analysis of the average vocabulary size (Heaps' law), we measured the standard deviation across different texts with the same length M . We found that the relative fluctuations (standard deviation divided by the mean) do not decay with M in contrast to simple sampling processes. We explained this observation using a simple stochastic process (Poisson usage of words) in which we account for topical aspects of written text, i.e., the frequency of an individual word is not treated as fixed across different documents. This heterogeneous dissemination of words across different texts leads to a reduction of the expected size of the vocabulary and to an increase in the variance. We have further shown the implications of these findings by proposing a practical measure of vocabulary richness that allows for a comparison of the vocabulary of texts with different lengths, including the quantification of statistical significance.

Our finding of anomalous fluctuation scaling implies that the vocabulary is a non-self-averaging quantity, meaning that the vocabulary of a single text is not representative of the whole ensemble. Here, we emphasized that topicality can be responsible for this effect. Although the existence of different topics is obvious for a collection of articles as broad in content as the Wikipedia, our analysis shows that we can apply the same reasoning for the Google-ngram data, in which case the frequency variation is measured at different times. This offers a new perspective on language change [44]: the difference in the vocabulary from different years can be seen as a shift in the topical content over time. Similarly, other systematic fluctuations (e.g., across different authors or in the parameters of the Zipf's law) can play a similar role as topicality.

Beyond linguistic applications, allometric scaling [4, 5] and other sublinear scalings similar to Heaps' law [28–33] have been observed in different complex systems. Our results show the importance of studying fluctuations around these scalings and provide a theoretical framework for the analysis.

Acknowledgements

We thank Diego Rybski for insightful discussion on fluctuation scaling.

Appendix A. Data

The Wikipedia database consists of the plain text of all 3, 743, 306 articles from a snapshot of the complete English Wikipedia [35]. The PlosOne database consists of all 76, 723 articles published in the journal PlosOne, which were accessible at the time of the data collection [36]. The Google-ngram database is a collection of printed books counting the number of times a word appears in a given year $t \in [1520–2008]$ [23]. We treat the collection of all books published in the the same year as a single document, yielding 393 observations for different t .

We apply the same filtering for each database: (i) we decapitalize each word (e.g., 'the' and 'The' are counted as the same word) and (ii) we restrict ourselves to words consisting uniquely of letters present in the alphabet of the English language. This is meant as a

conservative approach to minimize the influence of foreign words, numbers (e.g. prices), or scanning problems present in the raw data (for details on the preprocessing see [19]).

Due to peculiarities of the individual databases the data (Data: μ , σ) in figure 1, i.e., the calculation of the curves $\mu(M)$ and $\sigma(M)$ conditioned on the textlength M , is constructed in a slightly different way in each case. In the Wikipedia data, we order all datapoints $N(M)$ (of the full article) according to textlength M and consider 1000 consecutive datapoints (in M), from which we calculate the average value of the textlength M , and the conditional mean, $\mu(M)$, and variance, $\sigma(M)$, of the vocabulary N . In the PlosOne data, the length of all articles is much more concentrated, which is why we consider the full trajectory $N(M)$ with $M = 1, 2, \dots, M_{\max}$ for each individual article. For an arbitrary value of M , we calculate $\mu(M)$ and $\sigma(M)$ from the ensemble of all articles with vocabulary N at the particular textlength M . In the Google-ngram data, we impose a logarithmic binning in M such that we can calculate $\mu(M)$ and $\sigma(M)$ from a finite number of samples in each bin.

Appendix B. Poisson null model

The number of different words in each realization of the Poisson process is given by

$$N(M) = \sum_r I[n_r(M, F_r)], \quad (\text{B.1})$$

in which n_r is the integer number of times the word r occurs in a Poisson process of length M with frequency F_r and $I[x]$ is an indicator-type function, i.e., $I[x = 0] = 0$ and $I[x \geq 1] = 1$. Averaging over realizations of the Poisson process requires the calculation of $\mathbb{E}[I[n_r(M, F_r)]] \equiv \langle I[n_r(M)] \rangle = 1 - e^{-MF_r}$, which is the probability that the word with rank r appears at least once in a text of length M . Considering all words, we obtain

$$\mathbb{E}[N(M)] = \sum_r \langle I[n_r(M)] \rangle = \sum_r 1 - e^{-MF_r}, \quad (\text{B.2})$$

$$\mathbb{V}[N(M)] \equiv \mathbb{E}[N(M)^2] - \mathbb{E}[N(M)]^2 \quad (\text{B.3})$$

$$= \sum_{r,r'} \langle I[n_r] I[n_{r'}] \rangle - \sum_{r,r'} \langle I[n_r] \rangle \langle I[n_{r'}] \rangle \quad (\text{B.4})$$

$$= \sum_r \langle I[n_r]^2 \rangle + \sum_{\substack{r,r' \\ r \neq r'}} \langle I[n_r] I[n_{r'}] \rangle - \sum_{r,r'} \langle I[n_r] \rangle \langle I[n_{r'}] \rangle \quad (\text{B.5})$$

$$= \sum_r \langle I[n_r] \rangle + \sum_{\substack{r,r' \\ r \neq r'}} \langle I[n_r] \rangle \langle I[n_{r'}] \rangle - \sum_{r,r'} \langle I[n_r] \rangle \langle I[n_{r'}] \rangle \quad (\text{B.6})$$

$$= \sum_r e^{-MF_r} - e^{-2MF_r} \quad (\text{B.7})$$

where we used that $I[x]^2 = I[x]$ and that Poisson processes of different words ($r \neq r'$) are independent of each other.

Appendix C. Calculation $\mathbb{E}_q [N(M)^2]$

$$\mathbb{E}_q [N(M)^2] = \left\langle N(M)^{(i,j)} N(M)^{(i,j)} \right\rangle_{i,j} \quad (\text{C.1})$$

$$= \left\langle \sum_{r,r'} I[n_r^{(i)}(M, F_r^{(j)})] I[n_{r'}^{(i)}(M, F_{r'}^{(j)})] \right\rangle_{i,j} \quad (\text{C.2})$$

$$= \sum_r \left\langle I[n_r^{(i)}(M, F_r^{(j)})]^2 \right\rangle_{i,j} + \left\langle \sum_r \sum_{r' \neq r} I[n_r^{(i)}(M, F_r^{(j)})] I[n_{r'}^{(i)}(M, F_{r'}^{(j)})] \right\rangle_{i,j} \quad (\text{C.3})$$

$$= \sum_r \left\langle I[n_r^{(i)}(M, F_r^{(j)})] \right\rangle_{i,j} + \sum_r \sum_{r' \neq r} \left\langle \left\langle I[n_r^{(i)}(M, F_r^{(j)})] \right\rangle_i \left\langle I[n_{r'}^{(i)}(M, F_{r'}^{(j)})] \right\rangle_j \right\rangle \quad (\text{C.4})$$

$$= \sum_r 1 - \left\langle e^{-MF_r} \right\rangle + \sum_r \sum_{r' \neq r} \left\langle (1 - e^{-MF_r})(1 - e^{-MF_{r'}}) \right\rangle \quad (\text{C.5})$$

where we used $I[x]^2 = I[x]$, equation (9), and that two Poisson process of different words ($r \neq r'$) with a given set of frequencies $F_r^{(j)}$ are independent of each other.

Appendix D. Adding texts

In this section, we show the calculation for the quenched averages of the mean and the variance of the vocabulary growth when considering a text of length M' from the concatenation of k different texts of length M_i with $M' = \sum_{i=1}^k M_i$. We will first focus on the case $k = 2$, i.e., $M' = M_1 + M_2$, from which we can easily generalize to arbitrary k .

We consider the vocabulary growth, $N(M')$, as a random variable in which we concatenate two independent realizations of the stochastic process introduced in section 3.2 indicated by subscript (1) and (2) respectively:

$$N(M' = M_1 + M_2) = \sum_r I[n_r^{(1)}(M_1, F_r^{(1)}) + n_r^{(2)}(M_2, F_r^{(2)})] \quad (\text{D.1})$$

$$= \sum_r I[n_r^{(1)}(M_1, F_r^{(1)})] + I[n_r^{(2)}(M_2, F_r^{(2)})] - I[n_r^{(1)}(M_1, F_r^{(1)})] I[n_r^{(2)}(M_2, F_r^{(2)})] \quad (\text{D.2})$$

in which the word r is counted as part of the vocabulary if it appears in either of the two concatenated realizations of the stochastic process. In the same spirit as in section 3.2, taking expectation values requires averaging over all realizations of the Poisson process (i_1, i_2) given the frequencies $F_r^{(j_1)}, F_r^{(j_2)}$ as well as averaging over all realizations of those frequencies (j_1, j_2) , which we denote by $\langle \cdot \rangle_{i_1, i_2, j_1, j_2}$. For the individual terms appearing in $N(M' = M_1 + M_2)$, we obtain

$$\left\langle I \left[n_r^{(i_1)} \left(M_1, F_r^{(j_1)} \right) \right] \right\rangle_{i_1, i_2, j_1, j_2} = 1 - \left\langle e^{-M_1 F_r^{(j_1)}} \right\rangle_{j_1} \quad (\text{D.3})$$

$$\left\langle I \left[n_r^{(i_2)} \left(M_2, F_r^{(j_2)} \right) \right] \right\rangle_{i_1, i_2, j_1, j_2} = 1 - \left\langle e^{-M_2 F_r^{(j_2)}} \right\rangle_{j_2}, \quad (\text{D.4})$$

$$\begin{aligned} & \left\langle I \left[n_r^{(i_1)} \left(M_1, F_r^{(j_1)} \right) \right] I \left[n_r^{(i_2)} \left(M_2, F_r^{(j_2)} \right) \right] \right\rangle_{i_1, i_2, j_1, j_2} \\ &= \left(1 - \left\langle e^{-M_1 F_r^{(j_1)}} \right\rangle_{j_1} \right) \left(1 - \left\langle e^{-M_2 F_r^{(j_2)}} \right\rangle_{j_2} \right), \end{aligned} \quad (\text{D.5})$$

in which we can separate the average over (i_1, j_1) and (i_2, j_2) , assuming that the two concatenated realizations (i_1, j_1) and (i_2, j_2) of the original stochastic process are independent. For the calculation of the expectation of $N(M' = M_1 + M_2)^2$, we get higher order terms for $r \neq r'$:

$$\begin{aligned} & \left\langle I \left[n_r^{(i_1)} \left(M_1, F_r^{(j_1)} \right) \right] I \left[n_{r'}^{(i_1)} \left(M_1, F_{r'}^{(j_1)} \right) \right] \right\rangle_{i_1, i_2, j_1, j_2} \\ &= \left\langle \left(1 - e^{-M_1 F_r^{(j_1)}} \right) \left(1 - e^{-M_1 F_{r'}^{(j_1)}} \right) \right\rangle_{j_1}. \end{aligned} \quad (\text{D.6})$$

From this, we can evaluate the mean and variance

$$\mathbb{E}_q[N(M' = M_1 + M_2)] = \sum_r 1 - \langle e^{-M_1 F_r} \rangle \langle e^{-M_2 F_r} \rangle \quad (\text{D.7})$$

$$\begin{aligned} \mathbb{V}_q[N(M' = M_1 + M_2)] &= \sum_r \langle e^{-M_1 F_r} \rangle \langle e^{-M_2 F_r} \rangle - \langle e^{-2M_1 F_r} \rangle \langle e^{-2M_2 F_r} \rangle \\ &+ \sum_{r, r'} \langle e^{-M_1 F_r} e^{-M_1 F_{r'}} \rangle \langle e^{-M_2 F_r} e^{-M_2 F_{r'}} \rangle \\ &- \langle e^{-M_1 F_r} \rangle \langle e^{-M_1 F_{r'}} \rangle \langle e^{-M_2 F_r} \rangle \langle e^{-M_2 F_{r'}} \rangle. \end{aligned} \quad (\text{D.8})$$

Generalizing to the concatenation of an arbitrary number of k texts can be treated in the very same way; however, we will only state the result for the case of adding k texts of equal length M such that $M' = k \times M$:

$$\mathbb{E}_q[N(M' = k \times M)] = \sum_r 1 - \langle e^{-MF_r} \rangle^k \quad (\text{D.9})$$

$$\begin{aligned} \mathbb{V}_q[N(M' = k \times M)] &= \sum_r \langle e^{-MF_r} \rangle^k - \langle e^{-2MF_r} \rangle^k \\ &+ \sum_{r,r'} \langle e^{-MF_r} e^{-MF_{r'}} \rangle^k - \langle e^{-MF_r} \rangle^k \langle e^{-MF_{r'}} \rangle^k. \end{aligned} \quad (\text{D.10})$$

Appendix E. Vocabulary growth for Gamma-distributed frequency and a double power-law

Assuming a gamma-distribution for the distribution of the frequency of single words across different texts [38]

$$P_r(F_r = x; a, b) = \frac{1}{\Gamma(a)} b^{-a} x^{a-1} e^{-x/b} \quad (\text{E.1})$$

we can calculate the quenched average

$$\langle e^{-MF_r} \rangle = \int dx P_r(F_r = x; a, b) e^{-Mx} = (1 + bM)^{-a}. \quad (\text{E.2})$$

If we assume that the distribution of frequencies for all words is given by the same shape-parameter a (e.g., $a = 1$ corresponds to an exponential distribution) and fix the mean of the distribution, given by $\langle F_r \rangle = ab$ we get $\langle e^{-MF_r} \rangle = (1 + M\langle F_r \rangle/a)^{-a}$. Assuming a double power-law for the average rank-frequency distribution [19] with parameters γ and \tilde{r} , i.e., $\langle F_r \rangle = Cr^{-1}$ for $r \leq \tilde{r}$ and $\langle F_r \rangle = C\tilde{r}^{\gamma-1}r^{-\gamma}$ for $r > \tilde{r}$, where $C = C(\tilde{r}, \gamma)$ is the normalization constant determined by imposing $\sum_r \langle F_r \rangle = 1$, we can calculate the vocabulary growth according to equation (4) analytically in the continuum approximation by substituting $x := \langle F_r \rangle$

$$\mathbb{E}_q[N(M)] = \sum_r 1 - (1 + M\langle F_r \rangle/a)^{-a} \quad (\text{E.3})$$

$$= -\int_0^1 dx \frac{dr}{dx} [1 - (1 + Mx/a)^{-a}] \quad (\text{E.4})$$

which can be expressed in terms of the ordinary hypergeometric function $H := {}_2F_1$ [45] yielding

$$\begin{aligned} \mathbb{E}_q[N(M)] &= \tilde{r} - C + \tilde{r} \left[H\left(a, -\frac{1}{\gamma}, 1 - \frac{1}{\gamma}, -\frac{CM}{a\tilde{r}}\right) - 1 \right] \\ &- C \left(1 + \frac{M}{a}\right)^{-a} \left[a \frac{\Gamma(1+a)}{\Gamma(2+a)} H\left(1, 1, 2+a, -\frac{a}{M}\right) - 1 \right] \\ &+ \tilde{r} \left(1 + \frac{CM}{a\tilde{r}}\right)^{-a} \left[a \frac{\Gamma(1+a)}{\Gamma(2+a)} H\left(1, 1, 2+a, -\frac{a\tilde{r}}{CM}\right) - 1 \right], \end{aligned} \quad (\text{E.5})$$

where the vocabulary growth $\mathbb{E}_q[N(M)]$ is parametrized by γ , \tilde{r} , and a .

In the limit $a \rightarrow \infty$, the Gamma distribution $P_T(F_r = x; a, b)$ with given mean $\langle F_r \rangle = ab = \text{const.}$ converges to a Gaussian with $\sigma^2 = \langle F_r \rangle^2/a$. For $a \rightarrow \infty$, $\sigma^2 \rightarrow 0$ and we recover the Poisson null model, equations (4), (5), in which the individual frequencies F_r are fixed (annealed average).

References

- [1] Clauset A, Shalizi C R and Newman M E J 2009 *SIAM Review* **51** 661
- [2] Mitzenmacher M 2004 *Internet Math.* **1** 226–51
- [3] Newman M E J 2005 *Contemp. Phys.* **46** 323–51
- [4] West G B 1997 *Science* **276** 122–6
- [5] Bettencourt L M A, Lobo J, Helbing D, Kühnert C and West G B 2007 *Proc. Natl Acad. Sci.* **104** 7301–6
- [6] Taylor L R 1961 *Nature* **189** 732–5
- [7] de Menezes M and Barabási A L 2004 *Phys. Rev. Lett.* **92** 028701
- [8] Eisler Z, Bartos I and Kertész J 2008 *Adv. Phys.* **57** 89–142
- [9] Zipf G K 1949 *Human Behaviour and the Principle of Least Effort* (Cambridge, MA: Addison-Wesley Press)
- [10] Herdan G 1958 *Biometrika* **45** 222–8
- [11] Heaps H S 1978 *Information Retrieval* (New York: Academic)
- [12] Mandelbrot B 1961 On the Theory of Word Frequencies and on Related Markovian Models of Discourse *Structure of Language and Its Mathematical Aspects: Proceedings of Symposia in Applied Mathematics* Vol XII (Providence, RI: American Mathematical Society)
- [13] van Leijenhorst D C and van der Weide T P 2005 *Inf. Sci.* **170** 263–72
- [14] Zanette D and Montemurro M 2005 *J. Quant. Linguist.* **12** 29–40
- [15] Serrano M A, Flammini A and Menczer F 2009 *PloS One* **4** e5372
- [16] Williams H E and Zobel J 2005 *Int. J. Dig. Lib.* **5** 99–105
- [17] Bernhardsson S, Correa da Rocha L E and Minnhagen P 2009 *New J. Phys.* **11** 123015
- [18] Lü L, Zhang Z and Zhou T 2010 *PloS One* **5** e14139
- [19] Gerlach M and Altmann E G 2013 *Phys. Rev. X* **3** 021006
- [20] Font-Clos F, Boleda G and Corral A 2013 *New J. Phys.* **15** 093033
- [21] Manning C D, Raghavan P and Schuetze H 2008 *Introduction to Information Retrieval* (New York: Cambridge University Press)
- [22] Baeza-Yates R and Navarro G 2000 *J. Am. Soc. Inf. Sci.* **51** 69–82
- [23] Michel J B *et al* 2011 *Science* **331** 176–82
- [24] Klein W 2013 Von Reichtum und Armut des Deutschen Wortschatzes *Reichtum und Armut der Deutschen Sprache* (Berlin: De Gruyter) pp 15–56
- [25] Wimmer G and Altmann G 1999 *J. Quant. Linguist.* **6** 1–9
- [26] Baayen R H 2001 *Word Frequency Distributions* (Dordrecht, Netherlands: Kluwer Academic Publishers)
- [27] Yasserli T, Kornai A and Kertész J 2012 *PLoS One* **7** e48386
- [28] Brainerd B 1982 *J Appl. Probab.* **19** 785–93
- [29] García Martín H and Goldenfeld N 2006 *Proc. Natl Acad. Sci.* **103** 10310–5
- [30] Cattuto C, Barrat A, Baldassarri A, Schehr G and Loreto V 2009 *Proc. Natl Acad. Sci.* **106** 10511–5
- [31] Krapivsky P L and Redner S 2013 *J. Stat. Mech.* **2013** P06002
- [32] Perotti J I, Jo H H, Schaigorodsky A L and Billoni O V 2013 *EPL (Europhysics Letters)* **104** 48005
- [33] Tria F, Loreto V, Servedio V D P and Strogatz S H 2014 *Sci. Rep.* **4** 5890
- [34] Gnedin A, Hansen B and Pitman J 2007 *Probab. Surv.* **4** 146–71
- [35] Wikimedia. Dump of the English Wikipedia from 02/06/2012 (<http://dumps.wikimedia.org/enwiki/>); accessed 26/06/2012
- [36] PlosApi. All articles published in the journal plosone (<http://api.plos.org/>); accessed 17/10/2013

- [37] Eliazar I 2011 *Physica A* **390** 3189–203
- [38] Church K W and Gale W A 1995 *Nat. Lang. Eng.* **1** 163–190
- [39] Montemurro M and Zanette D 2010 *Adv. Complex Syst.* **13** 135
- [40] Altmann E G, Pierrehumbert J B and Motter A E 2011 *PloS One* **6** e19009
- [41] Blei D M 2012 *Commun. ACM* **55** 77
- [42] Blei D M, Ng A Y and Jordan M I 2003 *J. Mach. Learn. Res.* **3** 993–1022
- [43] Řehůřek R and Sojka P 2010 Software Framework for Topic Modelling with Large Corpora *Proc. LREC 2010 Workshop on New Challenges for NLP Frameworks* (Valletta, Malta: ELRA) pp 45–50
- [44] Baronchelli A, Loreto V and Tria F 2012 *Adv. Complex Syst.* **15** 1203002–1
- [45] Abramowitz M and Stegun I A (ed) 1972 *Handbook of Mathematical Functions* (New York: Dover Publications)