

Similarity of Symbol Frequency Distributions with Heavy Tails

Martin Gerlach,^{1,*} Francesc Font-Clos,^{2,3} and Eduardo G. Altmann¹

¹Max Planck Institute for the Physics of Complex Systems, D-01187 Dresden, Germany

²Centre de Recerca Matemàtica, Edifici C, Campus Bellaterra, E-08193 Barcelona, Spain

³Departament de Matemàtiques, Facultat de Ciències, Universitat Autònoma de Barcelona, E-08193 Barcelona, Spain

(Received 2 October 2015; published 15 April 2016)

Quantifying the similarity between symbolic sequences is a traditional problem in information theory which requires comparing the frequencies of symbols in different sequences. In numerous modern applications, ranging from DNA over music to texts, the distribution of symbol frequencies is characterized by heavy-tailed distributions (e.g., Zipf's law). The large number of low-frequency symbols in these distributions poses major difficulties to the estimation of the similarity between sequences; e.g., they hinder an accurate finite-size estimation of entropies. Here, we show analytically how the systematic (bias) and statistical (fluctuations) errors in these estimations depend on the sample size N and on the exponent γ of the heavy-tailed distribution. Our results are valid for the Shannon entropy ($\alpha = 1$), its corresponding similarity measures (e.g., the Jensen-Shanon divergence), and also for measures based on the generalized entropy of order α . For small α 's, including $\alpha = 1$, the errors decay slower than the $1/N$ decay observed in short-tailed distributions. For α larger than a critical value $\alpha^* = 1 + 1/\gamma \leq 2$, the $1/N$ decay is recovered. We show the practical significance of our results by quantifying the evolution of the English language over the last two centuries using a complete α spectrum of measures. We find that frequent words change more slowly than less frequent words and that $\alpha = 2$ provides the most robust measure to quantify language change.

DOI: 10.1103/PhysRevX.6.021009

Subject Areas: Complex Systems,
Interdisciplinary Physics,
Statistical Physics

I. INTRODUCTION

Quantifying the similarity of symbolic sequences is a classical problem in information theory [1] with modern applications in linguistics [2], genetics [3], and image processing [4]. The availability of large databases of texts sparked a renewed interest in the problem of similarity of the vocabulary of different collections of texts [5–9]. For instance, Fig. 1 shows the word-frequency distribution in three large collections of English texts: from 1850, 1900, and 1950. We see that the distribution itself remains essentially the same, a heavy-tailed Zipf distribution [11],

$$p(r) \propto r^{-\gamma}, \quad (1)$$

where p is the frequency of the r th most frequent word and $\gamma \gtrsim 1$. Changes are seen in the frequency p (or rank) of specific words; e.g., *ship* lost and *genetic* won popularity. Measures that quantify such changes are essential to answer

questions such as the following: Is the vocabulary from 1900 more similar to the one from 1850 or to the one from 1950? How similar are two vocabularies (e.g., from different years)? Are the two finite-size observations compatible

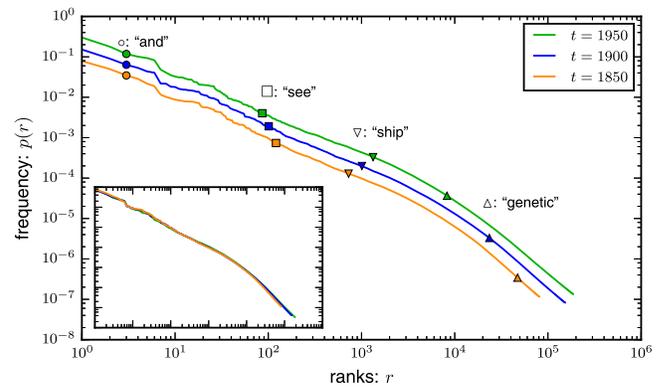


FIG. 1. The English vocabulary in different years. Rank-frequency distribution $p(r)$ of individual years t for $t = 1850$, 1900, and 1950 of the Google-ngram database [10], multiplied by a factor of 1, 2, and 4, respectively, for better visual comparison. The inset shows the original untransformed data (same axis), highlighting that the rank-frequency distributions are almost indistinguishable. Individual words (e.g., “and,” “see,” “ship,” and “genetic”) show changes in rank and frequency (symbols), where words with larger ranks (i.e., smaller frequencies) show larger change.

*Corresponding author.
gerlach@pks.mpg.de

Published by the American Physical Society under the terms of the Creative Commons Attribution 3.0 License. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

with a finite sample of the same underlying vocabulary? How similar are the vocabularies of different authors or disciplines? How fast is the lexical change taking place?

Heavy-tailed and broad distributions of symbol frequencies such as Eq. (1) are typical in natural languages [11–16] and appear also in the DNA (n-grams of base pairs for large n) [17], in gene expression [18], and in music [19]. The slow decay observed in a broad range of frequencies implies that there is no typical frequency for words, and therefore, relevant changes can occur in different ranges of the p spectrum, from the few large-frequency words all the way to the many low-frequency words. This imposes a challenge to define similarity measures that are able to account for this variability and that also yield accurate estimations based on finite-size observations.

In this paper, we quantify the vocabulary similarity using a spectrum of measures D_α based on the generalized entropy of order α [$D_{\alpha=1}$ recovers the usual Jensen-Shannon divergence]. We show how varying α magnifies differences in the vocabulary at different scales of the (heavy-tailed) frequency spectrum, thus providing different information on the vocabulary change. We then compute the accuracy (bias) and precision (variance) of estimations of D_α based on sequences of size N and find that in heavy-tailed distributions, the convergence is much slower than in non-heavy-tailed distributions (it often scales as $1/N^\beta$ with $\beta < 1$). Finally, we come back to the problem of comparing the English vocabulary in the last two centuries in order to illustrate the relevance of our general results.

II. DEFINITION

Consider the probability distribution $\mathbf{p} = (p_1, p_2, \dots, p_S)$ of a random variable over a discrete, countable set of symbols $i = 1, \dots, S$ (where later we include the possibility for $S \rightarrow \infty$). From an information theory standpoint, a natural measure to quantify the difference between two such probability distributions \mathbf{p} and \mathbf{q} is the Jensen-Shannon divergence (JSD) [20]

$$D(\mathbf{p}, \mathbf{q}) = H\left(\frac{\mathbf{p} + \mathbf{q}}{2}\right) - \frac{1}{2}H(\mathbf{p}) - \frac{1}{2}H(\mathbf{q}), \quad (2)$$

where H is the Shannon entropy [21],

$$H(\mathbf{p}) = -\sum_i p_i \log p_i. \quad (3)$$

This definition has several properties that are useful in the interpretation as a distance: (i) $D(\mathbf{p}, \mathbf{q}) \geq 0$, where the equality holds if and only if $\mathbf{p} = \mathbf{q}$, (ii) $D(\mathbf{p}, \mathbf{q}) = D(\mathbf{q}, \mathbf{p})$ (it is a symmetrized Kullback-Leiber divergence [20]), (iii) $\sqrt{D(\mathbf{p}, \mathbf{q})}$ fulfills the triangle inequality and thus is a metric [22], and (iv) $D(\mathbf{p}, \mathbf{q})$ equals the mutual information of variables sampled from \mathbf{p} and \mathbf{q} [3]; i.e., $D(\mathbf{p}, \mathbf{q})$ equals the average amount of information in one randomly sampled

word token about which of the two distributions it was sampled from [23]. The JSD is widely used in the statistical analysis of language [2], e.g., to automatically find individual documents that are (semantically) related [5,6] or to track the rate of evolution in the lexical inventory of a language over historical time scales [7,8].

Here, we also consider the generalization of JSD in which H in Eq. (3) is replaced by the generalized entropy of order α [24],

$$H_\alpha(\mathbf{p}) = \frac{1}{1-\alpha} \left(\sum_i p_i^\alpha - 1 \right), \quad (4)$$

yielding a spectrum of divergence measures D_α parametrized by α , first introduced in Ref. [25]. The usual JSD is retrieved for $\alpha = 1$. In (nonextensive) statistical mechanics, Eq. (4) was first proposed in Ref. [26] and D_α is sometimes called the Jensen-Tsallis divergence. While similar generalizations can be achieved with other formulations of generalized entropies such as the Renyi entropy [4,27], the corresponding divergences can become negative. In contrast, D_α is strictly non-negative, and it was recently shown that $\sqrt{D_\alpha(\mathbf{p}, \mathbf{q})}$ is a metric for any $\alpha \in (0, 2]$ [28]. For heavy-tailed distributions, Eq. (1), $H_\alpha < \infty$ for $\alpha > 1/\gamma$.

We define a normalized version of D_α as

$$\tilde{D}_\alpha(\mathbf{p}, \mathbf{q}) = \frac{D_\alpha(\mathbf{p}, \mathbf{q})}{D_\alpha^{\max}(\mathbf{p}, \mathbf{q})}, \quad (5)$$

where

$$D_\alpha^{\max}(\mathbf{p}, \mathbf{q}) = \frac{2^{1-\alpha} - 1}{2} \left(H_\alpha(\mathbf{p}) + H_\alpha(\mathbf{q}) + \frac{2}{1-\alpha} \right) \quad (6)$$

is the maximum possible D_α between \mathbf{p} and \mathbf{q} obtained by assuming that the set of symbols in each distribution (i.e., the support of \mathbf{p} and \mathbf{q}) are disjoint. The main motivation for using the measure (5) is that $\tilde{D}_\alpha(\mathbf{p}, \mathbf{q}) \in [0, 1]$, while the range of admissible values of D_α depends on α . This allows for a meaningful comparison of the divergences $\tilde{D}_\alpha(\mathbf{p}, \mathbf{q})$ and $\tilde{D}_{\alpha'}(\mathbf{p}, \mathbf{q})$ for $\alpha \neq \alpha'$ and therefore also for the full spectrum of α 's. In general, the metric properties of D_α are not preserved by \tilde{D}_α . An exception is the case in which the rank-frequency distribution $p(r)$ underlying all \mathbf{p} 's and \mathbf{q} 's is invariant (see Fig. 1). Noting that Eq. (6) is independent of the symbols, we obtain that $D_\alpha^{\max}(\mathbf{p}, \mathbf{q})$ is a constant for all \mathbf{p} 's and \mathbf{q} 's, and therefore, the metric property is preserved for \tilde{D}_α .

III. INTERPRETATION

In order to clarify the interpretation of D_α , it is useful to consider a toy model. As in Fig. 1, we consider two distributions \mathbf{p} and \mathbf{q} that have exactly the same rank-frequency distribution $p(r)$ but differ in (a subset of) the symbols they use. For simplicity, we consider that symbols

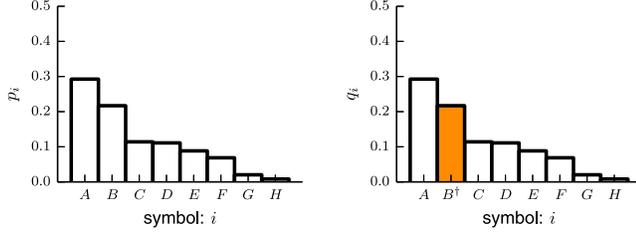


FIG. 2. Illustration of our toy model, where \mathbf{p} (left diagram) and \mathbf{q} (right diagram) have the same rank-frequency distribution but differ in the probability for individual symbols. In this example, \mathbf{p} and \mathbf{q} are the same ($p_i = q_i$) for $i \in \{A, C, D, E, F, G, H\}$, while the symbol $i = B$ in \mathbf{p} is replaced by $i = B^\dagger$ in \mathbf{q} with $p_{i=B} = q_{i=B^\dagger}$ and $p_{i=B^\dagger} = q_{i=B} = 0$.

that differ in the two cases appear only in one of the distributions. More precisely, denoting by $I_p = \{A, B, C, D, E, \dots\}$ the set of symbols in \mathbf{p} , we define the set of replaced symbols as $I^* \subset I_p$. The set of symbols in \mathbf{q} is chosen as $I_q = \{i | i \in I_p \setminus I^*\} \cup \{i^\dagger | i \in I^*\}$ with probabilities $p_i = q_i$ for $i \in I_p \setminus I^*$ and $p_i = q_{i^\dagger}$ for $i \in I^*$ (see Fig. 2 for one example).

For a given distribution \mathbf{p} and a set of replaced symbols I^* , we compute $D_\alpha(\mathbf{p}, I^*) \equiv D_\alpha(\mathbf{p}, \mathbf{q})$ as

$$D_\alpha(\mathbf{p}, I^*) = c_\alpha \sum_{i \in I^*} p_i^\alpha, \quad (7)$$

where $c_\alpha = (2^{(1-\alpha)} - 1)/(1 - \alpha)$. The maximum is given by

$$D_\alpha^{\max}(\mathbf{p}, I^*) = c_\alpha \sum_{i \in I_p} p_i^\alpha \quad (8)$$

such that

$$\tilde{D}_\alpha(\mathbf{p}, I^*) = \frac{\sum_{i \in I^*} p_i^\alpha}{\sum_{i \in I_p} p_i^\alpha}. \quad (9)$$

This shows that each symbol $i \in I^*$ that is replaced by a new symbol contributes p_i^α to D_α . It is thus clear that by varying α , the contribution of different frequencies becomes magnified (e.g., for $\alpha \gg 1$, large frequencies are enhanced, while for $\alpha < 0$ low frequencies contribute more to D_α than large frequencies).

In particular, for $\alpha = 0$, $\tilde{D}_{\alpha=0}(\mathbf{p}, I^*) = [(|I^*|)/(|I_p|)]$ is the fraction of symbols (types) that are different in \mathbf{p} and \mathbf{q} . Each symbol i counts the same, irrespective of their probabilities p_i . For $[(|I^*|)/(|I_p|)] \ll 1$, $\tilde{D}_{\alpha=0}(\mathbf{p}, I^*) = 1 - J(I_p, I_q)$, where $J(I_p, I_q) = [(|I_p \cap I_q|)/(|I_p \cup I_q|)]$ is the Jaccard coefficient between the two sets I_p and I_q , an *ad hoc* defined similarity measure widely used in information retrieval [2]. For $\alpha = 1$, $\tilde{D}_{\alpha=1}(\mathbf{p}, I^*) = \sum_{i \in I^*} p_i$, showing that each replaced symbol is weighted by its probability p_i and thus that $\tilde{D}_{\alpha=1}$ measures the distance in terms of tokens.

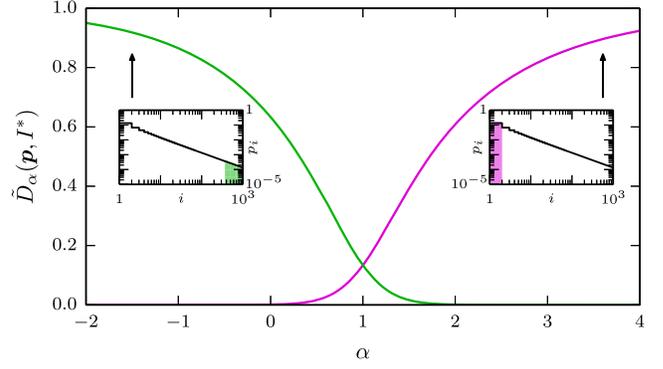


FIG. 3. The spectrum $\tilde{D}_\alpha(\mathbf{p}, I^*)$ for two different changes. The lines correspond to Eq. (9) with $p_i \propto i^{-1}$ with $i = 1, 2, \dots, 1000$ and two different sets of replaced symbols I_1^*, I_2^* . Right inset: $I_1^* = \{1\}$; i.e., only the symbol with the highest probability, $p_{i=1} \approx 0.13$, is changed. Left inset: $I_2^* = \{368, \dots, 1000\}$; i.e., the symbols with small probability are replaced. The choice of I_2^* was made such that $\sum_{i \in I_2^*} p_i \approx p_{i=1}$ and therefore $\tilde{D}_{\alpha=1}(\mathbf{p}, I_1^*) \approx \tilde{D}_{\alpha=1}(\mathbf{p}, I_2^*)$.

The full spectrum \tilde{D}_α offers information on changes in all frequencies, a point which is particularly important for the case of heavy-tailed distributions because word frequencies vary over many orders of magnitude. Figure 3 illustrates how different values of α are able to capture changes at different regions in the frequency spectrum. In particular, it shows that \tilde{D}_α grows (decays) with α when the modified symbols have high (low) frequency. Furthermore, the comparison between two given changes allows us to conclude about which change was more significant at different regions of the word-frequency spectrum. In the example of the figure, both changes (the two lines) are equally significant from the point of view of the modified tokens (\tilde{D}_1 are the same); the change in the left affects more types (\tilde{D}_0 is larger), and the change in the right affects more frequent words (\tilde{D}_α is larger for $\alpha \gg 1$).

IV. FINITE-SIZE ESTIMATION

In this section, we turn to the estimation of \tilde{D}_α from data. Even if \tilde{D}_α is defined with respect to distributions \mathbf{p} and \mathbf{q} , Eq. (5), in practice, these distributions are estimated from sequences with finite-size N (total number of symbols or word tokens) yielding finite-size estimates of the distributions $\hat{\mathbf{p}}$ and $\hat{\mathbf{q}}$. The main obstacle in obtaining accurate estimates of \tilde{D}_α is that it requires the estimation of entropies for which, in general, unbiased estimators do not exist [29]. Accordingly, even if $\mathbf{p} = \mathbf{q}$, in practice, $H_\alpha(\hat{\mathbf{p}}) \neq H_\alpha(\hat{\mathbf{q}})$ and $\tilde{D}_\alpha(\hat{\mathbf{p}}, \hat{\mathbf{q}}) > 0$ are measured not only in single realizations but also on average (the bias). Besides the bias, we are also interested in the expected fluctuation (standard deviation) of the estimations of H_α and \tilde{D}_α and how both depend

on the sequence size N for large N . In heavy-tailed distributions such as Eq. (1), these estimations are based on an observed vocabulary V (number of different symbols) that grows sublinearly with N as [30–32]

$$V(N) \propto N^{1/\gamma}. \quad (10)$$

This implies that the entropies in Eq. (4) are estimated based on a sum of $V \rightarrow \infty$ terms (for $N \rightarrow \infty$). In practice, γ and the precise functional form of the heavy-tailed distribution are unknown, and therefore, besides \tilde{D}_α , the estimation of H_α is also of interest (see Ref. [33] for the case in which a power-law form of \mathbf{p} is assumed to be known *a priori*).

A. Analytical calculations

Here, we extend previous results [34–37] and generalize them to arbitrary α . Given a probability distribution \mathbf{p} and the measured probabilities $\hat{\mathbf{p}}$ from a finite sample of N word tokens, we expand $H_\alpha(\hat{\mathbf{p}})$ around the true probabilities p_i up to second order as

$$H_\alpha(\hat{\mathbf{p}}) \approx H_\alpha(\mathbf{p}) + \sum_{i: \hat{p}_i > 0} (\hat{p}_i - p_i) \frac{\alpha}{1-\alpha} p_i^{\alpha-1} - \frac{1}{2} \sum_{i: \hat{p}_i > 0} (\hat{p}_i - p_i)^2 \alpha p_i^{\alpha-2}, \quad (11)$$

where we use $(\partial H_\alpha / \partial p_i) = \alpha / (1-\alpha) p_i^{\alpha-1}$ and $(\partial^2 H / \partial p_i \partial p_j) = -\alpha p_i^{\alpha-2} \delta_{i,j}$. We then calculate $\mathbb{E}[H_\alpha(\hat{\mathbf{p}})]$ by averaging over the different realizations of the random variables \hat{p}_i by assuming that the absolute frequency of each symbol i is drawn from an independent binomial with probability p_i such that $\mathbb{E}[\hat{p}_i] = p_i$ and $\mathbb{V}[\hat{p}_i] = p_i(1-p_i)/N \approx p_i/N$, yielding

$$\mathbb{E}[H_\alpha(\hat{\mathbf{p}})] \approx H_\alpha(\mathbf{p}) - \frac{\alpha}{2N} \sum_{i \in V} p_i^{\alpha-1} = H_\alpha(\mathbf{p}) - \frac{\alpha V^{(\alpha)}}{2N}, \quad (12)$$

which defines the vocabulary size of order α ,

$$V^{(\alpha)} \equiv \sum_{i \in V} p_i^{\alpha-1}. \quad (13)$$

From Eq. (12), we see that the bias in the entropy estimation $|H_\alpha(\mathbf{p}) - \mathbb{E}[H_\alpha(\hat{\mathbf{p}})]|$ depends only on $V^{(\alpha)}$ and N . Similar calculations (see Appendix B) show that the large- N behavior of the bias and the fluctuations (variance) of H_α , D_α , and \tilde{D}_α can be written as simple functions of $V^{(\alpha)}$ and N , as summarized in Table I.

We now focus on the dependence of $V^{(\alpha)}$ on N . The sum $\sum_{i \in V}$ in Eq. (13) indicates that in N samples, on average, $V = V(N) \equiv V^{(\alpha=1)}$ different symbols are observed. If for $N \rightarrow \infty$ the vocabulary V converges to a finite value, $V^{(\alpha)}$ in Eq. (13) also converges and the bias scales as $1/N$. A more interesting scenario happens when V grows with N . For the heavy-tailed case of interest here, V grows as $N^{1/\gamma}$,

TABLE I. Scaling of the bias $|\mathbb{E}[\hat{X}] - X|$ and the fluctuations $\mathbb{V}[X] \equiv \mathbb{E}[\hat{X}^2] - \mathbb{E}[\hat{X}]^2$ of estimations \hat{X} . The results are valid for large sequence sizes N and depend on the vocabulary of order α , $V^{(\alpha)}$ as in Eqs. (13) and (14). Results are shown for $X = H_\alpha$ [order α entropy, Eq. (4)], D_α [generalized divergence], and \tilde{D}_α [normalized divergence, Eq. (5)] (see Appendix B for the derivations). For \tilde{D}_α , we approximate $\tilde{D}_\alpha \approx D_\alpha / \mathbb{E}[D_\alpha^{\max}]$.

	H_α	$D_\alpha, \tilde{D}_\alpha(\mathbf{p} \neq \mathbf{q})$	$D_\alpha, \tilde{D}_\alpha(\mathbf{p} = \mathbf{q})$
Bias:	$V^{(\alpha)}/N$	$V^{(\alpha)}/N$	$V^{(\alpha)}/N$
Fluctuations:	$V^{(2\alpha)}/N$	$V^{(2\alpha)}/N$	$V^{(2\alpha-1)}/N^2$

Eq. (10), and we obtain (in Appendix C) that $V^{(\alpha)}$ scales for large N as

$$V^{(\alpha)} \propto \begin{cases} N^{-\alpha+1+1/\gamma} & \alpha < 1 + 1/\gamma \\ \text{constant} & \alpha > 1 + 1/\gamma, \end{cases} \quad (14)$$

where $\gamma > 1$ is the Zipf exponent defined in Eq. (1) and α is the order of the entropy in Eq. (4).

From the combination of Eq. (14) and Table I, we obtain the scalings with sequence size N of the estimators of H_α , D_α , and \tilde{D}_α in a heavy-tailed distribution with exponent γ . These scalings are summarized in Table II. Three scaling regimes can be identified for the bias and for the fluctuations. (i) For large α , the decay is $1/N$ (except when $\mathbf{p} = \mathbf{q}$, where the fluctuations decay even faster as $1/N^2$) as in the case of a finite vocabulary and short-tailed distributions. (ii) For intermediate α , a sublinear decay with N is observed. This regime appears exclusively in heavy-tailed distributions and has important consequences in real applications, as shown below. From the exponents of the sublinear decay, we see that the bias decays more slowly than the fluctuations. (iii) For small α , $\alpha < 1/\gamma$, $H_\alpha(\mathbf{p})$ is not defined; thus, the estimator for the mean of H_α and D_α diverges. The growth of H_α (and therefore D_α^{\max}) and D_α with N has the same scaling, and therefore, they cancel each other for \tilde{D}_α , in which case a convergence to a well-defined value is found (the fluctuation of \tilde{D}_α still decays in this regime).

B. Numerical simulations

Here, we perform numerical estimations of the normalized divergence spectrum \tilde{D}_α that illustrate the regimes derived above, we confirm the validity of the approximations used in their derivations, and we show that the same scalings are observed for different entropy estimators. We sample N symbols (tokens) twice from the same distribution ($\mathbf{p} = \mathbf{q}$), and therefore, $\tilde{D}_\alpha = 0$ and the expected value $\mathbb{E}[\tilde{D}_\alpha(\hat{\mathbf{p}}, \hat{\mathbf{q}})]$ is the bias. (The fact that the bias shows a slower decay with N than the fluctuations makes these two effects distinguishable also in this $\tilde{D}_\alpha = 0$ case because $\mathbb{E}[\tilde{D}_\alpha(\hat{\mathbf{p}}, \hat{\mathbf{q}})] \gg \mathbb{V}[\tilde{D}_\alpha(\hat{\mathbf{p}}, \hat{\mathbf{q}})]$ for large N).

TABLE II. Summary of finite-size scaling for distributions with heavy tails. Mean (\mathbb{E}) and variance (\mathbb{V}) of the plug-in estimator of H_α , D_α , and \tilde{D}_α for samples $\hat{\mathbf{p}}$ and $\hat{\mathbf{q}}$ each of size N drawn randomly from \mathbf{p} and \mathbf{q} with power-law rank-frequency distributions with exponent $\gamma > 1$, Eq. (1). The results are obtained by combining Table I with Eq. (14) (for details see Appendixes B and C). The constant c depends on α and has a different value in each case but is independent of N . The limit $\gamma \rightarrow \infty$ corresponds to the case in which both \mathbf{p} and \mathbf{q} have short tails.

	$\mathbb{E}[H_\alpha(\hat{\mathbf{p}})]$	$\mathbb{E}[D_\alpha(\hat{\mathbf{p}}, \hat{\mathbf{q}})]$		$\mathbb{E}[\tilde{D}_\alpha(\hat{\mathbf{p}}, \hat{\mathbf{q}})]$	
$\alpha_1^{\mathbb{E}}$	$1/\gamma$	$1/\gamma$		$1/\gamma$	
$\alpha_2^{\mathbb{E}}$	$1 + 1/\gamma$	$1 + 1/\gamma$		$1 + 1/\gamma$	
$\alpha < \alpha_1^{\mathbb{E}}$	$cN^{-\alpha+1/\gamma}$	$cN^{-\alpha+1/\gamma}$		c	
$\alpha_1^{\mathbb{E}} < \alpha < \alpha_2^{\mathbb{E}}$	$H_\alpha(\mathbf{p}) + cN^{-\alpha+1/\gamma}$	$D_\alpha(\mathbf{p}, \mathbf{q}) + cN^{-\alpha+1/\gamma}$		$\tilde{D}_\alpha(\mathbf{p}, \mathbf{q}) + cN^{-\alpha+1/\gamma}$	
$\alpha > \alpha_2^{\mathbb{E}}$	$H_\alpha(\mathbf{p}) + cN^{-1}$	$D_\alpha(\mathbf{p}, \mathbf{q}) + cN^{-1}$		$\tilde{D}_\alpha(\mathbf{p}, \mathbf{q}) + cN^{-1}$	
	$\mathbb{V}[H_\alpha(\hat{\mathbf{p}})]$	$\mathbb{V}[D_\alpha(\hat{\mathbf{p}}, \hat{\mathbf{q}})]$		$\mathbb{V}[\tilde{D}_\alpha(\hat{\mathbf{p}}, \hat{\mathbf{q}})]$	
		$\mathbf{p} \neq \mathbf{q}$	$\mathbf{p} = \mathbf{q}$	$\mathbf{p} \neq \mathbf{q}$	$\mathbf{p} = \mathbf{q}$
$\alpha_1^{\mathbb{V}}$	$1/(2\gamma)$	$1/(2\gamma)$	$1/(2\gamma)$	$1/\gamma$	$1/\gamma$
$\alpha_2^{\mathbb{V}}$	$\frac{1}{2}(1 + 1/\gamma)$	$\frac{1}{2}(1 + 1/\gamma)$	$1 + 1/(2\gamma)$	$\frac{1}{2}(1 + 1/\gamma)$	$1 + 1/(2\gamma)$
$\alpha < \alpha_1^{\mathbb{V}}$	$cN^{-2\alpha+1/\gamma}$	$cN^{-2\alpha+1/\gamma}$	$cN^{-2\alpha+1/\gamma}$	$cN^{-1/\gamma}$	$cN^{-1/\gamma}$
$\alpha_1^{\mathbb{V}} < \alpha < \alpha_2^{\mathbb{V}}$	$cN^{-2\alpha+1/\gamma}$	$cN^{-2\alpha+1/\gamma}$	$cN^{-2\alpha+1/\gamma}$	$cN^{-2\alpha+1/\gamma}$	$cN^{-2\alpha+1/\gamma}$
$\alpha > \alpha_2^{\mathbb{V}}$	cN^{-1}	cN^{-1}	cN^{-2}	cN^{-1}	cN^{-2}

We start with the most important prediction of our analytical calculations above: the existence in heavy-tailed distributions of a regime for which the bias and fluctuations of \tilde{D}_α decay with N more slowly than $1/N$. This holds already for $\alpha = 1$, i.e., for the usual Jensen-Shannon divergence, previously shown for the bias of $H_{\alpha=1}$ in Ref. [37]. One potential limitation of our analytical calculations is that they are based on the plug-in estimator obtained from replacing the p_i 's in the generalized entropies, Eq. (4), by the measured frequencies (i.e., $p_i \mapsto \hat{p}_i = N_i/N$, with N_i being the number of observed word tokens of type i). To test the generality of our results, in the numerical simulations, we use four different estimators of the Shannon entropy (i.e., $\alpha = 1$): (i) the plug-in estimator, (ii) Miller's estimator [34], which takes into account the approximation obtained from the expansion in Eq. (12), (iii) Grassberger's estimator [38], and (iv) a recently proposed Bayesian estimator described in Ref. [39], which is an extension of the approach by Nemenman *et al.* [40] to the case where the number of possible symbols is unknown or even countably infinite [41]. The numerical results in Fig. 4 show that the different estimators are indeed able to reduce the bias of the estimation but that the scaling of the bias with N remains the same. In particular, the transition from the short-tailed to the heavy-tailed distribution leads to the predicted transition from N^{-1} (N^{-2}) to the slower $N^{-1+1/\gamma}$ ($N^{-2+1/\gamma}$) decay for the bias (fluctuations) for all estimators. The only exception is in the bias of the Bayesian estimator for the exact Zipf's law (1), but since this estimator shows a bad performance for the fluctuation and for the real data, we

conclude that the slower scaling should be expected, in general, also for this elaborated estimator. These results confirm the generality of our finding that the bias and fluctuation in $\tilde{D}_{\alpha=1}$ decay more slowly than $1/N$ in heavy-tailed distributions. The consequence of this result to applications will be discussed in the next section.

We now consider the estimation of \tilde{D}_α for $\alpha \neq 1$ in the case of heavy-tailed distributions (1). The numerical results in Fig. 5 confirm the existence of the three scaling regimes discussed after Eq. (14) and in Table II. Panels (b) and (d) show the relative reduction in the bias and fluctuations achieved when the sequence size is doubled. For many different α 's, the relative reduction is larger than 0.5 (0.25) for the bias (fluctuations), a consequence of the slow decay with N that shows the difficulty in obtaining a good estimation of \tilde{D}_α . In practice, the exponent γ of the distribution is unknown, such that the critical values of α that separate these regimes (e.g., $\alpha_1^{\mathbb{E}} = 1/\gamma$ and $\alpha_2^{\mathbb{E}} = 1 + 1/\gamma$ for the bias) cannot be determined *a priori*. Yet, since $\gamma > 1$, we know that (i) $\alpha_1^{\mathbb{E}}, \alpha_1^{\mathbb{V}} < 1$, and therefore D_α for $\alpha \geq 1$ is such that $D_\alpha(\mathbf{p}, \mathbf{p}) = 0$ for $N \rightarrow \infty$; (ii) $\alpha_2^{\mathbb{E}}, \alpha_2^{\mathbb{V}} < 2$, and therefore the bias and fluctuations of D_α for $\alpha \geq 2$ decay as $1/N$ (or $1/N^2$ for the fluctuations in the case of $\mathbf{p} = \mathbf{q}$). This suggests $D_{\alpha=2}$ as a pragmatic choice for empirical measurements because any further increase in α will not lead to a faster convergence.

V. APPLICATION TO REAL DATA

In this section, we show the significance of the general results of the previous section to specific problems.

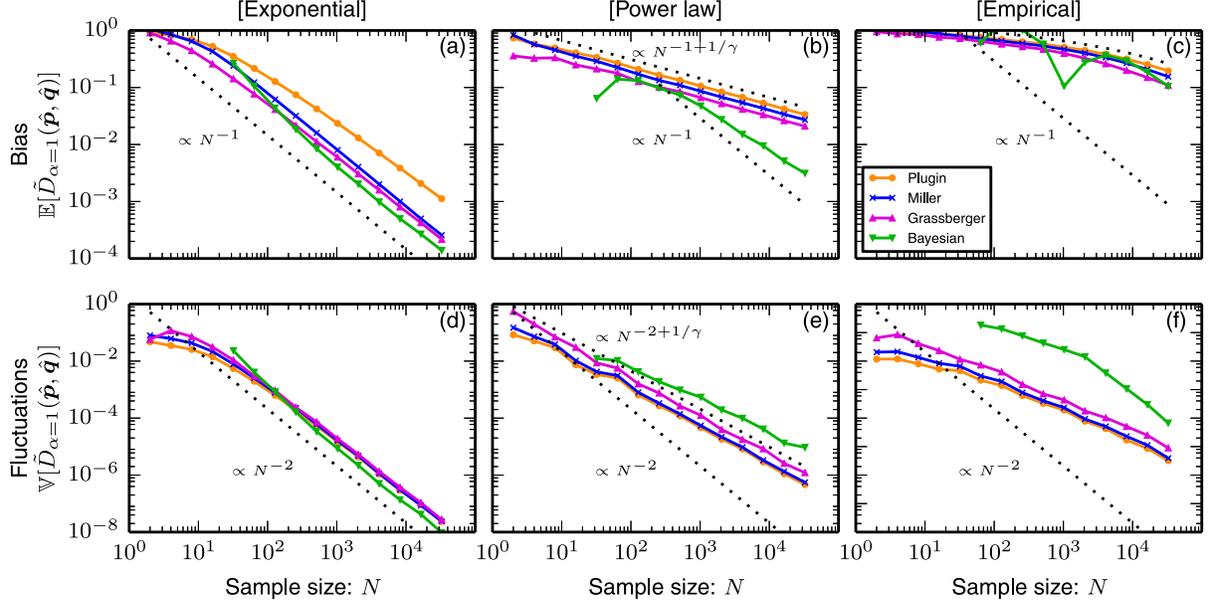


FIG. 4. Finite-size estimation of the normalized Jensen-Shannon divergence $\tilde{D} = \tilde{D}_{\alpha=1}$. (a)–(c) Estimation of $\mathbb{E}[\tilde{D}(\hat{\mathbf{p}}, \hat{\mathbf{q}})]$ between two sequences of size N drawn from the same distribution [i.e., $D(\mathbf{p}, \mathbf{q}) = 0$] using four different estimators of the entropy (see text) for three representative distributions: (a) Exponential (short-tailed) distribution $p_i \propto e^{-ai}$ for $i = 0, 1, \dots$ with $a = 0.1$, (b) power-law (heavy-tailed) distribution $p_i \propto i^{-\gamma}$ for $i = 1, 2, \dots$ with $\gamma = 3/2$, and (c) empirical Zipf distribution of word frequencies, i.e., rank-frequency distribution $p(r)$ from the complete Google-ngram data, $p_i = f(i = r)$ for $i = 1, \dots, 4623568$, which is well described by a double power law [11]. (d)–(f) The same as (a)–(c) for the fluctuations $\mathbb{V}[\tilde{D}(\hat{\mathbf{p}}, \hat{\mathbf{q}})]$. The dotted lines show the expected scalings from Table II for short-tailed distributions, i.e., N^{-1} (N^{-2}), and power-law distributions, i.e., $N^{-1+1/\gamma}$ ($N^{-2+1/\gamma}$), for the bias (fluctuations). In (c), we show the expected scaling for the bias, $V_{\text{emp}}(N)/N$, where $V_{\text{emp}}(N)$ is the expected number of different symbols in a random sample of size N from the empirical distribution [32]. Averages are taken over 1000 realizations.

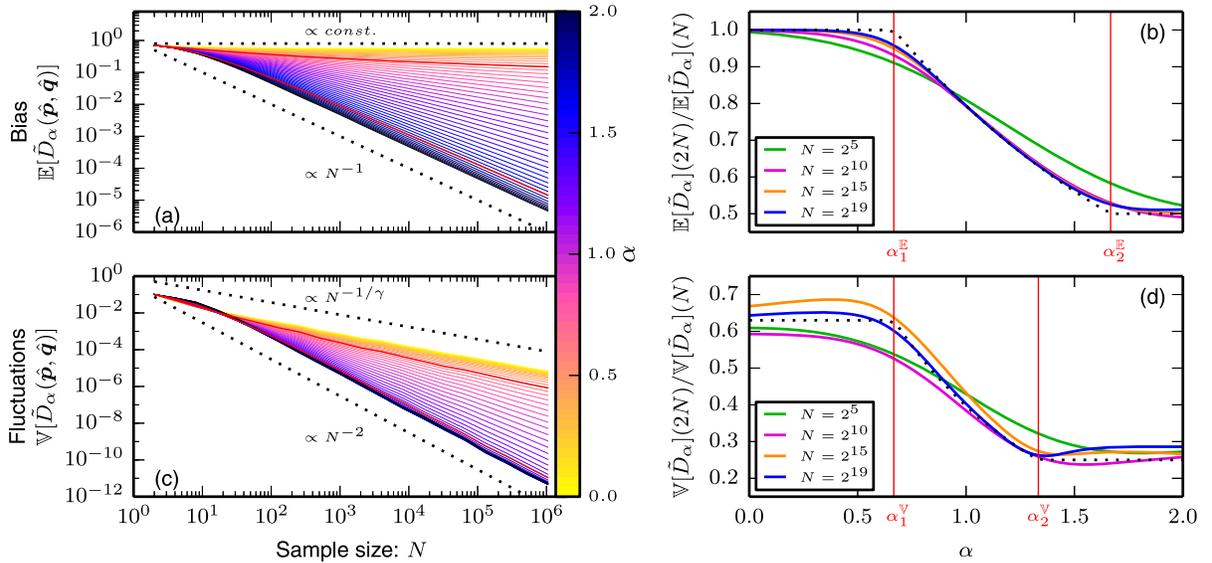


FIG. 5. Bias (a,b) and fluctuations (c,d) in the finite-size estimation of \tilde{D}_{α} . Estimation of $\mathbb{E}[\tilde{D}_{\alpha}(\hat{\mathbf{p}}, \hat{\mathbf{q}})]$ between two sequences, each of size N , drawn numerically from the same power-law distribution $p_i \propto i^{-\gamma}$ for $i = 1, 2, \dots, V \rightarrow \infty$ with $\gamma = 3/2$ using the plug-in estimator ($p_i \mapsto \hat{p}_i$) for the entropies of order α . (a) Scaling of the bias with N for different α . (b) Decrease of the bias in \tilde{D}_{α} when the sample size is doubled ($N \mapsto 2N$) for different values of N as a function of α . (c,d) The same as (a) and (b) for the fluctuations $\mathbb{V}[\tilde{D}_{\alpha}(\hat{\mathbf{p}}, \hat{\mathbf{q}})]$. Red lines in all plots indicate the borders between the regimes, $\alpha_1^E = 1/\gamma = 2/3$, $\alpha_2^E = 1 + 1/\gamma = 5/3$ [for the bias in (a,b)] and $\alpha_1^V = 1/\gamma = 2/3$, $\alpha_2^V = 1 + 1/(2\gamma) = 4/3$ [for the fluctuations in (c,d)]. Dotted lines indicate the predictions based on Table I for $\alpha < \alpha_1^E$, α_1^V and $\alpha > \alpha_2^E$, α_2^V [in (a,c)] and all values of α [in (b,d)]. Averages are taken over 1000 realizations.

A problem that appears in different contexts is to test whether two finite-size N sequences, described by their empirical distributions \hat{p} and \hat{q} , have a common source (null hypothesis). This involves the computation of a single divergence $\tilde{D}_\alpha(\mathbf{p}, \mathbf{q})$, which is then compared to the divergence $\tilde{D}_\alpha(\mathbf{p}', \mathbf{p}')$ between two finite-size (random) samplings of a single (properly chosen) distribution \mathbf{p}' (e.g., $\mathbf{p}' = 0.5\mathbf{p} + 0.5\mathbf{q}$). The probability of observing $\tilde{D}_\alpha(\mathbf{p}', \mathbf{p}') \geq \tilde{D}_\alpha(\mathbf{p}, \mathbf{q})$ is then reported as a p value [3]. Besides applications in language, e.g., comparing the distribution of word frequencies, this problem appears in the identification of coding and noncoding regions in DNA [42]. The significance of our results for finite-size estimations in Sec. IV is that for the case of heavy-tailed distribution, the expected $\tilde{D}_\alpha(\mathbf{p}, \mathbf{q})$ of the null model may be much larger than the predicted value based on a $1/N$ decay (as observed in short-tailed distributions). If the slower convergence in N is ignored, e.g., by applying standard tests [3] to heavy-tailed distributions, one rejects the null hypothesis (low p value) even if the data are drawn from the same source because the measured distance will be much larger. The example in Fig. 4(c) shows that, even when the size of both sequences is on the order of $N \approx 10^5$, the expected \tilde{D}_1 (JSD) is $\mathbb{E}[\tilde{D}_{\alpha=1}(\hat{p}, \hat{q})] \approx 10^{-1}$. This is 2 orders of magnitude larger than for the exponential distribution in Fig. 4(a), where $\mathbb{E}[\tilde{D}_{\alpha=1}(\hat{p}, \hat{q})] \approx 10^{-3}$.

The next problems we consider appear in the analysis of historical data and in the quantification of language change [43]. These problems are representative of problems that involve the comparison of two or more divergences $\tilde{D}_\alpha(\mathbf{p}, \mathbf{q})$, obtained from different distributions $\mathbf{p} \neq \mathbf{q}$ and

$\alpha' \neq \alpha$. As depicted in Fig. 1, the different distributions are obtained based on individual years ($t \in \{1850, 1900, 1950\}$), and we calculate the normalized spectrum $\tilde{D}_\alpha(\mathbf{p}_{t_1}, \mathbf{p}_{t_2})$ between pairs of years (t_1, t_2). As argued in Sec. II and Appendix A, $\tilde{D}_\alpha(\mathbf{p}, \mathbf{q})$ is meaningful even if the sequences used to estimate \mathbf{p} and \mathbf{q} have different sizes $N_p \neq N_q$. We summarize our results in Fig. 6, from which different conclusions can be drawn:

- Temporal change.* The change of English from 1850 to 1950 was larger than the change from 1850 to 1900 and from 1900 to 1950, as seen from the fact that the curve of $\tilde{D}_\alpha(\mathbf{p}_{1850}, \mathbf{p}_{1950})$ in Fig. 6(a) lies above the two other curves for all α . This intuitive result (evolutionary dynamics show no recurrences) confirms that the divergence spectrum $\tilde{D}_\alpha(\mathbf{p}_{t_1}, \mathbf{p}_{t_2})$ is a meaningful quantification of language change. The average dependency of $\tilde{D}_\alpha(\mathbf{p}_{1850}, \mathbf{p}_{1950})$ on $\Delta t = |t_2 - t_1|$, shown in Fig. 6(b), can thus be used as a quantification of the speed of language change. We observe an approximate relationship $\tilde{D}_\alpha(\Delta t) \approx \tilde{D}_\alpha^{(i)} + \tilde{D}_\alpha^{(ii)} \Delta t^2$ for $\Delta t \gg 1$ [see inset of Fig. 6(b)], where $\tilde{D}_\alpha^{(i)}$ and $\tilde{D}_\alpha^{(ii)}$ are constants and can be related to words that change because of fluctuations (finite sampling or topical dependencies), which are independent of Δt , and words that show a systematic increase or decrease over Δt , respectively (see Appendix D for a detailed discussion).
- Dependence on α .* All observed divergences $\tilde{D}_\alpha(\mathbf{p}_{t_1}, \mathbf{p}_{t_2})$ decay with α [e.g., the three curves in Fig. 6(a)]. As discussed in Sec. III, this shows that for

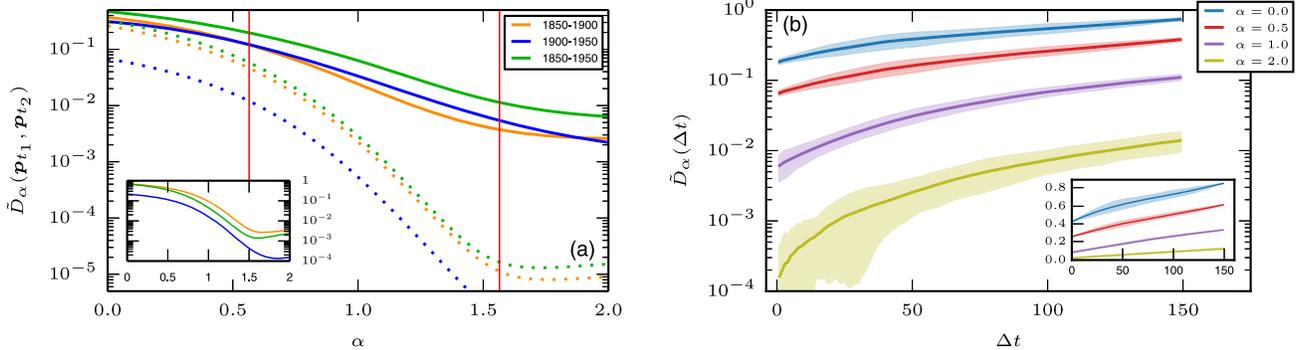


FIG. 6. Measuring change in the usage of language on historical time scales. (a) $\tilde{D}_\alpha(\mathbf{p}_{t_1}, \mathbf{p}_{t_2})$ as a function of α for pairs of word-frequency distributions of the Google-ngram database obtained from the yearly corpora t_1 and t_2 with $(t_1, t_2) \in \{(1850, 1900), (1900, 1950), (1850, 1950)\}$ (solid lines). The dotted lines with the same colors show the results of a null model in which samples of the same size as the ones in t_1 and t_2 are randomly drawn from the *same* distribution (obtained from combining the corpora in t_1 and t_2) mimicking a minimum distance that can be observed because of finite-size effects. The vertical lines show the three regimes $\alpha < 1/\gamma$, $1/\gamma < \alpha < 1 + 1/\gamma$, and $\alpha > 1 + 1/\gamma$ in the convergence of $\tilde{D}_\alpha(\mathbf{p}_{t_1}, \mathbf{p}_{t_2})$ with N (see Sec. IV), obtained using $\gamma = 1.77$ [11]. Inset: Ratio $\tilde{D}_\alpha(\mathbf{p}_{t_1}, \mathbf{p}_{t_2})/\tilde{D}_\alpha(\mathbf{p}_{t_1}, \mathbf{p}_{t_2})$. (b) Average divergence as a function of $\Delta t \equiv |t_2 - t_1|$, calculated as $\tilde{D}_\alpha(\Delta t) = (1/N_{\Delta t}) \sum_{t_1=1805}^{2000-\Delta t} \tilde{D}_\alpha(\mathbf{p}_{t_1}, \mathbf{p}_{t_1+\Delta t})$ for four different α (solid lines). Shaded areas represent the standard deviation associated with the average $\tilde{D}_\alpha(\Delta t)$. Inset: $\sqrt{\tilde{D}_\alpha(\Delta t)}$ as a function of Δt , highlighting the approximate relationship $\tilde{D}_\alpha(\Delta t) \sim \Delta t^2$ for $\Delta t \gg 1$.

words with a high (low) frequency, the distributions are more (less) similar and thus the change is slower (faster). This result is consistent with previous works on the evolution of individual words on historical time scales reporting that frequent words tend to be more stable [44,45]. This dependence on α is essential when comparing the change from 1850 to 1900 to the change from 1900 to 1950 [Fig. 6(a)]. While the earlier change was smaller if counted on a token basis, $\tilde{D}_{\alpha=1}(\mathbf{p}_{1850}, \mathbf{p}_{1900}) < \tilde{D}_{\alpha=1}(\mathbf{p}_{1900}, \mathbf{p}_{1950})$, it becomes larger if one focuses on the more frequent words [$\tilde{D}_{\alpha=2}(\mathbf{p}_{1850}, \mathbf{p}_{1900}) > \tilde{D}_{\alpha=2}(\mathbf{p}_{1900}, \mathbf{p}_{1950})$].

- (c) *Role of finite-size scalings.* Our finding that the scalings (of the bias and of the fluctuations) in \tilde{D}_α with sample size N depend on α allows for a deeper understanding of the $\tilde{D}_\alpha(\mathbf{p}_{t_1}, \mathbf{p}_{t_2})$ measurements discussed above. The expected \tilde{D}_α 's for random sampling of the same distribution [null model shown as dashed line in Fig. 6(a)] are of the same order as the empirical distance for small α [i.e., $\tilde{D}_\alpha(\mathbf{p}_{t_1}, \mathbf{p}_{t_2}) \approx \tilde{D}_\alpha(\mathbf{p}_{t_1}, \mathbf{p}_{t_2})$], and it is only for $\alpha > 1$ that the null model divergence becomes negligible compared to the empirical divergence [i.e., $\tilde{D}_\alpha(\mathbf{p}_{t_1}, \mathbf{p}_{t_2}) \ll \tilde{D}_\alpha(\mathbf{p}_{t_1}, \mathbf{p}_{t_2})$]. This implies that even though the size of the individual corpora is of the order of $N \approx 10^9$ word tokens, the empirically measured \tilde{D}_α is still strongly influenced by finite-size effects over a wide range of values for α , in agreement with our analysis in Sec. IV. In particular, the bias for the Jensen-Shannon divergence ($\alpha = 1$) is important even for the case of the (extremely large) Google-ngram database [e.g., the inset of Fig. 6(a) shows that the bias is $\approx 10\%$].
- (d) $\alpha = 2$ as a pragmatic choice. The slow decay of bias and fluctuations with database size suggests that $\tilde{D}_{\alpha=2}$ is a pragmatic choice in reducing such finite-size effects when the exponent γ in the power-law distribution is not known. This conclusion is further corroborated in the analysis of the dependence of \tilde{D}_α with Δt [Fig. 6(b)]. While $\tilde{D}_\alpha(\Delta t = 0) = 0$ by construction, \tilde{D}_α does not converge to zero for $\Delta t \rightarrow 0$ when extrapolating from $\tilde{D}_\alpha(\Delta t > 0)$, but instead it seems to saturate, i.e., $\tilde{D}_\alpha(\Delta t \rightarrow 0) \approx \tilde{D}_\alpha^{(i)} > 0$. For small values of α , $\tilde{D}_\alpha^{(i)}$ is of the same order of magnitude of the expected bias [e.g., shown as a dashed line in Fig. 6(a)] and even of the same order of magnitude of the divergence $\tilde{D}_\alpha(\Delta t = 100)$ between two corpora separated by 100 years. For small α and Δt , it is thus difficult to distinguish between finite-size effects ($\tilde{D}_\alpha^{(i)}$) and actual language change. Results for $\alpha = 2$ show the largest relative variation with Δt and

are therefore statistically more suited to quantify language change over time.

VI. CONCLUSIONS

In summary, we investigated the use of generalized entropies H_α to quantify the difference between symbolic sequences with heavy-tailed frequency distributions. In particular, we introduced a normalized spectrum of a generalized divergence, $\tilde{D}_\alpha(\mathbf{p}, \mathbf{q})$ in Eq. (5), that allows for a comparison between the different distributions \mathbf{p} and \mathbf{q} and also for different α 's. Increasing α , \tilde{D}_α attributes higher weights to high-frequency symbols. The more complete characterization given by the full spectrum \tilde{D}_α is particularly important in the case of heavy-tailed distributions because in this case, symbols do not have a characteristic frequency but instead show frequencies on a broad range of values.

Our main analytical finding is how the systematic (bias) and statistical (fluctuations) errors of finite-size (N) estimations of H_α and \tilde{D}_α scales with N (see Table II). The existence of regimes in which these scalings decay slower than $1/N$ shows that large uncertainties should be expected in H_α and \tilde{D}_α , estimated even for very large databases. This should be taken into account when comparing two or more \tilde{D}_α 's and when estimating the probability of two sequences having the same source. The fact that for large α we recover the usual scaling (decay with $1/N$) suggests $\tilde{D}_{\alpha=2}$ as a pragmatic choice in applications involving heavy-tailed distributions. Previous works using information theoretic measures in language used $\alpha = 1$ [5–8] and did not take into account the effect of (finite) database size. Our results show that the bias and fluctuations are significant even in the extremely large Google-ngram database. It is therefore essential to clarify what the role is of finite-size effects in the reported conclusions, in particular, in the (typical) case that database sizes change over time.

Our main empirical findings on language change are as follows: (i) least-frequent words contribute more to the total vocabulary change, (ii) the answer to the question of whether the speed of language change is accelerating depends on the emphasis that is given to either low-frequency or high-frequency words, and (iii) the quantification of the speed of vocabulary change in time, Δt , which roughly shows a dependence $\tilde{D}_\alpha(\Delta t) \approx \tilde{D}_\alpha^{(i)} + \tilde{D}_\alpha^{(ii)} \Delta t^2$, where $\tilde{D}_\alpha^{(i)}$ ($\tilde{D}_\alpha^{(ii)}$) quantifies the degree to which words change due to fluctuations independent of time (systematic increase or decrease of the frequency over time). More generally, our spectrum \tilde{D}_α opens the possibility of studying language change at different resolution, combining aspects from the analysis on the level of individual words (e.g., Refs. [44,45]) and the full vocabulary of a language (e.g., Refs. [7,9]).

Our results are also of interest beyond the cases treated here. First, the finite-size scaling we derive appears already in the entropy, and therefore, the same scalings are expected in any entropy-based measure, including those based on conditional entropies such as the Kullback-Leibler divergence [1]. Second, the analysis is not necessarily restricted to the word level; it can also be straightforwardly extended to n-grams of words that also show heavy-tailed distributions [46]. Third, the spectrum of divergences $\tilde{D}_\alpha(\mathbf{p}, \mathbf{q})$ offers a unifying framework that can be applied to problems involving different partitions of texts by varying the parameter α . For example, while in document classification [2] one tries to identify topical words (suggesting the use of low values of α), in applications of authorship attribution [47], it has been shown that the comparison of the most-frequent (function) words yields the best results (suggesting the use of large values of α). Fourth, heavy-tailed distributions appear in different problems involving symbolic sequences (e.g., in the DNA [17], in gene expression [18], and in music [19]), and the significance of our results is that they can be applied in all these cases.

ACKNOWLEDGMENTS

We thank Peter Grassberger for insightful discussions.

APPENDIX A: DOCUMENTS WITH DIFFERENT LENGTHS

Here, we discuss how to proceed if the JSD is computed from finite datasets with different finite lengths N , i.e., when \mathbf{p} (\mathbf{q}) is estimated from a sequence of length N_p ($N_q \neq N_p$).

1. Different weights

A possible way to extend Eq. (2), taking into account the unequal contribution $N_p \neq N_q$, is to consider weights π as [3]

$$D_\alpha^\pi(\mathbf{p}, \mathbf{q}) = H_\alpha(\pi_p \mathbf{p} + \pi_q \mathbf{q}) - \pi_p H_\alpha(\mathbf{p}) - \pi_q H_\alpha(\mathbf{q}), \quad (\text{A1})$$

with $\pi_p = N_p/N$ and $\pi_q = N_q/N$ such that $\pi_p + \pi_q = 1$ with $N = N_p + N_q$ (denoted as natural weights in the following). Obviously, if $N_p = N_q$ then $\pi_p = \pi_q = 1/2$ and D_α is recovered. The normalized distance (5) becomes

$$\tilde{D}_\alpha^\pi(\mathbf{p}, \mathbf{q}) = \frac{D_\alpha^\pi(\mathbf{p}, \mathbf{q})}{D_\alpha^{\pi, \max}(\mathbf{p}, \mathbf{q})}, \quad (\text{A2})$$

where

$$D_\alpha^{\pi, \max}(\mathbf{p}, \mathbf{q}) = (\pi_p^\alpha - \pi_p)H_\alpha(\mathbf{p}) + (\pi_q^\alpha - \pi_q)H_\alpha(\mathbf{q}) + \frac{1}{1-\alpha}(\pi_p^\alpha + \pi_q^\alpha - 1). \quad (\text{A3})$$

Our main results for the finite-size scaling of D_α summarized in Table II remain valid for the weighted divergences.

The approach above follows Ref. [3], which introduced weights to the usual JSD (non-normalized, $\alpha = 1$) and showed that the natural weights $\pi_p = N_p/N$ and $\pi_q = N_q/N$ imply certain useful properties for the JSD, e.g., that the bias does not depend on the relative size of the two samples. While their main motivation was to compare the statistical significance of a single measurement of the JSD in the identification of stationary subsequences (of possibly different lengths) in a nonstationary symbolic sequence, here, we are mainly interested in comparing two (or more) measured distances. In this case, choosing weights that depend on the size of the individual samples becomes problematic when the sequences are of different lengths. The demonstration that $\sqrt{D_\alpha(\mathbf{p}, \mathbf{q})}$ is a metric for any $\alpha \in (0, 2]$ [28] is valid for fixed weights $\pi_p = \pi_q = 1/2$. More generally, the measure D_α^π itself depends on the weights π such that D_α^π and $D_\alpha^{\pi'}$ constitute different measures when $\pi \neq \pi'$. It is therefore not meaningful to compare $D_\alpha^\pi(\mathbf{p}, \mathbf{q})$ and $D_\alpha^{\pi'}(\mathbf{p}', \mathbf{q}')$ if $N_p/N_q \neq N_{p'}/N_{q'}$ because this would imply that $\pi' \neq \pi$.

2. Equal weights

In the previous section, we argued that it is essential to choose fixed weights π when comparing different distances. The choice of equal weights $\pi_p = \pi_q = 1/2$ can, however, still be interpreted in the framework of natural weights ($\pi_p = N_p/N, \pi_q = N_q/N$) as the distance between undersampled versions of the sequences. For given \mathbf{p} and \mathbf{q} with $N_p \neq N_q$, we choose equal weights $\pi_p = \pi_q = 1/2$

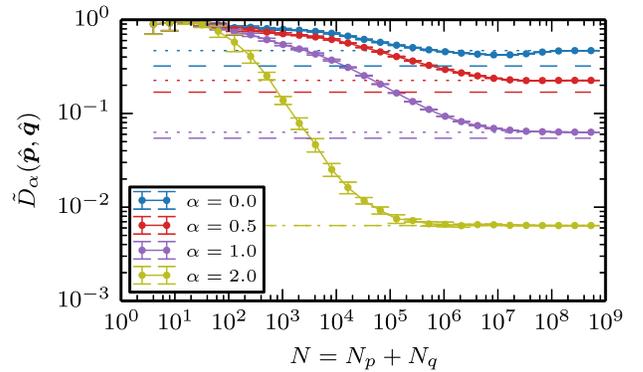


FIG. 7. JSD- α for sequences of different lengths. Measurement of $\tilde{D}_\alpha(\hat{\mathbf{p}}, \hat{\mathbf{q}})$ between sequences $\hat{\mathbf{p}}, \hat{\mathbf{q}}$ of size $N_p' = N_q'$ sampled randomly from the empirical distribution of the Google-ngram of the years $t \in \{1850, 1950\}$ with different sizes, i.e., $\mathbf{p} = \mathbf{p}_{t=1850}$ and $\mathbf{q} = \mathbf{p}_{t=1950}$ with $N_p \neq N_q$, as a function of the sample size $N' = N_p' + N_q'$ for different values of α . The dotted (dashed) lines show $\tilde{D}_\alpha^\pi(\mathbf{p}, \mathbf{q})$ between the full distributions \mathbf{p} and \mathbf{q} with equal (natural) weights, i.e., $\pi_p = \pi_q = 1/2$ [$\pi_p = N_p/(N_p + N_q) \approx 0.22$ and $\pi_q = N_q/(N_p + N_q) \approx 0.78$ corresponding to the relative size of \mathbf{p} and \mathbf{q}].

yielding a distance $D_\alpha^{1/2}(\mathbf{p}, \mathbf{q})$. If we randomly draw samples \mathbf{p}' and \mathbf{q}' of size $N_{p'} = N_{q'}$ from the distributions \mathbf{p} and \mathbf{q} , (by construction) the natural weights coincide with the equal weights, i.e., $\pi_p = \pi_q = N'_p/N = N'_q/N = 1/2$, and $\lim_{N'_p=N'_q \rightarrow \infty} D_\alpha^{1/2}(\mathbf{p}', \mathbf{q}') = D_\alpha^{1/2}(\mathbf{p}, \mathbf{q})$.

In Fig. 7, we show the difference in $\tilde{D}_\alpha^\pi(\mathbf{p}, \mathbf{q})$ between two empirical distributions from the Google-ngram with different sizes ($N_p \neq N_q$) when choosing equal and natural weights. Using equal weights corresponds to the case in which we draw samples $\hat{\mathbf{p}}$ and $\hat{\mathbf{q}}$ that are of equal length ($N_{p'} = N_{q'}$) such that equal and natural weights coincide, and taking the limit $N_{p'}, N_{q'} \rightarrow \infty$.

APPENDIX B: FINITE-SIZE ESTIMATION OF H_α , D_α , AND \tilde{D}_α

In this section, we present the calculations on the mean (i.e., the bias) and the fluctuations in finite-size estimates of H_α , D_α , and \tilde{D}_α . The starting point is a finite sample $\hat{\mathbf{p}} = (n_1/N, n_2/N, \dots, n_V/N)$ of size N (where n_i is the number of times symbol i was observed), which we assume is obtained from N identical and independent draws from the distribution \mathbf{p} giving an estimator for H_α :

$$H_\alpha(\hat{\mathbf{p}}) = \frac{1}{1-\alpha} \left(\sum_{i: \hat{p}_i > 0} \hat{p}_i^\alpha - 1 \right). \quad (\text{B1})$$

In order to take the corresponding expectation values, we expand \hat{p}_i^α around the true probabilities p_i up to second order

$$\hat{p}_i^\alpha \approx p_i^\alpha + (\hat{p}_i - p_i)\alpha p_i^{\alpha-1} + \frac{1}{2}(\hat{p}_i - p_i)^2 \alpha(\alpha-1) p_i^{\alpha-2} \quad (\text{B2})$$

and average over the realizations of the random variables \hat{p}_i^α by assuming that each symbol is drawn independently from a binomial with probability p_i such that $\langle (\hat{p}_i - p_i) \rangle = 0$ and $\langle (\hat{p}_i - p_i)^2 \rangle = p_i(1-p_i)/N \approx p_i/N$ yielding [37]

$$\langle \hat{p}_i^\alpha \rangle \approx p_i^\alpha + \frac{1}{2N} \alpha(\alpha-1) p_i^{\alpha-1}. \quad (\text{B3})$$

1. H_α

Combining Eqs. (B1) and (B3), we obtain for the mean

$$\begin{aligned} \mathbb{E}[H_\alpha(\hat{\mathbf{p}})] &\equiv \langle H_\alpha(\hat{\mathbf{p}}) \rangle = \frac{1}{1-\alpha} \left(\sum_{i \in \langle V_{\hat{\mathbf{p}}} \rangle} \langle \hat{p}_i^\alpha \rangle - 1 \right) \\ &= \frac{1}{1-\alpha} \left(\sum_{i \in \langle V_{\hat{\mathbf{p}}} \rangle} p_i^\alpha - 1 \right) - \frac{\alpha}{2N} \sum_{i \in \langle V_{\hat{\mathbf{p}}} \rangle} p_i^{\alpha-1} \\ &= \frac{1}{1-\alpha} (V_{\hat{\mathbf{p}}}^{(\alpha+1)} - 1) - \frac{\alpha}{2N} V_{\hat{\mathbf{p}}}^{(\alpha)}, \end{aligned} \quad (\text{B4})$$

where we introduce the notation $\sum_{i \in \langle V_{\hat{\mathbf{p}}} \rangle}$ indicating that we average only over the expected number of observed symbols $\langle V_{\hat{\mathbf{p}}} \rangle$ in samples $\hat{\mathbf{p}}$.

For the variance, we get

$$\begin{aligned} \mathbb{V}[H_\alpha(\hat{\mathbf{p}})] &\equiv \mathbb{E}[H_\alpha(\hat{\mathbf{p}})^2] - \mathbb{E}[H_\alpha(\hat{\mathbf{p}})]^2 \\ &= \frac{1}{(1-\alpha)^2} \sum_{i \in \langle V_{\hat{\mathbf{p}}} \rangle} \sum_{j \in \langle V_{\hat{\mathbf{p}}} \rangle} (\langle \hat{p}_i^\alpha \hat{p}_j^\alpha \rangle - \langle \hat{p}_i^\alpha \rangle \langle \hat{p}_j^\alpha \rangle) \\ &= \frac{\alpha^2}{(1-\alpha)^2 N} \sum_{i \in \langle V_{\hat{\mathbf{p}}} \rangle} p_i^{2\alpha-1} - \frac{\alpha^2}{4N^2} \sum_{i \in \langle V_{\hat{\mathbf{p}}} \rangle} p_i^{2\alpha-2} \\ &= \frac{\alpha^2}{(1-\alpha)^2} \frac{V_{\hat{\mathbf{p}}}^{(2\alpha)}}{N} - \frac{\alpha^2}{4} \frac{V_{\hat{\mathbf{p}}}^{(2\alpha-1)}}{N^2}, \end{aligned} \quad (\text{B5})$$

where we used the fact that two different symbols $i \neq j$ are independently drawn; thus, $\sum_{i,j} \langle \hat{p}_i^\alpha \hat{p}_j^\alpha \rangle = \sum_{i \neq j} \langle \hat{p}_i^\alpha \rangle \langle \hat{p}_j^\alpha \rangle + \sum_i \langle \hat{p}_i^{2\alpha} \rangle$.

2. D_α

For D_α , we have two samples $\hat{\mathbf{p}}$ and $\hat{\mathbf{q}}$ each of size N randomly sampled from the distributions \mathbf{p} and \mathbf{q} such that we can express the mean and the variance from the expectation values of the corresponding individual entropies.

Introducing the notation $\mathbf{P} \equiv \frac{1}{2}(\mathbf{p} + \mathbf{q})$, we get for the mean

$$\begin{aligned} \mathbb{E}[D_\alpha(\hat{\mathbf{p}}, \hat{\mathbf{q}})] &= \mathbb{E}[H_\alpha(\hat{\mathbf{P}})] - \frac{1}{2} \mathbb{E}[H_\alpha(\hat{\mathbf{p}})] - \frac{1}{2} \mathbb{E}[H_\alpha(\hat{\mathbf{q}})] \\ &= \frac{1}{1-\alpha} \left\{ V_{\hat{\mathbf{P}}}^{(\alpha+1)} - \frac{1}{2} V_{\hat{\mathbf{p}}}^{(\alpha+1)} - \frac{1}{2} V_{\hat{\mathbf{q}}}^{(\alpha+1)} \right\} \\ &\quad + \frac{\alpha}{2N} \left\{ \frac{1}{2} V_{\hat{\mathbf{P}}}^{(\alpha)} + \frac{1}{2} V_{\hat{\mathbf{q}}}^{(\alpha)} - \frac{1}{2} V_{\hat{\mathbf{p}}}^{(\alpha)} \right\}, \end{aligned} \quad (\text{B6})$$

where $V_{\hat{\mathbf{P}}}^{(\alpha)}$ denotes the generalized vocabulary, Eq. (13), for the combined sequence $\hat{\mathbf{P}} = \frac{1}{2}(\hat{\mathbf{p}} + \hat{\mathbf{q}})$, which is of length $2N$.

For the variance, we get

$$\begin{aligned} \mathbb{V}[D_\alpha(\hat{\mathbf{p}}, \hat{\mathbf{q}})] &\equiv \mathbb{E}[D_\alpha(\hat{\mathbf{p}}, \hat{\mathbf{q}})^2] - \mathbb{E}[D_\alpha(\hat{\mathbf{p}}, \hat{\mathbf{q}})]^2 \\ &= \mathbb{V}[H_\alpha(\hat{\mathbf{P}})] + \frac{1}{4} \mathbb{V}[H_\alpha(\hat{\mathbf{p}})] + \frac{1}{4} \mathbb{V}[H_\alpha(\hat{\mathbf{q}})] \\ &\quad - \text{Cov}[H_\alpha(\hat{\mathbf{P}}), H_\alpha(\hat{\mathbf{p}}) + H_\alpha(\hat{\mathbf{q}})], \end{aligned} \quad (\text{B7})$$

where $\text{Cov}[X, Y] \equiv \mathbb{E}[XY] - \mathbb{E}[X]\mathbb{E}[Y]$. We evaluate the covariance term in two different ways, i.e.,

$$\begin{aligned}
(1-\alpha)^2 \text{Cov}[H_\alpha(\hat{\mathbf{P}}), H_\alpha(\hat{\mathbf{p}}) + H_\alpha(\hat{\mathbf{q}})] &= \left\langle \sum_{i: \hat{p}_i + \hat{q}_i > 0} \hat{P}_i^\alpha \left(\sum_{j: \hat{p}_j > 0} \hat{p}_j^\alpha + \sum_{j: \hat{q}_j > 0} \hat{q}_j^\alpha \right) \right\rangle - \left\langle \sum_{i: \hat{p}_i + \hat{q}_i > 0} \hat{P}_i^\alpha \right\rangle \\
&\quad \times \left(\left\langle \sum_{j: \hat{p}_j > 0} \hat{p}_j^\alpha \right\rangle + \left\langle \sum_{j: \hat{q}_j > 0} \hat{q}_j^\alpha \right\rangle \right) \\
&= \left\langle \sum_{i: \hat{p}_i + \hat{q}_i > 0} \hat{P}_i^\alpha \sum_{j: \hat{p}_j + \hat{q}_j > 0} (\hat{p}_j^\alpha + \hat{q}_j^\alpha) \right\rangle - \left\langle \sum_{i: \hat{p}_i + \hat{q}_i > 0} \hat{P}_i^\alpha \right\rangle \left\langle \sum_{j: \hat{p}_j + \hat{q}_j > 0} (\hat{p}_j^\alpha + \hat{q}_j^\alpha) \right\rangle \\
&= \sum_{i \in \langle V_{\hat{p}} \rangle} \{ \langle \hat{P}_i^\alpha (\hat{p}_i^\alpha + \hat{q}_i^\alpha) \rangle - \langle \hat{P}_i^\alpha \rangle (\langle \hat{p}_i^\alpha \rangle + \langle \hat{q}_i^\alpha \rangle) \} \tag{B8}
\end{aligned}$$

and

$$\begin{aligned}
(1-\alpha)^2 \text{Cov}[H_\alpha(\hat{\mathbf{P}}), H_\alpha(\hat{\mathbf{p}}) + H_\alpha(\hat{\mathbf{q}})] &= \left\langle \sum_{i: \hat{p}_i + \hat{q}_i > 0} \hat{P}_i^\alpha \sum_{j: \hat{p}_j > 0} \hat{p}_j^\alpha \right\rangle - \left\langle \sum_{i: \hat{p}_i + \hat{q}_i > 0} \hat{P}_i^\alpha \right\rangle \left\langle \sum_{j: \hat{p}_j > 0} \hat{p}_j^\alpha \right\rangle \\
&\quad + \left\langle \sum_{i: \hat{p}_i + \hat{q}_i > 0} \hat{P}_i^\alpha \sum_{j: \hat{q}_j > 0} \hat{q}_j^\alpha \right\rangle - \left\langle \sum_{i: \hat{p}_i + \hat{q}_i > 0} \hat{P}_i^\alpha \right\rangle \left\langle \sum_{j: \hat{q}_j > 0} \hat{q}_j^\alpha \right\rangle \\
&= \sum_{i \in \langle V_{\hat{p}} \rangle} \{ \langle \hat{P}_i^\alpha \hat{p}_i^\alpha \rangle - \langle \hat{P}_i^\alpha \rangle \langle \hat{p}_i^\alpha \rangle \} + \sum_{i \in \langle V_{\hat{q}} \rangle} \{ \langle \hat{P}_i^\alpha \hat{q}_i^\alpha \rangle - \langle \hat{P}_i^\alpha \rangle \langle \hat{q}_i^\alpha \rangle \}. \tag{B9}
\end{aligned}$$

Similarly as in Eq. (B3), we can approximate

$$\begin{aligned}
\langle \hat{P}_i^\alpha \rangle &\approx P_i^\alpha + \frac{\alpha(\alpha-1)}{4N} P_i^{\alpha-1}, \\
\langle \hat{P}_i^\alpha \hat{p}_i^\alpha \rangle &\approx P_i^\alpha p_i^\alpha + \frac{\alpha}{4N} (3\alpha-1) P_i^{\alpha-1} p_i^\alpha + \frac{\alpha}{2N} (\alpha-1) P_i^\alpha p_i^{\alpha-1}, \\
\langle \hat{P}_i^\alpha \hat{q}_i^\alpha \rangle &\approx P_i^\alpha q_i^\alpha + \frac{\alpha}{4N} (3\alpha-1) P_i^{\alpha-1} q_i^\alpha + \frac{\alpha}{2N} (\alpha-1) P_i^\alpha q_i^{\alpha-1}. \tag{B10}
\end{aligned}$$

From this result, we get for the variance of D_α

$$\begin{aligned}
\mathbb{V}[D_\alpha(\hat{\mathbf{p}}, \hat{\mathbf{q}})] &= \sum_{i \in \langle V_{\hat{p}} \rangle} \left\{ \frac{\alpha^2}{(1-\alpha)^2} \frac{1}{2N} P_i^{\alpha-1} \left[P_i^\alpha - \frac{1}{2} (p_i^\alpha + q_i^\alpha) \right] - \frac{\alpha^2}{16N^2} P_i^{\alpha-1} [P_i^{\alpha-1} - (p_i^{\alpha-1} + q_i^{\alpha-1})] \right\} \\
&\quad + \frac{1}{2} \sum_{i \in \langle V_{\hat{p}} \rangle} \left\{ \frac{\alpha^2}{(1-\alpha)^2} \frac{1}{2N} P_i^\alpha [p_i^{\alpha-1} - P_i^{\alpha-1}] - \frac{\alpha^2}{8N^2} P_i^{\alpha-1} [p_i^{\alpha-1} - P_i^{\alpha-1}] \right\} \\
&\quad + \frac{1}{2} \sum_{i \in \langle V_{\hat{q}} \rangle} \left\{ \frac{\alpha^2}{(1-\alpha)^2} \frac{1}{2N} q_i^\alpha [q_i^{\alpha-1} - P_i^{\alpha-1}] - \frac{\alpha^2}{8N^2} q_i^{\alpha-1} [q_i^{\alpha-1} - P_i^{\alpha-1}] \right\}. \tag{B11}
\end{aligned}$$

Now we can see that for $\mathbf{p} = \mathbf{q} = \mathbf{P}$, we get

$$\mathbb{V}[D_\alpha(\hat{\mathbf{p}}, \hat{\mathbf{q}})]_{\mathbf{p}=\mathbf{q}} = \sum_{i \in \langle V_{\hat{p}} \rangle} \frac{1}{16N^2} \alpha^2 p_i^{2\alpha-2} = \frac{\alpha^2}{16N^2} V_{\hat{\mathbf{P}}}^{(2\alpha-1)}. \tag{B12}$$

While for arbitrary \mathbf{p} and \mathbf{q} the variance of the D_α contains the variances of the individual entropies (e.g., $V_{\hat{\mathbf{P}}}^{(2\alpha)}/N$) and a covariance term, (only) in the special case $\mathbf{p} = \mathbf{q}$ do all

first-order terms ($1/N$) vanish, yielding a qualitatively different behavior $V_{\hat{\mathbf{P}}}^{(2\alpha-1)}/N^2$.

3. \tilde{D}_α

The finite-size estimation of \tilde{D}_α can be obtained approximately by

$$\tilde{D}_\alpha(\hat{\mathbf{p}}, \hat{\mathbf{q}}) = \frac{D_\alpha(\hat{\mathbf{p}}, \hat{\mathbf{q}})}{D_\alpha(\hat{\mathbf{p}}, \hat{\mathbf{q}})_{\max}} \approx \frac{D_\alpha(\hat{\mathbf{p}}, \hat{\mathbf{q}})}{\mathbb{E}[D_\alpha^{\max}(\hat{\mathbf{p}}, \hat{\mathbf{q}})]} \tag{B13}$$

such that

$$\begin{aligned}\mathbb{E}[\tilde{D}_\alpha(\hat{p}, \hat{q})] &\approx \frac{\mathbb{E}[D_\alpha(\hat{p}, \hat{q})]}{\mathbb{E}[D_\alpha^{\max}(\hat{p}, \hat{q})]}, \\ \mathbb{V}[\tilde{D}_\alpha(\hat{p}, \hat{q})] &\approx \frac{\mathbb{V}[D_\alpha(\hat{p}, \hat{q})]}{\mathbb{E}[D_\alpha^{\max}(\hat{p}, \hat{q})]^2}.\end{aligned}\quad (\text{B14})$$

The mean of D_α^{\max} is given according to Eq. (6) as a linear combination of the individual entropies of \hat{p} and \hat{q} ,

$$\mathbb{E}[D_\alpha^{\max}(\hat{p}, \hat{q})] = \frac{2^{1-\alpha} - 1}{2} \left(\mathbb{E}[H_\alpha(\hat{p})] + \mathbb{E}[H_\alpha(\hat{q})] + \frac{2}{1-\alpha} \right).\quad (\text{B15})$$

APPENDIX C: DERIVATION OF EQ. (14)

In this section, we derive the scaling of the generalized vocabulary $V^{(\alpha)}$ defined in Eq. (13) assuming that p is a power law of the form $p_i \propto i^{-\gamma}$, Eq. (1). Instead of looking at the probability of individual symbols i , we consider the distribution of frequencies n , which in this case yields $p(n) \propto n^{-1-1/\gamma}$ [48]. Consider the sum $\sum_{i \in V} p_i = (1/N) \sum_{i \in V} n_i = (1/N) S_V(\gamma)$, where $S_V(\gamma)$ corresponds to the sum of V independent and identically distributed random variables n_i ($i = 1, \dots, V$) drawn from the distribution $p(n)$. It can be shown that [49]

$$S_V(\gamma) \propto \begin{cases} V^\gamma & \gamma > 1 \\ V & \gamma < 1 \end{cases}.\quad (\text{C1})$$

The case $\gamma = 1$ includes additional logarithmic corrections but is not of relevance for the discussion; therefore, for the sake of simplicity, we do not discuss it here. In the same way, we can treat $\sum_{i \in V} p_i^\mu = (1/N^\mu) \sum_{i \in V} n_i^\mu = (1/N^\mu) S_V(\gamma\mu)$ by noting that $S_V(\gamma\mu)$ can be interpreted as the sum of V independent and identically distributed random variables n_i ($i = 1, \dots, V$), where $n_i \sim \tilde{p}(n)$ with $\tilde{p}(n) \propto n^{-1-1/(\gamma\mu)}$ such that we get

$$S_V(\gamma\mu) \propto \begin{cases} V^{\gamma\mu} & \mu < 1/\gamma \\ V & \mu > 1/\gamma \end{cases}.\quad (\text{C2})$$

By setting $\mu = \alpha - 1$ in Eq. (13) and noting that for $p_i \propto i^{-\gamma}$, Eq. (1), the number of different symbols scales as $V \propto N^{1/\gamma}$, Eq. (10), we obtain Eq. (14).

APPENDIX D: TEMPORAL EVOLUTION OF $\tilde{D}_\alpha(\Delta t)$

We are interested in understanding the dependence of $\tilde{D}_\alpha(\Delta t) \equiv \tilde{D}_\alpha(t_0, t_0 + \Delta t)$ on Δt [we assume $\tilde{D}_\alpha(\Delta t)$ is the same for all t_0]. The triangle inequality implies that

$$\begin{aligned}\tilde{D}_\alpha(\Delta t) &\leq \left(\sqrt{\tilde{D}_\alpha(\Delta t - 1)} + \sqrt{\tilde{D}_\alpha(\Delta t = 1)} \right)^2 \\ &\leq (\Delta t)^2 \tilde{D}_\alpha(\Delta t = 1).\end{aligned}\quad (\text{D1})$$

In order to consider the origin of different Δt dependencies within the general bound given by Eq. (D1), we consider two classes of words subject to frequency change in Δt : (i) words that show fluctuations (e.g., finite sampling or topical dependencies), which do not depend on Δt and (ii) words that show a systematic increase or decrease over all t . If we assume that all words that change fall in one of these classes, we can use the fact that \tilde{D}_α is defined as a sum of word types and decompose the total change \tilde{D}_α as $\tilde{D}_\alpha = \tilde{D}_\alpha^{(i)} + \tilde{D}_\alpha^{(ii)}$, where $\tilde{D}_\alpha^{(i,ii)}$ is the divergence of all words falling into class (i) or (ii), respectively. For category (ii), the changes between consecutive years are independent; thus, the equality case of the triangle inequality is obtained: $\sqrt{\tilde{D}_\alpha(t_0, t_2)} = \sqrt{\tilde{D}_\alpha(t_0, t_1)} + \sqrt{\tilde{D}_\alpha(t_1, t_2)}$ for all t_1 , with $t_0 \leq t_1 \leq t_2$. Therefore, we obtain a quadratic dependence on Δt as

$$\tilde{D}_\alpha(\Delta t) = \tilde{D}_\alpha^{(i)} + \tilde{D}_\alpha^{(ii)}(\Delta t = 1)(\Delta t)^2.\quad (\text{D2})$$

-
- [1] S. Kullback, *Information Theory and Statistics* (Wiley, New York, 1959).
 - [2] C. D. Manning and H. Schütze, *Foundations of Statistical Natural Language Processing* (MIT Press, Cambridge, MA, 1999).
 - [3] I. Grosse, P. Bernaola-Galván, P. Carpena, R. Román-Roldán, J. Oliver, and H. E. Stanley, *Analysis of Symbolic Sequences Using the Jensen-Shannon Divergence*, *Phys. Rev. E* **65**, 041905 (2002).
 - [4] Y. He, A. Ben Hamza, and H. Krim, *A Generalized Divergence Measure for Robust Image Registration*, *IEEE Trans. Signal Process.* **51**, 1211 (2003).
 - [5] K. W. Boyack, D. Newman, R. J. Duhon, R. Klavans, M. Patek, J. R. Biberstine, B. Schijvenaars, A. Skupin, N. Ma, and K. Börner, *Clustering More Than Two Million Biomedical Publications: Comparing the Accuracies of Nine Text-Based Similarity Approaches*, *PLoS One* **6**, e18029 (2011).
 - [6] A. P. Masucci, A. Kalampokis, V. M. Eguíluz, and E. Hernández-García, *Wikipedia Information Flow Analysis Reveals the Scale-Free Architecture of the Semantic Space*, *PLoS One* **6**, e17333 (2011).
 - [7] V. Bochkarev, V. Solovyeu, and S. Wichmann, *Universals versus Historical Contingencies in Lexical Evolution*, *J. R. Soc. Interface* **11**, 20140841 (2014).
 - [8] E. A. Pechenick, C. M. Danforth, and P. S. Dodds, *Is Language Evolution Grinding to a Halt?: Exploring the Life and Death of Words in English Fiction*, [arXiv: 1503.03512](https://arxiv.org/abs/1503.03512).

- [9] G. Cocho, J. Flores, C. Gershenson, C. Pineda, and S. Sánchez, *Rank Diversity of Languages: Generic Behavior in Computational Linguistics*, *PLoS One* **10**, e0121898 (2015).
- [10] J.-B. Michel, Y. K. Shen, A. P. Aiden, A. Veres, M. K. Gray, J. P. Pickett, D. Hoiberg, D. Clancy, P. Norvig, J. Orwant, S. Pinker, M. A. Nowak, and E. L. Aiden, *Quantitative Analysis of Culture Using Millions of Digitized Books*, *Science* **331**, 176 (2011).
- [11] M. Gerlach and E. G. Altmann, *Stochastic Model for the Vocabulary Growth in Natural Languages*, *Phys. Rev. X* **3**, 021006 (2013).
- [12] G. K. Zipf, *The Psycho-Biology of Language* (Routledge, London, 1936).
- [13] R. H. Baayen, *Word Frequency Distributions* (Kluwer Academic Publishers, Dordrecht, Netherlands, 2001).
- [14] R. Ferrer i Cancho, *The Variation of Zipf's Law in Human Language*, *Eur. Phys. J. B* **44**, 249 (2005).
- [15] M. A. Serrano, A. Flammini, and F. Menczer, *Modeling Statistical Properties of Written Text*, *PLoS One* **4**, e5372 (2009).
- [16] S. K. Baek, S. Bernhardsson, and P. Minnhagen, *Zipf's Law Unzipped*, *New J. Phys.* **13**, 043004 (2011).
- [17] R. N. Mantegna, S. V. Buldyrev, A. L. Goldberger, S. Havlin, C.-K. Peng, M. Simons, and H. E. Stanley, *Linguistic Features of Noncoding DNA Sequences*, *Phys. Rev. Lett.* **73**, 3169 (1994).
- [18] C. Furusawa and K. Kaneko, *Zipf's Law in Gene Expression*, *Phys. Rev. Lett.* **90**, 088102 (2003).
- [19] J. Serrà, A. Corral, M. Boguñá, M. Haro, and J. L. Arcos, *Measuring the Evolution of Contemporary Western Popular Music*, *Sci. Rep.* **2**, 521 (2012).
- [20] J. Lin, *Divergence Measures Based on the Shannon Entropy*, *IEEE Trans. Inf. Theory* **37**, 145 (1991).
- [21] T. M. Cover and J. A. Thomas, *Elements of Information Theory* (Wiley-Interscience, New York, 2006).
- [22] D. M. Endres and J. E. Schindelin, *A New Metric for Probability Distributions*, *IEEE Trans. Inf. Theory* **49**, 1858 (2003).
- [23] M. A. Montemurro and D. H. Zanette, *Towards the Quantification of the Semantic Information Encoded in Written Language*, *Adv. Compl. Syst.* **13**, 135 (2010).
- [24] J. Havrda and F. Charvát, *Quantification Method of Classification Processes: Concept of Structural A-Entropy*, *Kybernetika* **3**, 30 (1967).
- [25] J. Burbea and C. Rao, *On the Convexity of Some Divergence Measures Based on Entropy Functions*, *IEEE Trans. Inf. Theory* **28**, 489 (1982).
- [26] C. Tsallis, *Possible Generalization of Boltzmann-Gibbs Statistics*, *J. Stat. Phys.* **52**, 479 (1988).
- [27] A. Rényi, *On Measures of Entropy and Information*, in *Berkeley Symposium on Mathematical Statistics and Probability* (University of California Press, Berkeley, CA, 1961).
- [28] J. Briët and P. Harremoës, *Properties of Classical and Quantum Jensen-Shannon Divergence*, *Phys. Rev. A* **79**, 052311 (2009).
- [29] T. Schürmann, *Bias Analysis in Entropy Estimation*, *J. Phys. A* **37**, L295 (2004).
- [30] G. Herdan, *Type-Token Mathematics* (Mouton, Den Haag, 1960).
- [31] H. S. Heaps, *Information Retrieval* (Academic Press, New York, 1978).
- [32] M. Gerlach and E. G. Altmann, *Scaling Laws and Fluctuations in the Statistics of Word Frequencies*, *New J. Phys.* **16**, 113010 (2014).
- [33] T. D. de Wit, *When Do Finite Sample Effects Significantly Affect Entropy Estimates?*, *Eur. Phys. J. B* **11**, 513 (1999).
- [34] G. A. Miller, *Note on the Bias of Information Estimates*, in *Information Theory in Psychology: Problems and Methods*, edited by H. Quastler (Free Press, Glencoe, IL, 1955).
- [35] G. P. Basharin, *On a Statistical Estimate for the Entropy of a Sequence of Independent Random Variables*, *Theory Probab. Appl.* **4**, 333 (1959).
- [36] B. Harris, *The Statistical Estimation of Entropy in the Non-Parametric Case*, *Colloq. Math. Soc. J. B* **16**, 323 (1975).
- [37] H. Herzel, A. O. Schmitt, and W. Ebeling, *Finite Sample Effects in Sequence Analysis*, *Chaos Solitons Fractals* **4**, 97 (1994).
- [38] P. Grassberger, *Entropy Estimates from Insufficient Samplings*, [arXiv:physics/0307138](https://arxiv.org/abs/physics/0307138).
- [39] E. Archer, I. M. Park, and J. W. Pillow, *Bayesian Entropy Estimation for Countable Discrete Distributions*, *J. Mach. Learn. Res.* **15**, 2833 (2014).
- [40] I. Nemenman, F. Shafee, and W. Bialek, *Entropy and Inference, Revisited*, in *Advances in Neural Information Processing Systems 14* (MIT Press, Cambridge, MA, 2002).
- [41] For small values of N , the Bayesian estimates of the individual entropies in some cases yield $\tilde{D}(\hat{p}, \hat{q}) < 0$ such that we consider averages over $|\tilde{D}(\hat{p}, \hat{q})|$.
- [42] P. Bernaola-Galván, I. Grosse, P. Carpena, J. L. Oliver, R. Román-Roldán, and H. E. Stanley, *Finding Borders between Coding and Noncoding DNA Regions by an Entropic Segmentation Method*, *Phys. Rev. Lett.* **85**, 1342 (2000).
- [43] We use the Google-ngram corpus containing the frequency of usage of millions of words from 4% of all books ever published in the English language with a yearly resolution [10].
- [44] M. Pagel, Q. D. Atkinson, and A. Meade, *Frequency of Word-Use Predicts Rates of Lexical Evolution throughout Indo-European History*, *Nature (London)* **449**, 717 (2007).
- [45] E. Lieberman, J.-B. Michel, J. Jackson, T. Tang, and M. A. Nowak, *Quantifying the Evolutionary Dynamics of Language*, *Nature (London)* **449**, 713 (2007).
- [46] T. Yasseri, A. Kornai, and J. Kertész, *A Practical Approach to Language Complexity: A Wikipedia Case Study*, *PLoS One* **7**, e48386 (2012).
- [47] E. Stamatatos, *A Survey of Modern Authorship Attribution Methods*, *J. Am. Soc. Inf. Sci. Tec.* **60**, 538 (2009).
- [48] M. E. J. Newman, *Power Laws, Pareto Distributions and Zipf's Law*, *Contemp. Phys.* **46**, 323 (2005).
- [49] J.-P. Bouchaud and A. Georges, *Anomalous Diffusion in Disordered Media: Statistical Mechanisms, Models and Physical Applications*, *Phys. Rep.* **195**, 127 (1990).