

Scaling laws and dynamics of hashtags on Twitter

Cite as: Chaos **30**, 063112 (2020); <https://doi.org/10.1063/5.0004983>

Submitted: 16 February 2020 . Accepted: 06 May 2020 . Published Online: 02 June 2020

Hongjia H. Chen , Tristram J. Alexander, Diego F. M. Oliveira , and Eduardo G. Altmann 



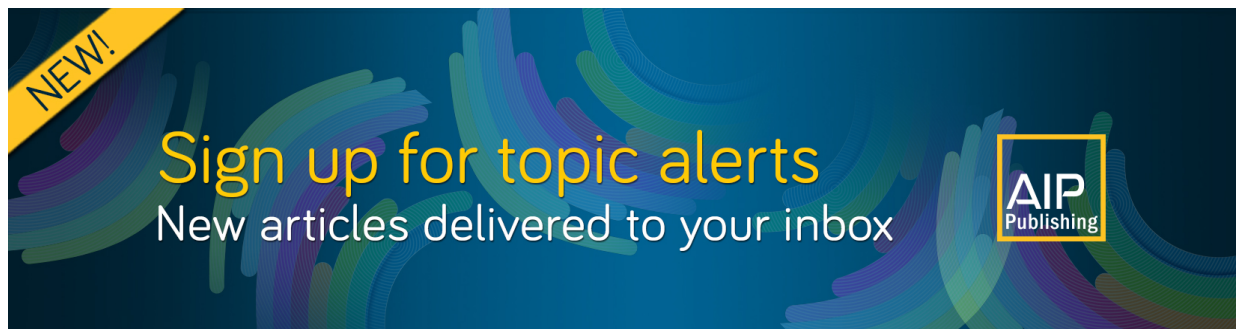
View Online



Export Citation



CrossMark



NEW!
Sign up for topic alerts
New articles delivered to your inbox
AIP
Publishing



Scaling laws and dynamics of hashtags on Twitter

Cite as: Chaos 30, 063112 (2020); doi: 10.1063/5.0004983

Submitted: 16 February 2020 · Accepted: 6 May 2020 ·

Published Online: 2 June 2020



View Online



Export Citation



CrossMark

Hongjia H. Chen,^{1,2}  Tristram J. Alexander,³ Diego F. M. Oliveira,^{4,5}  and Eduardo G. Altmann^{1,a)} 

AFFILIATIONS

¹School of Mathematics and Statistics, University of Sydney, Sydney, 2006 NSW, Australia

²Department of Mathematics, University of Auckland, 1010 Auckland, New Zealand

³School of Physics, University of Sydney, Sydney, 2006 NSW, Australia

⁴U.S. Army Research Laboratory, 2800 Powder Mill Rd., Adelphi, Maryland 20783, USA

⁵Network Science and Technology Center, Rensselaer Polytechnic Institute, 335 Materials Research Center 110 8th St., Troy, New York 12180, USA

Note: This article is part of the Focus Issue, Dynamics of Social Systems.

a) Author to whom correspondence should be addressed: eduardo.altmann@sydney.edu.au

ABSTRACT

In this paper, we quantify the statistical properties and dynamics of the frequency of hashtag use on Twitter. Hashtags are special words used in social media to attract attention and to organize content. Looking at the collection of all hashtags used in a period of time, we identify the scaling laws underpinning the hashtag frequency distribution (Zipf's law), the number of unique hashtags as a function of sample size (Heaps' law), and the fluctuations around expected values (Taylor's law). While these scaling laws appear to be universal, in the sense that similar exponents are observed irrespective of when the sample is gathered, the volume and the nature of the hashtags depend strongly on time, with the appearance of bursts at the minute scale, fat-tailed noise, and long-range correlations. We quantify this dynamics by computing the Jensen–Shannon divergence between hashtag distributions obtained τ times apart and we find that the speed of change decays roughly as $1/\tau$. Our findings are based on the analysis of 3.5×10^9 hashtags used between 2015 and 2016.

Published under license by AIP Publishing. <https://doi.org/10.1063/5.0004983>

The mathematical study of social systems is only possible because similar processes exist in seemingly different social configurations. Two examples from dynamical systems are rich-get-richer processes—responsible for the appearance of fat-tailed distributions—and evolutionary processes—controlling the dynamics of memes. Data from the microblogging platform Twitter allow us to study these two generic processes with an unprecedented quantitative accuracy. Here, we view hashtags as memes and quantify emerging properties of the collective interaction between these memes, including the appearance of scaling laws and the different time scales involved in their dynamics.

I. INTRODUCTION

Hashtags (“#”) have proven to be one of the most successful innovations in social-media language. They were originally introduced on Twitter to identify topical content in tweets,¹ essentially serving as topic markers to facilitate search and retrieval² in the face of an overwhelming amount of information. For instance, the hashtag “#DynamicsOfSocialSystems” could be used in social-media

messages to help users identify comments and papers relevant to this topic. In parallel to this, hashtags also provide a means for users to enhance social ties² and conduct a metacommentary distinct from other tweet content.³ Users exposed to a hashtag are invited to use (or modify) the hashtag, starting an imitation⁴ and mutation process that leads to a fat-tailed distribution^{5,6} of hashtag frequencies⁷ and that is typical of evolutionary dynamics observed more generally (e.g., in language and in memes).^{8,9} Hashtags are thus convenient—can be easily identified and traced—and generic—show behavior seen in various systems (e.g., language, social media, etc.)—creating thus an ideal scenario for a data-driven study of the dynamics of social systems.

Previous works examining the dynamical processes underpinning hashtag use have focused on the role of the connections between users on the resulting dynamics.^{7,10} As such, these works form part of a more general area of research exploring the nature of user driven dynamics on social media.^{11,12} For instance, models of user behavior have been able to explain the appearance of a fat-tailed distribution in the distribution of tweets^{11,13–15} and bursty behavior in the attention of specific topics in Twitter.¹² Other works have focused on specific hashtag dynamics, for instance, on the response

to an external event,¹⁶ for the purposes of ease of analysis while developing data-mining methods¹⁷ or while studying the competition behind diffusion processes.^{18–20} Instead, here we are interested not in the dynamics of specific hashtags, but rather in the general statistical behavior of all hashtags used during a particular time window. By looking at all hashtags simultaneously, we account for interactions between different hashtags and we provide an overall statistical characterization of the dynamics of hashtag usage. This is done by repeating classical analyses done in quantitative linguistics for word frequencies.^{21–27} This approach is justified not only because hashtags can be seen as special types of words but also because similar dynamical (evolutionary) processes affect the frequency of word usages (albeit at different scales).

The main findings of our manuscript are that hashtags follow statistical laws similar to the linguistic laws observed for words—such as Zipf's and Heaps' laws—but that differences appear due to the dynamics of the hashtags. We identify two main aspects of the dynamics of hashtags that differ from natural language: (i) extremely bursty behavior in the usage of hashtags over time leads to larger than expected fluctuations around the statistical laws, as characterized by an unusual scaling exponent of Taylor's law and (ii) hashtag usage evolves rapidly with time τ . We quantify the latter using the (generalized) Jensen–Shannon distance between hashtag observations separated by time τ ,²⁸ and we find a scaling law which characterizes the change in hashtag usage as a function of τ .

This paper is divided as follows. In Sec. II, we describe our data and we show relevant time scales of the dynamics. In Secs. III and IV, we focus on the distribution and scaling behavior of hashtag frequencies, comparing them to results for word frequencies. In Sec. V, we investigate how fast the hashtag distributions change, reporting a new scaling law for the dynamics of hashtags.

II. TIME SERIES OF TYPES AND TOKENS

Our database consists of all hashtags used in a 10% sample of all tweets published between November 1, 2015 and November 30, 2016. For a given time interval $([t, t + b])$ around time t and of (bin) size b , we count how many hashtags were used in our database. Here it is important to distinguish between hashtag types (i.e., unique hashtags) and hashtag tokens (i.e., the repetitive usage of potentially the same hashtags). For instance, in our complete database ($b = 392$ days), the hashtag type “#mtvstars” is the most frequently used hashtag (rank $r = 1$), responsible for the appearance of $M_{r=1} = 49M$ hashtag tokens. Next, we have “#kca” with $M_{r=2} = 28M$ and “#iheartawards” with $M_{r=3} = 26M$. Overall, we have $N = 57$ 876 308 types and $M = \sum_{r=1}^N M_r = 3\,492\,300\,357$ tokens in our database. We denote $N(t)$ and $M(t)$ as the number of hashtag types and tokens, respectively, in an interval of size b starting at time t .³³

Figure 1 shows how the number of hashtag tokens M and types N change in time t at different time scales. The time series of tokens $M(t)$ shows a more noisy behavior than the time series $N(t)$ of types: $M(t)$ shows pronounced bursts and spikes while $N(t)$ reflects more clearly the weekly and daily oscillations of Twitter usage. We see a weekly minimum in activity on a Sunday, while the daily maximum occurs around 1600 GMT. At short time scales, both time series have peaks at the first minute of each hour and each half hour, suggesting that a large number of pre-programmed tweets are being launched

at regular patterns. The main peak in $M(t)$ highlighted in this figure is mostly due to the hashtag “#iheartawards,” which was used during a music awards show that took place in the USA on April 3, 2016 and has rank $r = 3$ in our complete database.

III. ZIPF'S LAW

We are interested in the share of total hashtag tokens obtained by the different hashtag types, which can be interpreted as the success rate of individual memes in attracting the attention of users.³⁰ This possibility of the “rich-getting-richer” element of hashtag use suggests that a fat-tailed distribution should be expected, because of the ubiquity of such a distribution type in data from natural and social systems.^{5–7} Possibly the best known example of such a distribution is Zipf's law, which states that the frequency $F_r = M_r/M$ (i.e., the fraction of all tokens) of the r -th most frequent word (type) decays with r as

$$F_r \sim r^{-\gamma}, \quad (1)$$

with $\gamma \gtrsim 1$.

In Fig. 2, we show a representation of the hashtag distribution. We observe that a similar distribution is observed for different time intervals, that the distribution spans many orders of magnitude—in agreement with the fat-tailed character of Eq. (1)—, and that the distribution shows a positive concavity (in the double-logarithmic plot) indicating a faster than Zipfian decay. All these observations have been reported for the frequency of words in a recent analysis of Zipf's law in a large dataset (Google n-grams)³³ and are consistent with previous analysis of hashtag frequencies.⁷

The observations above motivate us to consider whether generalizations of Zipf's law proposed to describe word frequencies are also describing hashtag frequencies. We considered the distributions and methodology proposed in Ref. 23 to determine which of the eight parameterizations of F_r best describes our hashtag data. Table I lists the different distributions, the best inferred parameters, and a measure of the agreement between data and the (best) distributions. The results show that the best generalized Zipf's law is obtained by a lognormal fit of the rank distribution,

$$F_r = Cr^{-1} \exp\left(-\frac{1}{2}(\ln(r) - m)^2/s^2\right), \quad (2)$$

where $C = C(m, s)$ is a normalization constant and m, s are free parameters such that $m < s^2$. The restriction in the parameter choice is necessary to ensure that Eq. (2) is monotonically decaying in the integers r . This is necessary because, by construction, F_r is monotonically decaying (a lognormal distribution in F_r does not imply that the number of hashtag types with a given frequency M_i/M is also lognormal). A further indication that the distribution (2) provides a good description of the data for different times t is the fact that the estimated parameters m and s do not strongly depend on the size of the database b (see the lower panel of Fig. 2). This is a different finding from the one reported for natural language, where a double power-law distribution provided a better fit.^{21,23} Differently from the case of language, in the case of hashtags, the double gamma distribution (with three free parameters) leads to a smaller likelihood \mathcal{L} (or larger $-\log \mathcal{L}$) than the lognormal. Moreover, the parameters of the double gamma in the hashtag distributions differ from the

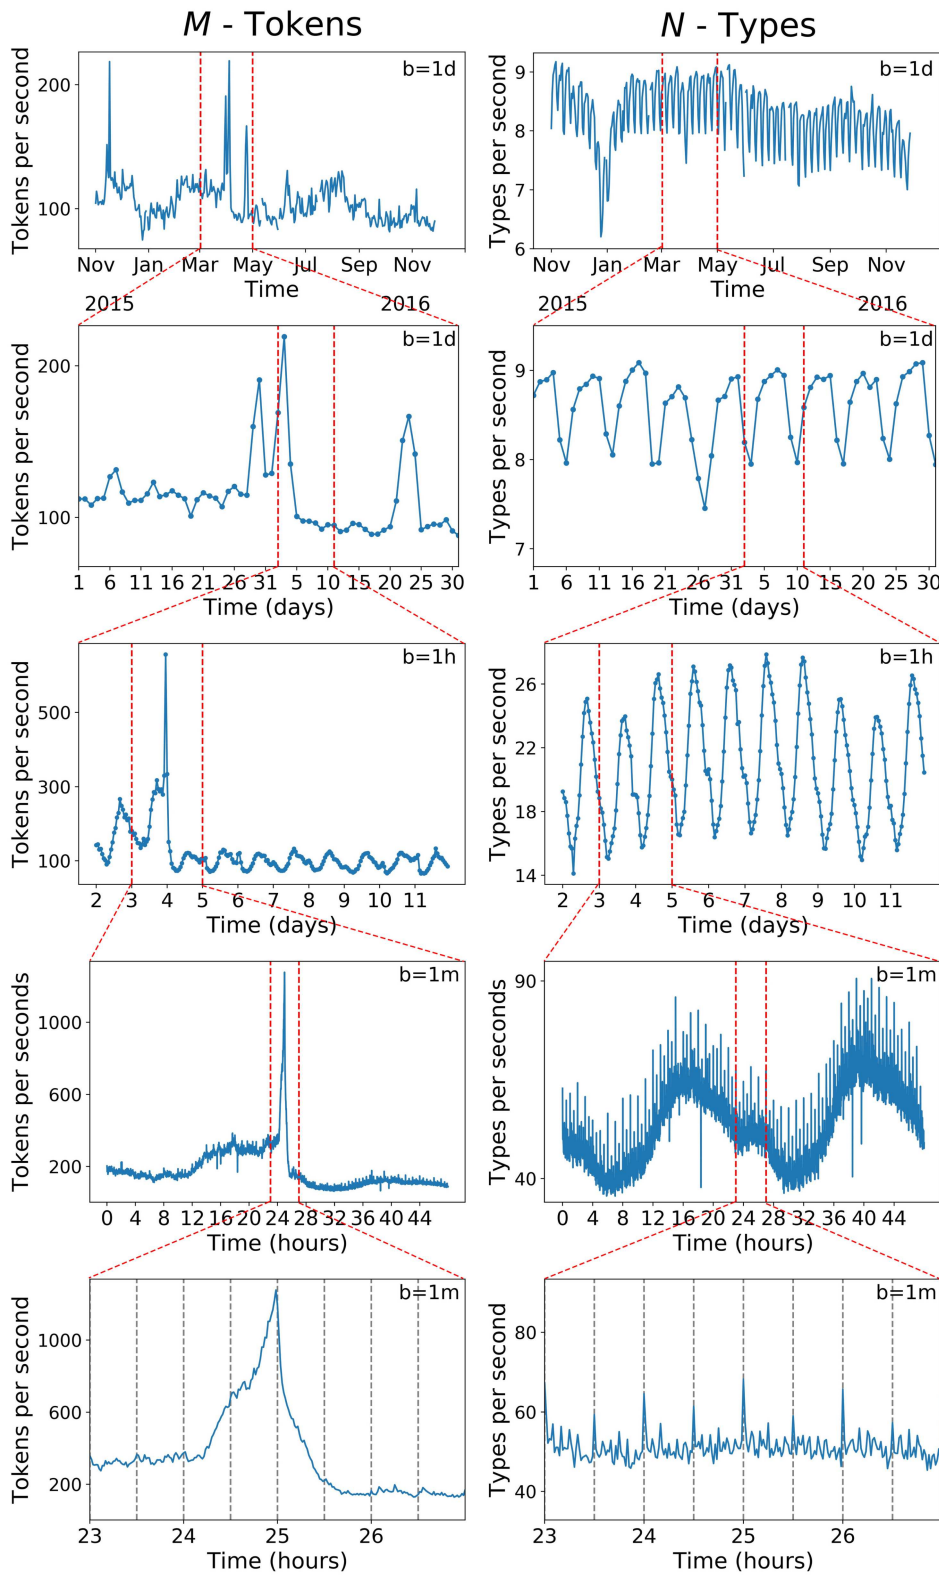


FIG. 1. Time series of hashtag tokens $M(t)$ (left column) and types $N(t)$ (right column). Reported (y axis) is the rate of usage (N/b and M/b with b measured in seconds). Each panel corresponds to a magnification in the time scale (x axis) of the panel immediately above it in the region indicated by vertical dashed (red) lines. In the two top panels, the data were aggregated at different scales b : $b = 1$ day for the two top panels, $b = 1$ h for the middle panels, and $b = 1$ min for the lower two panels.³³ Time corresponds to GMT.

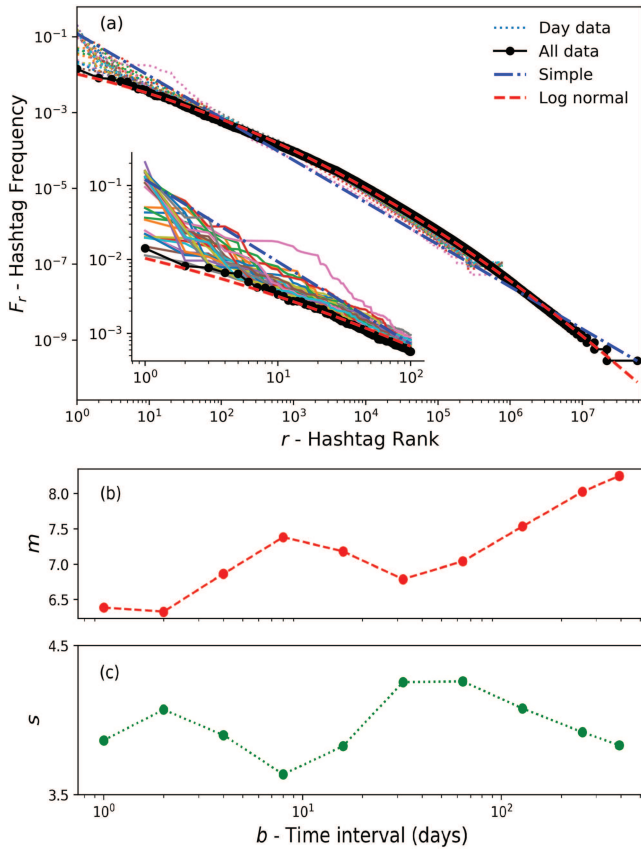


FIG. 2. Fat-tailed distribution of hashtag frequencies. (a) The solid (black) curve shows the frequency F_r (y axis) of the r th (x axis) most frequent hashtag in our complete database $b = 392$ days. The thin (colored) lines show the results obtained for 30 different cases with $b = 1$ day. The inset shows a magnification for small r . In the main panel, the dotted line corresponds to Zipf’s law (1) with the maximum-likelihood parameter $\hat{\gamma} = 1.13$ inferred from the data and the dashed (red) line correspond to the best generalized Zipf’s law given by Eq. (2) with estimated parameters $\hat{m} = 8.25$ and $\hat{s} = 3.83$. Panels (b) and (c) show how the parameters \hat{m} and \hat{s} vary with the size of the database b used in the estimate, i.e., we used all hashtags in the time interval $[t, t + b]$ with t fixed (the minimum) and varying b .

case of language: while for language the first exponent was $\gamma = 1$ (as originally proposed by Zipf), in the case of hashtags, the first exponent is $\gamma \approx 0.8 < 1$. Altogether, in comparison to word frequencies, hashtags have a slower initial decay of F_r (i.e., the top ranked hashtags have a more similar frequency) and a faster asymptotic decay of F_r (which is faster than a power law but slower than an exponential).

IV. HEAPS’ AND TAYLOR’S LAWS

The Zipfian-type behavior of hashtag frequencies motivates us to consider also other statistical laws proposed in quantitative linguistics.^{21–24,26,27} We start with Heaps’ law, which states that the

number of types N and tokens M scale nonlinearly as

$$N \sim M^\lambda, \tag{3}$$

where $\lambda < 1$ and the symbol \sim indicates that the ratio of the left and right sides tend to a constant for large M . To perform this analysis, we compute $M(t)$ and $N(t)$ at different time intervals $[t, t + b]$, for different t ’s and b ’s as above. We then consider averages $\langle \dots \rangle$ over all times t for a fixed b and compute the expected value and standard deviation of these quantities as

$$\mu_M = \langle M(t) \rangle, \sigma_M = \sqrt{\langle M^2(t) \rangle - \langle M(t) \rangle^2}, \tag{4}$$

$$\mu_N = \langle N(t) \rangle, \sigma_N = \sqrt{\langle N^2(t) \rangle - \langle N(t) \rangle^2}. \tag{5}$$

By varying b from minutes to months, we effectively vary the size of the database over many orders of magnitude, allowing us to explore the scaling between these quantities.

In our case, Heaps’ law (3) is interpreted as the relation between how μ_N (the expected number of types N) scales with μ_M (the expected number of tokens M). The results in Fig. 3 reveal a striking scaling law over more than four decades, with an estimated exponent $\lambda \approx 0.73$. In this plot, we also show the results obtained after shuffling the series at the scale of $b = 1$ min. We observe that μ_N is increased in the randomized data, reflecting the existing correlation between the hashtags used in neighboring time intervals.²⁵ However, the same Heaps law scaling is observed for the shuffled data, in agreement with the previous demonstrations that Heaps’ law can be obtained from a random sampling of Zipf’s law.²⁵

We now investigate how the fluctuations σ scale with the mean μ as

$$\sigma \sim \mu^\beta. \tag{6}$$

Reference 31 provides a review of this scaling, known as Taylor’s law, showing its appearance and significance in various complex systems. The exponent $\beta = 1/2$ is obtained if we consider that the quantity of interest (M in our case) is obtained as the sum of random quantities sampled independently from a distribution with a well-defined second moment. In our case, we can think that the values M in a time interval $[t, t + b]$ are obtained as the sum of the number of hashtags at smaller scales. The case $\beta = 1$ reflects the lack of mixing in the terms being summed.^{25,31} Nontrivial values, $0.5 < \beta < 1$, are obtained in the presence of long-range correlations (in time t) or if the underlying distribution from which samples are taken does not have a second moment (large fluctuations of M in small time intervals). In natural language, $\beta = 1$ was observed for the case of word types N ²⁵ and $0.5 < \beta < 1$ was reported for the fluctuation of individual words.²⁷

The results for our hashtag dataset are reported in Fig. 4 and indicate that the exponent $\beta \approx 0.84$ is clearly within the range of non-trivial values (i.e., clearly different from $\beta = 0.5$ and $\beta = 1.0$). In order to clarify the origin of this non-trivial exponent, we repeat the analysis after randomizing the time series $M(t)$. As expected, the exponent after the randomization $\beta_R \approx 0.6$ is smaller than the original exponent. The fact that this exponent is still larger than $1/2$ indicates that the origin of the non-trivial Taylor’s law in the hashtag frequencies is due to both long-range correlation in $M(t)$ and

TABLE I. Generalized Zipf's law for hashtag frequencies. Different models for the rank-frequency distribution $F_r \equiv F(r|\Omega)$ were fitted to the empirical distribution F_r using the maximum-likelihood methods proposed in Ref. 23. The parameters Ω that maximize the likelihood \mathcal{L} are reported together with the negative log-likelihood per token $-\log \mathcal{L}/M$ (at the given parameters). The model with maximum likelihood (minimum $-\log \mathcal{L}$) is the lognormal model (value in boldface in the last column).

Model	$F_r \equiv F(r \Omega)$	Parameter estimates	$-\log \mathcal{L}/M$
Simple	$Cr^{-\gamma}$	$\gamma = 1.11$	11.544
Shifted power law	$C(r+a)^{-\gamma}$	$\gamma = 1.25, a = 119.8$	11.205
Exponential cutoff	$C \exp(-ar)r^{-\gamma}$	$\gamma = 0.96, a = 1.11$	11.195
Naranan	$C \exp(-a/r)r^{-\gamma}$	$\gamma = 1.16, a = 5.0931$	11.347
Weibull	$C \exp(-ar^{-\gamma})r^{\gamma-1}$	$\gamma = -0.24, a = 4.51$	12.175
Lognormal	$Cr^{-1} \exp(-\frac{1}{2}(\ln(r) - m)^2/s^2)$	$m = 8.25, s = 3.83$	11.075
Double power law	$C \begin{cases} r^{-1} & r \leq a \\ a^{\gamma-1} r^{-\gamma} & r > a \end{cases}$	$\gamma = 1.57, a = 352288.8$	11.186
Double gamma	$C \begin{cases} r^{-\gamma_1} & r \leq a \\ a^{\gamma_2-\gamma_1} r^{-\gamma_2} & r > a \end{cases}$	$\gamma_1 = 0.8083, \gamma_2 = 1.4079, a = 18145.1$	11.091

sampling from an underlying fat-tailed distribution (with diverging second moment). The latter point is consistent with the bursty behavior of $N(t)$ reported in Fig. 1 above and also with the results of Ref. 12.

V. HASHTAG DYNAMICS

So far, we have concentrated on general statistical characterizations of hashtag frequencies that remain roughly invariant over time t , finding a Zipfian-like distribution and different scales between the total number of hashtag types N and tokens M . Underlying these relationships, there is a rich dynamical process of the usage $N_i(t)$ of individual hashtags. Our goal here is to quantify the extent into which, collectively, the frequency of all M hashtags change over time.

We use an information theoretic measure to quantify the similarity of two (normalized) frequency distributions, $F_r(t_1)$ and $F_r(t_2)$, following the approach used in Ref. 28 for language.

For each hashtag type $n = 1, \dots, N$, we define the frequency at time t_s as $p_n = M_n(t_s)/M(t_s)$. We consider the frequencies p_n to be an estimate of the probability of using this hashtag and $\mathbf{p} = (p_1, p_2, \dots, p_N)$ the probability distribution over all hashtags. The α entropy of \mathbf{p} is defined as

$$H_\alpha(\mathbf{p}) = \frac{1}{1-\alpha} \left(\sum_i p_i^\alpha - 1 \right) \tag{7}$$




FIG. 3. The expected number of types μ_N and tokens μ_M scale non-linearly as described by Heaps' law (3). The hashtag data—■ with (blue) line—were obtained using time intervals b ranging from $b = 1$ min to $b = 256$ days. The error bars correspond to σ_M (x axis) and σ_N (y axis). The results obtained after shuffling the temporal order of $M(t)$ and $N(t)$ obtained the scale of $b = 1$ min are shown as ● with (red) line. The scaling exponents λ indicated in the legend were obtained from a linear regression of the average results (dashed lines).

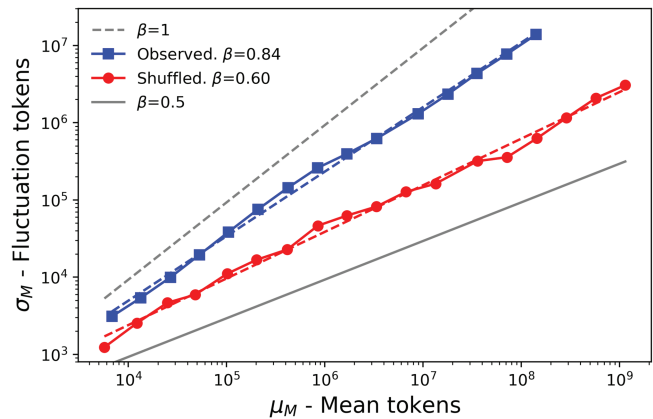


FIG. 4. Fluctuations and average number of tokens scale non-linearly as described by Taylor's law (6). The average μ_M (x axis) and standard deviation σ_M (y axis) of the hashtag data—■ with (blue) line—were obtained for intervals ranging from $b = 1$ min to $b = 256$ days. The results obtained after shuffling the temporal order of $M(t)$ and $N(t)$ are shown as ● with (red) line. The shuffling performed at the scales of $b = 1$ min, $b = 1$ h, and $b = 1$ day all showed the same scaling with a different pre-factor. The plot shows a combined curve obtained after re-scaling the curves for $b = 1$ h and $b = 1$ day by a constant factor so that they agree with the $b = 1$ min curve. The scaling exponents β indicated in the legend were obtained by linear regression (dashed lines).

and the similarity between two distributions, \mathbf{p} and \mathbf{q} , can be quantified using the α -generalized Jensen–Shannon divergence

$$D_\alpha(\mathbf{p}, \mathbf{q}) = H_\alpha\left(\frac{\mathbf{p} + \mathbf{q}}{2}\right) - \frac{1}{2}H_\alpha(\mathbf{p}) - \frac{1}{2}H_\alpha(\mathbf{q}). \quad (8)$$

For $\alpha = 1$, we recover the usual Shannon entropy $H(\mathbf{p}) = -\sum_i p_i \log p_i$ and Jensen–Shannon divergence, which can be viewed as a symmetrized Kullback–Leibler divergence. Finally, we quantify the similarity between distributions by taking the square root of the divergence

$$J_\alpha \equiv \sqrt{D_\alpha}. \quad (9)$$

As J_α has metric properties for $0 \leq \alpha \leq 2$, it is a natural choice to measure distance. We use $\alpha = 1$ and $\alpha = 2$ to obtain different

perspectives on the dynamics of hashtags: larger values of α give more weight to high-frequency hashtags.²⁸ Moreover, the statistical estimators of J_α converge very slowly with sample size M for data with Zipfian frequency distribution^{28,29} and are better for $\alpha = 2$ when compared to the usual $\alpha = 1$.

The results obtained for our hashtag data are reported in Fig. 5 and show rich dynamics. The growth of $\langle J_\alpha \rangle$ with τ indicates that the measures J_α (9) are able to quantify the changes in hashtag frequencies we are interested in. Weekly oscillations are clearly visible in J_1 but not in J_2 , indicating that there are a large number of hashtags that are not among the top ranked ones but are used repeatedly in the same day of the week (e.g., “#MondayMotivation”). The overall growth of J_α is slowing down with τ (i.e., the change in hashtag frequencies is larger for smaller τ 's). Our main empirical finding is that this slowdown follows an orderly pattern, described by the scaling law

$$d\langle J \rangle / d\tau \sim 1/\tau^\eta, \quad (10)$$

with $\eta \gtrsim 1$. This suggests that there is no characteristic time scale for the change of hashtag frequencies in Twitter but that instead it slows down in a self-similar fashion.

VI. CONCLUSIONS

In summary, we provided a general statistical characterization of the frequency of hashtags on Twitter. We found that the frequency distribution follows a Zipfian pattern with a faster decay than a simple power law. We found that this distribution is well described by the lognormal rank-frequency distribution (2). The type-token relationship shows a scaling law characteristic of Heaps' law, with non-trivial large fluctuations around expected values that follow a fluctuation scaling relationship (Taylor's law). These large fluctuations are due to the very noisy dynamics of hashtag tokens, which shows fat-tailed fluctuations and long-range temporal correlations. We also quantified the collective dynamics due to the change in the frequency p_i of individual hashtags i using a generalized Jensen–Shannon divergence. We found that the distance between hashtag distributions separated by time τ grows with τ , showing weak oscillations (i.e., distributions at the same day of the week are more similar to each other) and that the velocity of the change decays with τ , following a newly discovered scaling law, $1/\tau^\eta$ with $\eta \gtrsim 1$.

A comparison of our findings to previous results for the frequency of words in large collections of texts is given in Table II. It reveals striking similarities but also notable differences due to the different dynamics of hashtag and word frequencies. While in texts, word tokens of the same word type cluster together, this happens in the middle of many high-frequency function words that permeate the texts with a more regular frequency. In contrast, the appearance of a (new) hashtag can trigger a large response of the usage of the same hashtag, leading to much wider fluctuations and correlations. In fact, the top ranked word in English (“the”) remains the same over centuries, showing a frequency $F_1 \approx 5\%$ that varies only slightly (between 4% and 6%) over 200 years (in the Google n-gram database). In contrast, the most frequent hashtag not only varies from day to day but also the frequency of the top ranked hashtag can vary dramatically. For instance, on the first day of our dataset




FIG. 5. Dynamics of hashtags. Panels (a) and (b) show the distance $\langle J_\alpha \rangle$ (y axis) between the hashtag usage separated by time τ (x axis). The symbols (error bars) correspond to the average (standard deviation) J_α computed over all times t for a fixed time separation τ . The α entropy in Eq. (7) was used to compute J_α in Eq. (9) for $\alpha = 1$ (a) and $\alpha = 2$ (b). Panel (c) shows the derivative $d\langle J_\alpha \rangle / d\tau$ of the curves in panels (a) and (b), revealing a scaling law between $d\langle J_\alpha \rangle / d\tau$ and τ . The maximum possible value of J_1 is $J_1 = \sqrt{\ln 2} \approx 0.83$.²⁸

TABLE II. Statistical laws for the frequency of hashtags in Twitter and for the frequency of words in texts.

	Words in texts (English)	Hashtags in Twitter
Zipf-like decay of frequency F_r	Yes, faster than r^{-1}	Yes, faster than r^{-1}
Best generalized Zipf's law	Double power law ^{21,23} $\begin{cases} r^{-1} & r \leq a \\ a^{\gamma-1} r^{-\gamma} & r > a \end{cases}$	Lognormal $Cr^{-1} \exp(-\frac{1}{2}(\ln(r) - m)^2/s^2)$
Heaps' law $M \sim N^\beta$	$\beta \in [0.52, 0.62]$ ²³	$\beta = 0.73$
Taylor's law $\mu \sim \sigma^\lambda$	$\lambda = 1$ ²⁵	$\lambda = 0.84$
Dynamics J_α	Linear growth over centuries, constant dJ/dt ²⁸	Sub-linear growth over months, $dJ/dt \sim 1/\tau^n$

(01-11-2015), the top ranked hashtag was “#pushawardskathniels” with a frequency of 4.3%, while the hashtag “#mtvstars” was ranked 143rd with a frequency of 0.05%. Two weeks later, the hashtags “#pushawardskathniels” and “#mtvstars” were ranked tenth and first, respectively, with frequencies of 0.7% and 10.9%.

A number of our statistical observations are similar to observations reported in isolation in earlier work, such as the burstiness of hashtags and high variability between hashtag volumes,³⁰ the appearance of fat-tailed distributions in the frequency of hashtags,⁷ and the steady evolution of social-media language with time.³² With the combined statistical laws articulated here, we hope to provide a framework for generative models to be compared with. Our findings provide statistical results that (modifications of) existing mechanistic models of social dynamics,^{4,20} language,^{22,23} and Twitter^{11,12} should reproduce. Next steps could be to verify in which extent previous models are able to reproduce our observations and to look in more detail at the nature of the hashtag evolution, e.g., to clarify whether certain types (sub-populations) of hashtags lead to different statistical features or whether the nature of hashtag usage changes more broadly at longer timescales.

ACKNOWLEDGMENTS

We thank Martin Gerlach for sharing the code that was used in the generalized Zipf's law analysis. H.H.C. was funded by a Denison fellowship, and E.G.A. and T.J.A. were funded by the CTDS-Incubator Scheme (Grant No. G5121), both from The University of Sydney. D.F.M.O. was supported by the ARL through ARO Grant No. W911NF-16-1-0524. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation here on.

DATA AVAILABILITY

The data that support the findings of this study are openly available in Zenodo at <https://doi.org/10.5281/zenodo.3673744>, Ref. 34 and <https://doi.org/10.5281/zenodo.3842680>, Ref. 35.

REFERENCES

- Hurlock and M. L. Wilson, “Searching Twitter: Separating the Tweet from the Chaff,” in *Proceedings of the Fifth International AAAI Conference on Weblogs and Social Media (AAAI)*, (2011).
- M. Zappavigna, “Searchable talk: The linguistic functions of hashtags,” *Soc. Semiotics* **25**, 274 (2015).
- M. Zappavigna, *Searchable Talk: Hashtags and Social Media Metadiscourse* (Bloomsbury Academic, London, 2018).
- J. P. Bagrow and L. Mitchell, “The quoter model: A paradigmatic model of the social flow of written information,” *Chaos* **28**, 075304 (2018).
- M. E. J. Newman, “Power laws, Pareto distributions and Zipf's law,” *Contemp. Phys.* **46**, 323 (2005).
- M. Mitzenmacher, “A brief history of generative models for power law and log-normal distributions,” *Internet Math.* **1**, 226 (2004).
- E. Cunha, G. Magno, G. Comarella, V. Giovanni, M. A. Gonçalves, and F. Benvenuto, “Analyzing the dynamic evolution of hashtags on Twitter: A language-based approach,” in *Proceedings of the Workshop on Languages in Social Media (Association for Computational Linguistics)*, (2011), p. 58.
- M. Naaman, H. Becker, and L. Gravano, “Hip and trendy: Characterizing emerging trends on Twitter,” *J. Am. Soc. Inform. Sci. Tech.* **62**, 902 (2011).
- D. M. Beskow, S. Kumar, and K. M. Carley, “The evolution of political memes: Detecting and characterizing internet memes with multi-modal deep learning,” *Inf. Process. Manag.* **57**, 102170 (2020).
- D. M. Romero, B. Meeder, and J. Kleinberg, “Differences in the mechanics of information diffusion across topics: Idioms, political hashtags, and complex contagion on Twitter,” in *Proceedings of the 20th International Conference on World Wide Web (Association for Computing Machinery)*, (2011), p. 695.
- J. P. Gleeson, K. P. O'Sullivan, R. A. Baños, and Y. Moreno, “Effects of network structure, competition and memory time on social spreading phenomena,” *Phys. Rev. X* **6**, 021019 (2016).
- M. De Domenico and E. G. Altmann, “Unraveling the origin of social bursts in collective attention,” *arXiv:1903.06588* (2019).
- L. Weng, A. Flammini, A. Vespignani, and F. Menczer, “Competition among memes in a world with limited attention,” *Sci. Rep.* **2**, 335 (2012).
- K. Lerman, R. Ghosh, and T. Surachawala, “Social contagion: An empirical study of information spread on Digg and Twitter follower graphs,” *arXiv:1202.3162* (2012).
- N. Daniele and C. Castellano, “Analytical study of quality-biased competition dynamics for memes in social media,” *Europhys. Lett.* **122**, 28002 (2018).
- M. Tremayne, “Anatomy of protest in the digital era: A network analysis of Twitter and occupy wall street,” *Soc. Mov. Stud.* **13**, 110 (2014).
- K. G. Kapanova and S. Fidanova, “Generalized nets: A new approach to model a hashtag linguistic network on Twitter,” in *Advanced Computing in Industrial Mathematics: 12th Annual Meeting of the Bulgarian Section of SIAM, Sofia, Bulgaria, 20–22 December 2017, Revised Selected Papers*, edited by K. Georgiev, M. Todorov, and I. Georgiev (Springer International Publishing, 2019), p. 211.
- H. Bingol, “Fame emerges as a result of small memory,” *Phys. Rev. E* **77**, 036118 (2008).

- ¹⁹J. Ratkiewicz, S. Fortunato, A. Flammini, F. Menczer, and A. Vespignani, "Characterizing and modeling the dynamics of online popularity," *Phys. Rev. Lett.* **105**, 158701 (2010).
- ²⁰D. F. M. Oliveira and K. S. Chan, "The effects of trust and influence on the spreading of low and high quality information," *Phys. A Stat. Mech. Appl.* **525**, 657–663 (2019).
- ²¹R. Ferrer-i-Cancho and R. V. Solé, "Two regimes in the frequency of words and the origins of complex lexicons: Zipf's law revisited," *J. Quant. Linguist.* **8**, 165 (2001).
- ²²D. Zanette and M. Montemurro, "Dynamics of text generation with realistic Zipf's distribution," *J. Q. Linguistics* **12**, 29 (2005).
- ²³M. Gerlach and E. G. Altmann, "Stochastic model for the vocabulary growth in natural languages," *Phys. Rev. X* **3**, 021006 (2013).
- ²⁴F. Font-Clos, G. Boleda, and A. Corral, "A scaling law beyond Zipf's law and its relation to Heaps' law," *New J. Phys.* **15**, 093033 (2013).
- ²⁵M. Gerlach and E. G. Altmann, "Scaling laws and fluctuations in the statistics of word frequencies," *New J. Phys.* **15**, 113010 (2014).
- ²⁶E. G. Altmann and M. Gerlach, "Statistical laws in linguistics," in *Creativity and Universality in Language*, Lecture notes in Morphogenesis (Springer, 2016), pp. 7–26.
- ²⁷K. Tanaka-Ishii and T. Kobayashi, "Taylor's law for linguistic sequences and random walk models," *J. Phys. Commun.* **3**, 089401 (2019).
- ²⁸M. Gerlach, F. Font-Clos, and E. G. Altmann, "On the similarity of symbol-frequency distributions with heavy tails," *Phys. Rev. X* **6**, 021009 (2016).
- ²⁹A. Koplenig, S. Wolfer, and C. Müller-Spitzer, "Studying lexical dynamics and language change via generalized entropies: The problem of sample size," *Entropy* **21**, 464 (2019).
- ³⁰O. Tsur and A. Rappoport, "What's in a hashtag? Content based prediction of the spread of ideas in microblogging communities," in *Proceedings of the Fifth ACM International Conference on Web Search and Data Mining* (Association for Computing Machinery, 2012), p. 643.
- ³¹Z. Eisler, I. Bartos, and J. Kertész, "Fluctuation scaling in complex systems: Taylor's law and beyond," *Adv. Phys.* **57**, 89–142 (2008).
- ³²J. Grieve, A. Nini, and D. Guo, "Analyzing lexical emergence in Modern American English online," *Eng. Lang. Linguistics* **21**, 99 (2017).
- ³³An important difference between $N(t)$ and $M(t)$ is that $M(t)$ is simply the sum of $M(t)$ of sub-intervals (of smaller size b) while for $N(t)$, this is not the case because we are interested in unique hashtags. This leads to a sub-linear relationship between types and tokens, which is investigated in the context of Heaps' law.
- ³⁴H. H. Chen, T. J. Alexander, D. F. M. Oliveira, and E. G. Altmann (2020), "Scaling laws and dynamics of hashtags on Twitter," Zenodo3673744. <https://doi.org/10.5281/zenodo.3673744>
- ³⁵E. G. Altmann and H. H. Chen (2020), "edugalt/TwitterHashtags: First public," Zenodo3842680, <https://doi.org/10.5281/zenodo.3842680>.