Neutral theory for competing attention in social networks

Carlos A. Plata^{1,2}, Emanuele Pigani¹, Sandro Azaele¹, Violeta Calleja-Solanas³, María J. Palazzi⁴,
Albert Solé-Ribalta^{4,5}, Javier Borge-Holthoefer⁴, Sandro Meloni³, and Samir Suweis^{1,6,7}

¹Dipartimento di Fisica “G. Galilei”, INFN and University of Padova, Via Marzolo 8, 35151 Padova, Italy

²Université Paris-Saclay, CNRS, LPTMS, 91405 Orsay, France

³Institute for Cross-Disciplinary Physics and Complex Systems, CSIC-UIB, 07122 Palma, Spain

⁴Internet Interdisciplinary Institute, Universitat Oberta de Catalunya, 08018 Barcelona, Catalonia, Spain

⁵URPP Social Networks, University of Zurich, Zurich, Switzerland

⁶Padova Neuroscience Center, University of Padova, Via Orus 2/B, 35131 Padova, Italy

⁷European Center of Living Technology, Dorsoduro 3911, 30123 Venezia, Italy

(Received 5 October 2020; accepted 21 December 2020; published 22 January 2021)

Online social networks have significantly changed global communication dynamics. Despite the complexity of such dynamics, emergent patterns from empirical microblogging data on hashtag usage, such as meme popularity power-law distributions, have been observed. However, going beyond empirical or simulation-based studies, highlighting and distinguishing fundamental mechanisms at the origin of such patterns are still open problems. Herein, inspired by the neutral theory of ecology, we provide an analytical formalism capable of capturing the competition for attention in online social systems and explaining several emergent patterns observed in the context of information ecosystems. Specifically, besides meme popularity distributions, we study the persistence distributions of memes within the network and the capacity of the system to sustain several coexisting diverse memes. Such properties have a clear analogy to extinction time and biodiversity in natural ecosystems. In the social network context, diversity of ideas is indeed an identifying characteristic of the “health” of the communication process. We show how our theoretical framework is successful in explaining real data collected from the Twitter online social network. The proposed framework introduced here provides both insight into the roles played by competition and network structure on meme diversity and persistence, and an applicable null model to be tested against real data.

DOI: [10.1103/PhysRevResearch.3.013070](https://doi.org/10.1103/PhysRevResearch.3.013070)

I. INTRODUCTION

An online social network (OSN) is a virtual social structure made of individuals using the Internet as a communication medium for interacting and sharing content and opinions. Online social networks allow hundreds of millions of Internet users worldwide to produce and consume content, providing access to a vast source of information on an unprecedented scale. Therefore, nowadays, OSNs constitute mainstream communication channels to interact, exchange opinions, and reach consensus.

In recent years, it has become increasingly evident that competition significantly shapes the structure and the dynamics of these information-driven platforms [1–3]: Users thrive for visibility, while memes can be thought of as entities that compete for users’ attention. Nevertheless, it is hard to disentangle the effects of limited attention from many concurrent factors [4] such as the structure of the underlying social network [5,6], the activity of users [7], the different degrees of

influence of information spreaders [8], the intrinsic quality of the information they spread [9], and the persistence of topics [10].

Through the analysis of OSNs data, emergent properties, such as the emergence of consensus [1], viral spreading [11], power-law distribution of memes popularity [12], and echo-chambers effects [13,14], have been observed even among different platforms. Nevertheless, the availability of massive streams of online data does not give *per se* theoretical insight into the complex interplay among the structure of OSNs, meme popularity, and users’ attention dynamics.

The aforementioned ubiquitous properties are signatures of the emergent simplicity characterizing complex systems. Therefore, it naturally calls for a statistical mechanics approach, i.e., the attempt to understand regularities at large scale as collective effects of the dynamics at the individual/meme scale. The bridge between these scales would also allow us to better understand how these macroscopic (systemwide) patterns are affected by the microscopic dynamics of the interacting elements (users and memes) forming the OSN.

One of the first attempts to unveil the emergence of fat-tailed power-law distributions for meme popularity starting from an interacting particle model is the work of Gleeson *et al.* [15]. They studied a stochastic model with simple microscopic rules to describe the evolution, that is, the spreading of memes

competing for the limited resource of users' attention, on a Twitter-like OSN. In real life, each user in the OSN pays attention to a finite number of memes constrained by the finite capacity of users. In the model, this picture is simplified by assuming that each user can just pay attention to a single meme, although further generalizations have already been considered [12].

The type of stochastic process introduced in [15] has been used for decades in population genetics [16] and in ecology [17]. In particular, these models, based on simple stochastic rules which neglect species (gene) fitness, represent null neutral models in which no intrinsic advantage is ascribed to a particular type of species (gene). Therefore, the fate of a species depends on its topological role in the network and on random demographic effects.

The concept of neutrality and the so-called neutral models have attracted a great deal of interest among statistical physicists [18–20]. Indeed, neutral theory has been proven to be very successful in describing ubiquitous emergent patterns, which do not depend on the specific details of the system under study [21–24].

The analogy between memes competing for attention in an OSN and species competing for resources in an ecosystem suggests that the approaches introduced in the framework of ecological neutral theories can be exploited to extract novel relevant information about the dynamics of users' attention. Different memes can be thought of as different species and, in addition to the spreading dynamics, new memes are introduced in the system by innovation events (speciation processes).

Such qualitative similarities between natural and information ecosystems have already been explored using quantitative methods. In this direction, Borge-Holthoefer *et al.* [25] have found evidence of nested structural signatures, i.e., a landmark feature in natural mutualistic assemblages [26–29], when analyzing time-resolved online communication discussions after external information shocks, e.g., breaking news. Similarly, Lorenz-Spreen *et al.* [30], using a mathematical model based on Lotka-Volterra equations, have been able to explain some empirical data patterns in different OSNs. Moreover, they have suggested that relatively fast loss of attention is driven by increasing production and consumption of content, leading to higher turnover rates and shorter collective attention for individual topics.

Following this line of research, here we propose an analytically tractable neutral theory for competing attention in OSNs. In particular, we derive equations for the evolution of users' attention to different memes. The approach starts from a master equation describing the stochastic evolution of users' attention to a meme. In the limit of large networks, the equation is well approximated by a Fokker-Planck equation which resembles the dynamics of species abundance in ecological neutral theories [31]. We thus analytically compute several quantities of interest to characterize communication dynamics such as the number of different coexisting memes, i.e., a measure of diversity of coexisting ideas in the OSN, the distribution of user attention to these memes, and the average persistence time of user attention to a meme. Moreover, to investigate how the emergent properties depend on the underlying network of user interactions, we compare our

theoretical predictions to numerical simulations of the model and finally test our approach against empirical stream data from the Twitter OSN.

The work is organized as follows. In the next section we introduce the interacting particle model describing meme propagation. In Sec. III we derive the evolution equation for competing attention in OSNs. Therein, we present both (exact) discrete and (approximated) continuous descriptions of the system, computing several relevant patterns of the OSN related to attention. As expected, the evolution of a user's attention depends on the OSN structure. Moreover, therein we compare (and thereby successfully validate) the results of the model with real data collected from Twitter. In Sec. IV we summarize and discuss the achievements of the proposed neutral model for the study of competing attention in social networks.

II. NEUTRAL MODEL FOR USER ATTENTION

Let us consider a directed OSN of N nodes where each of them represents a user. The network is solely characterized by its out-degree distribution p_k , that is, a random user has k followers with probability p_k . Each node can be thought of as the screen of the individual displaying the meme of current interest for that user. For the sake of simplicity, we assume that each screen has the capacity for only one meme, although some generalizations of this aspect are possible [12,15]. Therefore, the state of the system at time $\tilde{t}$ is given by the list of memes appearing in all the nodes.

The dynamics is introduced in discrete time steps representing the subsequent times where an action is carried out by any of the users. During each time step one node is picked at random and the user either (i) (re)tweets the meme currently on the screen to all their followers with probability $1 - \tilde{\mu}$ while their own screen remains unchanged or (ii) innovates with probability $\tilde{\mu}$, generating a brand-new meme that appears on their screen and is tweeted to all their followers. An illustrative visualization of one time step of the dynamics is shown in Fig. 1.

The configurational state of the system in a given time is described by the corresponding list of memes the users are paying attention to. We define the set of attention variables x_m as the fraction of nodes in the system which are paying attention to meme m . Consequently, the normalization constraint $\sum_m x_m = 1$ holds for all time.

Another relevant variable in the simulated OSN, as originally introduced by Gleeson *et al.* [15], is the meme popularity. The popularity n_m of meme m is defined as the number of times that meme m has been broadcast. Because of this definition, n_m is a nondecreasing function of time, which increases by one every time m is (re)tweeted. On the contrary, the evolution of attention x_m can either increase or decrease until the meme disappears from the system. Note that, in our model, once a meme disappears from the OSN, there is no mechanism to reintroduce it in the system (innovation always introduces brand-new memes). Therefore, once a meme goes extinct, its popularity remains constant for the remaining simulation time.

In the aforementioned work by Gleeson *et al.* [15], they semianalytically found, and compared to numeric simulations,

FIG. 1. Sketch of the model. Each node represents a user in the network, with the color denoting the different memes on the user's screen. The directed edges stand for the relations of followers. Specifically, our convention is that the edge starts in the node which is followed and ends in the follower. The node chosen to spread the meme is highlighted with a dotted circle. A time step is represented with its two possible outcomes, either an innovation event (with probability $\tilde{\mu}$) or a spreading event (with probability $1 - \tilde{\mu}$).

asymptotic power-law behavior for the popularity distribution. In particular, they have investigated the cases where $p_k = \delta_{k,k^*}$ (constant number of followers for all users) and $p_k \propto k^{-\gamma}$, i.e., a power-law distribution for the number of followers. In both cases, they found that meme popularity has a fat-tail distribution, but with different exponents [15]. Popularity can be measured from real Twitter data streams as the number of (re)tweets since a given hashtag appears in the OSN. Although a systematic analysis of meme popularity distributions is still lacking, there are some results suggesting that indeed they display asymptotic power-law behavior [12], thus pointing out that the proposed model is a good candidate to elucidate at least this pattern observed in empirical OSNs. In the next section we show how in fact, exploiting neutral theory, we can go beyond the meme popularity distribution and analytically study also the diversity of coexisting ideas in the OSN and the distribution of user attention to such coexisting memes, as well as the average persistence time of a user's attention to a meme.

III. NEUTRAL EMERGENT PATTERNS IN USER ATTENTION

Because of the neutrality hypothesis, we assume complete symmetry among memes (i.e., we neglect meme identities or labels) and we thus consider the system as an effective two-meme system with large but finite N . We call A the meme we are focusing on, whereas B simply represents the collection of memes that are non- A . The variable $x \equiv x_A$ describes the fraction of users that are paying attention to meme A (i.e., the meme is on their screen) and hence $x_B = 1 - x$. It comes in handy to define the total number of users paying attention to meme A as v , i.e., $v \equiv xN$.

The dynamics, as described in the preceding section, is defined in discrete time steps and obeys the Markovian property,

i.e., the transition probability to the next state depends only on the current state. Consequently, we can properly formulate the time evolution of the probability $P(v, \tilde{t})$ of having v users paying attention to meme A at the time step $\tilde{t}$ by means of a discrete master equation, which reads

$$\begin{aligned} P(v, \tilde{t} + 1) - P(v, \tilde{t}) &= \sum_{k=0}^{N-1} \sum_{j=0}^k \{ W_{k,j}^{A,s}(v-j)P(v-j, \tilde{t}) \\ &+ W_{k,j}^{A,i}(v+j+1)P(v+j+1, \tilde{t}) \\ &+ W_{k,j}^B(v+j)P(v+j, \tilde{t}) \\ &- [W_{k,j}^{A,s}(v) + W_{k,j}^{A,i}(v) + W_{k,j}^B(v)]P(v, \tilde{t}) \}. \end{aligned} \quad (1)$$

In the master equation (1) we have taken into account that there are three main different classes of transformations with different transition probabilities W . Namely, the node selected to act in a particular time step may be either A or B ; this is indicated by the first superscript in the transition probabilities W . The second superscript denotes the type of event, i.e., $\{s\}$ if the selected node spreads the current meme with probability $1 - \tilde{\mu}$ or $\{i\}$ if it first innovates with probability $\tilde{\mu}$ and then spreads the new meme. Note that, since B is defined as non- A , innovation of B generates also a B meme and thus there is no need to differentiate between spread or innovation when the selected node contains a B meme. The subscripts give information about the connected nodes of the selected node (its followers). We define $k \in \{0, \dots, N-1\}$ as the number of followers of the selected node, whereas $j \in \{0, \dots, k\}$ is the number of followers having a meme on their screen that is different from the meme spreading from the i node. Thus, the dependence on the properties of the network is carried in the specific form of the transition rates. The detailed form of the transition rates W and the physical meaning of the master equation (1) are provided in Appendix A.

Assuming that the innovation rate scales as $\tilde{\mu} = \mu/N$ and defining the continuous-time variable $t = \tilde{t}/N^2$, it is possible to perform a diffusive approximation of the master equation (1). The Kramers-Moyal expansion [24,32] yields the Fokker-Planck equation

$$\partial_t P(x, t) = -\partial_x [A(x)P(x, t)] + \frac{1}{2} \partial_x^2 [B(x)P(x, t)], \quad (2)$$

with drift and diffusion coefficients

$$A(x) = -ax, \quad a = \mu(1 + \langle k \rangle), \quad (3a)$$

$$B(x) = bx(1-x), \quad b = (1 + \langle k \rangle)\langle k \rangle + \sigma_k^2, \quad (3b)$$

respectively. The drift term is negative and drives the memes to extinction with a rate that increases with the innovation rate and the mean number of followers in the network $\langle k \rangle = \sum_k k p_k$. The dependence of the diffusion term on $x(1-x)$ [33] indicates that there is a limit (a carrying capacity, in the language of ecology) in the attention that both competing memes can collect, as we are considering a finite number of users in the system. The constant b increases the (demographic) fluctuations with $\langle k \rangle$ and the variability of the network through the variance $\sigma_k^2 = \sum_k (k - \langle k \rangle)^2 p_k$. A detailed derivation of the Fokker-Planck equation (2) associated with the master equation (1) is provided in Appendix A.

Due to the master equation (1) and the Fokker-Planck equation (2), it is possible to study the emergent patterns in users' attention in the proposed OSN neutral model.

A. Mean persistence time for attention

The dynamics of the users' attention has an absorbing boundary at $x = 0$, i.e., the attention to a given meme will eventually go to zero. The time that a meme persists in the OSN receiving attention is a relevant quantity to study to better understand communication dynamics, as it gives insight into the virality of that meme, i.e., the longer it remains "active" the more probable it may go viral. In the ecological context, species persistence times [23] or lifetimes [34] have been thoroughly studied in different ecological neutral models and proved to be able to describe and thus elucidate the ecological mechanisms behind the power-law shape (with exponential cutoff) of species persistence patterns observed in real ecosystems [35]. Due to the stochastic nature of the neutral dynamics of the memes, these persistence times are random variables identically distributed among memes. Herein, we focus on the computation of the mean persistence time of a meme and its variance.

Previously, we derived the forward evolution equations that govern the probability density function $P(x, t)$ of the attention. The probability density function $Q(x_0, t)$ of reaching the absorbing boundary at zero at time t , departing from initial attention x_0 , obeys the backward Fokker-Planck equation [36]

$$\partial_t Q(x_0, t) = A(x_0) \partial_{x_0} Q(x_0, t) + \frac{1}{2} B(x_0) \partial_{x_0}^2 Q(x_0, t), \quad (4)$$

with the same A and B defined in Eq. (3).

Multiplying (4) by t and integrating over all times, one obtains the differential equation that governs the mean persistence time $\tau(x_0)$ to reach zero departing from x_0 . That equation can be analytically solved with the proper boundary conditions (see Appendix B for details), giving the result

$$\begin{aligned} \tau(x_0) = & - \left\{ (1 - x_0) {}_2F_1 \left(1, \frac{2a}{b}; 1 + \frac{2a}{b}; 1 - x_0 \right) \right. \\ & \left. + \frac{2a}{b} \left[\gamma + \psi \left(\frac{2a}{b} \right) + \ln x_0 \right] \right\} \left[a \left(1 - \frac{2a}{b} \right) \right]^{-1}, \end{aligned} \quad (5)$$

where ${}_2F_1$ stands for the ordinary hypergeometric function, γ is the Euler-Mascheroni constant, ψ denotes the digamma function, and the pair $\{a, b\}$ is defined in Eq. (3).

Having an analytical form for $\tau(x_0)$ allows us to study how the persistence of attention varies for the OSN's innovation rate and structure. We find that τ is a monotonically decreasing function of μ , $\langle k \rangle$, and σ_k^2 , taking the rest of the parameters as fixed.

We find that, as expected, a higher innovation rate favors the demise of living memes, and thus decreases the average persistence time of the meme's attention. Similarly, higher connectivity of the OSN, measured by the average network degree, also leads, on average, to a faster decay of attention to existing memes. This result can be understood by observing that large $\langle k \rangle$ increases the diffusivity of the memes spreading in the systems, thus accelerating the time at which the attention to the meme will reach zero. Finally,

we find that for fixed average connectivity, the heterogeneity of the network, as measured by σ_k^2 , disfavors, on average, long persistent attention to a given meme. In other words, we find that the presence of large hubs (influencers) coexisting with many poorly connected nodes (users) drives the communication dynamics towards a state dominated by fast, short, persistent tweets, which disappear relatively quickly from the OSN.

Note that (5) is an approximated result, since it comes from the diffusive continuum limit. To go beyond this approximation, we write the backward master equation for the discrete system, describing the evolution of the probability $Q(v_0, \tilde{t})$ for a meme to reach zero users' attention at time $\tilde{t}$, if at the initial time the meme has v_0 users following it. This equation reads

$$\begin{aligned} Q(v_0, \tilde{t} + 1) - Q(v_0, \tilde{t}) &= \sum_{k=0}^{N-1} \sum_{j=0}^k \{ W_{k,j}^{A,s}(v_0) Q(v_0 + j, \tilde{t}) \\ &+ W_{k,j}^{A,i}(v_0) Q(v_0 - j - 1, \tilde{t}) \\ &+ W_{k,j}^B(v_0) Q(v_0 - j, \tilde{t}) \} - Q(v_0, \tilde{t}). \end{aligned} \quad (6)$$

Multiplying by $\tilde{t}$ the above backward master equation and summing over all times, we obtain a system of linear equations whose solution is the vector of mean persistence times $\tilde{\tau}$, whose components are $\tilde{\tau}(v_0)$ with $v_0 \in \{1, \dots, N\}$. This equation, written in matrix form, yields

$$\tilde{\tau} = \mathbb{M} \tilde{\tau} + \mathbf{1}, \quad (7)$$

where the components of the matrix $\mathbb{M}$ are appropriate combinations of the transition probabilities W (see Appendix C for the details) and $\mathbf{1}$ is an N -dimensional vector of ones. Therefore, the exact result for the mean persistence times is given by

$$\tilde{\tau} = (\mathbb{I} - \mathbb{M})^{-1} \mathbf{1}, \quad (8)$$

with $\mathbb{I}$ the identity matrix in dimension N . Furthermore, one can obtain higher-order moments using (6), multiplying it by powers of $\tilde{t}$, and summing over all times. Specifically, we have done so in order to obtain the meme attention variance $\sigma_{\tilde{\tau}}^2 = \langle \tilde{t}^2 \rangle - \tilde{\tau}^2$. The calculation of the mean square attention persistence time results

$$\langle \tilde{t}^2 \rangle = (\mathbb{I} - \mathbb{M})^{-1} (2\tilde{\tau} - \mathbf{1}). \quad (9)$$

Detailed calculations are provided in Appendix C.

In Fig. 2 we compare the results obtained from numerical simulation of the model regarding the persistence time of memes with our theoretical predictions. We have done simulations for a system with $N = 10^3$ users. The two OSN structures used throughout this work to study the role of the network structure are (i) homogeneous networks, i.e., $p_k = \delta_{k,k^*}$ with constant k^* and thus all users having the same number of followers, and (ii) scale-free networks with $p_k \propto k^{-\alpha}$ for $k \geq k_{\min}$. For further details on the simulation algorithm, see Appendix E.

The value $k^* = 10$ has been chosen to avoid falling in a regime very far from the diffusive approximation. As can be seen in Fig. 2(a), the exact results perfectly agree with the simulations for both mean value τ and standard deviation

FIG. 2. Mean persistence time for meme attention for both (a) homogeneous and (b) scale-free networks. The simulations (circles) agrees perfectly with the exact solution (blue line) given by (8) in both panels. Moreover, the diffusive approximation in Eq. (4) (red line) gives also a quite good estimation in the homogeneous network, whereas that is not the case in the scale-free network. We have considered a system with $N = 10^3$, an out-degree distribution $p_k = \delta_{k,10}$, and $\mu = 15$ in (a), while we have used $p_k \propto k^{-2.5}$ for $k \geq 4$ and $\mu = 100$ in (b). In the insets, the variance of the persistence time shows perfect agreement with the exact solution (9).

σ_τ . Remarkably, in spite of the diffusive approximation, the prediction given by (5) still gives a very good estimation of the persistence time of the memes. This relative success of the diffusive approximation is reasonable in a system where there are not big chances of having burst events. Nevertheless, we expect the diffusive theory to fail when the distribution allows such burst events. In fact, if k^* is very large, then in a single step it is likely that most of the users in the network change their attention, and thus the dynamics regime for the attention is bursty rather than diffusive. This is the case, for example, in scale-free OSNs, where p_k presents long tails, i.e., few nodes have very large degree centrality.

Indeed, the most interesting case is when users have a distribution that is strongly uneven as in real OSNs. To investigate the impact of such heterogeneous structure on the dynamics, we have considered a power-law out-degree distribution which also has been found empirically [12] in some OSNs. More specifically, we assume that p_k is zero for $k < k_{\min}$ and $p_k \propto k^{-\alpha}$ for $k \geq k_{\min}$ [15]. In particular, in the following we consider $k_{\min} = 4$ and $\alpha = 2.5$ so that we have an average number of followers $\langle k \rangle \approx 10$ close to the previous case $k^* = 10$, but with much higher variance. In Fig. 2(b) we observe clearly that the result of the exact approach using the backward master equation matches accurately the values found in the numerical simulations. The diffusive approximation instead, as expected because of the presence of bursting events, fails to quantitatively describe the observed patterns.

B. Diversity patterns in information ecosystems

Biodiversity in natural ecological systems plays a fundamental role in preserving the ecosystem's functioning and health. Similarly, diversity of memes is a crucial aspect if we want to maintain “healthy” information in social systems, where a plurality of ideas circulates, in opposition to echo-chambered low polarized amplification of ideologies in the OSN.

In this section we focus specifically on two different aspects of meme diversity: the number of different memes coexisting in the system and how user attention is distributed among such memes, i.e., how memes are distributed on users' screens.

1. Number of coexisting memes

The capacity of attention of users in an OSN is finite and affects meme spreading [9,37,38]. Therefore, a given network cannot sustain an arbitrarily large number of memes. Interestingly, due to both neutrality and the assumption of a constant innovation rate (μ), it is possible to exactly calculate the probability P_S of having S different memes coexisting in the OSN [39]. In fact, the stationary solution for P_S , reached in the long time limit, is a Poisson distribution

$$P_S = \frac{e^{-\langle S \rangle} \langle S \rangle^S}{S!}, \quad (10)$$

with an average number of memes given by

$$\langle S \rangle = \tilde{\mu} \sum_{v_0=1}^N p_{v_0-1} \tilde{\tau}(v_0), \quad (11)$$

which, for the particular case of a δ -peaked out-degree distribution $p_k = \delta_{k,k^*}$, simply reduces to $\langle S \rangle = \tilde{\mu} \tilde{\tau}(k^* + 1)$. Interestingly, the Poisson shape of the distribution of the meme diversity is “universal” regardless of the details of the underlying network. The specific effect of the network structure is in determining the average number $\langle S \rangle$ of the coexisting memes in the systems.

This result was originally derived in the ecological context for a coarse-grained birth-death model for species diversity [39]. However, in our model a meme may appear with any v_0 initial number of users paying attention to it, with a probability given by the out-degree distribution p_{v_0-1} . We can thus see how the diversity of ideas circulating in the network is promoted by long mean persistence times, and thus scale-free networks are detrimental to a plurality of ideas in the OSN and promote strong polarization (and visibility) of few ideologies.

FIG. 3. Histogram of the number of different memes present in the whole system for both (a) homogeneous and (b) scale-free networks. Numerical simulations agree with the theoretically predicted Poisson distributions, especially for the homogeneous network. Note that there is no fitting parameter in the distribution since the mean is given by (11). We have considered the same model parameters as those in Fig. 2. The inset shows the mean number of memes (species) $\langle S \rangle$ for different values of the innovation rate μ . Apart from μ , the rest of the parameters are the same as those in the main figure.

Note that, equivalently, we can formulate Eq. (11) using the framework of continuous time. It simply reads $\langle S \rangle = N\mu \int_0^1 dx_0 \tau(x_0) p(x_0)$, where we have used that, in the continuous description, innovation is a Poisson process with rate $N\mu$ and $p(x_0)$ is the continuous counterpart of the out-degree distribution, giving the initial fraction of users' attention when the meme appears in the OSN. Consistently, the fraction $\langle S \rangle / N$ depends only on the parameters of the continuous model.

In Fig. 3 we have compared the theoretically predicted Poisson distribution with the histogram obtained in numerical simulations for long times in the systems studied above, that is, we consider $N = 10^3$ and either $\mu = 15$ and $p_k = \delta_{k,10}$ or $\mu = 100$ and $p_k \propto k^{-2.5}$ for $k \geq 4$. Moreover, in the homogeneous network, we have considered in the inset different values for the innovation rate μ to study the dependence of the average number of memes. On the one hand, we obtain perfect agreement when we use the exact calculation for $\bar{\tau}$ as given by (8), especially for the homogeneous network. A Poisson distribution explains quite successfully the pattern observed in the scale-free network, although the quality of the agreement is slightly lower, but still quite good. On the other hand, although we did not expect quantitative agreement (because of the bursty nature of the communication dynamics), also the prediction provided by the diffusive approximation (5) remarkably shares the same qualitative behavior. As expected, both theories and simulations converge to $\langle S \rangle = 1$ for the limit case $\mu \rightarrow 0$ since, in the absence of innovation, fluctuations eventually drive the system to the absorbing state and thus there is only one meme surviving at the end.

Finally, motivated by the good agreement between theory and numerical simulations, we decided to test our predictions in a real scenario against data from the Twitter OSN. To do so, we collected samples of Twitter activity freely available from the Twitter Streaming application programming interface (API), spanning a period of four years between 2015 and 2018 and comprising more than 135×10^6 tweets (see Appendix F for details on the data collection).

For the comparison, we assume that each meme in the model represents a hashtag on Twitter. Its popularity is traceable through the number of times it has been observed in the data set. However, given obvious privacy restrictions in the data that allow us to see when a hashtag has been tweeted but not when it appeared on a screen, some quantities such as the attention paid to a hashtag at a given time cannot be directly measured. For these reasons we decided to limit the comparison only to observables that can be estimated directly from the data, such as the distribution P_S of different memes present in the system and the mean number of coexisting memes $\langle S \rangle$ as a function of $\bar{\mu}$ (see Appendix F for the estimation of $\bar{\mu}$ and S from the data).

In Fig. 4 we show information analogous to that presented in Fig. 3, but using the data extracted from Twitter. In this case, since each sample in the data has a different $\bar{\mu}$ and different size, the theoretical distribution (blue line in Fig. 4) has been obtained as a weighted average of different Poisson distributions given by Eq. (10), one for each value of $\bar{\mu}$ observed in the data (see Appendix F). In the inset of Fig. 4, we repeated exactly the same analysis done for Fig. 3. Remarkably, given the correlations and possible biases in the data that could contradict the assumptions of our model, the good agreement demonstrates that, at least at the macroscopic scale, such emergent human communication patterns in OSNs can be explained successfully by the proposed neutral model.

2. Relative meme attention

Equations (10) and (11) characterize the diversity of memes at stationarity, but give no information on the relative attention that users give to such memes. We are thus interested in this new pattern that can be described by the probability P_{RMA} that a meme receives attention from a fraction x of users in the OSN. We call this pattern relative meme attention (RMA), which is completely analogous to the relative species abundance (RSA) pattern in ecology [18,21,24]. This quantity is clearly relevant in communication dynamics, as it allows

FIG. 4. Comparison between theoretical predictions and Twitter data (without a fit). The histogram of the normalized number of different memes present in the whole system is obtained considering the full data set. For the Twitter data S has been evaluated as the number of active hashtags at the end of each time slice. The theoretical distribution (blue line) has been calculated as a weighted average of Poisson distributions, each for the value of $\bar{\mu}$ observed in the data. The inset shows the mean number of memes (species) divided by the systems size $\langle S \rangle / N$ for different values of the innovation rate $\bar{\mu}$. For the theoretical curve, we have assumed a power-law out-degree distribution with exponent 2.5 as the one presented in this work.

one to understand how users' attention is distributed among memes. For example, when the RMA is peaked around a specific value, each meme attracts attention from a characteristic number of users. On the other hand, if RMA is fat tailed, then users' attention is extremely unevenly distributed among tweets, with some meme attracting attention of the majority of users.

Let us define $\phi(x)$ as the density function of the average number of memes receiving attention x . Note that the only difference between $\phi(x)$ and $P_{\text{RMA}}(x)$ is the normalization

constant, the latter being normalized to unity and the former to the total number of different memes S . In the stationary state, the memes that contribute to the RMA are those which have been generated by innovation a time t before and they are still present in the OSN. They thus contribute to $\phi(x)$ an amount $\int_0^1 dx_0 p(x_0) P(x, t|x_0)$, which is integrated over all possible initial configurations x_0 , weighted by the out-degree distribution. We have explicitly included the initial condition in the solution of the forward Fokker-Planck equation (2). Since the number of species generated in a small time interval dt is simply $\mu N dt$, the RMA finally reads [18]

$$\phi(x) = \mu N \int_0^\infty dt \int_0^1 dx_0 p(x_0) P(x, t|x_0). \quad (12)$$

The application of Eq. (12) is not straightforward, since we need the full time solution for $P(x, t|x_0)$. Although the Fokker-Planck equation (2) cannot be exactly solved, we have an approximated solution that can be obtained by neglecting the quadratic term in x in the diffusion coefficient (3b), i.e., the diffusive term is simply proportional to x [31]. This approximation is particularly good if one can assume that, during the time evolution, the invasion of most of the users' screens by one single meme is unlikely.

Plugging the full transient solution of $P(x, t|x_0)$ [31] into (12) is not immediate and an explicit analytical integration is not possible. However, if x_0 is small enough, a Taylor expansion of $P(x, t|x_0)$ in x_0 is suitable and makes the integration possible. Once the integral is carried out, we just need to take into account the normalization of P_{RMA} to finally obtain

$$P_{\text{RMA}}(x) = \frac{\exp(-2\frac{a}{b}x)}{x\psi(\frac{2a}{Nb})}, \quad (13)$$

where we recall that ψ denotes the digamma function. The normalization has been imposed from $1/N$, which is the minimum nonvanishing value for an attention fraction up to infinity. Note that our assumptions legitimate the change of one by infinity in the upper bound. For a detailed derivation of (13), see Appendix D.

FIG. 5. Relative meme attention for both (a) homogeneous and (b) scale-free networks. The probability $P_{\text{RMA}}(x)$ of finding a meme which receives a fraction x of attention is plotted. The theoretical prediction (red line) given by (13) reproduces remarkably well the tail of the pattern obtained by numerical simulations (circles). We have considered the same model parameters as in Fig. 2.

We find that the outcome of our prediction for the RMA is a logarithmic-series distribution, rekindling a classical result for the RSA in ecological systems [40]. Such RMA displays a power-law behavior $p(x) \sim x^{-\beta}$ with an exponent $\beta = 1$. Therefore, as expected, our model predicts that users' attention is extremely heterogeneously distributed among memes, with most of them attracting a negligible number of users, with a few (the viral ones) catching the attention of almost all of the OSN's users.

We compare the logarithmic series with the relative meme attention obtained from numerical simulations of the interacting particle model in Fig. 5. As before, the results for both homogeneous and scale-free networks are shown. On the one hand, in spite of the strong approximation we have carried out, the logarithmic series reproduces very well the simulated pattern for the homogeneous network. On the other hand, as expected, we find that the approximated diffusive theory cannot reproduce such a simulated pattern, which is the result of a bursty dynamics. Interestingly, in the latter case P_{RMA} is compatible with a power law with exponent 1.5 in a broad range of attention regime.

IV. DISCUSSION AND CONCLUSIONS

Having high diversity in an information ecosystem is just as important as having a rich biodiversity in natural ecosystems. In fact, diversity of ideas in OSNs is a crucial aspect for the quality of the deliberative process [41]. Online spaces dominated by one or few visions, i.e., with very low biodiversity, represent biased information ecosystems where phenomena such as fake news, filter bubbles, and echo chambers crystallize beliefs and annihilate diverse opinions [42]. Therefore, being able to characterize diversity patterns in information ecosystems is not only a theoretically intriguing problem, but also an important aspect to measure their health.

In this work we have proposed an analytically tractable neutral theory to describe the dynamics of user attention to competing memes in OSNs. In particular, we have shown that we are able to compute several quantities of interest such as the number of coexisting memes, the distribution of user attention among these memes, and the average persistence time for attention to a meme. All these emergent properties have an ecological analogy in natural systems, suggesting that an ecological approach to study information ecosystems can provide an opportunity for understanding the dynamics of memes in OSN.

By comparing our theoretical predictions to numerical simulations of the model, we have shown that the continuous approximation provides accurate results when the dynamics of the user attention x is diffusive rather than bursty. This in turn depends on the underlying architecture of the user-user network. In fact, if the network is characterized by a scale-free degree distribution, then the diffusive approximation fails to quantitatively reproduce the biodiversity patterns, although the qualitative behavior is still well described. Using a backward semianalytical master-equation approach allows us to overcome this limitation and to correctly predict the mean number of different coexisting memes and their mean persistence attention time. We note that having an analytical theory to calculate such quantities is of paramount impor-

tance, as it allows us to easily understand the effect of the system parameters (such as the network connectivity or the innovation rate) on the persistence of active memes in the OSN and their diversity. Moreover, we note that, because of very strong finite-size effects and fluctuations in the dynamics, simulations may lead to wrong conclusions especially when considering (as they typically are) large and heterogeneous systems.

Just as for the numerical simulations, the excellent agreement between the theory and the patterns extracted from the Twitter samples has important implications, both theoretical and practical. First of all, it indicates that our framework can provide information on the status of real information ecosystems. From few samples, it could be possible to estimate the plurality and health of online discourse. Moreover, our results demonstrate that OSNs can be studied through the lens of neutral theory. Such a result draws a direct connection between information and a natural ecosystem and allows the use of all the machinery developed in theoretical ecology [24] for the study of online human interactions.

We would like to stress that the success of our theory in explaining and reproducing real patterns does not translate into the claim that our model is capturing all the different mechanisms driving social interactions and meme spreading. Instead, our work shows that we have put forward a simple enough model to explain and reproduce successfully some real emergent patterns in OSNs, hinting that for such properties the dominating driver is competition.

Finally, we note that, differently from popularity measures, user attention to competing memes is not a feature that can be directly measured from data, but can only be evaluated through some proxy. In this regard, a future prospect for this work would be to connect the mean persistence attention time and RMA to actual measurable proxies. Doing so, we would also understand if and when the neutrality of the dynamics is broken. For instance, strong external events, such as breaking news, may have an impact on the attention dynamics that cannot be described in terms of demographic stochasticity, but calls for incorporating in this framework also environmental noise [43] and non-neutral effects [44].

ACKNOWLEDGMENTS

This project was funded by the Cariparo Foundation Visiting Program. S.S. and C.A.P. acknowledge UNIPD STARS_Stg grant BioReACT. S.M. and V.C.-S. also acknowledge support from the Spanish Ministry of Science and Innovation, the AEI, and FEDER (EU) through PACSS (Grant No. RTI2018-093732-B-C22) and the Maria de Maeztu program for Units of Excellence in R&D (Grant No. MDM-2017-0711). M.J.P., A.S.-R., and J.B.-H. acknowledge support from the Spanish MICINN Project No. PGC2018-096999-A-I00. M.J.P. acknowledges as well support from the Universitat Oberta de Catalunya through a doctoral grant. S.S. acknowledges DFA and UNIPD for SUWE_BIRD2020_01 grant. C.A.P. acknowledges the support provided by grant No. ANR-18-CE30-0013 from the Agence Nationale de la Recherche. We acknowledge CloudVeneto computing service, INFN, Padova section.

APPENDIX A: TRANSITION PROBABILITIES AND DIFFUSIVE APPROXIMATION

We put forward explicitly, in this Appendix, the functional form of the transition probabilities assumed in our model of an OSN. As explained in the main text, there are three different classes of transformations that correspond to (i) the spread of a node that holds meme A and has k followers, j of the which hold meme B , where the probability of such a transition in a system with ν nodes carrying meme A is

$$W_{k,j}^{A,s}(\nu) = \frac{\nu}{N}(1 - \tilde{\mu})p_k \frac{\binom{\nu-1}{k-j}\binom{N-\nu}{j}}{\binom{N-1}{k}}; \quad (\text{A1})$$

(ii) the innovation of a node that holds meme A and has k followers, j of the which hold meme A , where the probability of such a transition in a system with ν nodes carrying meme A is

$$W_{k,j}^{A,i}(\nu) = \frac{\nu}{N}\tilde{\mu}p_k \frac{\binom{\nu-1}{j}\binom{N-\nu}{k-j}}{\binom{N-1}{k}}; \quad (\text{A2})$$

and (iii) the action of a node that holds meme B and has k followers, j of the which hold meme A . The probability of such a transition in a system with ν nodes carrying meme A is

$$W_{k,j}^B(\nu) = \frac{N - \nu}{N}p_k \frac{\binom{\nu}{j}\binom{N-\nu-1}{k-j}}{\binom{N-1}{k}}. \quad (\text{A3})$$

All expressions are obtained as the product of four probability factors. The first one is the probability of choosing the node that acts. The second one is the probability of the type of event being either spread or innovation. The third one is the probability of having k followers. The last one is the hypergeometric probability coming from the distribution of the possible states of the k followers. Note that the “second” factor in Eq. (A3) is a factor 1 that is not explicitly written.

The master equation (1) is simply a balance of gained and lost probability due to transitions between states. Note that $\nu - \Delta\nu$ appears as the argument of P , with $\Delta\nu$ being the change in the number of users paying attention to node A .

In order to perform the diffusive approximation, we consider the limit $N \rightarrow \infty$, with $x = \nu/N$. In such a limit, we can approximate the hypergeometric distribution with a binomial one, which yields the continuous description of the transition rates

$$W_{k,j}^{A,s}(x) = x(1 - \tilde{\mu})p_k x^j (1 - x)^{k-j}, \quad (\text{A4a})$$

$$W_{k,j}^{A,i}(x) = x\tilde{\mu}p_k (1 - x)^j x^{k-j}, \quad (\text{A4b})$$

$$W_{k,j}^B(x) = (1 - x)p_k x^j (1 - x)^{k-j}. \quad (\text{A4c})$$

Carrying out an expansion of the master equation (1) in powers of N^{-1} , introducing transition probabilities as in Eq. (A4), and assuming scaling of time $t = \tilde{t}/N^2$ and innovation rate $\mu = N\tilde{\mu}$ as introduced in the main text, we achieve the diffusive approximation given by the Fokker-Planck equation (2) with the coefficient reported in Eq. (3).

APPENDIX B: BACKWARD FOKKER-PLANCK EQUATION AND LIFETIME DISTRIBUTION

We have obtained the forward Fokker-Planck equation

$$\partial_t P(x, t) = \partial_x [axP(x, t)] + \frac{1}{2}\partial_x^2 [bx(1 - x)P(x, t)]. \quad (\text{B1})$$

The backward version of this equation,

$$\partial_t Q(x_0, t) = -ax_0\partial_{x_0}Q(x_0, t) + \frac{bx_0(1 - x_0)}{2}\partial_{x_0}^2 Q(x_0, t), \quad (\text{B2})$$

governs the probability $Q(x_0, t)$ of reaching the absorbing boundary placed at $x = 0$ departing from x_0 after an evolution of time t . Consequently, the differential equation that holds for the probability of reaching the absorbing point at any time $\Pi(x_0) = \int_0^\infty dt Q(x_0, t)$ is

$$0 = -ax_0\partial_{x_0}\Pi(x_0) + \frac{bx_0(1 - x_0)}{2}\partial_{x_0}^2 \Pi(x_0), \quad (\text{B3})$$

where we have used that $Q(x_0, 0) = \lim_{t \rightarrow \infty} Q(x_0, t) = 0$. Since the only absorbing boundary is that at $x = 0$, the meaningful solution for Eq. (B3) is $\Pi(x_0) = 1$. In other words, all the memes will eventually reach extinction.

To study the mean lifetime $\tau(x_0)$ of the memes we need to multiply the backward equation (B2) by t and then integrate for all time. After assuming regular behavior for the boundary terms and taking into account that $\Pi(x_0) = 1$, we get

$$1 = -ax_0\partial_{x_0}\tau(x_0) + \frac{bx_0(1 - x_0)}{2}\partial_{x_0}^2 \tau(x_0). \quad (\text{B4})$$

Imposing the boundary conditions $\tau(0) = 0$ and $\lim_{x_0 \rightarrow 1} \tau(x_0) < \infty$, we obtain the solution (5) presented in the main text.

APPENDIX C: BACKWARD MASTER EQUATION AND LIFETIME DISTRIBUTION

The backward master equation that rules the dynamics of the probability $Q(v_0, \tilde{t})$ of reaching extinction departing from the initial attention of ν_0 after $\tilde{t}$ time steps is written in Eq. (6). The probability in the next time step is a linear combination of the probability in the current one. Therefore, we can write a vector equation of dimension $N + 1$ using matrix notation

$$\hat{Q}(\tilde{t} + 1) = \hat{\mathbb{M}}\hat{Q}(\tilde{t}), \quad (\text{C1})$$

with $\hat{Q}(\tilde{t})$ the vector with components $Q(v_0, \tilde{t})$, for $v_0 = \{0, \dots, N\}$, and $\hat{\mathbb{M}}$ the $(N + 1) \times (N + 1)$ matrix with elements

$$M_{v_0, v'_0} = \sum_{k=0}^{N-1} [W_{k, v'_0 - v_0}^{A,s}(v_0) + W_{k, v'_0 - v'_0 - 1}^{A,i}(v_0) + W_{k, v'_0 - v'_0}^B(v_0)]. \quad (\text{C2})$$

In this equation, to avoid cluttering our formulas, we assume that transition probabilities are equal to zero when the index involves a transformation with no physical meaning, e.g., the second subindex does not belong to the interval $[0, k]$. Probability conservation is reflected in the property

$$\sum_{v'_0=0}^N M_{v_0, v'_0} = 1. \quad (\text{C3})$$

The absorbing boundary condition at $v = 0$ makes that $Q(0, \tilde{t}) = \delta_{\tilde{t},0}$. Hence, in dimension N , we have that

$$\mathbf{Q}(\tilde{t} + 1) = \mathbb{M}\mathbf{Q}(\tilde{t}) + \mathbf{M}_0\delta_{\tilde{t},0}, \quad (\text{C4})$$

where $\mathbf{Q}(\tilde{t})$ is the vector $\hat{\mathbf{Q}}(\tilde{t})$ after removing the first component corresponding to $v_0 = 0$, $\mathbb{M}$ is the submatrix of $\hat{\mathbb{M}}$ obtained after removing the rows and columns corresponding to value 0, and $\mathbf{M}_0$ is a column vector that contains the transition probabilities to the extinction state.

We define the total probability of arriving at 0 departing from n as the time sum

$$\mathbf{\Pi} = \sum_{\tilde{t}=0}^{\infty} \mathbf{q}(\tilde{t}), \quad (\text{C5})$$

where the components of $\mathbf{\Pi}$ are $\Pi(v_0)$. Summing the backward master equation (C4) for all time and taking into account that $\mathbf{Q}(0) = \lim_{\tilde{t} \rightarrow \infty} \mathbf{Q}(\tilde{t}) = 0$, we get

$$\mathbf{\Pi} = \mathbb{M}\mathbf{\Pi} + \mathbf{M}_0. \quad (\text{C6})$$

Since the property (C3) holds, we have that the solution for (C6) is simply $\mathbf{\Pi} = \mathbf{1}$, where we have used the same notation for the vector of dimension N full of ones as in the main text. Consistently with our results in the continuous description and with our physical understanding, we get, in this exact framework, that all the memes will reach extinction eventually.

The mean first-passage time $\tilde{\tau}(v_0)$ to reach the absorbing point starting from the state v_0 is given by

$$\tilde{\tau} = \sum_{\tilde{t}=0}^{\infty} \tilde{t}\mathbf{Q}(\tilde{t}). \quad (\text{C7})$$

Multiplying (C4) by $\tilde{t}$ and summing for all time, we obtain the equation for the mean persistence time (7) presented in the main text with its corresponding solution (8).

It is also possible to compute the average of the square passage time, which is useful for the variance. We define

$$\langle \tilde{t}^2 \rangle = \sum_{\tilde{t}=0}^{\infty} \tilde{t}^2 \mathbf{Q}(\tilde{t}). \quad (\text{C8})$$

Multiplying (C4) by $\tilde{t}^2$ and summing for all time, we obtain

$$\langle \tilde{t}^2 \rangle - 2\tilde{\tau} + \mathbf{1} = \mathbb{M}\langle \tilde{t}^2 \rangle, \quad (\text{C9})$$

which has the solution (9) presented in the main text.

APPENDIX D: DERIVATION OF THE RELATIVE MEME ATTENTION

Our starting point here is the approximated Fokker-Planck equation after neglecting the squared term in the diffusion coefficient, that is,

$$\partial_t P(x, t|x_0) = \partial_x [axP(x, t|x_0)] + \frac{1}{2}\partial_x^2 [bxP(x, t|x_0)]. \quad (\text{D1})$$

From now on, we consider that the equation is defined in the region $x > 0$. We understand the neglect of the squared term as a rescaling of the attention that moves the boundary from $x = 1$ to infinity. The solution of (D1) submitted to an absorbing boundary at $x = 0$ and the initial condition

$P(x, t|x_0) = \delta(x - x_0)$ is [31]

$$P(x, t|x_0) = \frac{2a}{b} \frac{1}{1 - e^{-at}} \exp\left(-\frac{2a}{b} \frac{x + x_0 e^{-at}}{1 - e^{-at}}\right) \times \left(\frac{x_0}{x} e^{-at}\right)^{1/2} I_1\left(\frac{4a}{b} \frac{\sqrt{x_0 x e^{-at}}}{1 - e^{-at}}\right). \quad (\text{D2})$$

A Taylor expansion of this solution up to linear order yields

$$P(x, t|x_0) \simeq \frac{4a^2 \exp\left[-\left(at + \frac{2a}{b} \frac{x}{1 - e^{-at}}\right)\right]}{b^2(1 - e^{-at})^2} x_0. \quad (\text{D3})$$

With this approximation, valid for small x_0 , we can carry out exactly the integration for all time, with the result

$$\int_0^\infty dt P(x, t|x_0) \propto \frac{\exp\left(-2\frac{a}{b}x\right)}{x}, \quad (\text{D4})$$

where we have explicitly written the dependence with x . Since the relative meme attention P_{RMA} is proportional to the integral in Eq. (D4), we just need to impose the normalization condition. Taking into account that the minimum nonvanishing value for the attention is $1/N$ and assuming, as said above, that the upper bound is at infinity, we end up with the P_{RMA} reported in Eq. (13).

APPENDIX E: SIMULATION DETAILS

In this Appendix we cover in detail the algorithm used to carry out the numerical simulations presented in this work. We have explicitly simulated the microscopic dynamics described in the main text. Specifically, we consider a system made of N nodes. Each of them carries just one meme. The dynamics is generated by repetition of the following recipe for each time step.

- (i) We randomly choose a node with homogeneous probability $1/N$.
- (ii) With probability $\tilde{\mu}$, the node innovates the meme.
- (iii) The degree of node k is chosen from the out-degree distribution p_k .
- (iv) From the available $N - 1$ nodes, we randomly choose the k followers.
- (v) The meme of the first chosen node spreads to all the followers.

Note that the connections in the network are not fixed; we work with a random network which is dynamic. This random dynamic feature guarantees that the approach given by the master equation used in the main text is the correct mathematical description of our system.

For each set of parameters, we have run $N_S = 100$ simulations of a total duration of $\tilde{t}_f = 5 \times 10^6$ time steps. The measurements that correspond to the persistence time come from the average of all registered times for those memes which reach extinction inside the simulation window. Recording the history of all memes, it is easy to recover the times to extinction starting from all possible initial attentions. We just need to subtract the time in which the meme has that initial attention to the extinction time of the corresponding meme. We have increased the statistics of the measurements of biodiversity, which are performed in the stationary state, taking different times of observation. We have taken not only $\tilde{t}_f$, but also $\tilde{t}_f - n\Delta\tilde{t}$, with $n = 1, \dots, 10$ and $\Delta\tilde{t} = 2.5 \times 10^5$ time

steps. The choice of Δt has to guarantee some requirements. On the one hand, $\Delta \tilde{t}$ has to be big enough to prevent correlation between the different observation times. On the other hand, $\Delta \tilde{t}$ has to be small enough to ensure that the system is close enough to the stationary state in the first observation time $\tilde{t}_f - 10\Delta \tilde{t}$.

In general, in the main text we reported all values using the continuous scales, which are related to the discrete dynamics through a factor N^2 for the timescale $t = \tilde{t}/N^2$ and a factor N for the attention $x = v/N$.

APPENDIX F: TWITTER DATA

In this Appendix we describe the details of the data collection from the Twitter OSN and the analyses used in the comparison with the theoretical predictions. In order to get an unbiased sample of Twitter activity, we collected data from the Twitter Streaming API. The data collection covers a period of four years from January 2015 to December 2018. Only geolocalized tweets originated in a rectangle area covering the UK have been requested. This latter condition allowed us to avoid the 1% total traffic limitations imposed by the Streaming API and to ensure that the majority of tweets and hashtags would be in English. Finally, to protect privacy, users' information and the tweet ID have been anonymized through a hash function before their storage and, in any case, for each tweet, only the timestamp and the hashtags contained in the text have been used in the analyses. The complete data set comprises approximately 135×10^6 tweets containing at least one hashtag.

To reduce the possible noise generated by misspelled hashtags or bots, we filtered out all the hashtags that have been tweeted less than, i.e., have a minimum popularity of, five times. We have explicitly checked that our results remain almost unchanged under different filters between 3 and 500. After the filtering and with the aim of different samples for the statistical analyses, we divided the final data set in 10^4 bins/samples, each of the size of around 12 000 tweets. Moreover, to test the robustness of our results with respect to the choice of the number and size of the bins, we also repeated the calculations for 10^3 , 2×10^4 , and 5×10^4 bins along with one test in which we considered one day in the data as one bin. In all the cases, we found robust results that behave qualitatively in the same way as those presented in Fig. 4.

Once we filtered and divided the data, for each bin/sample, we estimate the innovation rate $\tilde{\mu}$ as the fraction of new

hashtags, i.e., used for the first time, observed in a bin over the total number of unique hashtags observed until that bin. Note that, therefore, the innovation rate does not remain constant. This dynamics is in contrast to our model. Nevertheless, each bin/sample can be thought of as a realization of our model with different innovation rates. For S/N , we take the number of unique hashtags surviving at the end of each bin normalized by the maximum value of $\langle S \rangle/N$ given by the theory. The histogram in Fig. 4 is made by just binning and counting these data along the S/N axis. In order to get better statistics to determine average values, we group the samples binning in $\tilde{\mu}$, where we denote by $n_{\tilde{\mu}}$ the number of samples considered within the bin with the center at $\tilde{\mu}$. Therefore, we obtain in this way an averaged number of coexistent memes $\langle S \rangle/N$ and its variance $\sigma_{S/N}^2$ for each innovation rate considered. The inset of Fig. 4 is constructed with these averages, where the size of the error bars is given by $\sigma_{S/N}$.

For the construction of the theoretical prediction in Fig. 4, we have used, in an auxiliary way, the full data set, that is, the full list of corresponding $(\tilde{\mu}, n_{\tilde{\mu}}, \langle S \rangle/N)$. Specifically, we have used a weighted combination of Poisson distributions, one for each bin of $\tilde{\mu}$ with its observed $\langle S \rangle/N$. In other words, we have defined a weight $w_{\tilde{\mu}} = n_{\tilde{\mu}} / \sum_{\tilde{\mu}'} n_{\tilde{\mu}'}$ and we have built the distribution

$$P_S = \sum_{\tilde{\mu}} w_{\tilde{\mu}} P_S^{(\tilde{\mu})} \quad (\text{F1})$$

as the linear combination of Poisson distributions $P_S^{(\tilde{\mu})}$ parametrized by the average $\langle S \rangle$ corresponding to that $\tilde{\mu}$. Note that, in order to combine the Poisson distributions and then rescale from S to S/N , the value of N is required. Therefore, we have exploited the Poisson theoretical prediction associating an effective size with the data set given by

$$N = \frac{\langle S \rangle/N}{\sigma_{S/N}^2}, \quad (\text{F2})$$

which is also averaged over bins.

For the theoretical prediction in the inset, since Twitter users' out-degree (follower) distribution is extremely fat tailed [12,45], in the model we assumed the same power-law degree distribution as that shown in the main text. We have used such parameters, also introduced in Ref. [15], in the absence of a better knowledge of the real network.

-
- [1] L. Weng, A. Flammini, A. Vespignani, and F. Menczer, Competition among memes in a world with limited attention, *Sci. Rep.* **2**, 335 (2012).
 - [2] E. Altman, P. Kumar, S. Venkatramanan, and A. Kumar, *Proceedings of the 2013 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining, Niagara, 2013* (ACM, New York, 2013), pp. 1352–1357.
 - [3] K.-K. Kleineberg and M. Boguñá, Competition between global and local online social networks, *Sci. Rep.* **6**, 25116 (2016).
 - [4] D. M. Romero, B. Meeder, and J. Kleinberg, *Proceedings of the 20th International Conference on World Wide Web, Hyderabad, 2011* (ACM, New York, 2011), pp. 695–704.
 - [5] D. Centola, The spread of behavior in an online social network experiment, *Science* **329**, 1194 (2010).
 - [6] J. Borge-Holthoefer, A. Rivero, I. García, E. Cauhé, A. Ferrer, D. Ferrer, D. Francos, D. Iniguez, M. P. Pérez, G. Ruiz *et al.*, Structural and dynamical patterns on online social networks: The Spanish May 15th movement as a case study, *PLoS One* **6**, e23883 (2011).
 - [7] K. Lerman and R. Ghosh, *Proceedings of the Fourth International AAAI Conference on Weblogs and Social Media, Washington, DC, 2010* (AAAI, Menlo Park, 2010).
 - [8] D. M. Romero, W. Galuba, S. Asur, and B. A. Huberman, in *Joint European Conference on Machine Learning and*

- Knowledge Discovery in Databases*, edited by D. Gunopulos, T. Hofmann, D. Malerba, and M. Vazirgiannis, Lecture Notes in Computer Science Vol. 6913 (Springer, Berlin, 2011), pp. 18–33.
- [9] X. Qiu, D. F. Oliveira, A. S. Shirazi, A. Flammini, and F. Menczer, Limited individual attention and online virality of low-quality information, *Nat. Human Behav.* **1**, 0132 (2017).
 - [10] S. Asur, B. A. Huberman, G. Szabo, and C. Wang, *Proceedings of the Fifth International AAAI Conference on Weblogs and Social Media, Barcelona, 2011* (AIAA, Menlo Park, 2011), pp. 434–437.
 - [11] L. Weng, F. Menczer, and Y.-Y. Ahn, Virality prediction and community structure in social networks, *Sci. Rep.* **3**, 2522 (2013).
 - [12] 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).
 - [13] M. Del Vicario, A. Bessi, F. Zollo, F. Petroni, A. Scala, G. Caldarelli, H. E. Stanley, and W. Quattrociocchi, The spreading of misinformation online, *Proc. Natl. Acad. Sci. USA* **113**, 554 (2016).
 - [14] F. Baumann, P. Lorenz-Spreen, I. M. Sokolov, and M. Starnini, Modeling Echo Chambers and Polarization Dynamics in Social Networks, *Phys. Rev. Lett.* **124**, 048301 (2020).
 - [15] J. P. Gleeson, J. A. Ward, K. P. O’Sullivan, and W. T. Lee, Competition-Induced Criticality in a Model of Meme Popularity, *Phys. Rev. Lett.* **112**, 048701 (2014).
 - [16] M. Kimura, *The Neutral Theory of Molecular Evolution* (Cambridge University Press, Cambridge, 1983).
 - [17] S. P. Hubbell, *The Unified Neutral Theory of Biodiversity and Biogeography (MPB-32)* (Princeton University Press, Princeton, 2001).
 - [18] M. Vallade and B. Houchmandzadeh, Analytical solution of a neutral model of biodiversity, *Phys. Rev. E* **68**, 061902 (2003).
 - [19] R. A. Blythe and A. J. McKane, Stochastic models of evolution in genetics, ecology and linguistics, *J. Stat. Mech.* (2007) P07018.
 - [20] M. Martinello, J. Hidalgo, A. Maritan, S. Di Santo, D. Plenz, and M. A. Muñoz, Neutral Theory and Scale-Free Neural Dynamics, *Phys. Rev. X* **7**, 041071 (2017).
 - [21] I. Volkov, J. R. Banavar, S. P. Hubbell, and A. Maritan, Neutral theory and relative species abundance in ecology, *Nature (London)* **424**, 1035 (2003).
 - [22] I. Volkov, J. Banavar, F. He, S. Hubbell, and A. Maritan, Patterns of relative species abundance in rainforests and coral reefs, *Nature (London)* **450**, 45 (2007).
 - [23] E. Bertuzzo, S. Suweis, L. Mari, A. Maritan, I. Rodríguez-Iturbe, and A. Rinaldo, Spatial effects on species persistence and implications for biodiversity, *Proc. Natl. Acad. Sci. USA* **108**, 4346 (2011).
 - [24] S. Azaele, S. Suweis, J. Grilli, I. Volkov, J. R. Banavar, and A. Maritan, Statistical mechanics of ecological systems: Neutral theory and beyond, *Rev. Mod. Phys.* **88**, 035003 (2016).
 - [25] J. Borge-Holthoefer, R. A. Baños, C. Gracia-Lázaro, and Y. Moreno, Emergence of consensus as a modular-to-nested transition in communication dynamics, *Sci. Rep.* **7**, 41673 (2017).
 - [26] J. Bascompte, P. Jordano, C. J. Melián, and J. M. Olesen, The nested assembly of plant-animal mutualistic networks, *Proc. Natl. Acad. Sci. USA* **100**, 9383 (2003).
 - [27] S. Suweis, F. Simini, J. R. Banavar, and A. Maritan, Emergence of structural and dynamical properties of ecological mutualistic networks, *Nature (London)* **500**, 449 (2013).
 - [28] S. Suweis, J. Grilli, J. R. Banavar, S. Allesina, and A. Maritan, Effect of localization on the stability of mutualistic ecological networks, *Nat. Commun.* **6**, 10179 (2015).
 - [29] C. Payrató-Borràs, L. Hernández, and Y. Moreno, Breaking the Spell of Nestedness: The Entropic Origin of Nestedness in Mutualistic Systems, *Phys. Rev. X* **9**, 031024 (2019).
 - [30] P. Lorenz-Spreen, B. M. Mønsted, P. Hövel, and S. Lehmann, Accelerating dynamics of collective attention, *Nat. Commun.* **10**, 1759 (2019).
 - [31] S. Azaele, S. Pigolotti, J. R. Banavar, and A. Maritan, Dynamical evolution of ecosystems, *Nature (London)* **444**, 926 (2006).
 - [32] N. G. Van Kampen, *Stochastic Processes in Physics and Chemistry* (North-Holland, Amsterdam, 1992).
 - [33] O. Al Hammal, H. Chaté, I. Dornic, and M. A. Muñoz, Langevin Description of Critical Phenomena with Two Symmetric Absorbing States, *Phys. Rev. Lett.* **94**, 230601 (2005).
 - [34] S. Pigolotti, A. Flammini, M. Marsili, and A. Maritan, Species lifetime distribution for simple models of ecologies, *Proc. Natl. Acad. Sci. USA* **102**, 15747 (2005).
 - [35] S. Suweis, E. Bertuzzo, L. Mari, I. Rodríguez-Iturbe, A. Maritan, and A. Rinaldo, On species persistence-time distributions, *J. Theor. Biol.* **303**, 15 (2012).
 - [36] S. Redner, *A Guide to First-Passage Processes* (Cambridge University Press, Cambridge, 2007).
 - [37] B. Gonçalves, N. Perra, and A. Vespignani, Modeling users’ activity on twitter networks: Validation of Dunbar’s number, *PLoS One* **6**, e22656 (2011).
 - [38] N. O. Hodas and K. Lerman, *Proceedings of the 2012 International Conference on Privacy, Security, Risk and Trust and 2012 International Conference on Social Computing* (IEEE, Piscataway, 2012), pp. 249–257.
 - [39] S. Suweis, A. Rinaldo, and A. Maritan, An exactly solvable coarse-grained model for species diversity, *J. Stat. Mech.* (2012) P07017.
 - [40] R. A. Fisher, A. S. Corbet, and C. B. Williams, The relation between the number of species and the number of individuals in a random sample of an animal population, *J. Anim. Ecol.* **12**, 42 (1943).
 - [41] N. Helberger, in *Digital Dominance: The Power of Google, Amazon, Facebook, and Apple*, edited by M. Moore and D. Tambini (Oxford University Press, Oxford, 2018), pp. 153–175.
 - [42] M. Del Vicario, G. Vivaldo, A. Bessi, F. Zollo, A. Scala, G. Caldarelli, and W. Quattrociocchi, Echo chambers: Emotional contagion and group polarization on Facebook, *Sci. Rep.* **6**, 37825 (2016).
 - [43] J. Hidalgo, S. Suweis, and A. Maritan, Species coexistence in a neutral dynamics with environmental noise, *J. Theor. Biol.* **413**, 1 (2017).
 - [44] P. V. Martín, J. Hidalgo, R. R. de Casas, and M. A. Muñoz, Eco-evolutionary model of rapid phenotypic diversification in species-rich communities, *PLoS Comput. Biol.* **12**, e1005139 (2016).
 - [45] V. Martha, W. Zhao, and X. Xu, *Proceedings of the 2013 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining, Niagara, 2013* (Ref. [2]), pp. 1405–1409.