

Zipf and Heaps laws from dependency structures in component systems

Andrea Mazzolini,¹ Jacopo Grilli,² Eleonora De Lazzari,³ Matteo Osella,¹
 Marco Cosentino Lagomarsino,^{3,4,5} and Marco Gherardi^{3,6,*}

¹*Dipartimento di Fisica and INFN, Università degli Studi di Torino, Via Pietro Giuria 1, 10125 Torino, Italy*

²*Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, New Mexico 87501, USA*

³*Sorbonne Universités, UPMC Univ Paris 06, UMR 7238, Computational and Quantitative Biology, 4 Place Jussieu, Paris, France*

⁴*CNRS, UMR 7238, Paris, France*

⁵*IFOM, Milan, Italy*

⁶*Dipartimento di Fisica, Università degli Studi di Milano, via Celoria 16, 20133 Milano, Italy*



(Received 2 March 2018; published 25 July 2018)

Complex natural and technological systems can be considered, on a coarse-grained level, as assemblies of elementary components: for example, genomes as sets of genes or texts as sets of words. On one hand, the joint occurrence of components emerges from architectural and specific constraints in such systems. On the other hand, general regularities may unify different systems, such as the broadly studied Zipf and Heaps laws, respectively concerning the distribution of component frequencies and their number as a function of system size. Dependency structures (i.e., directed networks encoding the dependency relations between the components in a system) were proposed recently as a possible organizing principles underlying some of the regularities observed. However, the consequences of this assumption were explored only in binary component systems, where solely the presence or absence of components is considered, and multiple copies of the same component are not allowed. Here we consider a simple model that generates, from a given ensemble of dependency structures, a statistical ensemble of sets of components, allowing for components to appear with any multiplicity. Our model is a minimal extension that is memoryless and therefore accessible to analytical calculations. A mean-field analytical approach (analogous to the “Zipfian ensemble” in the linguistics literature) captures the relevant laws describing the component statistics as we show by comparison with numerical computations. In particular, we recover a power-law Zipf rank plot, with a set of core components, and a Heaps law displaying three consecutive regimes (linear, sublinear, and saturating) that we characterize quantitatively.

DOI: [10.1103/PhysRevE.98.012315](https://doi.org/10.1103/PhysRevE.98.012315)

I. INTRODUCTION

Several physical, biological, and artificial systems are modular and can be naturally partitioned into well-defined basic components. These components can be considered as the building blocks that constitute each element of the system. For example, books can be regarded as sets of words, each word possibly occurring multiple times in the same book; likewise, operating systems such as Linux are sets of packages, and genomes are sets of genes, where each gene belongs to an evolutionary family [1–5]. In general, a system with this simple modular structure, called a “component system,” can be considered as a specific collection of partitions of elementary components. For instance, bacterial genomes are composed of genes drawn from a common pool of gene families, but each genome possesses them in different multiplicities. The same description applies to books in a linguistic corpus, where different books are written using words from the same vocabulary. These partitions, and their statistical properties, are generally an extension of the classic partitions extensively studied in probability [6] as well as in statistical mechanics, for example, related to the equilibrium

statistics of particles in energy states [7], to nonequilibrium site occupancy in driven diffusive systems [8], or resulting from stochastic processes based on duplication-innovation mechanisms [9,10].

The representation of complex systems as component systems may reveal information on the system’s generative mechanisms and architecture. Notably, a variety of quantitative laws have been discovered, which, intriguingly, are often conserved across very different systems. Prominent examples are Zipf’s law, concerning the power-law distribution of component frequencies [11–16], and the sublinear scaling of the number of different component classes with system size, often referred to as Heaps’s law [10,17–19]. The analysis of these laws has a long tradition in quantitative linguistics [4,11]. Several mechanisms of text generation, based on different hypotheses on the fundamental structures or principles of natural language, have been proposed to explain these statistical patterns [20–22]. Analogously, in genomics, the coarse-grained view of genomes as component systems reveals emergent quantitative invariants pointing to relevant underlying evolutionary and architectural properties of genomes [1,3,23–26].

Scale invariance (and universality), suggested by the presence of Zipf’s law, may be a natural consequence of criticality [27], either due to evolutionary tuning or self-organization [28,29], as is well understood via the renormalization group

*Corresponding author: marco.gherardi@mi.infn.it

in statistical mechanics. However, more and more often scale-free features in different contexts have been recognized as a possible consequence of stochastic processes based on specific features, such as preferential attachment [12,13] or a history-dependent entropic reduction of the accessible states [30,31]. The plethora of possible mechanisms that can generate these pervasive statistical patterns raises questions concerning their true origin and their robustness [32].

Typically, these proposed mechanisms do not explicitly take into account the functional specificities of single components, their functional dependencies, or the possible synergy and conflicts between components. These interactions are essential in most empirical systems. For example, the importance of this aspect is clear in operating systems (or other large software projects), where a package performs a specific function and typically requires the presence of other components for functioning. Indeed, models based on dependency structures have emerged recently as a promising framework to rationalize component dependencies and analyze their consequences in terms of component statistics [3]. Similar dependency structures also emerge in preference prediction [33] or for addressing causality in financial data [34].

More precisely, a dependency structure is a directed graph (most often, but not necessarily, acyclic), whose nodes are the components (e.g., Linux packages, or genes) and whose links are the dependency relations occurring between them (e.g., requirement constraints or regulatory pathways). A component depends on another if it is not functional unless the latter is present. A simple mechanism to build a component system compatible with a given dependency structure identifies a system realization (e.g., a genome or a book) with the choice of a node and all its direct and indirect dependencies [3]. This model allows us to link quantitative laws of the component statistics to topological properties of the dependency structure (and hence to the generative processes sculpting it). For instance, a broad ensemble of dependency structures has the property that the total number of nodes that ultimately depend on a given node has scale-free distribution in the limit of large system size. Notice that this number does not coincide with the degree of the node, since it counts both direct and indirect dependencies, whereas degree only counts the direct ones. Actually, a scale-free distribution of total dependencies is a much weaker condition than the power-law distribution of degrees (i.e., of direct dependencies), and therefore includes a much broader class of graphs than that of “scale-free networks,” as per the usual definition.

This topological property explains the fat-tailed distribution of component occurrences across realizations, both in genomes and operating systems [3]. However, since this simple model was inspired by software packages, it was developed in the restrictive case of binary (presence or absence) abundances of components. While a piece of software is either installed or absent, in most empirical component systems a component type can be present in many instances. For example, a gene is typically present in multiple copies (called paralogs) in a genome, and the same word is typically used several times in a single book.

Here, to capture this large class of component systems, we consider an extension of the model proposed in Ref. [3]

to the cases in which components can appear with arbitrary abundances. This extension is able to address the question of how dependency structures affect abundance-related features, such as Zipf’s laws and in particular Heaps’s law. The same question is not defined in the case of binary component abundances, where, for example, Heaps’s law is trivially linear.

II. MODEL

We consider a simplified description of complex functional architectures as unordered sets of modular components. Let \mathcal{U} be the set of all unique components (the *universe*), and let $U = |\mathcal{U}|$ denote its cardinality. A *realization* of this “component system” is a set $r = \{c_i\}$ of components $c_i \in \mathcal{U}$, with $i = 1, \dots, N$, where N is the *size* of the realization. The rules for constructing a realization are specified by a network of dependencies, as we explain below.

A *dependency structure* is a directed acyclic graph \mathcal{G} on \mathcal{U} , which encodes the dependencies between the components. An edge $i \rightarrow j$ between two nodes i and j represents the relation “ i depends on j .” A component i is said to depend on another component j if i is not functional without j . In empirical cases, such a relation can be more or less strict depending on the system; for instance, it is enforced in software operating systems, where a package cannot function unless all its dependencies are installed, but not in metabolic networks, where alternative pathways can be followed to the same metabolite [26,35]. The model assumes strict unbroken dependencies. Notice that acyclicity of \mathcal{G} is not stringently necessary; however, as will be clear in the following, a cycle in \mathcal{G} would behave as a single node in the model. Our convention in the direction of the arrow describing the dependence $i \rightarrow j$ may appear unnatural. We made this choice to be consistent with previous works [3], and because this is the actual direction of pointers in software systems.

The topology of the dependency structure is conceptually separated from the procedure that generates realizations satisfying the dependency constraints. Here we use the ensemble of dependency structures introduced in Ref. [3], and we define a novel method to build the realizations. Specifically, as sketched in Fig. 1, the growth process that creates the dependency network is an incremental node-addition process generating structures with power-law-distributed sizes of direct and indirect dependencies (such property is crucial to reproduce Zipf’s law, see below). More specifically, let us fix an average out-degree $D \geq 1$, i.e., an average number of direct dependencies of a given component. Starting with an initial graph consisting of a single node, the full graph is built node by node by attaching the new node to $d + 1$ randomly chosen existing nodes (possibly with repetitions), where d is a Poissonian random variable of mean $D - 1$. The process is stopped when the network reaches the predetermined size U . A graph assembled with these rules is directed and acyclic, and hence a good dependency structure, as can be seen by labeling each node by the time $t = 1, \dots, U$ it was added to the network, and noticing that there can be no links $t \rightarrow t'$ with $t < t'$. Given a node c , the set $\wedge(c) \subseteq \mathcal{U}$ is defined as the set of all nodes c' such that there exists at least one directed path in \mathcal{G} starting from c and arriving at c' . We will call the set $\wedge(c)$ the *forward cone* of c . Similarly, we define the *backward cone*

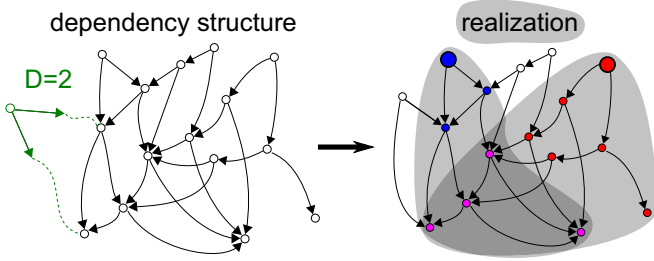


FIG. 1. Illustration of the model. The model is defined by two steps. The first step creates a dependency structure by an incremental node-addition process with mean out-degree D (left panel; the new node on the left attaches to two existing nodes, indicated by dotted lines). The second step builds realizations drawing at random k precursor nodes (larger circles in figure, where $k = 2$) and taking all components belonging to their forward “cones” of dependency (shaded in gray). This second step assigns multiplicities to the components appearing in each realization: The abundance of component i is the number of precursors whose forward cones contain i . In the realization illustrated in the right panel, the nodes (components) in the light-gray regions have multiplicity 1, while those in the dark-gray region have multiplicity 2.

$\wedge(c)$ of c as the set of all nodes c' such that there exists a path from c' to c . In other words, $\wedge(c)$ is the set of all components required by the (direct and indirect) dependencies of c , whereas $\vee(c)$ is the set of the nodes that depend (directly or indirectly) on c . We remark that networks grown by this process do not have fat-tailed degree distributions. Instead, it is the size of the backward cones that is approximately power-law distributed (see Sec. III A).

Once a dependency structure \mathcal{G} is established, realizations of the model, i.e., sets of components, are generated by the following procedure (see also Fig. 1). Let us fix a positive integer k , which represents the number of “precursor” components determining a realization. The k precursors $\{p_j\}$, $j = 1, \dots, k$, are chosen randomly and independently among the nodes of \mathcal{G} . Then the corresponding realization is produced by taking all components belonging to the forward cones of the precursors. To complete the model specification one needs a rule to choose the multiplicities of the components. We allow a component belonging to multiple cones to appear in multiple copies. Let us imagine that precursors are added one at a time to the realization. At the j th step the existing realization r_{j-1} (possibly empty, when $j = 1$) is extended by the addition of elements from the cone $\wedge(p_j)$ of the precursor p_j :

$$r_j = r_{j-1} \cup \Delta_j, \quad \Delta_j \subseteq \wedge(p_j). \quad (1)$$

The choice of the incrementing set Δ_j must be done so as to satisfy the dependency relations, i.e., $r_j \supseteq \wedge(p_j)$. Doing this at every step ensures that the final realization r_k will not have any unsatisfied dependency. Other than that, Δ_j is in principle unconstrained, and it may be a random variable even at fixed p_j and r_{j-1} . Here we make the simplest choice

$$\Delta_j = \wedge(p_j). \quad (2)$$

This makes the process Markovian, in the sense that $r_j \setminus r_{j-1} \equiv \Delta_j$ is independent of r_{j-1} . The case $k = 1$, when a

realization is specified by a single precursor, reproduces the model of Ref. [3].

III. RESULTS

Our description separates the relational constraints existing between the components, such as dependency and incompatibility, from their functional correlations, such as synergy, co-occurrence, interchangeability, conflict, and so on. For what concerns the functional correlations, our model is the simplest least-constrained model where no correlations between components are dictated, other than those arising from the dependencies. An advantage of this model is its analytical tractability. The form of Zipf’s law is basically the same as in the binary model, and the mean-field analysis parallels that in Ref. [3]. The main additional output of our extension is a nontrivial Heaps law, which is derived analytically below.

A. Zipf’s law and the occurrence-abundance relation

We now set out to compute the distribution of component abundance emerging from a set of realizations in this model, equivalent to the empirical Zipf’s law measured from a set of texts.

Given a set of R realizations of a component system, the “popularity” of a component i can be measured in two ways: by its *abundance* a_i and by its *occurrence* o_i . The abundance counts the number of times that i appears in all realizations (with multiplicities):

$$a_i = \frac{1}{kR} \sum_r \sum_{c \in \mathcal{R}} \delta_{c,i}. \quad (3)$$

In the model, the maximum abundance of a component i corresponds to drawing i each time a cone is selected, for each realization. In such a case, the double sum in (3) is kR . Therefore, the abundance a_i is normalized so that $0 \leq a_i \leq 1$ (we will call a_i the *relative abundance* when it is important to stress that it is an intensive quantity). The occurrence o_i measures the fraction of realizations containing the component, regardless of its abundance:

$$o_i = \frac{1}{R} \sum_r \left[1 - \prod_{c \in \mathcal{R}} (1 - \delta_{c,i}) \right]. \quad (4)$$

With this definition, the occurrence is normalized so that $0 \leq o_i \leq 1$.

Zipf’s law is a statement about the rank-frequency relation of components. Specifically, the frequency of a component decreases as a power of its rank r , $f_r \propto r^{-\gamma}$ (where the rank is 1 for the most frequent component, 2 for the second most frequent, and so on) [4]. The Zipf relation is expected to be independent of the number of cones k (at least for large systems). In fact, the abundance a_i of a given component in R realizations constructed with k cones each has the same distribution as that in kR single-cone realizations, since the choices of the cones are independent. a_i can be estimated as the probability of choosing a cone that contains i , which is proportional to the size $|\vee(i)|$ of the backward cone of i :

$$a_i = \frac{|\vee(i)|}{U}. \quad (5)$$

Let us call $\text{rank}(i)$ the rank of component i when all components are ranked by their abundance. Following Ref. [3], an approximate relation can be derived between $|\vee(i)|$ and $\text{rank}(i)$, which will allow to obtain an analytical estimate of Zipf’s plot. The t th node in the network (the one added at the t th step of the construction, when a network of size $t - 1$ has been already generated) has approximately $(U/t)^D$ nodes that depend on it. This result can be obtained by writing an equation stating that the backward cone of the t th node is the union of the backward cones of all the nodes that, at later times t' , will directly attach to the t th node. Neglecting the intersections between these cones allows one to write the recursion

$$|\vee(t)| = 1 + \sum_{t'=t+1}^U \frac{D}{t'} |\vee(t')|, \quad (6)$$

where the factor D/t' estimates the probability that the t' th node attaches to the t th node [with a slight abuse of notation, we write $\wedge(t)$ and $\vee(t)$ for the forward and backward cones of the t th node]. By approximating the sum by an integral and taking a derivative with respect to t , one obtains a differential equation that is solved by $|\vee(t)| = (U/t)^D$.

For small t , however, $(U/t)^D$ is greater than the size of the network U . In fact, the relation can hold only down to a cutoff t_{\min} , which can be estimated by the condition that the whole network depends on the t_{\min} th node, i.e., $(U/t_{\min})^D = U$, which gives $t_{\min} = U^{1-1/D}$. For any node below t_{\min} , the size of its backward cone is $\approx U$:

$$|\vee(t)| \approx \begin{cases} U & t < U^{1-1/D} \\ (U/t)^D & t \geq U^{1-1/D}. \end{cases} \quad (7)$$

Equations (5) and (7) imply that if node i is the t th node in the network growth process, then $t = \text{rank}(i)$. (This identification does not hold for the first $U^{1-1/D}$ components, but this does not influence the result since the size of their backward cones are equal in this approximation.) Therefore, one obtains

$$a_i \approx \begin{cases} 1 & \text{rank}(i) < U^{1-1/D} \\ \text{rank}(i)^{-D} U^{D-1} & \text{rank}(i) \geq U^{1-1/D}. \end{cases} \quad (8)$$

This relation has the form of a Zipf power law (with exponent $-D$) with an initial “core set” consisting of $U^{1-1/D}$ components having similar abundances. This is reminiscent of a result found in natural language showing that the empirical double scaling in Zipf’s law can indeed be explained by the presence of a finite set of specific core words [22,36]. For the particular growth process chosen here, the core set coincides with the oldest nodes, i.e., the ones appearing earlier in the growth of the network. For more general ensembles of topologies, the core can be defined as the set of all nodes whose backward cone has size of order U , for large U . Figure 2(a) compares the analytical form (8) with the results of simulations, showing good accord, especially in the behavior of the fat tail. The transition between the core and the tail, instead, is less sharp than predicted. This is due to the fact that the relation $|\vee(t)| = (U/t)^D$ starts to break down before reaching U and saturates more smoothly than in the approximation made above. Importantly, the relation between rank and relative abundance does not depend on the number of cones k , in accord with the foregoing prediction.

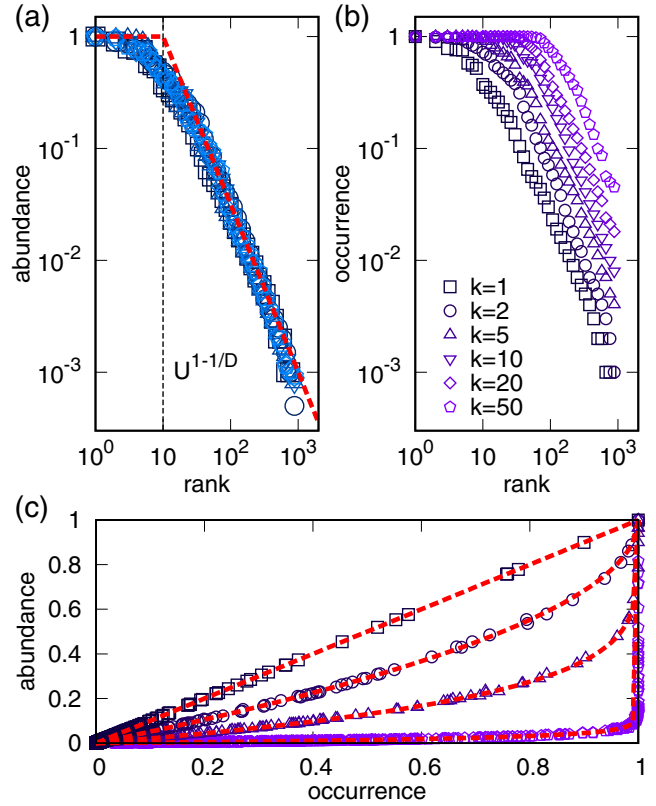


FIG. 2. Simulations match the analytical prediction for the distribution of total abundance in a set of realizations (empirical Zipf’s law) and for the abundance-occurrence relation. (a) Rank-plot of component abundances. (b) Rank-plot of component occurrences. (c) Scatterplot of the abundance (y axis) versus the occurrence (x axis) of each component. Colored symbols are results of simulations (1000 realizations) with varying number of precursors k , and the dashed lines are the analytical prediction. (Parameters: $U = 1000$, $D = 1.5$.)

Contrary to the binary case discussed in Ref. [3], where abundance and occurrence are the same quantity [as shown by Eq. (9)], the rank-occurrence and rank-abundance relations in our model are different in general. Figure 2(b) shows that the Zipfian plot for the occurrence has a similar appearance to the usual one but with a larger and larger core as k increases. When plotted as a probability distribution function in linear scale, the rank-occurrence relation has a characteristic U shape [37,38] that highlights a subset constituted by a large number of rare components and a subset of components shared by a large fraction of realizations, usually referred to as the “cloud” and “core” sets, respectively [25]. The specificity of the core set in our model emerges from the dependency structure. How these features are related to Zipf’s law in a scenario of random sampling is explored in detail in Ref. [5].

The occurrence-abundance relation predicted by our model turns out to be universal (or “null”), meaning that it is insensitive to the explicit form of Zipf’s law, to the detailed structure of the network, and even to the size U of the component universe. In fact, we show here that a simple probabilistic argument gives a relation that is consistent with simulations of the full model. In the limit of large R , we can assume that the occurrence of a component i is equal to the probability of choosing i at least

once in a single realization: $o_i = 1 - (1 - a_i)^k$, hence

$$a_i = 1 - (1 - o_i)^{1/k}. \quad (9)$$

For realizations with a single precursor ($k = 1$), abundance and occurrence are equal. As k increases, more and more components (with larger and larger occurrence) assume small abundances. In the large- k limit, all components have zero (relative) abundance, except those with occurrence 1. Figure 2(c) shows that a scatterplot of abundance versus occurrence in simulations perfectly matches the theoretical prediction (9). The figure shows results for a single choice of D and U , but we checked that these parameters have no effect on the curves.

B. Analytical form of Heaps's law

We now ask about the change of the typical number of distinct components with realization size. The calculation of the number $F(N)$ of unique components in a realization of size N can be performed in a mean-field approximation, where the correlations between nodes are neglected. We consider a process where a realization is generated by extracting N nodes independently. The full model adds entire dependency cones at once. However, one expects that the size N of a realization scales linearly with the number of cones k , and we have verified with simulations that this is indeed the case. The assumption of independent extractions makes our mean-field approach to evaluate Heaps's law analogous to the "Zipfian ensemble" of quantitative linguistics, in which realizations are generated through random extractions of components with probabilities proportional to their frequencies [5,39,40] or through a Poisson process in which component frequencies establish their arrival rates [41]. In this framework, Heaps's law is the natural result of the heterogeneity in component frequencies described by Zipf's law.

The probability

$$p(t) = \frac{1}{\Omega} |\vee(t)| \quad (10)$$

of drawing the node t is proportional to the size $|\vee(t)|$ of the node's backward cone. In a continuous approximation, the normalization Ω can be fixed by the condition $\int_0^\infty p(t) dt = 1$, which yields

$$\Omega = U \frac{DU^{1-1/D} - 1}{D - 1}. \quad (11)$$

Note that $\Omega > U$ whenever $U > 1$ and $D > 1$.

Let $p_1(t, n)$ be the probability that the t th node in the network is drawn for the first time when the system being constructed has size n :

$$p_1(t, n) = p(t)[1 - p(t)]^{n-1}. \quad (12)$$

A mean-field estimate of F can then be obtained as

$$F(N) = \sum_{n=1}^N \sum_{t=1}^U p_1(t, n) \approx \int_0^U dt \sum_{n=1}^N p_1(t, n). \quad (13)$$

The geometric sum in n gives simply the probability $1 - [1 - p(t)]^N$ that the t th node has been drawn at least once after N steps. The mean-field expression for Heaps's law is then given

by the following integral:

$$\begin{aligned} F(N) &= \int_0^U dt \{1 - [1 - p(t)]^N\} \\ &= U - U^{1-1/D} \left(1 - \frac{U}{\Omega}\right)^N - \mathcal{I}(N), \end{aligned} \quad (14)$$

where the first term (U) comes from the integral of 1, and the second and third terms are the contributions of the two regions in (7). The remaining integral,

$$\mathcal{I}(N) = \int_{U^{1-1/D}}^U dt \left[1 - \frac{1}{\Omega} \left(\frac{U}{t}\right)^D\right]^N, \quad (15)$$

can be evaluated with the change of variables $z = (U/t)^D/\Omega$, which gives

$$\mathcal{I}(N) = \frac{U}{D} \Omega^{-1/D} \int_{1/\Omega}^{U/\Omega} (1-z)^N z^{-1-1/D} dz. \quad (16)$$

By remembering that the primitive of $(1-z)^\alpha z^\beta$ is $z^{\beta+1} {}_2F_1(-\alpha, \beta+1, \beta+2, z)/(\beta+1)$, where ${}_2F_1$ is the Gauss hypergeometric function, one finally obtains

$$\begin{aligned} F(N) &= U - U^{1-1/D} \left(1 - \frac{U}{\Omega}\right)^N \\ &\quad - {}_2F_1\left(-N, -\frac{1}{D}, 1 - \frac{1}{D}, \frac{1}{\Omega}\right) U \\ &\quad + {}_2F_1\left(-N, -\frac{1}{D}, 1 - \frac{1}{D}, \frac{U}{\Omega}\right) U^{1-1/D}. \end{aligned} \quad (17)$$

The hypergeometric function ${}_2F_1(-N, \cdot, \cdot, \cdot)$ is real when N is an integer. Figure 3 shows that the analytical mean-field expression (17) nicely matches the results of numerical simulations of the model. A similar functional form to Eq. (17) was found in Appendix E of Ref. [42], by assuming a Gamma distribution for the frequencies of words in texts.

C. Linear, sublinear, and saturation regimes of Heaps's law

If a realization is constructed by incremental addition of randomly chosen components, then one expects $F(N)$ to be approximately linear for small N , as it is unlikely to draw the same component twice. Intuitively, the probability to do so increases with N , up to a point where approximately all components in the universe will have been included, and $F(N)$ will saturate to U . For instance, this is the behavior observed empirically for the species-area relationship in ecology [43–45]. This behavior becomes apparent by plotting $F(N)$ in log-log scale [see Fig. 3(a)]. There emerge three distinct regimes: a linear increase for small N , a saturation to U for large N , and an intermediate regime where the sublinear increase of $F(N)$ appears to be well described by a power law. Two transition points can be identified, N_c and N_s , respectively at the crossover between the linear and the sublinear regimes, and at the onset of saturation. We collect here a few analytical estimates and observations.

It is clear from expression (14) that, since $p(t) > 0$ for a finite universe,

$$\lim_{N \rightarrow \infty} F(N) = U. \quad (18)$$

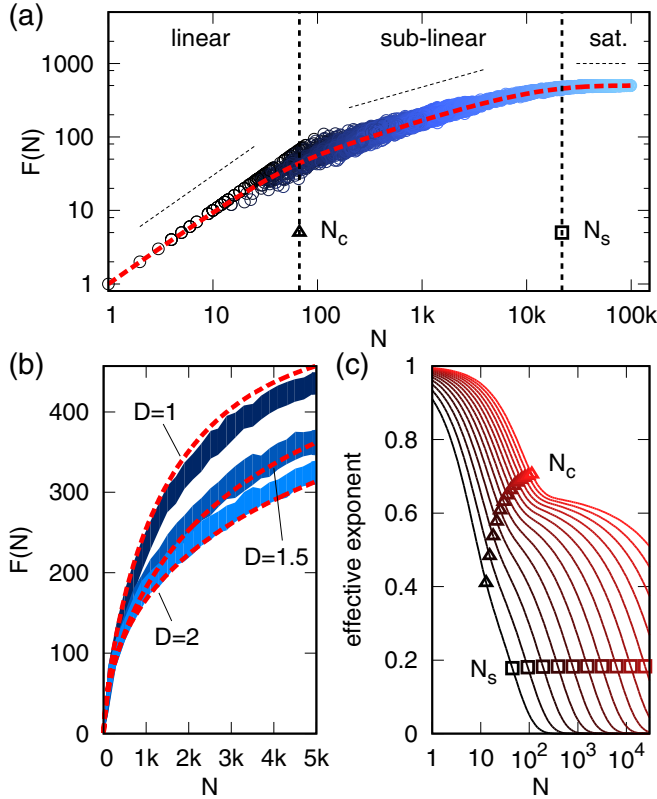


FIG. 3. The model reproduces the characteristic three-regimes structure of Heaps's law. (a) The number of distinct components $F(N)$ of a realization as a function of realization size N (in log-log scale). The colored circles are simulations, and lighter shades of blue correspond to larger numbers k of precursors (k ranges from 1 to 3000). The red dashed lines are the analytical prediction (17). The area of the solid curves represents the 90% variability interval. (Parameters: $U = 500$, $D = 2$, 13 000 realizations.) (b) The same plot of $F(N)$, in linear scale, for three values of the mean out-degree D : 1 (dark blue), 1.5 (blue), and 2 (light blue). (c) Effective exponent versus realization size for increasing values of U (black- to red-colored lines, ranging logarithmically from $U = 10$ to $U \approx 6000$). Triangles correspond to the analytical estimate of N_c , which captures well the transition to a plateau region. The analytical estimate of N_s (squares) corresponds to an approximately constant effective exponent.

This is a consequence of the definition of the model, whereby $F(N)$ is monotonic by construction and $F(N) \leq U$. However, this limit is not apparent from the final formula (17). What happens is that the (essential) singularities of the two hypergeometric functions cancel out in the large- N limit. This makes it difficult to compute values of $F(N)$ numerically in this regime (see below).

An estimate of the point N_s where the saturation regime sets in can be obtained from (14). The term with $(1 - U/\Omega)^N$ is significantly different from zero when $U/\Omega \lesssim 1/N$, i.e., when $N \lesssim \Omega/U$. The integral $\mathcal{I}(N)$, instead, can be evaluated for large N in a saddle-point approximation. The integrand [Eq. (15)] attains its minimum at $t = U$, where it is equal to $(1 - 1/\Omega)^N$; hence, it is significantly different from zero when $N \lesssim \Omega$. Therefore, both N -dependent terms in (14) are negligible when $N \gtrsim N_s = \Omega$, where Ω is given by (11).

The small- N behavior at finite U can be obtained in principle from Eq. (14) as well by expanding in N before performing the integral in $\mathcal{I}(N)$. However, it is easier to analyze the onset of sublinearity in the large- U regime by expanding Eq. (17) in powers of $1/U^{1-1/D}$. This can be done by using the definition of the hypergeometric function:

$${}_2F_1(a, b, c, z) = \sum_{k=0}^{\infty} \frac{(a)_k (b)_k}{(c)_k k!} z^k, \quad (19)$$

where $(\alpha)_k = \alpha(\alpha+1)(\alpha+2)\cdots(\alpha+k-1)$ is the Pochhammer symbol. The two terms of the form ${}_2F_1(\dots)$ in (17) can be expanded in powers of z ; Eq. (19) shows that the term of order z^j in such an expansion is a polynomial of order j in N . The same property holds for the small- z expansion of the first N -dependent term in (17), of the form $(1-z)^N$. It is then easy to see, by keeping track of the analytical and nonanalytical powers of U , that in the limit $U \rightarrow \infty$ the only nonvanishing term is N , and Heaps's law reduces to the identity

$$\lim_{U \rightarrow \infty} F(N) = N. \quad (20)$$

A linear onset is expected for small N even when U is finite. Performing explicitly the expansion to first order in $U^{-1+1/D}$ yields

$$F(N) \approx N - \frac{1}{2}N(N-1) \frac{2(D-1)^2}{D(2D-1)} U^{-1+1/D}. \quad (21)$$

The crossover point N_c separating the linear and sublinear regimes can be estimated by the point where (21) reaches its maximum:

$$N_c = \frac{D(2D-1)}{2(D-1)^2} U^{1-1/D}. \quad (22)$$

This expression is expected to become inaccurate when $D \approx 1$ (where in fact it diverges), because all terms of order $U^{-j+j/D}$ with $j > 1$, which are neglected in (21), approach constants for $D \rightarrow 1$.

In order to identify more visually the transition points from the data, one can plot the *effective exponent*

$$\gamma_{\text{eff}}(N) = \frac{d \log F(N)}{d \log N}, \quad (23)$$

which is easily computed from numerical data as a discrete derivative. $\gamma_{\text{eff}}(N)$ measures the apparent exponent that is obtained by approximating the function $F(N)$ locally by a power-law $N^{\gamma_{\text{eff}}}$. Figure 3(c) shows the effective exponent for a range of values of U . For small U , the regimes are somewhat intertwined, and no sharp transitions appear. For larger U , γ_{eff} shows three plateaus, corresponding to $\gamma_{\text{eff}} = 1$, $\gamma_{\text{eff}} = 0$, and an intermediate value $\gamma_{\text{eff}} = \gamma$. This intermediate plateau, which becomes more and more extended by increasing U , indicates the robustness of the effective power-law behavior with the nontrivial exponent γ . While the figure suggests that there may be an abrupt regime shift in U between the absence and the presence of the intermediate plateau, a more refined analysis (not reported here) shows a very smooth expansion of the plateau for increasing values of U . An approximate value for γ can be obtained by making the approximation

$(1 - \epsilon)^N \approx \exp(-N\epsilon)$ in Eq. (15), where $\epsilon = (U/t)^D/\Omega$. The integral then becomes

$$\mathcal{I}(N) \approx \frac{U}{D} \left(\frac{N}{\Omega}\right)^{1/D} \Gamma\left(-\frac{1}{D}, x\right) \Big|_{x=NU/\Omega}^{x=N/\Omega}. \quad (24)$$

In the intermediate regime $N_c \ll N \ll N_s$ one can use the recurrence relation of the upper incomplete gamma function Γ and its asymptotic expansion, finding

$$F(N) \approx \Gamma\left(1 - \frac{1}{D}\right) U \left(\frac{N}{\Omega}\right)^{1/D}, \quad (25)$$

where the one-argument Γ is the (complete) Euler gamma. The equation then provides the power-law exponent: $\gamma = 1/D$. This same exponent can be derived in the framework of the Zipfian ensemble computations by a simple scaling argument by assuming a pure power-law behavior of $F(N)$ [39,41,46].

Figure 3(c) also shows that the transition points computed above, i.e., N_c given by (22), and

$$N_s = U \frac{DU^{1-1/D} - 1}{D - 1}, \quad (26)$$

are reasonable estimates of the sizes where the two regime shifts occur. Surprisingly, the estimate of N_s turns out to correspond to an approximately U -independent value of the effective exponent.

D. Stretched exponential saturation as a phenomenological expression of Heaps's law

As pointed out above, the asymptotically flat behavior of Heaps's law $F(N)$ in this model results from the cancellation of two infinite contributions in the analytical formula. This subtlety makes it numerically challenging to evaluate $F(N)$, especially for large U and N . Such a difficulty prevents the use of Eq. (17) to estimate the parameters by fitting against empirical data. However, the analytical expression (16) suggests a simple phenomenological expression, which can be useful in fits. Since the integration variable z is small for large U , one can attempt to approximate the integrand in $\mathcal{I}(N)$ by $z^{-1-1/D} \exp(-zN) dz$. In this form, the integral is similar to a representation of the stretched exponential function $\psi_{\gamma,a}(x) = \exp(-ax^\gamma)$ in terms of exponential decays as a Laplace transform,

$$\psi_{\gamma,a}(x) = \int_0^\infty P_{\gamma,a}(z) e^{-zx} dz. \quad (27)$$

The asymptotic behavior of $P_{\gamma,a}(z)$ is known to be

$$P_{\gamma,a}(z) \sim z^{\gamma+1} \quad (28)$$

for large z , and an exponential decrease for small z [47]. This suggests the following phenomenological expression:

$$F_{\text{ph}}(N) = U[1 - \exp(-aN^\gamma)]. \quad (29)$$

Figure 4 shows that the stretched-exponential saturation, Eq. (29), is a remarkably good approximation of the simulated data. The log-linear scale reveals that the agreement is tight on the whole range of N . However, the phenomenological expression fails to capture the transient linear increase at small sizes (see Sec. III C). Indeed, the small- N behavior

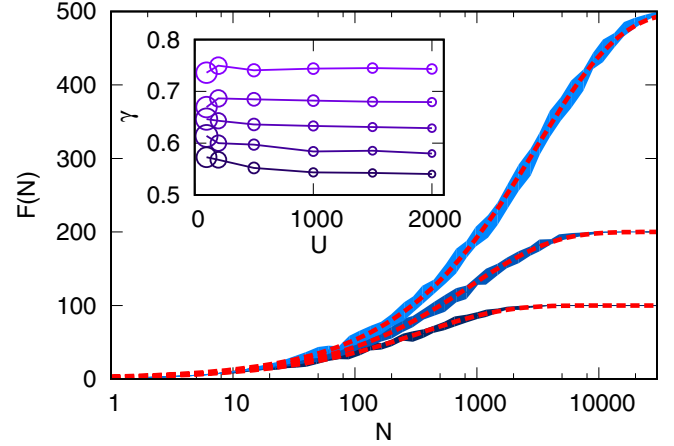


FIG. 4. The stretched exponential is a good approximation of the Heaps plot generated by the model. Number of unique components (y axis, in linear scale) as a function of realization size (x axis, in logarithmic scale), for three different values of U , the size of the universe. Simulations are the solid curves (dark blue $U = 100$, blue $U = 200$, and light blue $U = 500$; line width is the 90% variability interval). The dashed red lines are the analytical prediction. (Parameters: $D = 1.5$, k ranging from 1 to 5000.) (Inset) The fitted stretched-exponential exponent γ as a function of U . Different shades correspond to $D = 1, 1.25, 1.5, 1.75, 2$, from top to bottom; point size represents errors.

of F_{ph} is $F_{\text{ph}}(N) \sim UaN^\gamma$. Interestingly, if one extracts γ by matching the large- z power-law scaling in Eq. (28) with the factor $z^{-1-1/D}$ in Eq. (16), one obtains $\gamma = 1/D$, which is the same exponent that was derived from Eq. (24). Note however, that the integration range in (16) is very different from the one in the integral representation (27) of $\psi_{\gamma,a}$, that is, $(0, \infty)$. As a consequence, the fitted exponents γ can deviate from the simple scaling relation $\gamma = 1/D$. Altogether, these considerations suggest that it is quite surprising that the stretched exponential can be such a good approximation.

IV. DISCUSSION

It is interesting to compare the results of this positive model, where the component statistics are determined by component dependencies, with more null views of Heaps's and Zipf's laws. We assume the usual notion of null models, as the least-constrained or maximally random scenarios, with which more refined and elaborate (i.e., positive) models are compared. Our analytical computations for the rank-abundance and rank-occurrence and Heaps's relations were all performed in an ensemble where correlations between different components are neglected. The agreement between the results obtained in such a mean-field approximation and the numerical simulations of the full model reflects the fact that the main statistics considered are not sensitive to the correlations, and they depend only weakly on the details of the network ensemble. Here we have tried to pinpoint the salient features that generate Zipf's and Heaps's laws, namely the fat tails of the size distribution of backward cones and the size of the core. Other details are expected to be relevant for other observables (see below).

Our results show that the null relation between Zipf's law and Heaps's law is still satisfied even in a strongly interacting

setting such as the one defined by dependency structures. At a finer resolution, the Heaps plot may display non-null deviations on top of the Zipfian ensemble prediction, perhaps carrying information about the underlying dependencies. However, the overall semiquantitative accord suggests that the classic observables related to frequency and occurrence are not good proxies of interdependency between components. In particular, the approximate stretched-exponential form of Heaps's law is expected to hold not only for this special positive model, but for any model yielding a Zipf law similar to Eq. (8). These observations are in agreement with the results of Ref. [5], where it is shown that building realizations of a component system by randomly sampling from the global frequency distribution of its components reproduces some empirical statistics of real systems to a surprising detail.

Understanding how the positive trends emerge from the random features is clearly crucial to determine the role played by dependencies in a given empirical system. We briefly discuss two such positive observables here, leaving a more detailed investigation for future work. First, the Zipf law in Eq. (8), which is the main input of a sampling procedure [5], is instead a positive output of the model studied here: fixing the network of dependencies between components constrains all the main statistics regarding their abundance and occurrence, including their frequency distribution. Second, the existence of dependency relations generates nontrivial correlations between the components. These correlations can be observed, e.g., by measuring the empirical distribution of the mutual information between the occurrences of pairs of components. Let $p_i(x)$ be the fraction of realizations in which the component i has state x , where $x = 1$ means it is present and $x = 0$ means it is absent, and let $p_{ij}(x, y)$ be the fraction of realizations in which the component i has state x and the component j has state y . The mutual information between the two components, measuring how much the presence of one informs on the presence of the other, is defined as

$$I(i, j) = \sum_{x,y} p_{ij}(x, y) \log \frac{p_{ij}(x, y)}{p_i(x)p_j(y)}. \quad (30)$$

If the probabilities are factorized, i.e., $p_{ij}(x, y) = p_i(x)p_j(y)$, as it is the case in the random sampling model, then $I(i, j) = 0$. If the empirical probabilities are computed from R samples, one expects $I(i, j)$ to converge to 0 in the large- R limit; yet it will have a nontrivial distribution for finite R . Figure 5 shows the distribution of the mutual information $I(i, j)$, averaged over all pairs i and j , for a set of realizations generated from random sampling, compared to the positive model based on dependency structures studied here. It is clear that, beyond a common bulk due to finite-size fluctuations, the two models have very different profiles in terms of pairwise correlations. The effect of the underlying dependency structure generates a high mutual information tail, populated by highly correlated pairs of components and emerging from a background of pairwise relationships compatible with the null random-sampling model. We also expect that dependency structures will create correlations that go beyond the pairwise relations between components.

The model presented here provides the simplest generative mechanism producing collections of components consistent

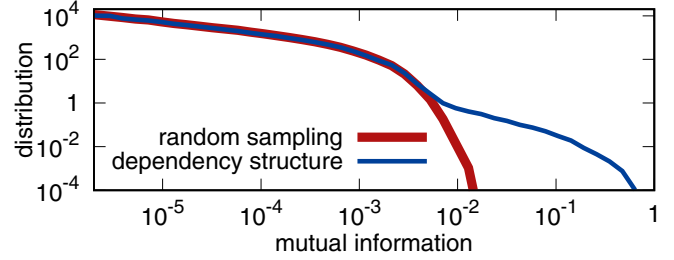


FIG. 5. Non-null signatures of component dependency. The distribution of mutual information (x axis), Eq. (30), for the random sampling protocol of Ref. [5] (thick red line) and for the model studied here (thin blue line). The frequency distribution of components was obtained by the positive model (with parameters $U = 1000$, $D = 1.5$, $k = 50$), and $I(i, j)$ was computed for $R = 1000$ realizations of both models. The distribution (y axis) is the normalized empirical distribution for all realizations and all pairs (i, j) .

with a given dependency structure. Our construction extends the one proposed in Ref. [3] to the case of components with nonbinary abundance. While the distribution of component abundances has a power-law tail with an initial core, like in the case of Ref. [3], the situation is more complex here, as the distribution of component abundances does not coincide with that of occurrences, due to the nonbinary nature of the multiplicities. More general families of models can be defined by specifying the rules for selecting the abundance of nodes belonging to the dependency cones of more than one precursor. The additive choice taken here, where each cone determines the addition of one copy of each component belonging to it, provides a minimal model that is memoryless, and therefore still accessible analytically.

Additionally, our model can be linked to other recent studies investigating the dynamics of innovations [48,49]. In such descriptions, Heaps's law is seen as informative of how the space of the "adjacent possible" is explored by a process. In particular, Ref. [49] describes the occurrence of novelties as a (noncausal) network exploration process (an edge-reinforcing random walk) showing the appearance of Heaps's law. In our model, the causal relationships between individual components encoded in the network affect the trend of Heaps's law and thus the probability of finding new components in a dependency cone added to a realization. The two transition points N_c and N_s define three nontrivially different regimes for the innovation dynamics in a typical realization.

Since a wide variety of systems can be represented as collections of components belonging to a common pool, our model has general applicability. A special case is that of genomes, which deserves some observations. A dependency structure between genes represents the recipes binding the functional roles of different protein families, thereby determining their usefulness in the same genome. For example, a gene could depend on another if it is found downstream in the same metabolic pathway [3]. The topology of such a dependency has not been fully characterized. Likely, it comprises both feedforward and feedback structures, as well as nondirected exclusion principles (whereby a gene might not be necessary or useful if another one is present). Concerning the additive choice discussed above, it is possible (and likely)

that the choice of a memoryless process is too restrictive in the context of genomics. For instance, not all gene families required by the presence of multiple precursors need to be present in multiple copies. Future investigations could aim at defining more stringently from data the minimal features of a model of dependency that could realistically describe genomes. This could be inferred by the correlation structure of domain abundances from sets of entirely sequenced genomes.

ACKNOWLEDGMENTS

We thank Michele Caselle, Amos Maritan, and Sergei Maslov for useful discussions and advice. J.G. was supported by an Omidyar Postdoctoral Fellowship at the Santa Fe Institute. M.O. is supported by the Departments of Excellence Grant awarded by the Italian Ministry of Education, University and Research (MIUR) (L. 232/2016).

-
- [1] E. van Nimwegen, *Trends Genet.* **19**, 479 (2003).
- [2] M. Gherardi, S. Mandrà, B. Bassetti, and M. C. Lagomarsino, *Proc. Natl. Acad. Sci. USA* **110**, 21054 (2013).
- [3] T. Y. Pang and S. Maslov, *Proc. Natl. Acad. Sci. USA* **110**, 6235 (2013).
- [4] E. G. Altmann and M. Gerlach, Statistical laws in linguistics, in *Creativity and Universality in Language*, edited by M. Degli Esposti, E. G. Altmann, and F. Pachet (Springer International Publishing, Cham, 2016), pp. 7–26.
- [5] A. Mazzolini, M. Gherardi, M. Caselle, M. Cosentino Lagomarsino, and M. Osella, *Phys. Rev. X* **8**, 021023 (2018).
- [6] J. Pitman, *Probab. Theory Relat. Fields* **102**, 145 (1995).
- [7] J. M. Leinaas and J. Myrheim, *Il Nuovo Cimento B (1971–1996)* **37**, 1 (1977).
- [8] M. R. Evans and T. Hanney, *J. Phys. A* **38**, R195 (2005).
- [9] A. Angelini, A. Amato, G. Bianconi, B. Bassetti, and M. Cosentino Lagomarsino, *Phys. Rev. E* **81**, 021919 (2010).
- [10] M. Cosentino Lagomarsino, A. Sellerio, P. Heijning, and B. Bassetti, *Genome Biol.* **10**, R12 (2009).
- [11] G. K. Zipf, *The Psycho-Biology of Language* (Houghton Mifflin, New York, 1935).
- [12] M. E. Newman, *Contemp. Phys.* **46**, 323 (2005).
- [13] M. Mitzenmacher, *Internet Math.* **1**, 226 (2004).
- [14] A. J. Lotka, *J. Wash. Acad. Sci.* **16**, 317 (1926).
- [15] M. Huynen and E. van Nimwegen, *Mol. Biol. Evol.* **15**, 583 (1998).
- [16] R. L. Axtell, *Science* **293**, 1818 (2001).
- [17] G. Herdan, *Type-Token Mathematics* (Mouton, The Hague, Netherlands, 1960), Vol. 4.
- [18] H. S. Heaps, *Information Retrieval: Computational and Theoretical Aspects* (Academic Press, New York, 1978).
- [19] D. M. Powers, in *Proceedings of the Joint Conferences on New methods in Language Processing and Computational Natural Language Learning* (Association for Computational Linguistics, Sydney, Australia, 1998), pp. 151–160.
- [20] R. F. i Cancho and R. V. Solé, *Proc. Natl. Acad. Sci. USA* **100**, 788 (2003).
- [21] D. Zanette and M. Montemurro, *J. Quant. Ling.* **12**, 29 (2005).
- [22] M. Gerlach and E. G. Altmann, *Phys. Rev. X* **3**, 021006 (2013).
- [23] E. V. Koonin, *PLoS Comput. Biol.* **7**, e1002173 (2011).
- [24] J. Grilli, B. Bassetti, S. Maslov, and M. Cosentino Lagomarsino, *Nucl. Acids Res.* **40**, 530 (2012).
- [25] A. E. Lobkovsky, Y. I. Wolf, and E. V. Koonin, *Genome Biol. Evol.* **5**, 233 (2013).
- [26] S. Maslov, S. Krishna, T. Y. Pang, and K. Sneppen, *Proc. Natl. Acad. Sci. USA* **106**, 9743 (2009).
- [27] T. Mora and W. Bialek, *J. Stat. Phys.* **144**, 268 (2011).
- [28] P. Bak, C. Tang, and K. Wiesenfeld, *Phys. Rev. Lett.* **59**, 381 (1987).
- [29] J. Hidalgo, J. Grilli, S. Suweis, M. A. Muñoz, J. R. Banavar, and A. Maritan, *Proc. Natl. Acad. Sci. USA* **111**, 10095 (2014).
- [30] R. Hanel and S. Thurner, *Europhys. Lett.* **96**, 50003 (2011).
- [31] B. Corominas-Murtra, R. Hanel, and S. Thurner, *Proc. Natl. Acad. Sci. USA* **112**, 5348 (2015).
- [32] D. Sornette, *Critical Phenomena in Natural Sciences: Chaos, Fractals, Selforganization and Disorder: Concepts and Tools* (Springer Science & Business Media, New York, 2006).
- [33] D. Heckerman, D. M. Chickering, C. Meek, R. Rounthwaite, and C. Kadie, *J. Mach. Learn. Res.* **1**, 49 (2000).
- [34] D. Y. Kenett, M. Tumminello, A. Madi, G. Gur-Gershgoren, R. N. Mantegna, and E. Ben-Jacob, *PLoS ONE* **5**, e15032 (2010).
- [35] M. A. Fortuna, J. A. Bonachela, and S. A. Levin, *Proc. Natl. Acad. Sci. USA* **108**, 19985 (2011).
- [36] R. F. i Cancho and R. V. Solé, *J. Quant. Ling.* **8**, 165 (2001).
- [37] E. V. Koonin and Y. I. Wolf, *Nucl. Acids Res.* **36**, 6688 (2008).
- [38] M. Touchon, C. Hoede, O. Tenaillon, V. Barbe, S. Baeriswyl, P. Bidet, E. Bingen, S. Bonacorsi, C. Bouchier, O. Bouvet *et al.*, *PLoS Genet.* **5**, e1000344 (2009).
- [39] D. C. van Leijenhorst and T. P. Van der Weide, *Inf. Sci.* **170**, 263 (2005).
- [40] F. Font-Clos and Á. Corral, *Phys. Rev. Lett.* **114**, 238701 (2015).
- [41] I. Eliazar, *Physica A (Amsterdam)* **390**, 3189 (2011).
- [42] M. Gerlach and E. G. Altmann, *New J. Phys.* **16**, 113010 (2014).
- [43] J. Rosindell and S. J. Cornell, *Ecol. Lett.* **10**, 586 (2007).
- [44] J. Grilli, S. Azaele, J. R. Banavar, and A. Maritan, *J. Theor. Biol.* **313**, 87 (2012).
- [45] J. P. O’Dwyer and S. J. Cornell, *Sci. Rep.* **8**, 10200 (2018).
- [46] L. Lü, Z.-K. Zhang, and T. Zhou, *PLoS One* **5**, e14139 (2010).
- [47] D. C. Johnston, *Phys. Rev. B* **74**, 184430 (2006).
- [48] F. Triá, V. Loreto, V. D. P. Servedio, and S. H. Strogatz, *Sci. Rep.* **4**, 5890 (2014).
- [49] I. Iacopini, S. Milojević, and V. Latora, *Phys. Rev. Lett.* **120**, 048301 (2018).