

Memory-cognizant generalization to Simon's random-copying neutral model

Joseph D. O'Brien  and James P. Gleeson 

MACSI, Department of Mathematics and Statistics, University of Limerick, Limerick V94 T9PX, Ireland



(Received 23 February 2021; accepted 16 September 2021; published 21 October 2021)

Simon's classical random-copying model, introduced in 1955, has garnered much attention for its ability, in spite of an apparent simplicity, to produce characteristics similar to those observed across the spectrum of complex systems. Through a discrete-time mechanism in which items are added to a sequence based upon rich-gets-richer dynamics, Simon demonstrated that the resulting size distributions of such sequences exhibit power-law tails. The simplicity of this model arises from the approach by which copying occurs uniformly over all previous elements in the sequence. Here we propose a generalization of this model which moves away from this uniform assumption, instead incorporating memory effects that allow the copying event to occur via an arbitrary age-dependent kernel. Through this approach, we first demonstrate the potential to determine further information regarding the structure of sequences from the classical model before illustrating, via analytical study and numeric simulation, the flexibility offered by the arbitrary choice of memory. Furthermore, we demonstrate how previously proposed memory-dependent models can be further studied as specific cases of the proposed framework.

DOI: [10.1103/PhysRevResearch.3.043057](https://doi.org/10.1103/PhysRevResearch.3.043057)

I. INTRODUCTION

Within natural systems from an assortment of domains there are underlying properties which are found to consistently appear. These occurrences result in researchers aiming to develop general theories which can capture aspects of these phenomena independent of the domain under consideration [1,2]. One such instance is the distribution underlying the abundance of a system's constituents, which is frequently described by a highly right-skewed distribution [3–5]. This apparently ubiquitous phenomena has been observed in a variety of domains, including baby-name popularity [6,7], citations to scientific literature [8–11], user-generated passwords [12], and the market share of different cryptocurrencies [13].

Motivated by this apparently ubiquitous property, researchers have investigated mechanisms which may reproduce such heavy-tailed size distributions. These models tend to create a population of elements, each of which has a certain identity or variant, who reproduce over time through a given process, the abundances of said variants are generally the quantities of interest. Typically these mechanisms also offer the possibility of mutation or the creation of a new type of variant in a given reproduction event. The most common of these frameworks are known as *neutral models* due to the assumption that each variant in the population has equal fitness [14,15]. This neutrality implies that there are no intrinsic

advantage offered to any type of variant within the model. For a detailed overview of the most common of these models, alongside their limitations, we refer the interested reader to Ref. [16].

In 1955 Herbert Simon, building upon the earlier work of G. U. Yule [17], famously introduced such a neutral model based upon random-copying (reproduction was viewed as a copying event) with mutation (which he referred to as an innovation event) framework that could reproduce the power-law distributions observed within empirical systems [4]. The fundamental property of this model is that the likelihood of the next element being of a certain variant is dependent upon the number of previous occurrences of this variant. Notwithstanding the model's apparent simplicity, it has been shown to accurately describe the distribution of abundances within a number of complex systems, including the citation dynamics of scientific literature [18,19], family names [20], and the growth of both the world wide web [21] and open-source software developments [22].

Similar models have appeared throughout the literature since Simon's original paper—in particular those placing emphasis upon the cumulative advantage or rich-get-richer mechanisms. Arguably the most famous such approach has been in the development of network growth models within the field of complex networks initiated by Barabási and Albert in 1999 [23] (although the first such framework was produced through Price's model of citation growth [8,24]). In spite of the misleading use of the word “advantage” within these frameworks it is important to highlight, as mentioned in Ref. [16], that these models are in fact neutral models as the nodes have no inherent fitness: their likelihood of selection is determined only by the number of times they have previously been selected, although there have been a number of

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

extensions proposed that move beyond from this neutral assumption [25,26].

As demonstrated through the preceding commentary, Simon's model has proven extremely adaptable in describing behaviors within a wide spectrum of domains. Remarkably, in spite of the model's supposed straightforwardness, the research community is still providing new insights into its behaviors. For example, analytical quantification of the first-mover advantage offered towards the initial variant in Simon's model was demonstrated in Ref. [19], while the effect of averaging across multiple realizations of a closely related model upon the distribution of larger abundances, denoted as peloton dynamics, have also been considered [27]. There have also been a number of studies which have incorporated memory effects within Simon's model such that the likelihood of a variant being copied is no longer uniform across all previous elements but rather is dependent upon how long ago the variant was last used. Such a framework requires the incorporation of a probability function $\phi(\tau, t)$ describing the likelihood of copying at time t an element which appeared at time τ . For example, in Refs. [28,29], the authors introduce a fat-tailed memory distribution to the model which could capture long-range correlations in times between copying events. Furthermore, the authors in Ref. [30] consider a finite-sized memory kernel such that if a certain time has passed since a variant has last been copied it becomes extinct. Reference [31] studies the scenario whereby the likelihood of copying is dependent on a power function of its abundance, with the asymptotic behavior of said abundance being analyzed via a branching process interpretation. Lastly, a specific age-dependent model described by an additional parameter is again studied in terms of its branching behavior in Ref. [32]. Questions remain however as to the effect of incorporating a more general memory function, depending on the time since an element appeared, to Simon's model.

With these questions in mind, in this paper, we consider a branching process approach [33,34] towards describing a generalized Simon's model. Such techniques have previously proven extremely useful in representing similar random-copying phenomena in online diffusion scenarios [35–40]. Initially, we take this approach in describing the classical Simon's model and demonstrate how the interpretation allows numerous quantities describing the distribution of abundances within the process to be obtained in an analytically tractable manner. Moreover, we obtain results regarding the temporal evolution of a given variant as a consequence of the model incorporating the time at which a given element first appears. We proceed to use a similar framework to study a general memory function which depends upon the time $t - \tau$ elapsed since an element appeared to produce a generalized Simon's model (GSM) and demonstrate how statistical properties of abundances arising through such a process may be obtained for an arbitrary choice of said memory function. We also highlight how the previous literature focusing on specific choices of memory function [28–30], which are viewed as particular scenarios of our proposed framework, may be further understood using the branching process interpretation proposed here.

The remainder of the paper is laid out as follows. In Sec. II, we introduce Simon's original model in detail before demon-

strating a branching process approach towards describing it from which classic, alongside new, results are derived. Having demonstrated the usefulness of our framework we proceed in Sec. III to introduce a generalization of Simon's model incorporating arbitrary age-dependent memory effects and present a thorough analysis of its features. In Sec. IV, we proceed to validate the GSM framework through extensive numerical simulations. We demonstrate how previous extensions to Simon's model may be viewed further studied using the framework presented here resulting in additional information regarding said extensions now being attainable in Sec. V before drawing our conclusions in Sec. VI.

II. RESULTS FROM SIMON'S CLASSICAL MODEL - NEW AND OLD

The model devised by Simon [4], like most neutral models, is remarkable for both its apparent simplicity and the variety of behaviors it exhibits. The original representation describes an author creating a body of text that is the population of interest, with every word used representing an element: each unique word corresponds to a variant in the neutral model framework. In the analysis to follow, we shall use the neutral model descriptors (e.g., population, element, variant). The system initializes at time $t = 1$ with a single element and then at each subsequent discrete time-step of length Δt another element is added. How the variant of this element is chosen occurs through a probabilistic framework whereby there may be a mutation (innovation event in the original model) with probability (w.p.) μ such that this new element is a new variant in the population. Alternatively, occurring w.p. $1 - \mu$, the new element is from a reproduction event implying a random replication from one of the previous elements. Our main focus in this paper is on considering how the new element chooses a previous element to copy from; initially however, we study the case considered by Simon whereby this choice is made uniformly over all previously used elements.

We now focus on the main quantity of interest which is the number of times n that a variant which first appeared at time τ has been used by time Ω . Of course, due to the stochastic nature of the model, this quantity is a random variable and so we define $q_n(\tau, \Omega)$ to describe the corresponding probability of it taking the value n . More specifically, we shall interact with this quantity through its corresponding probability generating function (PGF) [41] given by

$$H(\tau, \Omega; x) = \sum_{n=1}^{\infty} q_n(\tau, \Omega) x^n, \quad (1)$$

with final condition $H(\Omega, \Omega; x) = x$, i.e., a variant which is created through a mutation event just as the process concludes would have an abundance of one. We now examine the functional form of the PGF that results from the mechanisms of Simon's model (see Fig. 1). In particular, we consider how the distribution of a variant's abundance may change over a short time interval $(\tau - \Delta t, \tau)$. Within this window there are three possible options from which the variant of a new element may be determined.

(i) First, the new element may be a result of a reproduction event w.p. $1 - \mu$ and the element which we are focusing on




FIG. 1. Schematic illustration of the classical Simon's model and its corresponding branching structure. At each discrete time interval of length Δt , a new element is created either via reproduction such that it copies the variant of a previous element or mutation where it is an entirely novel variant. The variants are represented here by colors, and we focus on the dynamics of the red variant that first appeared via mutation at time τ . A reproduction event occurs at time t_1 whereby the new element randomly chooses from all those which had appeared previously (indicated by the dashed line) to determine its type of variant, which is the red element of interest. This results in a new branch being created in the branching process interpretation (bottom plot). The branch may itself reproduce in the future, as occurs at time t_2 while the original element may also still produce progeny, as occurs at time t_3 . The quantity of interest is the abundance of this variant at time Ω : in the realization shown, the final abundance (total number of red elements in the sequence) is four.

$$H(\tau - \Delta t, \Omega; x) = (1 - \mu) \left(\frac{\Delta t}{\tau} \right) [H(\tau, \Omega; x)]^2 + (1 - \mu) \left(\frac{\tau - \Delta t}{\tau} \right) H(\tau, \Omega; x) + \mu H(\tau, \Omega; x), \quad (2)$$

and if we proceed to consider the scenario in which $\Delta t \ll \tau$ we may move to a continuum time limit that proves more tractable for mathematical analysis:

$$-\frac{dH}{d\tau} = \frac{1 - \mu}{\tau} (H^2 - H). \quad (3)$$

This is an ordinary differential equation under the assumption of fixed Ω , where $H = H(\tau, \Omega; x)$. This equation is in fact of the Riccati form which, when combined with the aforementioned final condition, may be solved to obtain an exact representation of the PGF's functional form by

$$H(\tau, \Omega; x) = \frac{x}{x \left[1 - \left(\frac{\tau}{\Omega} \right)^{-(1-\mu)} \right] + \left(\frac{\tau}{\Omega} \right)^{-(1-\mu)}}. \quad (4)$$

With this solution for the PGF at hand we may proceed to conduct analysis regarding the underlying distribution of variant abundance.

A. Analytically obtained statistical properties

The advantage of having obtained a functional representation regarding the dynamics of Simon's model is that it provides information regarding the entire distribution of abundances and is also extremely amenable to quantitative analysis. We now proceed to highlight some, although far from all, of these possible analyses regarding the statistical behavior of the variants' abundance.

is chosen w.p. $\Delta t/\tau$, this being the time during which the element was present Δt divided by the entire time τ of the process until this point. The reproduction event results in a further element of the variant which itself may reproduce in the future thus contributing $[H(\tau, \Omega; x)]^2$ to the PGF, where we have used the fact that both elements are now identical due to the variant abundance being the only factor influencing future copying events alongside the property that the PGF of the sum of two random variables is simply the product of their corresponding PGFs [41].

(ii) Second, the new element may again be a result of reproduction w.p. $1 - \mu$ but with the source of the reproduction (copying) being another element rather than the one under consideration. Such an copying event occurs w.p. $(\tau - \Delta t)/\tau$, i.e., the fraction of time the other elements were present for. In this scenario a further branch is not created and as such the contribution to the PGF of interest is simply $H(\tau, \Omega; x)$, i.e., a single branch that may result in future copying events.

(iii) Finally, the new element may be as a result of a mutation, w.p. μ . This of course results in a different variant to that under consideration and as such again contributes $H(\tau, \Omega; x)$ to the PGF.

Taking these terms and their corresponding probabilities together we produce the following difference equation describing the time evolution of the PGF:

1. Distribution of variant abundance for given seed time

First of all we consider the probability distribution of a variant's abundance for a given seed-time. Specifically we focus on the previously defined probability $q_n(\tau, \Omega)$ which may be accessed through the PGF H via the well-known property

$$q_n(\tau, \Omega) = \frac{1}{(n-1)!} \left. \frac{\partial^n}{\partial x^n} \left[\frac{H(\tau, \Omega; x)}{x} \right] \right|_{x=0}, \quad (5)$$

such that knowledge of the n th derivative of Eq. (4) is required. Generally this distribution is completed numerically or, to avoid the inaccuracies caused from numerical differentiation, via transformation to the complex plane from which Cauchy's theorem may be used [38,42]. Fortunately, in this specific case, it can be seen that the PGF in fact describes a geometric distribution (with an additional power of x) such that the probability mass function can be written explicitly as

$$q_n(\tau, \Omega) = \left(\frac{\tau}{\Omega} \right)^{1-\mu} \left[1 - \left(\frac{\tau}{\Omega} \right)^{1-\mu} \right]^{n-1}. \quad (6)$$

An array of analysis is possible with the analytical distribution at hand, for example we may consider the likelihood that a variant which first appeared at time τ does not appear again by time Ω which is given by

$$q_1(\tau, \Omega) = \left(\frac{\tau}{\Omega} \right)^{1-\mu}. \quad (7)$$

Furthermore, we may readily calculate the moments of abundance size directly from the distribution, however, with the analysis to come later in this paper in mind, we will now focus on obtaining such quantities directly from the PGF.

2. Moments of variant abundance distribution

We first consider the average number of times $m(\tau, \Omega)$ that a variant that first appeared at time τ has appeared across all elements by time Ω . Taking advantage of the fact that the distribution's moments are readily attainable through repeated differentiation of Eq. (4) with respect to x , the average or first moment is obtained through the following calculation

$$m(\tau, \Omega) = \left. \frac{\partial H(\tau, \Omega; x)}{\partial x} \right|_{x=1} = \left(\frac{\tau}{\Omega} \right)^{-(1-\mu)}. \tag{8}$$

Interestingly, we note that this quantity is the inverse of the probability that the variant had only appeared once in the population given by Eq. (7). This expression highlights a number of facets within Simon's model, most obviously the "early mover" advantage suggesting that variants with earlier seed times τ will have, on average, appeared more times in a given realization of the process. We specifically highlight that the mean abundance of a variant in fact exhibits a power-law relationship with the seed time.

One final point we comment on regarding this moment is the average number of appearances for the variant describing the introductory element seeded at time $\tau = 1$, which is given by $(1/\Omega)^{-(1-\mu)}$. This expression is in general agreement with the analysis found in [19] (assuming that $\mu \ll 1$ such that $1/\Gamma(2 - \mu) \approx 1$) demonstrating the intrinsic advantage offered to the first variant to appear in a realization of Simon's model.

Of course we may obtain further moments of the distribution through similar calculations to those shown above. For example, the variance of variant abundance $v(\tau, \Omega)$ may be found by observing that

$$v(\tau, \Omega) = \frac{\partial^2 H(\tau, \Omega; x)}{\partial x^2} + \frac{\partial H(\tau, \Omega; x)}{\partial x} - \left[\frac{\partial H(\tau, \Omega; x)}{\partial x} \right]_{x=1}^2 = \left(\frac{\tau}{\Omega} \right)^{-1-\mu} \left[\left(\frac{\tau}{\Omega} \right)^{-1-\mu} - 1 \right], \tag{9}$$

which allows analysis regarding the dispersal of the abundance size to be readily calculated.

3. Distribution of variant abundances across all seed times

Simon's model originally found fame for its ability to generate power-law distributed popularity values when all variants within a given sequence are considered, regardless of their ages. This distribution $\tilde{q}_n(\Omega)$ thus focuses on all seed times rather than the distribution of variants with a given seed time as our preceding analysis has considered. The classic result is however directly obtainable from our analysis by taking the distribution for a given seed time from Eq. (6) and averaging over all seed times, which are uniformly distributed

over the interval $[0, \Omega]$, to obtain

$$\begin{aligned} \tilde{q}_n &= \frac{1}{\Omega} \int_0^\Omega q_n(\tau, \Omega) d\tau \\ &= \frac{1}{\Omega} \int_0^\Omega \left(\frac{\tau}{\Omega} \right)^{1-\mu} \left[1 - \left(\frac{\tau}{\Omega} \right)^{1-\mu} \right]^{n-1} d\tau \\ &\sim \frac{1}{1-\mu} \beta \left[\frac{2-\mu}{1-\mu}, n \right], \quad \text{for large } n, \end{aligned} \tag{10}$$

where β represent the beta distribution and by consequently making use of the approximation $\beta(a, b) \sim b^{-a}$ for large b , we obtain

$$\tilde{q}_n \sim \frac{1}{1-\mu} n^{-\left(\frac{2-\mu}{1-\mu}\right)}, \tag{11}$$

the power-law distribution with exponent larger than two governing the entire population sequence as originally demonstrated by Simon.

B. Numerical simulations

We now provide results from numerical simulations of the classical Simon's model with the aim of validating the branching process approach described in the preceding section. Specifically, we compute an ensemble of 10^6 realizations of the process with length Ω before considering the quantities derived above. In general, when considering a given seed time τ we force a variant to be created at this time but otherwise the simulations are exactly as described by Simon in his original work. We first consider the overall distribution of variant abundance originally shown to result in a power-law distribution for larger abundances and as described through Eq. (11). The corresponding distribution is shown in Fig. 2(a) where simulations have been performed with $\Omega = 5 \times 10^5$ and $\mu = 0.01$.

We now consider the first original result obtained as a consequence of our approach. Specifically, the distribution of abundances for a variant seeded at time τ given by Eq. (6) is considered for a number of different seed times $\tau = 5 \times 10^2, 10^3, 5 \times 10^3, \text{ and } 10^4$. We then observe the distribution of abundances observed after 2.5×10^4 time units in each scenario such that each seed time is considered with the same age at observation time $\Omega = \tau + 2.5 \times 10^4$. The results of said simulations are shown in Fig. 2(b) where we observe the 'early mover' advantage offered by Simon's model which has been a motivating factor in similar preferential attachment based models. Figure 2(c) considers a similar framework but instead focuses on how the probability of a variant having abundance n varies as a function of seed time, where in each case $\Omega = 10^5$. We again observe the early mover advantage here with those large n values being more likely for the early seed times, we also show the probability of the variant not being used again after creation given by Eq. (7). Lastly, Fig. 2(d) shows the average popularity of a variant seeded at time τ and observed at time $\Omega = 10^5$ along with the theoretical calculation given by Eq. (8), demonstrating the power-law relationship between the two quantities. We conclude by commenting on the generally excellent agreement between simulations and the theoretical results offered by the branching process interpretation of Simon's classic model.

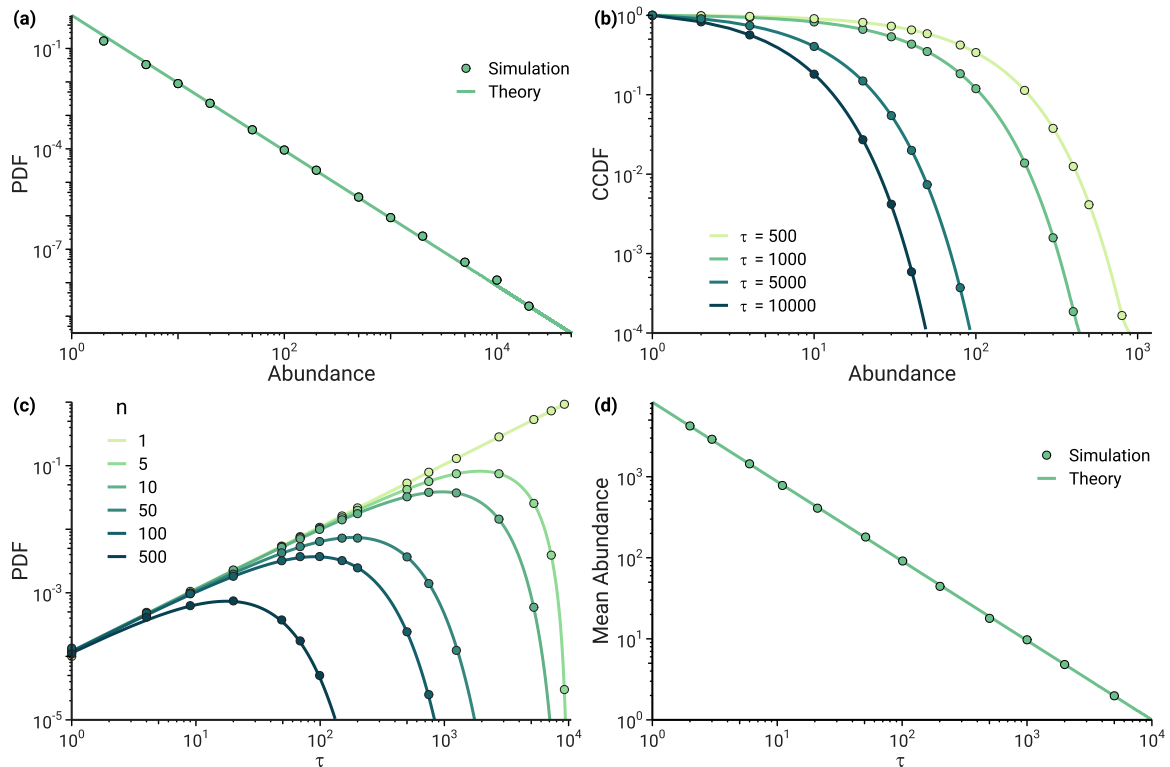


FIG. 2. Simulations and corresponding analytical results for the classical Simon's model. (a) Probability distribution across all abundances of variants within realizations of Simon's model with $\Omega = 5 \times 10^5$ and $\mu = 0.01$. (b) The complementary cumulative distribution of variant abundance based upon their seed time ($\tau = 5 \times 10^2$, 10^3 , 5×10^3 , and 10^4) where in each case the observation time is 2.5×10^4 time units after the seed time. The early mover advantage is observed: those variants with earlier seed times have larger likelihood of larger abundance despite all having the same age. (c) Probability distribution function of variant abundance dependent upon the variant's seed time for a number of different abundance values with $\Omega = 10^5$. (d) The average popularity of a variant with $\Omega = 10^5$ in each case for multiple seed times is shown. The power-law relationship with seed time is demonstrated, alongside theoretical values given by Eq. (8). Each of the above simulations are performed over 10^6 ensembles.

III. A GENERALIZED SIMON'S MODEL

One of the enticing properties of Simon's original model is the simple manner in which previously used elements are chosen from when a reproduction event occurs. Specifically, the choice is made over all previously used elements such that the likelihood of choosing a certain variant is proportional to the number of times that variant has appeared prior to the choice, resulting in a preferential attachment form of dynamics. While such properties have resulted in an extensive array of literature in more recent times the concept of temporal memory—more formally known as the burstiness [43–45]—underlying how one makes decisions when copying has become increasingly important from a modeling perspective.

Accordingly, here we focus on such aspects by proposing a generalized Simon's model. In this model, when a new element from a reproduction event is choosing their variant type from the previously used elements the likelihood of copying from an element which was present t time units prior is proportional to a given function $\phi(t)$. This probability may be viewed as a memory function through assuming the property $\int_0^\infty \phi(w) dw = 1$. The quantity of interest is, as before, the probability that a variant which first appeared at time τ has

been used n times by time Ω , denoted by $q_n(\tau, \Omega)$. Incorporation of the temporal element of memory requires some further considerations when performing calculations for this quantity. Specifically, we start by considering the PGF $H(\tau, a, \Omega; x)$ of the random variable describing the number of times that an element, seeded at time τ due to a reproduction event from an element which had age a at the time (i.e., first appeared itself at time $\tau - a$), has caused further elements to have the same variant type. Now, as in the classical case of Sec. II, we consider how this PGF changes over the time interval $(\tau - \Delta t, \tau)$. If Δt is sufficiently small such that at most one event occurs in said time interval, there are three possible events.

(1) First, there may be a reproduction event w.p. $1 - \mu$ where the new element chooses to copy the element of interest's variant by looking back to the time which they appeared w.p. $\phi(a)\Delta t$. This results in a new element of age $a = 0$ with the same variant that may generate further progeny, the total size of which would be added to the size of the previously seeded element and as such this event contributes $H(\tau, a, \Omega; x)H(\tau, 0, \Omega; x)$ to the PGF $H(\tau - \Delta t, a - \Delta t, \Omega; x)$.

(2) There may again be a reproduction event w.p. $1 - \mu$ but in this instance the new element copies its variant from

an alternative element w.p. $1 - \phi(a)\Delta t$, which does not affect the abundance resulting from the element of interest, thus the contribution to its PGF is simply $H(\tau, a, \Omega; x)$.

(3) Finally, there may be a mutation event rather than a reproduction w.p. μ which again does not change the abundance

$$H(\tau - \Delta t, a - \Delta t, \Omega; x) = (1 - \mu)\phi(a)H(\tau, a, \Omega; x)[H(\tau, 0, \Omega; x) - 1]\Delta t + H(\tau, a, \Omega; x). \tag{12}$$

Taking, as in Sec. II, the case where $\Delta t \ll \tau$, and using a two-dimensional Taylor approximation gives the following partial differential equation:

$$\frac{\partial H(\tau, a, \Omega; x)}{\partial \tau} + \frac{\partial H(\tau, a, \Omega; x)}{\partial a} = (1 - \mu)\phi(a)H(\tau, a, \Omega; x)[1 - H(\tau, 0, \Omega; x)], \tag{13}$$

which may be solved via the method of characteristics (we refer the interested reader to [46] for a similar calculation) to obtain

$$\tilde{H}(t, a; x) = x \exp \left\{ (1 - \mu) \int_0^t \phi(u + a) [\tilde{H}(t - u, 0; x) - 1] du \right\}, \tag{14}$$

where $\tilde{H}(\Omega - \tau, a; x) = H(\tau, a, \Omega; x)$, i.e., the functional form of this PGF now depends only upon the time interval $t = \Omega - \tau$ between seeding and observation.

This PGF describes the distribution of abundance as a result of a new element's variant being chosen by copying from a previous element of age a . We are generally interested in the first such seeding i.e., by an element with age $a = 0$, to obtain information regarding the entire tree size, and as such we may drop the explicit use of a and obtain

$$\tilde{H}(t; x) = x \exp \left\{ (1 - \mu) \int_0^t \phi(u) [\tilde{H}(t - u; x) - 1] du \right\}. \tag{15}$$

We note the similarity between this distribution and Eq. (8) of [46], indicating that the GSM may be viewed as a specific version of the self-exciting component within a Hawkes process [47] for a given variant. As such, previous results regarding the predictability of future popularity size of said processes are readily applicable to abundance size in the case of GSM. This apparent equivalence is rather intriguing and demonstrates a deep underlying relationship between age-dependent branching processes and a range of random-copying continuous processes. We do note however that this equivalence arises from the branching process assumption, which describes the Hawkes process exactly, but in the case of Simon's model is an approximation as during each time step only one event should occur. This violates the independence assumption of the branching process, although in Appendix B we demonstrate that the errors arising from this approximation in the classical model are minor.

A. Analysis

Having obtained the functional form of the PGF through Eq. (15) for the GSM we may now proceed to consider some possible analytical quantities obtainable regarding the process.

1. Mean popularity

As in Sec. IIA2, the mean abundance of a given variant after t time units have passed since it first appeared may be

size resulting from the element under consideration here thus contributes $H(\tau, a, \Omega; x)$ to the PGF.

Taking these terms alongside their corresponding probabilities leads to the following difference equation describing the time evolution of the PGF:

obtained directly from Eq. (15) through the following calculation

$$\begin{aligned} m(t) &= \left. \frac{\partial \tilde{H}(t; x)}{\partial x} \right|_{x=1} \\ &= 1 + (1 - \mu) \int_0^t \phi(w) m(t - w) dw, \end{aligned} \tag{16}$$

which does not, in general, have an analytical solution. We may, however, conduct analysis in the Laplace space to determine properties of this quantity as follows:

$$\hat{m}(s) = \frac{1}{s} + (1 - \mu) \hat{\phi}(s) \hat{m}(s), \tag{17}$$

where hatted quantities represent their Laplace transform $\hat{\phi}(s) = \int_0^\infty \phi(t) e^{-st} ds$. We proceed to solve this equation for the Laplace-transformed average popularity to obtain

$$\hat{m}(s) = \frac{1}{s[1 - (1 - \mu)\hat{\phi}(s)]}. \tag{18}$$

From this expression we may immediately calculate, by means of the final-value theorem, the average abundance in the infinite time limit as follows:

$$\begin{aligned} m(\infty) &= \lim_{t \rightarrow \infty} m(t) = \lim_{s \rightarrow 0} s \hat{m}(s) \\ &= \frac{1}{\mu}. \end{aligned} \tag{19}$$

More generally (aside from specific distributions, e.g., the exponential, which may be handled analytically) the finite-time behavior of this quantity is obtained via numerical inversion of the Laplace transform given by Eq. (16). This is readily performed using the efficient Talbot algorithm [39,48,49].

One final consideration we describe here is the relationship between the mean memory-time $T = \int_0^\infty t \phi(t) dt$ and the corresponding large-time mean abundance. This is studied by considering the small- s behavior of Eq. (18) in particular

noting that

$$\begin{aligned} \hat{\phi}(s) &\approx \hat{\phi}(0) + s \frac{d\hat{\phi}(0)}{ds} \\ &= 1 - sT, \end{aligned} \tag{20}$$

where we have used the fact that $\hat{\phi}(0) = \int_0^\infty \phi(t) dt = 1$ and $\frac{d\hat{\phi}(0)}{ds} = \int_0^\infty -t\phi(t) dt = -T$. Using this approximation in Eq. (18) allows the mean abundance of a given variant t time units after first appearing to be written explicitly as

$$m(t) \approx \frac{1}{\mu} [1 - e^{-\mu t / (1-\mu)T}], \quad t \gg 1. \tag{21}$$

Furthermore, in the case of an exponential memory kernel with mean time β , i.e., $\phi(t) = \frac{1}{\beta} e^{-t/\beta}$, we may exactly determine the mean abundance through inversion of Eq. (18) to obtain

$$m(t) = \frac{1}{\mu} [1 - (1 - \mu)e^{-(\mu/\beta)t}]. \tag{22}$$

Indeed, this expression for the exponential memory kernel may be expanded to first order in the case where $t \ll \beta/\mu$ giving

$$m(t) \approx 1 + \frac{(1 - \mu)}{\beta} t, \tag{23}$$

i.e., a linear function of time with an inverse dependence upon the mean memory-time of the distribution.

2. Infinite-age distribution

Analytically obtaining the probability distribution associated with the PGF given by Eq. (15), unlike that for the classical model, generally proves infeasible as a consequence of its convoluted nature. We may, however, conduct asymptotic analysis regarding the distribution in the case of those variants with large age [38]. Specifically we proceed to assume that $\tilde{H}(t;x) \rightarrow \tilde{H}_\infty(x)$, independent of t , as $t \rightarrow \infty$ and make use of the fact that $\int_0^\infty \phi(w) dw = 1$ in Eq. (15) to obtain

$$\tilde{H}_\infty(x) = x \exp \{ (1 - \mu) [\tilde{H}_\infty(x) - 1] \}. \tag{24}$$

The large- n behavior of the underlying distribution $q_n(\infty)$ may be studied by letting $x = 1 - w$ and $\tilde{H}_\infty(x) = 1 - \gamma(w)$ before proceeding to make use of the small- w and small- γ asymptotics of Eq. (24). Through this analysis (detailed in Appendix C) it may be shown that the large- n abundance distribution is described by a power law with exponential cutoff

$$q_n(\infty) \sim A n^{-3/2} e^{-n/\kappa}, \quad \text{as } n \rightarrow \infty, \tag{25}$$

where the prefactor A is given by

$$A = \sqrt{\frac{1 - \mu}{2\pi}}, \tag{26}$$

and the cutoff scale κ is

$$\kappa = \frac{2(1 - \mu)^2}{\mu^2}. \tag{27}$$

We note that this heavy-tailed behavior is distinct from that found in Simon's original analysis, as in the original case Simon considered an abundance distribution across all variants within the system whereas we focus on a single variant. The present result however suggests that the incorporation of a decaying memory within the model results in a heavier tail describing the large- n behavior of the abundance distribution even in the case of a single variant. A similar result has been shown in Ref. [30] for a specific memory function (see Sec. V) which also resulted in a power law with an exponent of $-3/2$ and exponential cutoff, the results presented here therefore demonstrate the generalization of this behavior to the case of an arbitrary memory kernel.

IV. SIMULATIONS

We now proceed to consider different forms of the GSM via numerical simulations alongside the analytical estimates of their behavior obtained in the preceding section. In each case we generate sequences of elements similar to those in Sec. II B but now, in the case of a replication event, the time from which to choose a target element for copying is a random variable drawn from the specified memory distribution. Due to the generally infinite support of the GSM kernel there are times in which the copying time may occur prior to the start time of the system; we consider such cases as mutation events due to elements from prior to the system's start time not being of the same variant as that in which we are interested.

First, we simulate processes of length 10^4 with an exponentially distributed memory kernel with mean time β and innovation probability $\mu = 0.01$. The average number of times the variant has been copied by the time it has age t , i.e., when there have been t further elements created, is shown in Fig. 3(a) for $\beta = 100, 300, 500$, along with the theoretical prediction given by Eq. (22). We point out the apparent inverse relationship between mean memory time and average abundance of a given variant, suggesting that shorter memory times favor the likelihood of a variant having larger popularity. We proceed to consider the distribution of abundances across these realizations in the case $\beta = 300$ at a number of different ages $t = 100, 500, 1000, 5000$ in Fig. 3(b), where the complementary cumulative distribution function is shown along with the corresponding theoretical estimates obtained from inversion of Eq. (15). Lastly we consider, as in Fig. 2(b), two variants seeded at different times $\tau = 500$ and 5000 and observe the abundance of the variant at $\Omega = \tau + 50000$, i.e., when they both have the same age. Now, unlike in the classical Simon's model, the distributions are equivalent such that there is no longer a early mover advantage for variants.

We now proceed to consider an alternative memory kernel to demonstrate the robustness of our GSM formalism, specifically we consider a gamma-distributed memory given by

$$\phi(t) = \frac{t^{k-1} e^{-t/\theta}}{\theta^k \Gamma(k)}, \tag{28}$$

where k is called the shape parameter and θ the scale parameter. This distribution offers a wide-range of behaviors, the effect of which upon the GSM may be observed. In particular, the average memory time is given by $T = k\theta$

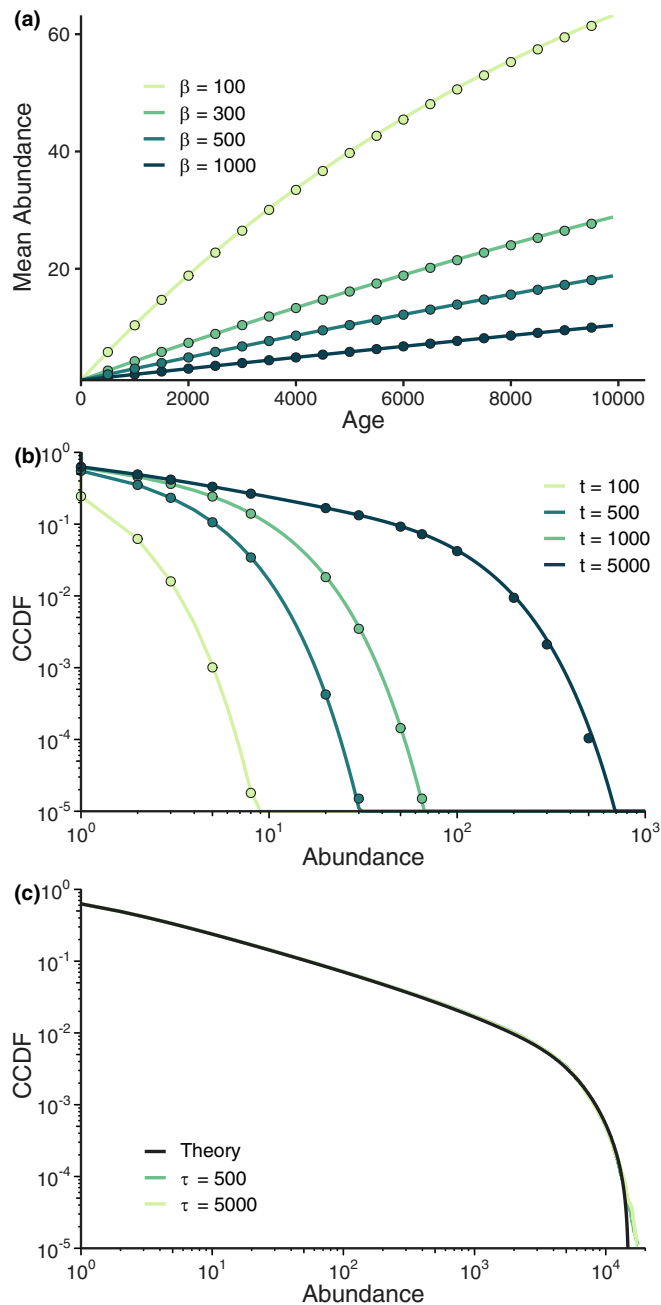


FIG. 3. Simulations and corresponding analytical results for the generalized Simon's model in the case of an exponential memory kernel. (a) Average abundance of a given variant as a function of its age obtained over 10^6 realizations for various β values and $\mu = 0.01$; the dots represent simulated values and curves are the theoretical results. (b) CCDF of variant abundance at a number of different ages with $\beta = 300$, alongside theoretical distributions from inversion of Eq. (15). (c) CCDF of abundance for two variants which have been seeded at different times ($\tau = 500$ and 5000) but have been observed after the same amount of time has passed after seeding, resulting in identical distributions. Also note the power-law along with exponential cutoff behavior as predicted by Eq. (25).

while the structure of the distribution can vary widely from exponential (when $k = 1$) to more symmetrically distributed about the mean when $k > \theta$. Figure 4 again shows the

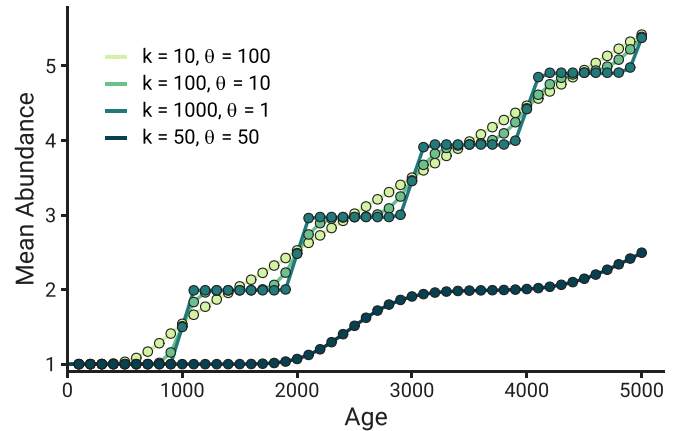


FIG. 4. Simulations and corresponding analytical results for the average abundance of the generalized Simon's model in the case of a Gamma distributed memory kernel. A number of parameter combinations for the distribution are used. We see how the general trend of the average abundance is represented through the mean $k\theta$ of the Gamma distribution while the temporal dynamics of the abundances fluctuate based upon its shape. In each case, the average is computed over 10^6 realizations with $\mu = 0.01$.

average abundance of a variant simulated with this process over 10^6 realizations. Two features are immediately noted, firstly we see that the general trend of the abundance is dependent upon the mean memory-time whereby the cases where $k = 10, 100, 1000$ and $\theta = 100, 10, 1$ all have an equivalent general trend while the process with larger mean memory-time ($k = \theta = 50$) demonstrates greater average abundance growth. The second interesting aspect here is the effect of the shape of the memory distribution upon the dynamics of the average abundance. In particular, the more skewed scenario ($k = 10, \theta = 100$) results in linear-like growth as observed in the exponentially distributed case. On the other hand, when the memory distribution is more symmetric the likelihood of a new element replicating the variant of interest is most likely at time units close to the average memory times, resulting in the almost stepwise behavior of the averages. Importantly, in this case, there is no closed-form solution for the theoretical expression of the expected abundance and the results are obtained via numerical inversion of Eq. (18).

V. SPECIAL CASES

In the preceding section, we proposed the incorporation of a general memory kernel within a Simon's model framework. Prior literature has considered specific kernels, including Refs. [28,29] where a hyperbolic memory was studied, and Ref. [30] which, like the original Simon's model, utilizes a uniform memory that is, however, bounded such that elements are not considered after a certain amount of time has passed. Here we demonstrate how our proposed framework can describe these generalizations (particularly in the case of the latter) in continuous time while also allowing the determination of additional information regarding the statistical properties of the abundance distribution.

A. Schaigorodsky *et al.*

In Ref. [30], Schaigorodsky *et al.* propose an extension to Simon's model which they coin bounded memory preferential growth (BMPG) whereby their choice of memory kernel is given by a step function

$$\phi(t) = \begin{cases} \frac{1}{\kappa} & \text{if } t \leq \kappa, \\ 0 & \text{if } t > \kappa. \end{cases} \quad (29)$$

This implies that the variant of a new element is uniformly chosen from the previous κ elements (with $\Delta t = 1$). We may proceed, as before, to consider the average popularity of a variant with age t by calculating the Laplace transform of Eq. (29) given by

$$\hat{\phi}(s) = \frac{1}{\kappa s} [1 - e^{-\kappa s}], \quad (30)$$

before substituting in Eq. (18) to obtain

$$\hat{m}(s) = \frac{1}{s - \frac{(1-\mu)}{\kappa} [1 - e^{-\kappa s}]}, \quad (31)$$

which gives the Laplace transform of the mean abundance size at a time t after the variant first appears in the BMPG model. While this quantity itself does not have an analytical inverse transform, it is still possible to invert it numerically as described in Sec. IIIA1. Furthermore, we may perform approximate analysis regarding the mean abundance, as per Eq. (21), which is a valid approximation here in the regime where $t \gg \kappa$. This calculation depends on the mean-memory time which in the case of BMPG is given by $T = \kappa/2$ and the corresponding expected abundance for a variant with age t is given by

$$m(t) \approx \frac{1}{\mu} \left[1 - e^{-\frac{2\mu}{(1-\mu)\kappa} t} \right]. \quad (32)$$

If we then consider the Taylor series of this function about $t = 0$, particularly focusing on the linear term, which is valid when $t \ll \frac{(1-\mu)}{2\mu} \kappa$, we obtain

$$\begin{aligned} m(t) &\approx \frac{1}{\mu} \left[1 - \left(1 - \frac{2\mu}{(1-\mu)\kappa} t \right) \right] \\ &\approx \frac{2}{(1-\mu)\kappa} t, \quad \kappa \ll t \ll \frac{(1-\mu)}{2\mu} \kappa. \end{aligned} \quad (33)$$

Furthermore, we readily obtain the entire distribution of abundances for variants of any age by substituting Eq. (29) into Eq. (15) and inverting numerically as before, from which moments of the distribution may also be calculated.

In order to validate the applicability of our GSM to this scenario, we proceed to conduct numerical simulations of the process, with results shown in Fig. 5. Each simulation is averaged over 10^6 realizations with the dots representing simulated values and the curves their theoretical equivalents. Figure 5(a) shows the average abundance of a variant for various κ values where, as per Ref. [30], in each case the innovation probability μ is the reciprocal of the memory length κ . The theoretical results are determined via numerical inversion of Eq. (31) and generally offer excellent agreement with simulation. We also observe that in this case the linear approximation offered by Eq. (33) holds within the time regime




FIG. 5. Analysis of the bounded memory preferential growth (BMPG) model alongside analytical estimates from the branching process description. (a) Average abundance of a given variant as a function of its age obtained over 10^6 realizations for a number of κ values representing the length of the memory kernel. Here $\mu = \kappa^{-1}$ and the theoretical curves are from Eq. (31). (b) PDF of variant abundance at a number of different ages in the case $\kappa = 100$, with theoretical distributions obtained from inversion of Eq. (15).

$\kappa \ll t \ll \kappa^2/2$, which supports the linear behavior observed in the mean abundance and furthermore gives an estimate for the slope of these functions to be approximately $2/\kappa$. The corresponding probability distribution functions of abundance size at a number of different ages are shown in Fig. 5(b) with $\kappa = 100$. We also observe the “peloton dynamics” reported in Ref. [27] whereby there are bulges within the distribution as it converges towards its large-time limit.

B. Cattuto *et al.*

A further study which analyses a memory-dependent Simon's model may be found in the works [28,29] by Cattuto *et al.* in which the memory kernel is defined as

$$\phi(t) = \frac{C}{\gamma + t}, \quad (34)$$

where γ represents a timescale characteristic to the system under consideration and C is simply a normalization factor determined by the condition $\int_0^{t_{\max}} \phi(w) dw = 1$, where t_{\max} is the current total time of the process, i.e., the memory distribution

is varying at each step. This specific time-dependent memory kernel is beyond the scope of our proposed framework in that we assume a fixed kernel defined such that $\int_0^\infty \phi(w) dw = 1$. We may, however, consider a similar constant kernel for which the support is bounded by some upper time limit t^* whereby C is determined to ensure $\int_0^{t^*} \phi(w) dw = 1$. Numerical simulations (omitted from the present paper) of the GSM with this memory kernel, like those studied above, offer perfect agreement with theoretical estimates of the abundance size distribution obtained from inversion of Eq. (15). Exact agreement with the original model proposed in Refs. [28,29] however would require the incorporation of a time-varying memory function within our framework which is beyond the scope of the present paper.

VI. CONCLUSIONS

In this paper, we have proposed a branching process framework to describe Simon's renowned random-copying model [4] and generalizations thereof. This mapping is determined by identifying the replication events of the original model as equivalent birth events in the branching process. Through this description we obtain a detailed understanding of the temporal evolution of a new variant's abundance via a knowledge of its probability distribution under the branching process assumption. Using this analysis, we reproduce classical results while also demonstrating the potential to determine further statistical properties. These findings are validated via extensive numerical simulations which ratify this approach in describing the model. Interestingly, in spite of the model's discrete nature, we find that the continuous time description suggested by our framework is extremely useful in further understanding properties of the original model.

Motivated by the branching process framework accurately describing Simon's original model (see Appendix B for a discussion on the accuracy of this interpretation) we proceeded to extend beyond the classical uniform random-copying probability to any arbitrary age-dependent memory kernel resulting in our GSM framework. Such considerations have previously been studied in the case of specific kernels [28–31] and we now extend this to the general scenario. Using these ideas we demonstrate the potential for rich behaviors in such models that are in strong agreement with numerical simulations. Indeed, as per the classical model, it is possible to determine statistical properties of the process via analytical (or semi-analytical) calculation rather than through extensive Monte Carlo simulations. Lastly, we highlight the fact that power-law distributions of abundance size, as originally highlighted in Simon's model, occur for a different reason in the GSM for a variant with large age, albeit now with a smaller exponent (3/2) and an exponential tail.

Given the ubiquitous use of Simon's model as a basis for a range of different frameworks, including growing networks [23,27] and in describing the abundance distributions in numerous empirical systems [20–22], we believe that the theory introduced in this paper offers the potential for much further exciting research. An important example is the incorporation of memory effects [39,45] within models for dynamical systems describing the behaviors of human activity. Lastly, while Simon's model is arguably the most notable neutral

model [16], there are many other similar models which we anticipate to benefit from our framework in further understanding the temporal dynamics of variant abundance size.

ACKNOWLEDGMENTS

Helpful comments from two anonymous reviewers are gratefully acknowledged. This work was supported by Science Foundation Ireland Grants No. 16/IA/4470, No. 16/RC/3918, and No. 12/RC/2289P2 (J.D.O'B. and J.P.G.). We acknowledge the DJEI/DES/SFI/HEA Irish Centre for High-End Computing (ICHEC) for the provision of computational facilities and support.

APPENDIX A: ALTERNATIVE DERIVATIONS OF EQ. (4)

We shall now consider alternative approaches to the PGF, given by Eq. (4) in the main text, from which the probability distribution associated with the classical Simon's model may be derived. First we define $q_n(t)$ to be the probability of a specific variant having abundance n at time t . We now consider the different ways in which the variant could have appeared n times by time t .

(1) The variant could have appeared $n - 1$ times by time t and then be chosen again as a result of a new element replicating from one of the previous elements, with probability $(1 - \mu)\binom{n-1}{t}$.

(2) The variant could have instead appeared n times by time t and a new element again chooses to replicate, but this time chose an alternative variant with probability $(1 - \mu)(1 - \frac{n}{t})$.

(3) Finally, the variant could have appeared n times by time t and the new element was a mutation with probability μ , thus not creating another copy of the variant in question.

Combining these possibilities gives us the following master equation:

$$\begin{aligned} q_n(t+1) &= (1 - \mu)\binom{n-1}{t}q_{n-1}(t) \\ &\quad + (1 - \mu)\left(1 - \frac{n}{t}\right)q_n(t) + \mu q_n(t) \\ &= q_n(t) + \frac{1 - \mu}{t+1}[(n-1)q_{n-1}(t) - nq_n(t)]. \end{aligned} \quad (\text{A1})$$

With this equation at hand, we now proceed to consider two mathematical approaches from which we may obtain the corresponding PGF.

1. Generating function approach

Defining $G(t; x) = \sum_n q_n(t)x^n$ to be the PGF of the random variable $q_n(t)$, we then multiply Eq. (A1) by x^n and sum over all n to obtain

$$G(t+1; x) = G(t; x) + \frac{1 - \mu}{t}(x^2 - x)\frac{\partial G(t; x)}{\partial x}, \quad (\text{A2})$$

and considering the large-time case such that the change in one time unit is relatively small we obtain

$$\frac{\partial G(t;x)}{\partial t} = \frac{1-\mu}{t}(x^2-x)\frac{\partial G(t;x)}{\partial x}. \quad (A3)$$

We now also note that the initial condition governing this first order homogeneous partial differential equation (PDE) is dependent on when the variant first appeared in the sequence, which we shall call time τ , such that $G(\tau;x) = x$. To solve this PDE, we use the method of characteristics noting that the two equations governing the shape of the characteristic curve are given by

$$\frac{dt}{dr} = \frac{t}{1-\mu}, \quad \frac{dx}{dr} = x(1-x), \quad (A4)$$

which are readily solved to give

$$r = (1-\mu)\ln(t), \quad r = \ln\left(\frac{x}{1-x}\right) + C, \quad (A5)$$

where C is the constant of integration. We now note that due to the PDE being homogeneous, $G(t;x)$ is a constant along the characteristic curve and we can deduce the general solution to be given by

$$G(t;x) = F(C) = F\left[(1-\mu)\ln(t) - \ln\left(\frac{x}{1-x}\right)\right]. \quad (A6)$$

As we are aware of the initial condition of the differential equation, i.e., $G(\tau;x) = x$, we may determine the functional form of F , which allows us to conclude that

$$G(t;x) = \frac{x}{x\left[1 - \left(\frac{\tau}{t}\right)^{-(1-\mu)}\right] + \left(\frac{\tau}{t}\right)^{-(1-\mu)}}, \quad (A7)$$

in agreement with the PGF given by Eq. (4) obtained via the branching process analysis.

2. Inductive approach

Another possible approach to obtaining this distribution is to take Eq. (A1) and consider the small time change without directly transforming to a PGF in order to obtain

$$q'_n(t) + n\left(\frac{1-\mu}{t}\right)q_n(t) = (n-1)\left(\frac{1-\mu}{t}\right)q_{n-1}(t), \quad (A8)$$

which implies

$$\frac{d}{dt}\left[t^{n(1-\mu)}q_n(t)\right] = t^{n(1-\mu)}\left(\frac{1-\mu}{t}\right)(n-1)q_{n-1}(t). \quad (A9)$$

Again we note the dependence on the seed time τ , as $q_1(\tau) = 1$ and $q_0(t) = 1, \forall t < \tau$. Solving this equation we find the probability of a variant which first appeared at time τ having popularity n at time t to be given by

$$q_n(t) = \left(\frac{\tau}{t}\right)^{1-\mu} \left[1 - \left(\frac{\tau}{t}\right)^{1-\mu}\right]^{n-1}, \quad (A10)$$

as per Eq. (6) determined via the branching process description.

APPENDIX B: ERRORS FROM THE BRANCHING PROCESS APPROXIMATION

Throughout Sec. II of the main text we describe the classical Simon's model via a branching process approximation. It is important to note that this approach is, as stated, an approximation and in particular we consider the regime in which the time-step is smaller than the observation time, i.e., $\Delta t \ll \Omega$, enabling a continuous time analysis to be conducted. Simon's model is, however, a discrete process by nature. In spite of this we demonstrate, via numerical simulations, that the approach appears to work extremely well in terms of capturing the behavior of the abundance of a given variant. Here we consider the branching process approximation but in a discrete-time setting describing variants which appear near the observation time, for which the exact probabilities in both the BP setting and the model itself may be calculated and compared to determine the error associated with the approximation.

We define $G_\tau(\Omega) = G(\tau, \Omega; x)$ to be the PGF for the probability mass function describing the abundance of a variant which appeared at time step τ observed at time Ω . Using similar arguments to those found in the main text, we obtain the following difference equation to describe the dynamics of the process

$$G_\tau(\Omega) = G_{\tau+1}(\Omega) + \frac{(1-\mu)}{\tau} \{[G_{\tau+1}(\Omega)]^2 - G_{\tau+1}(\Omega)\}, \quad (B1)$$

where the squared term arises from the fact that a successful copying event can be construed as two separate trees beginning due to each element having the same probability of being copied in the classical model. Furthermore, we are aware of the final condition of this equation namely $G_\Omega(\Omega) = x$, i.e., a variant seeded at the observation time must have only appeared once. This equation, while due to its nonlinear nature being analytically troublesome, may be comfortably studied symbolically.

The approach taken above may seem extremely close to that underlying Simon's original model however the branching process description is not exact due to the fact that each tree created as a consequence of a successful copying event must be treated independently of one another from the branching process assumption. In the model itself, on the other hand, these trees are not independent as only one event may happen in a time step, i.e., if one element is successfully copied the other can not be.

In order to study the effect of this assumption in the branching process model we consider a variant which appeared at time $\tau = \Omega - 2$ and we are concerned with the probability mass function of its abundance at time Ω , i.e., after three potential copying events. Exact calculation of the quantities in the classical model, $q_n^{\text{exact}}(\Omega - 2, \Omega)$ requires one to consider the possible sample paths for a variant to be copied $n - 1$ times, for example, we now demonstrate how one may calculate $q_2^{\text{exact}}(\Omega - 2, \Omega)$:

TABLE I. Analysis of errors associated with the branching process approximation of Simon's classical model. The distribution of variant abundance from a mutation event at time $\Omega - 2$ observed at time Ω is determined. The probabilities are calculated via both the branching process model from Eq. (B1) and the exact values via the sample paths of the true process, alongside the error calculated as the difference between the two.

Abundance n	$q_n^{\text{exact}}(\Omega - 2, \Omega)$	$q_n^{\text{BP}}(\Omega - 2, \Omega)$	Error
1	$\frac{(\Omega - 3 + \mu)(\Omega - 2 + \mu)}{(\Omega - 2)(\Omega - 1)}$	$\frac{(\Omega - 3 + \mu)(\Omega - 2 + \mu)}{(\Omega - 2)(\Omega - 1)}$	0
2	$\frac{(1 - \mu)(2\Omega - 6 + 3\mu)}{(\Omega - 2)(\Omega - 1)}$	$\frac{(1 - \mu)(\mu^2 + 2\Omega^2 + 3\mu\Omega - 5\mu - 8\Omega + 7)}{(\Omega - 2)(\Omega - 1)^2}$	$\frac{-(1 - \mu)^3}{2(1 - \mu)^3}$
3	$\frac{2(1 - \mu)^2}{(\Omega - 2)(\Omega - 1)}$	$\frac{2(1 - \mu)^2(\Omega + \mu - 2)}{(\Omega - 2)(\Omega - 1)^2}$	$\frac{-(1 - \mu)^3}{2(1 - \mu)^3}$
4	0	$\frac{(1 - \mu)^3}{(\Omega - 2)(\Omega - 1)^2}$	$\frac{-(1 - \mu)^3}{(\Omega - 2)(\Omega - 1)^2}$

$$q_2^{\text{exact}}(\Omega - 2, \Omega) = \left\{ \underbrace{\frac{1 - \mu}{\Omega - 2}}_{\text{Successful copy at time } \Omega - 1} \times \underbrace{\left[\mu + (1 - \mu) \left(\frac{\Omega - 3}{\Omega - 1} \right) \right]}_{\text{Unsuccessful copy at time } \Omega} \right\} + \left\{ \underbrace{\left[\mu + (1 - \mu) \left(\frac{\Omega - 3}{\Omega - 2} \right) \right]}_{\text{Unsuccessful copy at time } \Omega - 1} \times \underbrace{\frac{1 - \mu}{\Omega - 1}}_{\text{Successful copy at time } \Omega} \right\}. \tag{B2}$$

Clearly, this approach itself becomes computationally expensive as we move the seeding event further backwards in time. Similarly, we iterate Eq. (B1) to obtain the equivalent probabilities for the branching process interpretation, $q_n^{\text{BP}}(\Omega - 2, \Omega)$, which are the coefficients of the resulting polynomial in x . Table I demonstrates the probability values obtained from both approaches alongside the error resulting from the branching process approximation. Importantly, the error scales inversely with the observation time to the power of the number of time steps—three in the above example—which, in the case of large observation time, explains the good agreement observed between the branching process approximation and numerical simulations.

APPENDIX C: DERIVATION OF LARGE POPULARITY BEHAVIOR

The large n behavior of $q_n(a)$ may be obtained in the $a \rightarrow \infty$ limit by asymptotic analysis of Eq. (24) around $x = 1$, via a similar approach as taken in Sec. S3 of Ref. [38]. Specifically, for x near 1, we let $w = 1 - x$ and suppose $\tilde{H}_\infty(x) \sim 1 - \psi(w)$ as $w \rightarrow 0$. Now Eq. (24) may be expressed as

$$1 - \psi = (1 - w) \exp \{ -(1 - \mu) \psi \}, \tag{C1}$$

Now let us consider the case where $0 < \mu \ll 1$, again taking $\tilde{H}_\infty(1 - w) = 1 - \psi(w)$ and Taylor expanding Eq. (24) about $w = 0$ and keeping terms of order w , ψ and ψ^2 but ignoring terms of order ψw we obtain

$$\frac{1}{2}(1 - \mu)^2 \psi^2 + \mu \psi - w = 0, \tag{C2}$$

which has solution

$$\psi = \frac{-\mu + \sqrt{\mu^2 + 2w(1 - \mu)^2}}{(1 - \mu)^2}, \tag{C3}$$

where we have taken the positive root as the PGF should be bounded above by one. Note that a branch point α exists in the

complex x plane at $x = 1 + \frac{\mu^2}{2(1 - \mu)^3}$. Now, using the Cauchy theorem, we may express the discrete probabilities in terms of the underlying generating function as

$$q_n = \frac{1}{2\pi i} \oint_C \frac{\tilde{H}_\infty(x)}{x^{n+1}} dx, \tag{C4}$$

where C is a closed contour in the complex plane which does not enclose any poles of \tilde{H} . We deform C such that it does not enclose the branch point α into the contour $C_R \cup l_2 \cup C_\epsilon \cup l_1$ as seen in Fig. 6.

It can be easily shown that the contributions to the integral from the circular contours C_R and C_ϵ limit to zero as $R \rightarrow \infty$




FIG. 6. The contour C in the complex x plane for the PGF inversion integral Eq. (C4). A branch cut extends from α to ∞ .

and $\varepsilon \rightarrow 0$ [50]. Also, the contribution from the analytical part of \tilde{H} will cancel due to the opposite directions of the line integrals. Therefore, as $w \rightarrow 0$, we have

$$q_n = \frac{-1}{2\pi i} \int_{l_1 \cup l_2} \frac{\sqrt{\mu^2 + 2(1-x)(1-\mu)^2}}{(1-\mu)^2 x^{n+1}} dx$$

$$= \frac{1}{2\pi i} \int_{l_1 \cup l_2} \frac{\sqrt{\mu^2 + 2w(1-\mu)^2}}{(1-\mu)^2 (1-w)^{n+1}} dw, \quad (C5)$$

now letting $w = 1 - e^\rho \approx -\rho$ as $\rho \rightarrow 0$, we obtain

$$q_n = \frac{-1}{2\pi i} \int_{l_1 \cup l_2} \frac{\sqrt{\mu^2 - 2\rho(1-\mu)^2}}{(1-\mu)^2} e^{-n\rho} d\rho,$$

and making the substitution $v = \mu^2 - 2(1-\mu)^2 \rho$ gives

$$q_n = \frac{1}{4\pi i(1-\mu)^4} \int_{l_1 \cup l_2} \sqrt{v} e^{-n\left(\frac{\mu^2-v}{2(1-\mu)^2}\right)} dv.$$

We now make the substitution $v = re^{i\theta}$ where along l_1 $\theta = \pi$, $v = -r$ and $\sqrt{v} = \sqrt{r} e^{i\frac{\pi}{2}}$ and along l_2 $\theta = -\pi$, $v = -r$

and $\sqrt{v} = \sqrt{r} e^{-i\frac{\pi}{2}}$, taking the limit as $R \rightarrow \infty$ we now consider the limits on the integral. In the x -plane l_1 varies from $1 + \frac{\mu^2}{2(1-\mu)^2}$ to ∞ and as such w varies from $-\frac{\mu^2}{2(1-\mu)^2}$ to $-\infty$, hence ρ takes values between $\frac{\mu^2}{2(1-\mu)^2}$ and ∞ , and finally v varies between 0 to $-\infty$. A similar line of thought allows one to see that along l_2 , v varies between $-\infty$ to 0. Thus our integral may now be expressed as

$$q_n = \frac{1}{2\pi(1-\mu)^4} e^{\frac{-n\mu^2}{2(1-\mu)^2}} \left[\frac{e^{\frac{i\pi}{2}} - e^{-i\frac{\pi}{2}}}{2i} \right] \int_0^\infty r^{1/2} e^{\left(\frac{-n}{2(1-\mu)^2}\right)r} dr, \quad (C6)$$

and lastly letting $t = \frac{n}{2(1-\mu)^2} r$, we obtain

$$q_n = \frac{\sqrt{2}}{\pi} (1-\mu)^{1/2} n^{-3/2} e^{\frac{-n\mu^2}{2(1-\mu)^2}} \int_0^\infty t^{1/2} e^{-t} dt, \quad (C7)$$

where the integral is now well-known with solution

$$q_n = \frac{(1-\mu)^{1/2}}{\sqrt{2\pi}} n^{-\frac{3}{2}} e^{\frac{-n\mu^2}{2(1-\mu)^2}}. \quad (C8)$$

-
- [1] T. Gisiger, Scale invariance in biology: Coincidence or footprint of a universal mechanism? *Biol. Rev. Camb. Philos. Soc.* **76**, 161 (2001).
 - [2] G. West, *Scale: The Universal Laws of Life and Death in Organisms, Cities and Companies* (Penguin Press, New York, 2017).
 - [3] J. C. Willis, *Age and Area: A Study in Geographical Distribution and Origin of Species* (Cambridge University Press, Cambridge, 1922).
 - [4] H. A. Simon, On a class of skew distribution functions, *Biometrika* **42**, 425 (1955).
 - [5] M. Newman, Power laws, Pareto distributions and Zipf's law, *Contemp. Phys.* **46**, 323 (2005).
 - [6] M. W. Hahn and R. A. Bentley, Drift as a mechanism for cultural change: An example from baby names, *Proc. R. Soc. Lond. B* **270**, S120 (2003).
 - [7] R. A. Bentley, M. W. Hahn, and S. J. Shennan, Random drift and culture change, *Proc. R. Soc. Lond. B* **271**, 1443 (2004).
 - [8] D. de Solla Price, Networks of scientific papers, *Science* **149**, 510 (1965).
 - [9] S. Redner, Citation statistics from 110 years of Physical Review, *Phys. Today* **58**(6), 49 (2005).
 - [10] M. V. Simkin and V. P. Roychowdhury, A mathematical theory of citing, *J. Am. Soc. Inf. Sci. Tech.* **58**, 1661 (2007).
 - [11] F. Radicchi, S. Fortunato, and C. Castellano, Universality of citation distributions: Toward an objective measure of scientific impact, *Proc. Natl. Acad. Sci. USA* **105**, 17268 (2008).
 - [12] D. Wang, H. Cheng, P. Wang, X. Huang, and G. Jian, Zipf's law in passwords, *IEEE Trans. Inform. Forensic Secur.* **12**, 2776 (2017).
 - [13] A. ElBahrawy, L. Alessandretti, A. Kandler, R. Pastor-Satorras, and A. Baronchelli, Evolutionary dynamics of the cryptocurrency market, *R. Soc. Open Sci.* **4**, 170623 (2017).
 - [14] 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).
 - [15] 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).
 - [16] A. M. Leroi, B. Lambert, J. Rosindell, X. Zhang, and G. D. Kokkoris, Neutral syndrome, *Nat. Hum. Behav.* **4**, 780 (2020).
 - [17] G. U. Yule, A mathematical theory of evolution based on the conclusions of Dr. J. C. Willis, F.R.S., *Philos. Trans. R. Soc. London, Ser. B* **88**, 433 (1925).
 - [18] M. Newman, The first-mover advantage in scientific publication, *Europhys. Lett.* **86**, 68001 (2009).
 - [19] P. S. Dodds, D. R. Dewhurst, F. F. Hazlehurst, C. M. Van Oort, L. Mitchell, A. J. Reagan, J. R. Williams, and C. M. Danforth, Simon's fundamental rich-get-richer model entails a dominant first-mover advantage, *Phys. Rev. E* **95**, 052301 (2017).
 - [20] D. H. Zanette and S. C. Manrubia, Vertical transmission of culture and the distribution of family names, *Physica A* **295**, 1 (2001).
 - [21] S. Bornholdt and H. Ebel, World wide web scaling exponent from Simon's 1955 model, *Phys. Rev. E* **64**, 035104(R) (2001).
 - [22] T. Maillart, D. Sornette, S. Spaeth, and G. von Krogh, Empirical Tests of Zipf's Law Mechanism in Open Source Linux Distribution, *Phys. Rev. Lett.* **101**, 218701 (2008).
 - [23] A.-L. Barabási and R. Albert, Emergence of scaling in random networks, *Science* **286**, 509 (1999).
 - [24] D. D. S. Price, A general theory of bibliometric and other cumulative advantage processes, *J. Am. Soc. Inf. Sci.* **27**, 292 (1976).
 - [25] G. Bianconi and A.-L. Barabási, Competition and multiscaling in evolving networks, *Europhys. Lett.* **54**, 436 (2001).
 - [26] G. Caldarelli, A. Capocci, P. De Los Rios, and M. A. Muñoz, Scale-Free Networks From Varying Vertex Intrinsic Fitness, *Phys. Rev. Lett.* **89**, 258702 (2002).
 - [27] L. Hébert-Dufresne, A. Allard, V. Marceau, P.-A. Noël, and L. J. Dubé, Structural preferential attachment: Stochastic process for

- the growth of scale-free, modular, and self-similar systems, *Phys. Rev. E* **85**, 026108 (2012).
- [28] C. Cattuto, V. Loreto, and V. P. Servedio, A Yule-Simon process with memory, *Europhys. Lett.* **76**, 208 (2006).
- [29] C. Cattuto, V. Loreto, and L. Pietronero, Semiotic dynamics and collaborative tagging, *Proc. Natl. Acad. Sci. USA* **104**, 1461 (2007).
- [30] A. L. Schaigorodsky, J. I. Perotti, N. Almeida, and O. V. Billoni, Short-ranged memory model with preferential growth, *Phys. Rev. E* **97**, 022132 (2018).
- [31] J. Bertoin, A version of Herbert A. Simon's model with slowly fading memory and its connections to branching processes, *J. Stat. Phys.* **176**, 679 (2019).
- [32] E. Baur and J. Bertoin, On a two-parameter Yule-Simon distribution, [arXiv:2001.01486](https://arxiv.org/abs/2001.01486).
- [33] T. E. Harris, *The Theory of Branching Processes* (Springer, Berlin, Heidelberg, 1963).
- [34] K. B. Athreya and P. E. Ney, *Branching Processes* (Springer, Berlin, Heidelberg, 1972).
- [35] R. van der Lans, G. van Bruggen, J. Eliashberg, and B. Wierenga, A viral branching model for predicting the spread of electronic word of mouth, *Marketing Science* **29**, 348 (2010).
- [36] J. L. Iribarren and E. Moro, Branching dynamics of viral information spreading, *Phys. Rev. E* **84**, 046116 (2011).
- [37] O. Yagan, D. Qian, J. Zhang, and D. Cochran, Conjoining speeds up information diffusion in overlaying social-physical networks, *IEEE J. Sel. Areas Commun.* **31**, 1038 (2013).
- [38] 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).
- [39] 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).
- [40] J. D. O'Brien, I. K. Dassios, and J. P. Gleeson, Spreading of memes on multiplex networks, *New J. Phys.* **21**, 025001 (2019).
- [41] H. S. Wilf, *Generatingfunctionology* (Academic Press, Boston, 1994).
- [42] J. Abate and W. Whitt, The Fourier-series method for inverting transforms of probability distributions, *Queueing Syst.* **10**, 5 (1992).
- [43] A.-L. Barabási, The origin of bursts and heavy tails in human dynamics, *Nature (London)* **435**, 207 (2005).
- [44] J.-C. Delvenne, R. Lambiotte, and L. E. C. Rocha, Diffusion on networked systems is a question of time or structure, *Nat. Commun.* **6**, 7366 (2015).
- [45] M. Karsai, H.-H. Jo, and K. Kaski, *Bursty Human Dynamics* (Springer, Cham, 2018).
- [46] J. D. O'Brien, A. Aleta, Y. Moreno, and J. P. Gleeson, Quantifying uncertainty in a predictive model for popularity dynamics, *Phys. Rev. E* **101**, 062311 (2020).
- [47] A. G. Hawkes, Spectra of some self-exciting and mutually exciting point processes, *Biometrika* **58**, 83 (1971).
- [48] A. Talbot, The accurate numerical inversion of laplace transforms, *IMA J. Appl. Math.* **23**, 97 (1979).
- [49] J. Abate and W. Whitt, A unified framework for numerically inverting laplace transforms, *INFORMS J. Comput.* **18**, 407 (2006).
- [50] J. P. Keener, *Principles of Applied Mathematics: Transformation and Approximation* (CRC Press, 2018).