Self-organization of dragon king failures

Yuansheng Lin, Keith Burghardt, Martin Rohden, Pierre-André Noël, and Raissa M. D'Souza

Phys. Rev. E **98**, 022127 — Published 27 August 2018

DOI: 10.1103/PhysRevE.98.022127

Self-Organization of Dragon King Failures

Yuansheng Lin,[1, 2, 3, *] Keith Burghardt,[4] Martin Rohden,[3] Pierre-André Noël,[3] and Raissa M. D'Souza[3, 5, 6]

[1] *School of Reliability and Systems Engineering, Beihang University, Beijing, China, 100191*
[2] *Beijing Jingdong Century Trade Co., Ltd., Beijing, China, 101111*
[3] *Department of Computer Science, University of California, Davis, California, USA, 95616*
[4] *Information Sciences Institute, University of Southern California, Marina del Rey, California, USA, 90292*
[5] *Department of Mechanical and Aerospace Engineering, University of California, Davis, California, USA, 95616*
[6] *Santa Fe Institute, Santa Fe, New Mexico, USA, 87501*

(Dated: August 9, 2018)

The mechanisms underlying cascading failures are often modeled via the paradigm of self-organized criticality. Here we introduce a simple network model where nodes self-organize to be either weakly or strongly protected against failure in a manner that captures the trade-off between degradation and reinforcement of nodes inherent in many network systems. If strong nodes cannot fail, any failure is contained to a single, isolated cluster of weak nodes and the model produces power-law distributions of failure sizes. We classify the large, rare events that involve the failure of only a single cluster as "Black Swans". In contrast, if strong nodes fail once a sufficient fraction of their neighbors fail, then failure can cascade across multiple clusters of weak nodes. If over 99.9% of the nodes fail due to this cluster hopping mechanism, we classify this as a "Dragon King", which are massive failures caused by mechanisms distinct from smaller failures. The Dragon Kings observed are self-organized, existing over a wide range of reinforcement rates and system sizes. We find that once an initial cluster of failing weak nodes is above a critical size, the Dragon King mechanism kicks in, leading to piggybacking system-wide failures. We demonstrate that the size of the initial failed weak cluster predicts the likelihood of a Dragon King event with high accuracy and we develop a simple control strategy that can dramatically reduce Dragon Kings and other large failures.

PACS numbers: 89.75.Da 02.30.Yy 05.65.+b

I. INTRODUCTION

Many natural and engineered systems exhibit rare, catastrophic events that lead to widespread failure of the system [1–13]. Two categories for these large failure cascades have been proposed: Black Swans, which are tail events in a power-law distribution, and Dragon Kings (DKs), which are outliers involving mechanisms absent in smaller events leading to more frequent failure than a power-law would predict. The power-law distribution necessary for Black Swans to exist is often explained by self-organized criticality (SOC): a tug-of-war that poises the system close to a critical point without any need for tuning of external parameters [5–8, 14, 15].

In this paper, we demonstrate that a simple cascading failure model inspired by engineered systems self-organizes to form DKs. Furthermore, the DKs occur for almost any parameter chosen, but can be predicted and controlled. We call this type of model a *Self-Organized Dragon King* model, in contrast to SOC models in which DK-size cascades occur as finite-size effects for super-critical parameters. Furthermore, because DKs are less common than small events, our model reveals that an observer of the dynamics for short or moderately long times might naively claim to be confident of the system's resiliency, and thus taken by surprise when a DK event inevitably occurs.

We model an idealized engineered system as a network in which nodes self-organize to be "weak" or "strong" (i.e., more susceptible or less susceptible to neighboring failures), and failures spread between neighboring nodes. The initial failure of a random weak node can then lead to a very large cascade of subsequent node failures, i.e., a failure cascade [16–19]. A weak node fails as soon as *one* of its neighbors fails, and a failed weak node has small probability, ϵ, to be reinforced and upgraded to a strong node upon repair. This captures the common practice of *failure-based resource allocation* in many engineered systems [20, 21]. Strong nodes independently degrade (i.e., become weak) at a slow rate.

If strong nodes *cannot* fail, we call the model the "*inoculation*" (IN) model [22]. This is akin to site percolation because a failure is contained to an individual cluster of adjacent weak nodes. Like other SOC models, the IN model self-organizes to a critical state with a power-law distribution in event sizes. Since failure is contained to a single cluster no matter the size of the failure, the same mechanism underlies all events and we classify the rare events as "Black Swans". If strong nodes fail as soon as *two* of their neighbors fail, we call the model the "*complex contagion*" (CC) model [23]. This CC model can lead to self-amplifying failures that cascade across multiple clusters of weak nodes, and it is similar in spirit to bootstrap percolation with two activation thresholds [24, 25]. If the failure cascades across multiple clusters of weak nodes and leads to the failure of over 99.9% of

* bhlys1106@gmail.com

nodes in the system, we classify the event as a "Dragon King".

Surprisingly, the small modification from IN to CC, allows this model to self-organize to a state that creates DKs. We will show that DKs are inevitable and generic in the CC model. They are predicted to occur for any finite-sized network and any non-trivial value of the one tuning parameter, ϵ, which has been numerically confirmed through extensive simulation. The CC model is furthermore the simplest model we are aware of that produces self-amplifying cascading failures, while the IN model provides a SOC null model for baseline behavior.

Our model is also novel in that it allows us to quantitatively define a DK event and also allows us to accurately predict and control DKs. Although DKs may be more destructive than Black Swans, differences between DK and Black Swan events allow for DKs to be predicted and controlled in ways that Black Swans can not. At best, Black Swans can be predicted a little better than random chance [26, 27], but usually predicting them is inherently difficult [28, 29] because they often have no associated length- and time-scales. DK events, on the other hand, can be predicted: there are typical places and times when DKs will and will not occur. This has been successfully applied to, for example, prediction of material failure and crashes of stock markets [1]. Our model allows us to predict DKs with a near-perfect true-positive rate.

Even though Black Swans are not very predictable, there are simple methods to push SOC systems away from criticality, thus reducing the size of Black Swans [8, 30, 31]. It has been an open problem, however, to control DKs in many situations and to elucidate the mechanisms underlying these often self-amplifying cascades [2]. Recent advances in controlling DKs have been based on low-dimensional models, such as coupled oscillators [32]. We develop a simple targeted-reinforcement control strategy, in which we turn a few fairly well-chosen weak-nodes into strong nodes to decrease the likelihood of DKs and other large failures by orders of magnitude. Importantly, our strategy can significantly reduce both large DK and non-DK events *without changing the reinforcement rate*.

Finally, we have an analytical understanding of DKs in our model. The main reason DKs appear in the CC model is because failure cascades hop from one weak cluster to another if the initial weak cluster is large enough. Once a sufficient number of nodes fail, there is a high likelihood for almost all nodes to fail, because each node will likely connect to the failed portion of the network and subsequently fail. We take advantage of this finding to predict whether a small initial failure will cascade into a DK event by showing the probability that a strong node, which bridges weak-node clusters, will have two neighbors in the initial failing cluster can be mapped onto a generalization of the birthday problem [33]. The birthday problem considers that, given c people, what is the probability that any of the people share the same birthday. Once the probability is significant for a strong node to have two neighbors in the failing cluster, then failures are likely to spread from the first weak-node cluster to subsequent weak-node clusters. More strong nodes are then likely to fail by piggybacking off of the previous failures. We can make a qualitative analogy to the gas-water phase transition in condensed matter, where droplets can nucleate. In both our model and in droplet nucleation, there is a critical size, above which the droplet or failed cluster grows almost without bound [34], although in the CC model, clusters of any size can form (*i.e.,* there is no analogous surface tension).

We organize the paper as follows. In Section II, we discuss the model in detail. We discuss the mechanism behind the CC model's DKs in Section III. In Sections IV & V, we discuss how the probability of DKs and large events, respectively, vary with ϵ and N. In Section VI, we discuss ways to predict DK events, and in Section VII, we discuss how to control DKs with a simple algorithm, before concluding the paper.

II. SELF-ORGANIZING MODELS

The dynamics of our models depend on two competing mechanisms: degradation and reinforcement. Degradation, which represents the aging of infrastructure or an increase of load placed on them, is modeled by slowly converting strong nodes into weak ones. Conversely, reinforcement converts weak nodes that fail into strong nodes at rate ϵ, representing the hardening of nodes in an attempt to prevent future failures. This repair strategy mimics modern-day power grid guidelines [20], where resources are allocated to places were failures happen more often. The trade-off between degradation and reinforcement drives the system to an SOC state.

We consider dynamics on 3-regular random networks with N nodes, where N is an even positive integer. The motivation for this network is two-fold. First, complex engineered systems, such as electrical grids, have low degree (approximately degree 3) and low degree heterogeneity [19, 35]. Second, this is one of the simplest networks we can create, which aids analytic understanding of the dynamics. Repeated edges and self-loops are allowed, but are rare when N is large. The system size N and the probability ϵ are the model's only parameters. We are particularly interested in large-but-finite N and small ϵ to simulate real systems, which may be large and expensive to upgrade.

Both the CC and IN models follow the same algorithm. We initialize all N nodes as weak and for each discrete time step $0 \leq t \leq t_{\text{stop}}$, the model does the following.

Degradation: Select a node uniformly at random. If that node is strong, make it weak and proceed to the beginning of the Degradation step with $t \leftarrow t + 1$. If the selected node is already weak, then it fails, and continue with the remaining three steps.

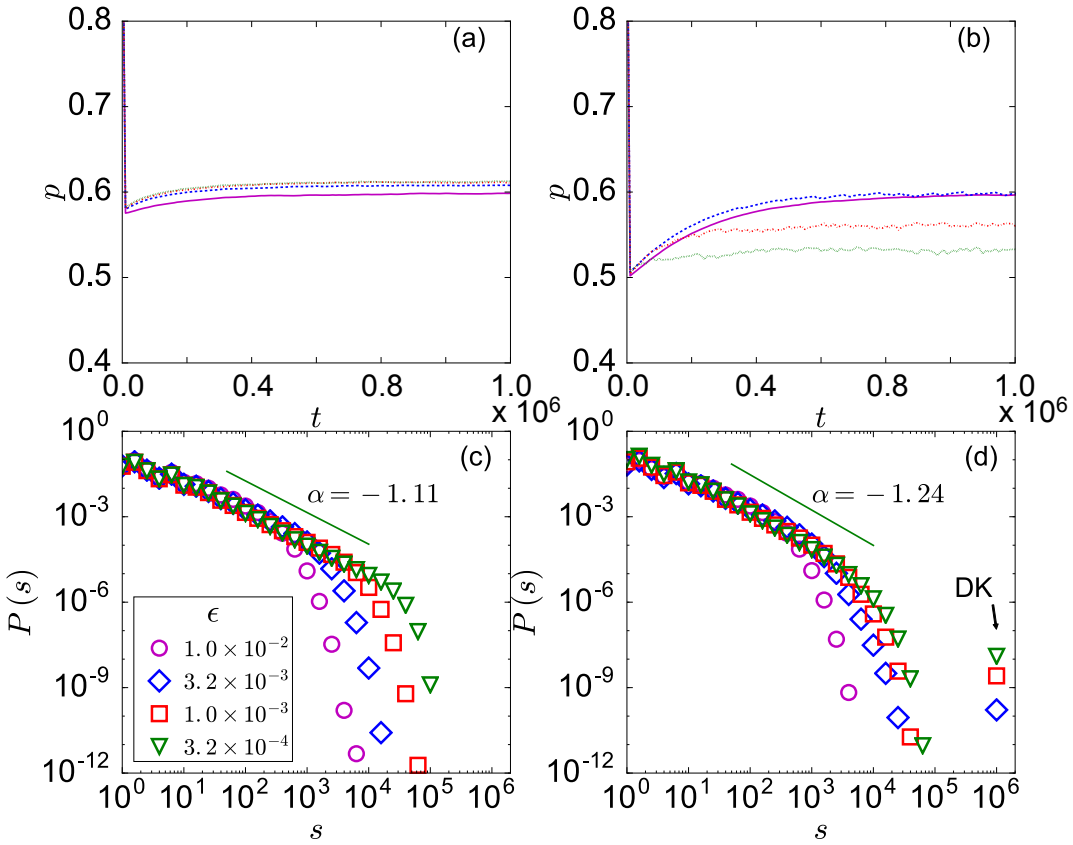

FIG. 1. (Color online) Self-organizing behavior and failure size. Top row: Fraction of weak nodes $p(t)$ vs. t for the (a) IN and (b) CC models over individual network realizations for $N = 10^6$ and varying ϵ ($\epsilon = 1.0 \times 10^{-2}$, solid line; $\epsilon = 3.2 \times 10^{-3}$, dashed line; $\epsilon = 1.0 \times 10^{-3}$, dash-dot line; $\epsilon = 3.2 \times 10^{-4}$, dotted line). Bottom row: Failure size distribution for the (c) IN and (d) CC models, with the DK outliers labeled. Symbols denote results of simulations on random 3-regular graphs averaged over ten network realizations and $15 \times N$ time steps.

Cascade: Apply the IN or CC failure-spreading mechanism until no more failures occur. Failed nodes remain failed for the duration of the cascade.

Repair: All failed nodes are un-failed (strong failed nodes become strong un-failed nodes, and weak failed nodes become weak un-failed nodes).

Reinforcement: Each weak node that failed at this time step has probability ϵ to become strong. Proceed back to the Degradation step with $t \leftarrow t + 1$.

The distinction between the two models is only in the cascade step of the algotirhm. Specifically, under the CC model, strong nodes fail if at least two of their neighbors fail, whereas in the IN model, strong nodes cannot fail. The IN model is similar to the SIRS model in epidemiology [36], except that failed (i.e., infected) nodes can directly become un-failed (i.e., susceptible) again. Many choices for initial conditions are possible, but our investigations show that the steady state behavior is independent of these choices (see Appendix XI). Because we currently initialize all nodes as weak, the sizes of the first few cascades are on the order of the system size, and numerous node upgrades take place before the system equilibrates. An important indicator that we have reached the relaxation time is the proportion of nodes that are weak at time t, $p(t)$, shown in Fig. 1. We wait until well after $p(t)$ stabilizes ($5 \times N$ timesteps) and then calculate failure sizes for a subsequent $15N$ timesteps. As shown in Appendix X, waiting longer, and varying the initial conditions produces quantitatively similar results. For the IN model, we find that $p(t)$ is almost independent of ϵ as $\epsilon \to 0$, but in the CC model, the steady-state value of $p(t)$ depends on ϵ. Importantly, the dynamics imply cascades occur before any reinforcement or repair, therefore the models assume that failure cascades occur on shorter time-scales than repair or reinforcement. This is similar to real systems, such as electric grids, where blackouts may occur over minutes, repairs occur over hours or days, and reinforcement may occur over months or years. This also motivates the small values of ϵ used throughout the paper.

Importantly, we notice that $p(t)$ is almost always above $p_c = 0.5$. If weak and strong nodes were distributed completely at random in the network, then there would be a spanning cluster of weak nodes, because the fraction

of weak nodes is above the critical percolation fraction of 0.5 [37]. With a non-zero probability, the failure of a weak node could lead to the failure of the spanning cluster, and therefore $O(N)$ failed nodes. The fact that we do not see this (e.g., Fig. 1c) suggests non-trivial correlations between weak and strong nodes. As we mentioned previously, this model is similar to bootstrap percolation with two thresholds. Importantly, however, bootstrap percolation literature assumes a finite fraction of nodes fail throughout the network before a first order transition [24], while the CC model described here is for a more realistic scenario in which a single failure, centered inside a cluster that is small compared to the network size, can create a near-total failure of the system. Furthermore, because node failure thresholds are not randomly distributed in the system, we will not be able to analyze this model using bootstrap percolation literature [24].

The failure size distribution, $P(s)$, is illustrated in Figs. 1c & 1d. The probability of large failures generally increases with decreasing ϵ for both the IN and CC models because, if less nodes are reinforced, cascades can more easily spread and affect larger portions of a network. For small enough ϵ, we find that the cascade size distributions for the IN and CC models exhibit a power-law with exponential decay, however the CC model also has a DK tail, where *over 99.9% of nodes fail* in each DK event (cf. Fig. 1d). Furthermore, the IN and CC model appear to have two different power-law exponents: $\alpha = -1.11$ for the IN model and $\alpha = -1.24$ for the CC model when $\epsilon = 3.2 \times 10^{-4}$ and $N = 10^6$ (traditional SOC models yield $\alpha = -1.5$ [38, 39]). In the following section, we will explore why the DK events exist in the CC model.

III. DRAGON KING MECHANISM

In this section we will give an overview of our analysis used to understand DKs in the network. Details of this analysis are left in the appendix.

In order for a DK to occur, strong nodes that bridge weak node clusters must fail. When the first weak-node fails, intuitively, all of its weak-node neighbors, and those neighbors' neighbors fail as well. In this "zeroth order" cascade, any weak node failure will trigger its connected weak-node cluster to fail (see Fig. 2a). Clearly, for a DK to occur, multiple clusters must fail, so there must also exist at least one strong node that fails and bridges two clusters, allowing for yet more clusters to fail (e.g., S_1 in Fig. 2a). This is what we call a *first-order* cascade. We derive analytical arguments about the likelihood of an S_1 node in Appendix X, which we summarize below.

We find through numerical simulations that the weak-node cluster is approximately a tree. Intuitively, this is because the likelihood for short cycles, such as triangles, in a random network is very low [40]. The number of links, $L_{\mathrm{w},1}$, from the first weak-node cluster to strong nodes is therefore approximately $C_{\mathrm{w},1} + 2$, where $C_{\mathrm{w},1}$ is the size of the weak-node cluster. Each link can

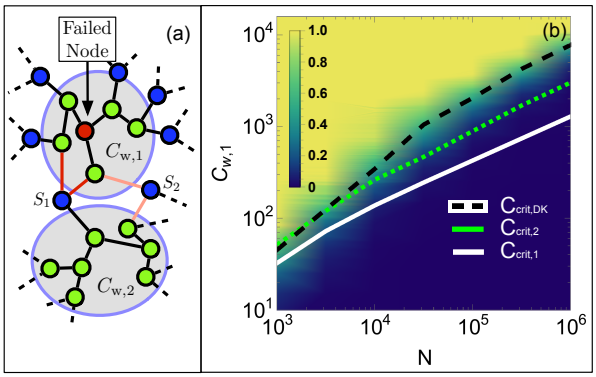

FIG. 2. (Color online) DKs form by cascading failures of weak-node clusters. (a) Weak-node clusters (circled) are surrounded by strong nodes. If one such strong node (S_1) has two links connecting to the same weak-node cluster, the failure of this cluster (with size $C_{\mathrm{w},1}$) will make the strong node fail, which we call a first-order cascade. The cascade may spread to other weak-node clusters (e.g., one of size $C_{\mathrm{w},2}$), and thus other strong nodes (e.g., S_2), which we call a second-order cascade. The failure of subsequent strong nodes can eventually lead to a DK event. (b) A heat map of the probability a DK occurs in the CC model conditioned on $C_{\mathrm{w},1}$, the size of the weak-node cluster that first fails, versus N and $C_{\mathrm{w},1}$ for $\epsilon = 10^{-3}$. Black dashed line is the simulation result for the critical $C_{\mathrm{w},1}$ to create DKs, $C_{\mathrm{crit,DK}}$, while the solid lines are the analytically derived critical $C_{\mathrm{w},1}$ to create a second-order cascade ($C_{\mathrm{crit},2}$, green line), and first-order cascade ($C_{\mathrm{crit},1}$, white line).

connect to a given strong node with a small probability ($\sim L_{\mathrm{w},1}/N$), but if $C_{\mathrm{w},1}$ is large, there are many opportunities for such a link to occur ($L_{\mathrm{w},1}$). The probability for two links to reach the same strong node is therefore $\sim \frac{L_{\mathrm{w},1}(L_{\mathrm{w},1}-1)}{N}$. The probability for an S_1-like node, which we define as a strong node with two links in one weak-node cluster and a third in a second cluster, is similar but with N rescaled. Although there are more complicated ways for a failure to spread between clusters, an S_1-like node is a simplest case. Interestingly, because the probabiliy two links connect to a given node is the same across all nodes, this problem can be exactly mapped onto a generalized birthday problem [33]. Applying results from previous analysis of the birthday problem [41], the probability that the failure of a weak-node cluster of size $C_{\mathrm{w},1}$ creates a first-order cascade is

$$C_{\mathrm{crit},1} = \sqrt{2\log(2) N_{\mathrm{eff}}}, \tag{1}$$

where

$$N_{\mathrm{eff}} = N \frac{(1 - \langle p \rangle) \langle k \rangle^2}{6q}. \tag{2}$$

In the above equation, q is the average fraction of strong nodes with three weak-node neighbors and $\langle k \rangle$ is the average number of weak nodes a strong node connects to (see

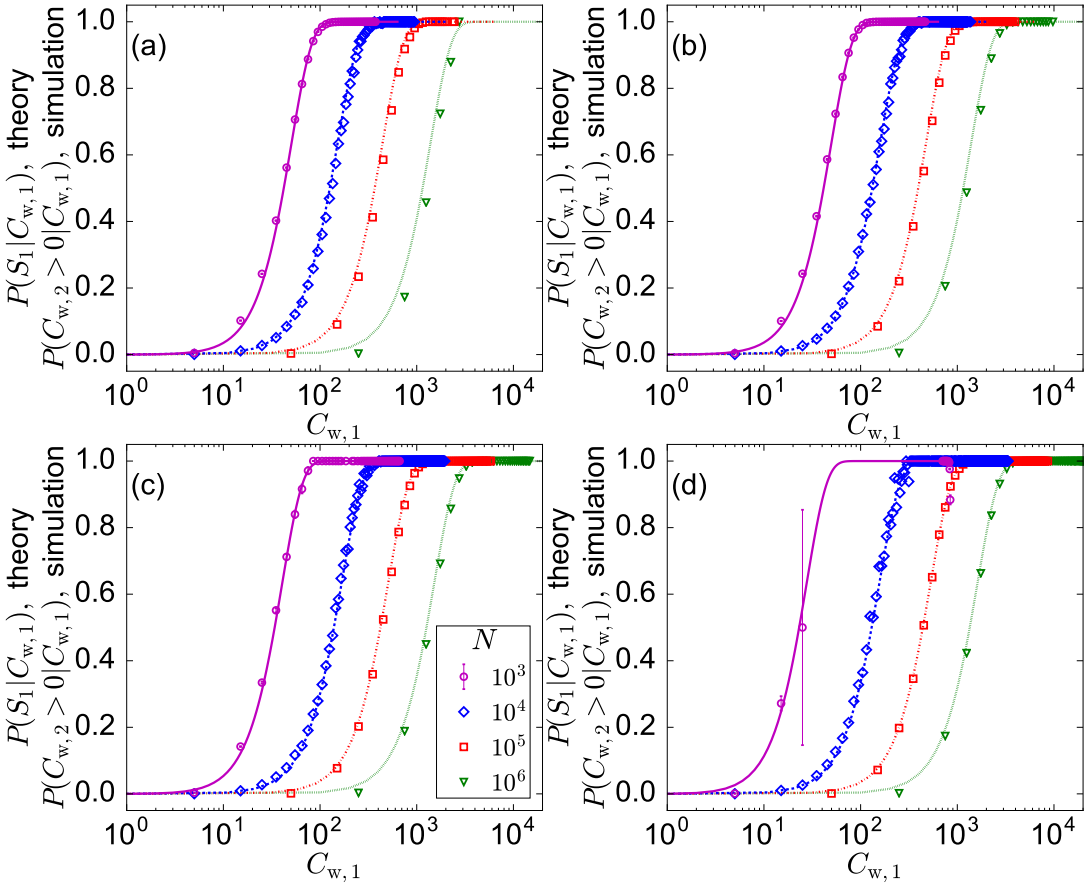

FIG. 3. (Color online) The probability of failure spreading out from the initial failed weak-node cluster. Open markers represent the probability that a failure spreads from the first weak-node cluster to any other weak-node clusters, based on model simulations. Lines represent theoretical results of $P(S_1|C_{\mathrm{w},1})$ ($N = 10^3$, solid line; $N = 10^4$, dashed line; $N = 10^5$, dash-dot line; $N = 10^6$, dotted line), which is the simplest way a failure can spread between weak-node clusters. Standard errors are smaller than marker sizes, except for $N = 10^3$ and $\epsilon = 1.0 \times 10^{-3}$. We show results for (a) $\epsilon = 1.0 \times 10^{-2}$;(b) $\epsilon = 3.2 \times 10^{-3}$;(c) $\epsilon = 1.0 \times 10^{-3}$;(d) $\epsilon = 3.2 \times 10^{-4}$.

Appendix X for details of the derivation). Intuitively, we only need $C_{\mathrm{w},1} \sim N^{1/2}$ for $\frac{L_{\mathrm{w},1}(L_{\mathrm{w},1}-1)}{N} \sim O(1)$, therefore $C_{\mathrm{w},1} \ll N$ for there to exist a first-order cascade.

The exact value of $C_{\mathrm{crit},1}$ is shown by the white line in Fig. 2b. The line, however, is as much as an order of magnitude smaller than the equivalent critical value empirically found for DKs, $C_{\mathrm{crit,DK}}$, making it a poor approximation for the DK mechanism. That said, we compare the first-order cascade theory (Eq. (14)) to first-order cascade simulations (Fig. 3), and find that theory and simulation results match well. The underlying assumptions of this theory are therefore valid.

Because this theory can be mapped onto the birthday problem, we also find that the number of S_1-like nodes is Poisson distributed with $\lambda = \frac{L_{\mathrm{w},1}(L_{\mathrm{w},1}-1)}{N_{\mathrm{eff}}}$ (see Appendix X). The initial cascade can therefore produce first order cascades that increase quadratically with the initial cascade size. Higher-order theory would predict even larger cascades (and eventually DKs). Here we explicitly calculate the likelihood of *second-order* cascades, in which weak-node clusters fail once a first-order cascade occurs (S_1' in Fig. 4b). Second-order cascades are only likely to occur when (a) a S_1'-like node bridges the first failed cluster and a neighboring failed cluster, or (b) it connects one or two neighboring failed clusters (see Fig. 5). The probability for a second-order cascade is therefore one minus the probability that neither (a) or (b) occurs. The details of the derivation are similar to the first-order cascade theory (see Appendix X). The probability of a second-order cascade can only be found numerically. We plot the size, $C_{\mathrm{crit},2}$ when this probability is 1/2 in Fig. 2b. Higher order theory is not analytically tractable after this point, but once a sufficient number of nodes fail, each node will likely connect to the failed portion of the network and subsequently fail, which leads to the failure of approximately N nodes (a DK event).

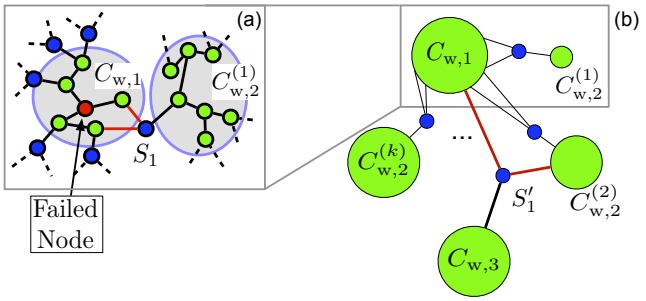

FIG. 4. (Color online) How second-order cascades occur. (a) Failures spread from a single weak-node to an entire connected cluster of weak-nodes. Strong nodes are usually resistant to failure, unless they connect to the same failed cluster (S_1). (b) Once enough weak-node clusters fail, there is a chance that failures can spread to yet more clusters (of size $C_{w,3}$) via strong nodes that bridge failed weak clusters (S'_1).

$$P(C_{w,3} > 0 | C_{w,1}) = 1 - P\left(\begin{array}{c} 2 \\ 1 \end{array} \; S'_1 \; 3 = 0 \right) P\left(\begin{array}{c} 2 \\ 1 \end{array} \; S'_1 \; 3 = 0 \right)$$

FIG. 5. (Color online) The probability of a second-order cascade. This is 1 minus the probability there are no S'_1 nodes from the first-order cascade clusters (labeled "2") to a subsequent weak-node cluster (labeled "3") times the probability there are no S'_1 nodes with one edge in the initial weak-node cluster (labeled "1"), the first-order cascade clusters, and a subsequent weak-node cluster.

IV. HOW THE PROBABILITY OF DRAGON KINGS VARIES WITH N AND EPSILON

We next consider how the probability of DKs varies with N and ϵ. We approximate the weak-node cluster size distribution, $P(C_w)$, as:

$$P(C_w) \sim C_w^{-\eta} e^{-\lambda C_w}. \tag{3}$$

We fit these coefficients via maximum likelihood estimation and find that $\eta = 2.26 \pm 0.03$, when $\epsilon = 3.2 \times 10^{-4}$. We also find that $\lambda = (0.39 \pm 0.01) \times \epsilon$ when $N = 10^5$ and 10^6 (cf. Appendix XII). This is consistent with our hypothesis that, in the dual limit that $N \to \infty$ and $\epsilon \to 0$, the model approaches a self-organized critical state, where the distribution of $C_{w,1}$ becomes a power-law.

The probability for the size of the first weak cluster that fails, $P(C_{w,1})$, is equivalent to picking a node in a cluster of size C_w, leading to

$$P(C_{w,1}) = \frac{C_w P(C_w)}{\langle C_w \rangle}, \tag{4}$$

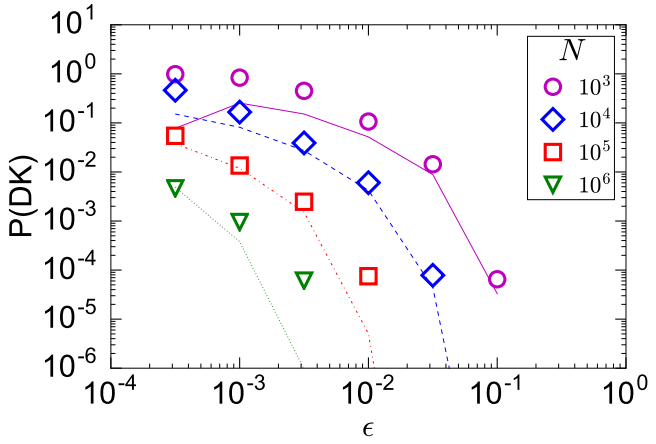

FIG. 6. (Color online) The probability of DKs as a function of epsilon. Regardless of N, we find close agreement between simulations-based probabilities of $P(DK)$ (open symbols) and theory (Eq. (9), $N = 10^3$, solid line; $N = 10^4$, dashed line; $N = 10^5$, dash-dot line; $N = 10^6$, dotted line), especially for moderate values of ϵ.

or

$$P(C_{w,1}) = \frac{C_w^{1-\eta} e^{\lambda C_w}}{E_{\eta-1}(\lambda)}, \tag{5}$$

where

$$E_n(z) = \int_1^N \frac{e^{-zt}}{t^n} dt. \tag{6}$$

Furthermore, we can approximate $P(\mathrm{DK}|C_{w,1})$ as a step function,

$$P(\mathrm{DK}|C_{w,1}) \approx H(C_{w,1} - C_{\mathrm{crit,DK}}). \tag{7}$$

With these assumptions, we can approximate $P(\mathrm{DK}|N, \epsilon)$ as

$$P(\mathrm{DK}|N, \epsilon) = \int_{C_{\mathrm{crit,DK}}}^{N} P(C_{w,1}) dC_{w,1}. \tag{8}$$

If $\epsilon > 0$, then we can further approximate the integral limit as $N \to \infty$. This implies that

$$P(\mathrm{DK}|N, \epsilon) = \begin{cases} \lambda^{\eta-1} \frac{\Gamma(1-\eta, C_{\mathrm{crit,DK}}\lambda)}{E_{\eta-1}(\lambda)} & \epsilon > 0 \\ \frac{N^\eta (C_{\mathrm{crit,DK}}^{2-\eta} - N^{2-\eta})}{N^\eta - N^2} & \epsilon \to 0. \end{cases} \tag{9}$$

We notice two interesting findings. First, we see that $P(\mathrm{DK})$ varies non-linearly with ϵ: a slight increase in repair frequency can dramatically reduce the number of system-wide failures. Agreement is strongest when ϵ is moderate, possibly because, when ϵ is too small for small N, the largest values of C_w are $O(N)$, therefore the model under-estimates how many DKs are possible. When ϵ is too large, however, we appear to underestimate $P(\mathrm{DK})$ again, potentially because the Heaviside approximation

breaks down for small probabilities. We nonetheless find good overall agreement with our simulations (see Fig. 6).

Another interesting implication of our model is that for $\epsilon \to 0$ and $N \to \infty$,

$$P(\mathrm{DK}|N,0) \sim C_{\mathrm{crit,DK}}^{\eta-2} - O(N^{2-\eta}). \tag{10}$$

Through simulations, we find that $C_{\mathrm{crit,DK}} \sim N^{0.59\pm0.03}$ (see Fig. 12), therefore

$$P(\mathrm{DK}|N,0) \sim N^{-\beta} - O(N^{2-\eta}), \tag{11}$$

where $\beta = 0.59(2-\eta) = 0.15 \pm 0.02$. In other words, we find that DK events are surprisingly common for finite systems: $P(\mathrm{DK})$ decreases by only a factor of 10 when the system increases in size by ten million.

V. HOW THE PROBABILITY OF LARGE FAILURES VARIES WITH N AND EPSILON

Figure 1 demonstrates that DKs in the CC model occur for various values of ϵ, and our analytic arguments show DKs occur for all values of $\epsilon < 1$ and any finite system. We find a superficially similar failure size distribution, with a bump of size $O(N)$, for the IN model (see Figs. 7a & 7b) but only when N and ϵ are sufficiently small, e.g., $N = 10^4$ for $\epsilon \leq 10^{-3}$. This is not due to cascading failures, however, because in the IN model only a single weak-node cluster fails. This bump exists because there are weak-node clusters that are $O(N)$, meaning we are in a parameter regime where there are super-critical percolating clusters.

This contrasts with the CC model, where we see over 99.9% of nodes fail almost independent of the values of ϵ and N (see Figs. 7c & 7d). The difference is due to cascading failures in the CC model, where the moment a cluster greater than a critical size fails, strong nodes begin to fail, which triggers more weak-node clusters to fail, etc., until almost all nodes fail (a DK event). Because the critical weak-node cluster size increases sub-linearly with N, and because weak-node clusters can be any size less than N, there is always a chance for weak-node clusters larger than the critical size to fail, triggering a DK-size event.

It is an open question in the field of DK theory how to further classify events such as the bump in the IN model, which has a heavier-than-power-law probability, yet shares the same underlying mechanism as small events (i.e., the mechanism being the failure of a single cluster of weak nodes in the IN model.)

VI. PREDICTING DRAGON KINGS

DKs are, in contrast to Black Swans [26–28], relatively predictable [1], although it may not be obvious what independent variables best indicate these events. For example, we find little correlation in the time between DKs (the autocorrelation is < 0.01 for $N = 10^6$, see Fig. 8), therefore, knowing the time-series of DKs will not tell us when another will necessarily occur. We analyze two different predictors. The first is the fraction of weak nodes present in the network. The rationale is that more weak nodes create larger initial failures, and therefore more DKs. The second predictor is the size of the first weak-node cluster, $C_{\mathrm{w},1}$. Both of these predictors are complementary, because the former would tell us *when* a DK might occur, while the latter would tell us *where* a DK might originate.

We model $P(\mathrm{DK}|p)$ versus p, and $P(\mathrm{DK}|C_{\mathrm{w},1})$ versus $C_{\mathrm{w},1}$, respectively, using logistic regression. Examples of Receiver Operating Characteristic (ROC) curves are illustrated in Appendix XIII. Unless ϵ is relatively large, p is a poor predictor as based on the area under the ROC curve (AUC, cf. Fig. 9) [42]. Thus, predicting when a DK would occur is inherently challenging. In contrast, by knowing $C_{\mathrm{w},1}$ alone, we can predict DKs with astounding accuracy, almost independent of N and ϵ. The high accuracy is due to the characteristic size of the initial failure that triggers a DK, $C_{\mathrm{crit,DK}}$. This is reminiscent of previous results on controlling DKs in a system of oscillators where a trajectory straying past a particular threshold is very likely to create a DK [32, 43]. Finding the weak-node cluster size, $C_{\mathrm{w},1}$ (which is much smaller than the system size), for each weak node requires only searching locally in the network. Similarly, to "tame" DKs, we can use a simple control mechanism that requires knowing the size of just a few weak-node clusters, as seen in the next section.

VII. CONTROLLING DRAGON KINGS

Because large weak-node cluster failures precede DKs, we can reasonably ask whether breaking up these clusters before they fail can reduce the prevalence of DKs. Assuming that the rate of node upgrades is proportional to the amount of money or effort allocated for repairing nodes, we create control strategies where this rate is kept the same on average as the non-controlled case, meaning $p(t)$ remains approximately constant. But, instead of randomly reinforcing failed weak nodes, we upgrade weak nodes in large clusters by picking r weak nodes and finding the size of the weak-node clusters to which they belong. Practically speaking, determining weak nodes would be equivalent to finding "old" nodes or nodes similar to those that failed in the past, indicating that this type of control does not require that we create failures. The largest of these weak-node clusters is selected and with probability $1 - p(t)$, a random node in that cluster is reinforced.

We find that, when $r = 1$, more DKs occur than the non-control reinforcement therefore random attempts to reduce the size of failures could actually make the failures substantially worse. The reason $r = 1$ is a poor control protocol is that one node in a weak-node cluster is re-

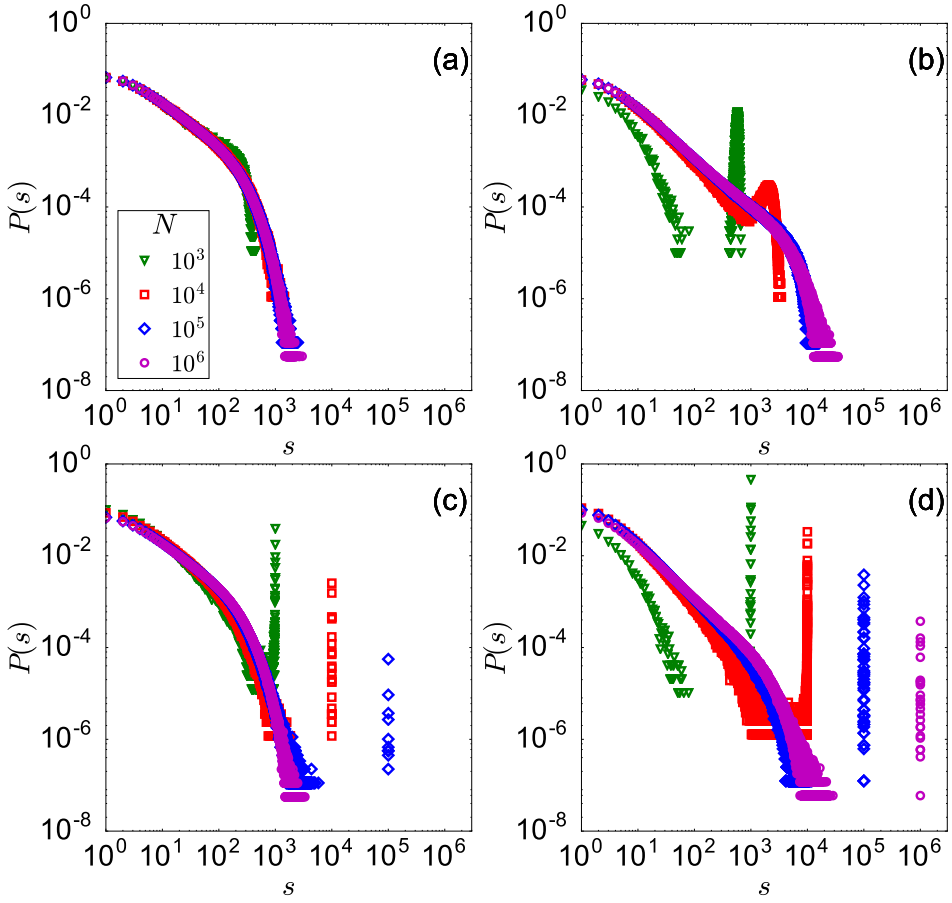

FIG. 7. (Color online) Failure size distribution for networks with different N and different reinforcement probability ϵ. (a) IN model, $\epsilon = 1.0 \times 10^{-2}$; (b) IN model, $\epsilon = 1.0 \times 10^{-3}$; (c) CC model, $\epsilon = 1.0 \times 10^{-2}$; (d) CC model, $\epsilon = 1.0 \times 10^{-3}$.

inforced with a probability $P(C_{\mathrm{w},1}) = \frac{C_{\mathrm{w}} P(C_{\mathrm{w}})}{\langle C_{\mathrm{w}} \rangle}$. In contrast, with non-controlled reinforcement, weak-node clusters fail with a probability $P(C_{\mathrm{w},1})$, and approximately $\epsilon C_{\mathrm{w},1}$ of the first failed weak-node cluster is reinforced.

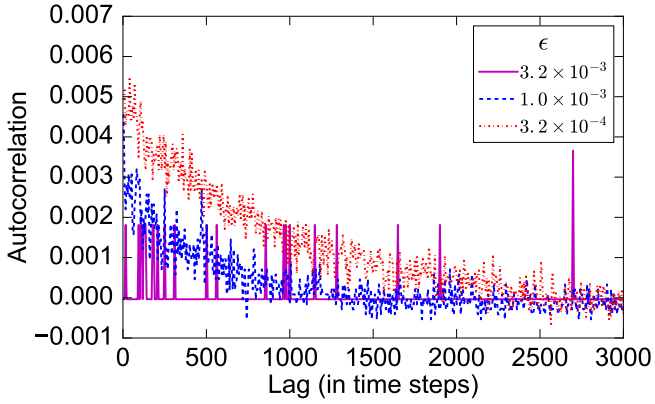

FIG. 8. DK autocorrelation versus lag time. We find that, regardless of the value of ϵ and regardless of the lag time, the autocorrelation of DKs over 1.5×10^7 timesteps is very low, especially as ϵ increases. We do not plot autocorrelations for $\epsilon \geq 1.0 \times 10^{-2}$ because there are few DKs for these values of ϵ in the time frame studied. The network size is $N = 10^6$.

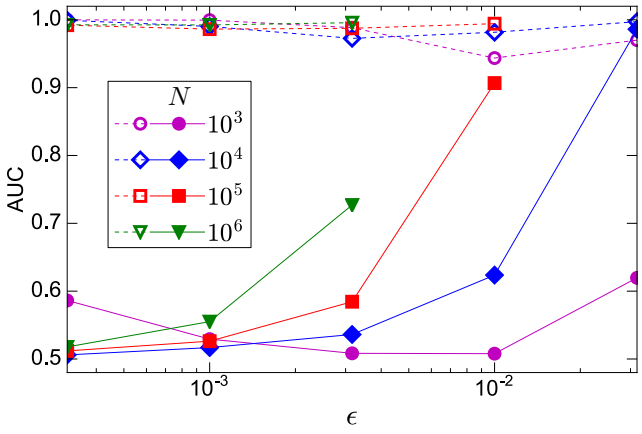

FIG. 9. (Color online) Predicting DKs. The area under the receiver operating characteristic (AUC) for logistic models of $P(\mathrm{DK}|p)$ (closed symbols and solid lines) and $P(\mathrm{DK}|C_{\mathrm{w},1})$ (open symbols and dashed lines) for varying N [42].

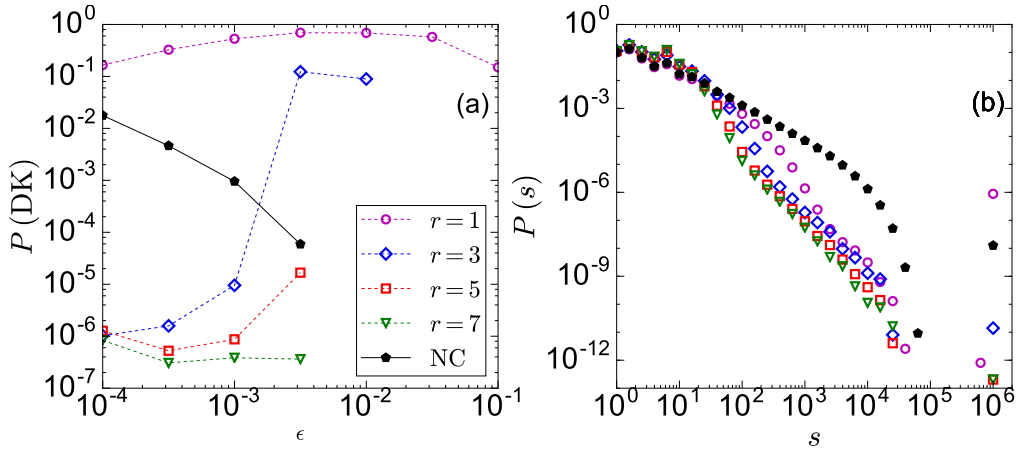

FIG. 10. (Color online) Controlling DKs. (a) The probability a DK occurs over time versus ϵ in both the non-controlled scenario (NC, pentagons), and in the controlled scenario with r weak-nodes chosen: $r = 1$ (circles), $r = 3$ (diamonds), $r = 5$ (squares) and $r = 7$ (triangles). Simulations are realized for $N = 10^6$, and standard errors are smaller than marker sizes. See main text for details of the control method. (b) Failure size distributions for different control strategies and non-controlled with $\epsilon = 3.2 \times 10^{-4}$.

This prioritizes reinforcement to large weak-node clusters, which are more likely to create DKs, while keeping the reinforcement rate fixed. Furthermore, larger r represents a better sampling of the cluster sizes, and more heavily prioritizes reinforcement to large clusters, which reduces the probability of DKs by orders of magnitude (cf. Fig. 10a), as well as large failures that are not DKs (cf. Fig. 10b). In Appendix XIII, we demonstrate that varying r does not significantly change the average fraction of weak nodes. Furthermore, the number of nodes we have to search through is only $r \times \langle C_{\mathrm{w},1} \rangle \ll N$ on average, which makes this technique applicable in systems where global knowledge of the network is lacking.

VIII. CONCLUSION

We have shown that DKs can self-organize in the CC model via failures that cascade across distinct weak-node clusters. Moreover, this mechanism allows for DKs to be easily predicted and controlled. Surprisingly, we find that reinforcing the network slightly more often, or selectively reinforcing nodes (the control strategy with $r > 1$), creates a significant percentage drop in the frequency of DKs. In contrast, naively reinforcing nodes at random (the control strategy with $r = 1$) dramatically increases the frequency of DKs. We believe that this model may describe many engineered systems which exhibit (1) nodes that fail when neighbors do (2) node degredation leading to a greater failure likelihood, and (3) failure-based resource allocation [20]. Although our model is highly idealized, it still makes practical suggestions to reduce large cascades: (1) increase (even slightly) the rate of repairs or (2) upgrade systems that are vulnerable rather than upgrade systems that have recently failed.

Due to the simplicity of the mechanism, we also believe this model may also explain other systems with DKs.

First, when extended to a directed network, the CC model might explain failure cascades on fault trees, which are used to explain the cause of failures on many complex engineered systems [44, 45], such as in nuclear power plants [9, 44], chemical processes, and aerospace systems [44]. Namely, fault trees assume that failures are caused by one or more effects (alike to the failure mechanism in the CC model), and they can cause multiple effects (common-cause failures [44, 45]). Failures therefore lead to further failures thus creating a failure cascade. Our model, however, is deterministic: the failure of a node always causes its neighboring weak nodes to fail, when, in reality, a failure might spread probabilistically. The CC model therefore produces a worst-case scenario, although even non-deterministic networks could still create DKs, just at lower probabilities. Importantly, the model assumes that components in the fault tree degrade (e.g., they may have failure rates that increase over some timescale) or are upgraded over time (based on whether they have failed in the past), which we believe is realistic in many scenarios.

Second, when extended to a social network this model may instead describe financial draw-downs in markets. Let us assume brokers buy or sell due to peer influence, a reasonable assumption which has been observed in stock market participation [46] and foreign exchange trading [47]. This is in contrast to previous models to explain large drawdowns (or drawups) in which agents make choices strictly due to their past behavior [48]. Some brokers will buy (sell) stock when any neighbor does ("weak" brokers) [49], while others buy (sell) stock only after a sufficient fraction of their neighbors do ("strong" brokers) [50]. Brokers can further become "complex" at a rate ϵ after they have adopted an idea (i.e. failed), which can be interpreted as agents exhibiting greater stubbornness to

new ideas. The CC model suggests that agents can self-organize to a state in which global adoption (DKs) occurs surprisingly often. This could explain, for example, why large financial drawdowns in stock markets appear to be DKs [1, 48].

Finally, there are several ways the CC model can be explored more deeply in the future. First, it can provide a concrete methodology to begin studying how DKs are driven by the interplay of heterogeneity (for example the variance of node degree and the diversity of thresholds for strong nodes) and coupling (e.g., average node degree) in a principled manner, which is still in its infancy [1]. Second, the CC model can be analyzed for different network topologies, such as assortative, clustered, or heavy-tailed degree distribution networks. Incorporating these features also creates additional degrees of freedom. For example, the failure dynamics could depend on a minimum number of neighboring agents failing [23], or minimum *fraction* of neighboring agents failing [51]. This distinction becomes especially important for heavy-tailed degree distributions.

IX. ACKNOWLEDGEMENTS

We gratefully acknowledge support from the US Army Research Office MURI award W911NF-13-1-0340 and Cooperative Agreement W911NF-09-2-0053; The Minerva Initiative award W911NF-15-1-0502; and DARPA award W911NF-17-1-0077; and financial support from the program of China Scholarships Council (No. 201506020065, Y.L.).

X. APPENDIX A: FAILURE CASCADE ANALYSIS

To establish a theoretical understanding of DKs, we first note that a failure in any part of a weak-node cluster makes that entire cluster fail. A necessary, but not sufficient, condition for a DK to occur is that strong nodes bridging the first failed weak-node cluster must also fail (cf. S_1 in Fig. 2a). We call the failure of weak-node clusters due to failure of such strong nodes a first-order cascade. If the initial weak-node cluster is large enough, however, then a first-order cascade can lead to even more strong nodes failing (cf. S_2 in Fig. 2a). If weak-node clusters fail due to S_2-like nodes, then we call it a second-order cascade. In the simplest case, only one strong node bridges two weak-node clusters. We first analyze the probability that the failure of a weak-node cluster, with size $C_{\mathrm{w},1}$, will lead to the failure of at least one bridging strong node, denoted by S_1, and find that S_1 nodes can accurately model the probability of multiple weak-node clusters failing (thus explaining how first-order cascades occur). Subsequent analysis for second-order cascades, however, gives us a much more accurate approximation of DKs, although it is not a sufficient mechanism to produce DKs. Third-order and higher-order theory does not appear to be analytically tractable, but we remark that the number of failed weak nodes in a first-order cascade scales as $(C_{\mathrm{w},1})^2$, therefore, once strong nodes fail, they lead to large numbers of subsequent failures. Once a sufficient number of nodes fail, each node will likely connect to the failed portion of the network and subsequently fail, which leads to the failure of $\sim N$ nodes (a DK event).

A. Mechanism for First-Order Cascades

We first present the mechanism behind first-order cascades. We assume there are N nodes, and $(1-\langle p\rangle)N$ strong nodes each of which may have between zero and three weak-node neighbors. In addition, a failure begins at a weak node within a weak-node connected cluster of size $C_{\mathrm{w},1}$. Based on simulations, the weak-node cluster is approximately a tree, therefore the number of links, $L_{\mathrm{w},1}$, from the first weak-node cluster to strong nodes should be $C_{\mathrm{w},1}+2$. Under an annealed network configuration model assumption, we sequentially connect links from the weak-node cluster to strong nodes. After m links are added, the probability that a subsequent link connects to a strong node with three weak-node neighbors, is

$$\begin{aligned}\rho_{m,j} &= \frac{3q(1-\langle p\rangle)N-j}{(1-\langle p\rangle)N\langle k\rangle - m}\\ &\approx \rho = \frac{3q}{\langle k\rangle},\end{aligned} \tag{12}$$

where q is the average fraction of strong nodes with three weak-node neighbors, j is the number of links already connected to strong nodes with three weak-node neighbors and $\langle k\rangle$ is the average number of weak nodes a strong node connects to. We assume $m, j \ll (1-\langle p\rangle)N\langle k\rangle$, and therefore drop second-order terms.

The probability for the $m+1^{\mathrm{th}}$ link from the weak-node cluster to connect to *any* strong node that has three weak-node neighbors with one neighbor already in the same weak-node cluster is

$$\begin{aligned}P(\text{new } S_1|m) &= \sum_{j=0}^{m}\left[\binom{m}{j}\rho^j(1-\rho)^{m-j}\frac{2j}{(1-\langle p\rangle)N\langle k\rangle - m}\right]\\ &= \frac{6mq}{\langle k\rangle}\frac{1}{(1-\langle p\rangle)N\langle k\rangle - m},\end{aligned} \tag{13}$$

where we average over j. The probability that two or more links from the weak-node cluster connect to *at least one* strong node with three weak-node neighbors (where we again average over all j) is

$$\begin{aligned} P(S_1|C_{\mathrm{w},1}) &= 1 - \prod_{m=1}^{L_{\mathrm{w},1}-1} (1 - P(\text{new } S_1|m)) \\ &\approx 1 - \prod_{m=1}^{L_{\mathrm{w},1}-1} \left(1 - \frac{m}{N_{\mathrm{eff}}}\right), \end{aligned} \tag{14}$$

where N_{eff} is Eq. 2.

Interestingly, the formula (14) is equivalent to a generalized birthday problem, where N_{eff} is the effective number of "days in a year" and $L_{\mathrm{w},1} = C_{\mathrm{w},1}+2$ is the number of "people". Following previous literature [41], the critical value of $C_{\mathrm{w},1}$, when the probability that S_1 will fail is 1/2, is Eq. (1).

B. Mechanism for Second-Order Cascades

Because a first-order cascade is a poor approximation of the DK mechanism, we will analytically study the probability of cascades producing yet more cascades. The next simplest case that we can study are second-order cascades.

In a second-order cascade, more strong nodes fail after an initial first-order cascade (for example, "S_2" in Fig. 2a). The probability of a second-order cascade acts as an upper bound on the probability of DKs. The simplest way this can occur is when a strong node with three weak-node neighbors fails (see S_1' in Fig. 4b), because it can bridge the failed cluster and a new weak-node cluster. In order to find this probability, we will need to find the probability that the S_1' node will connect to two failed nodes. First, to determine how many nodes failed, we first recall that the probability for the first strong nodes to fail in the cascade, S_1 in Fig. 4a, is a generalized birthday problem, therefore the probability for k nodes like S_1 to fail is Poisson distributed, for large N_{eff} and $L_{\mathrm{w},1}$:

$$P_{S_1}(k) = e^{-\lambda_{\mathrm{w},1}} \frac{\lambda_{\mathrm{w},1}^k}{k!}, \tag{15}$$

where

$$\lambda_{\mathrm{w},1} = \frac{L_{\mathrm{w},1}(L_{\mathrm{w},1}-1)}{2N_{\mathrm{eff}}}. \tag{16}$$

This finding implies that the size of secondary failures increases *quadratically* with the size of the initial failure, therefore, initial failures lead to unexpectedly large secondary failures. Intuitively, the statistics are Poisson because each event (the failure of S_1-like nodes) is statistically independent, and occurs with a low probability ($P(C_{\mathrm{w},1} > C_{\mathrm{crit,DK}})$ is small), and there are lots of opportunities for the event to occur (N_{eff} is large). Let the size of a cluster, i, be $C_{\mathrm{w},2}^{(i)}$, and let $C_{\mathrm{w},2} = \sum_{i=1}^{k} C_{\mathrm{w},2}^{(i)}$, then the probability $C_{\mathrm{w},2} > 0$ is

$$P(S_1|C_{\mathrm{w},1}) = 1 - e^{-\lambda_{\mathrm{w},1}}. \tag{17}$$

This can also be derived from Eq. (14), by taking

$$\begin{aligned} \prod_{m=1}^{L_{\mathrm{w},1}-1} \left(1 - \frac{m}{N_{\mathrm{eff}}}\right) &= e^{\log\left(\prod_{m=1}^{L_{\mathrm{w},1}-1}\left(1-\frac{m}{N_{\mathrm{eff}}}\right)\right)} \\ &= e^{\sum_{m=1}^{L_{\mathrm{w},1}-1}\log\left(1-\frac{m}{N_{\mathrm{eff}}}\right)} \\ &\approx e^{-L_{\mathrm{w},1}(L_{\mathrm{w},1}-1)/(2N_{\mathrm{eff}})}. \end{aligned} \tag{18}$$

Using this approximation, we can directly solve for when $P(S_1|C_{\mathrm{w},1}) = 1/2$ (Eq. (1)).

Now that we know the distribution of weak-node clusters that first fail (Eq. (15)), we need to know how many nodes are in each weak-node secondary cluster fail. For each cluster of size $C_{\mathrm{w},2}^{(i)}$, there are $L_{\mathrm{w},2}^{(i)} = C_{\mathrm{w},2}^{(i)} + 2$ edges connected to strong nodes. Furthermore, we find empirically that the size distribution of weak-node clusters is

$$P(C_{\mathrm{w}}) \sim C_{\mathrm{w}}^{-\eta}, \tag{19}$$

where $\eta = 2.26 \pm 0.03$ (measured for $\epsilon = 3.2 \times 10^{-4}$ and $N = 10^6$, see Appendix XII). Note, however, that the size distribution of $C_{\mathrm{w},2}^{(i)}$ is not Eq. (19), because the number of opportunities to connect to a node of size $C_{\mathrm{w},2}^{(i)}$ is proportional to the number of links emanating from the cluster, which increases as $L_{\mathrm{w},2}^{(i)}$. Namely, S_1 is a strong node with three weak-node neighbors, and there are on average $qL_{\mathrm{w},2}^{(i)}$ strong nodes with three weak-node neighbors connected to a cluster of size $C_{\mathrm{w},2}^{(i)}$. Therefore, the probability to connect to *any* cluster of size $C_{\mathrm{w},2}^{(i)}$, $P_{\mathrm{w},2}(C_{\mathrm{w},2}^{(i)})$, is

$$\begin{aligned} P_{\mathrm{w},2}(C_{\mathrm{w},2}^{(i)}) &= \frac{P(C_{\mathrm{w}} = C_{\mathrm{w},2}^{(i)}) q L_{\mathrm{w},2}^{(i)}}{\langle q L_{\mathrm{w},2}^{(i)} \rangle} \\ &= \frac{P(C_{\mathrm{w},2}^{(i)}) L_{\mathrm{w},2}^{(i)}}{\langle L_{\mathrm{w},2}^{(i)} \rangle}. \end{aligned} \tag{20}$$

This is intuitively similar to "excess degree" seen in random network literature (Eq. (22) in [52]).

S_1'-like nodes can only appear through one of two conditions: (1) an S_1' node is created from two links in the non-initial clusters (left probability in Fig. 5), or (2) an S_1' node spans the initial weak-node cluster, and a newer cluster (right probability in Fig. 5). Therefore $P(C_{\mathrm{w},3} > 0|C_{\mathrm{w},1})$ is 1 minus the probability that both conditions do not occur.

Condition (1) is the simpler of the two to calculate, because the derivation for the equation is very similar to that of equation (17). For each secondary cluster, the

number of links is $L'^{(i)}_{\mathrm{w},2} = C^{(i)}_{\mathrm{w},2} + 2 - 1$, where we subtract one because one link is used to connect to the initial cluster, therefore $L'_{\mathrm{w},2} = \sum_{i=1}^{k} L'^{(i)}_{\mathrm{w},2} = \sum_{i=1}^{k} C^{(i)}_{\mathrm{w},2} + k$, and

$$P(C_{\mathrm{w},3} = 0|\text{condition (1)}) = e^{-\lambda_{\mathrm{w},2}}, \tag{21}$$

where

$$\lambda_{\mathrm{w},2} = \frac{L'_{\mathrm{w},2}(L'_{\mathrm{w},2} - 1)}{2N_{\mathrm{eff}}}. \tag{22}$$

To solve for condition (2), note that there are $L'_{\mathrm{w},1} = C_{\mathrm{w},1} + 2 - 2k = C_{\mathrm{w},1} - 2(k-1)$ links available from the initial cluster, because $2k$ links were used to connect to S_1-like nodes. There are $qL'_{\mathrm{w},1}$ links from the initial weak-node cluster to strong nodes with three weak-node neighbors, therefore, there are $2qL'_{\mathrm{w},1}$ opportunities for strong nodes with three weak-node neighbors to connect to any secondary failed weak-node clusters. The probability for one link from the initial cluster to connect to any of the secondary failed weak-node clusters is

$$P(1 \text{ link} \to C_{\mathrm{w},2}) = \frac{qL'_{\mathrm{w},2}}{qpN}, \tag{23}$$

where the denominator is the total number of links from weak-node clusters to strong nodes with three weak-node neighbors. This implies that the probability at least one strong node has a link in the initial and secondary clusters is

$$P(C_{\mathrm{w},3} > 0|C_{\mathrm{w},1} \leftrightarrow C_{\mathrm{w},2}) = 1 - \left(1 - \frac{qL'_{\mathrm{w},2}}{qpN}\right)^{2qL'_{\mathrm{w},1}}. \tag{24}$$

Similar logic from the perspective of the secondary weak-node clusters suggests that

$$P(C_{\mathrm{w},3} > 0|C_{\mathrm{w},2} \leftrightarrow C_{\mathrm{w},1}) = 1 - \left(1 - \frac{L'_{\mathrm{w},1}}{pN}\right)^{2qL'_{\mathrm{w},2}}. \tag{25}$$

We therefore have a paradox. We expect that the probability for a strong node to span initial and secondary clusters should be independent of the order we choose to connect them (the probability of the secondary clusters connecting to the initial cluster should be the same as the initial to the secondary clusters). However, if we make an ansatz that $1 \ll L'_{\mathrm{w},1} \ll N$ and $1 \ll L'_{\mathrm{w},2} \ll N$, we can take the Taylor series of either equation and approximate the sum as

$$\begin{aligned} P(C_{\mathrm{w},3} > 0|C_{\mathrm{w},1} \leftrightarrow C_{\mathrm{w},2}) &\approx P(C_{\mathrm{w},3} > 0|C_{\mathrm{w},2} \leftrightarrow C_{\mathrm{w},1}) \\ &\approx 1 - e^{-\lambda_{\mathrm{w},1\leftrightarrow 2}}, \end{aligned} \tag{26}$$

or

$$P(C_{\mathrm{w},3} = 0|\text{condition (2)}) \approx e^{-\lambda_{\mathrm{w},1\leftrightarrow 2}}, \tag{27}$$

where

$$\lambda_{\mathrm{w},1\leftrightarrow 2} = \frac{2qL'_{\mathrm{w},1}L'_{\mathrm{w},2}}{pN}. \tag{28}$$

This equation is also order-independent, as we expect. We can therefore write $P(C_{\mathrm{w},3} > 0|C_{\mathrm{w},1})$ as the probability that neither condition (1) nor condition (2) occurs (i.e., Fig. 5). Recall that this is the probability that at least one S'_1-like node occurs over all values of $C^{(1)}_{\mathrm{w},2}$, $C^{(2)}_{\mathrm{w},2}$, ..., $C^{(k)}_{\mathrm{w},2}$. The probability of k different S_1-like nodes is $P_{S_1}(k)$, while $P_{\mathrm{w},2}(C^{(1)}_{\mathrm{w},2}, C^{(2)}_{\mathrm{w},2}, ..., C^{(k)}_{\mathrm{w},2}) = P_{\mathrm{w},2}(C^{(1)}_{\mathrm{w},2}) \times P_{\mathrm{w},2}(C^{(1)}_{\mathrm{w},2}) \times ... \times P_{\mathrm{w},2}(C^{(k)}_{\mathrm{w},2})$, where $P_{\mathrm{w},2}(C^{(i)}_{\mathrm{w},2})$ is Eq. (20). Therefore, if we approximate $C^{(i)}_{\mathrm{w},2}$ as a continuous variable, $P(C_{\mathrm{w},3} > 0|C_{\mathrm{w},1})$ can be written as

$$\begin{aligned} &P(C_{\mathrm{w},3} > 0|C_{\mathrm{w},1}) = \\ &\sum_{k=1}^{\infty} P_{S_1}(k) \int_{C^{(1)}_{\mathrm{w},2}=1}^{\infty} \int_{C^{(2)}_{\mathrm{w},2}=1}^{\infty} ... \int_{C^{(k)}_{\mathrm{w},2}=1}^{\infty} P(C_{\mathrm{w},3} > 0|C_{\mathrm{w},1}, C^{(1)}_{\mathrm{w},2}, ..., C^{(k)}_{\mathrm{w},2}) P_{\mathrm{w},2}(C^{(1)}_{\mathrm{w},2}) \times P_{\mathrm{w},2}(C^{(1)}_{\mathrm{w},2}) \times ... \times P_{\mathrm{w},2}(C^{(k)}_{\mathrm{w},2}), \end{aligned} \tag{29}$$

where

$$P(C_{\mathrm{w},3} > 0|C_{\mathrm{w},1}, C^{(1)}_{\mathrm{w},2}, ..., C^{(k)}_{\mathrm{w},2}) = 1 - e^{-\lambda_{\mathrm{w},2} - \lambda_{\mathrm{w},1\leftrightarrow 2}}. \tag{30}$$

In Eq. (29), we notice that the lower bound of each integral is 1 simply because $C^{(i)}_{\mathrm{w},2} \geq 1$, i.e., there must be at least one node. Using the above findings, Eq. (29) can be written more compactly as

$$P(C_{\mathrm{w},3} > 0|C_{\mathrm{w},1}) = \sum_{k=1}^{\infty} P_{S_1}(k) \left(1 - \int_{C^{(1)}_{\mathrm{w},2}, C^{(2)}_{\mathrm{w},2}, ... C^{(k)}_{\mathrm{w},2}=1}^{\infty} e^{-\lambda_{\mathrm{w},2} - \lambda_{\mathrm{w},1\leftrightarrow 2}} \prod_{i=1}^{k} P_{\mathrm{w},2}(C^{(i)}_{\mathrm{w},2}) dC^{(i)}_{\mathrm{w},2}\right). \tag{31}$$

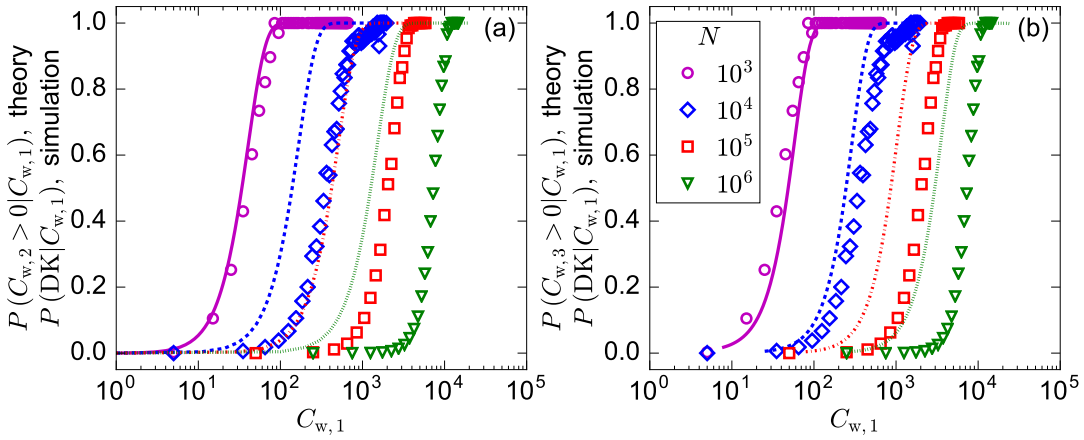

FIG. 11. (Color online) Probability of first-order cascades, second-order cascades, and DKs versus $C_{\rm w,1}$. Lines are the theoretical probabilities of (a) first-order cascades and (b) second-order cascades ($N = 10^3$, solid line; $N = 10^4$, dashed line; $N = 10^5$, dash-dot line; $N = 10^6$, dotted line), and plot markers are the simulation-based probabilities of DKs versus $C_{\rm w,1}$ for $\epsilon = 1.00 \times 10^{-3}$.

As a sanity check, if $\lambda_{\rm w,2} = \lambda_{\rm w,1\leftrightarrow 2} \to \infty$, which is the unrealistic condition that $C_{\rm w,3} > 0$ whenever $C_{\rm w,2} > 0$, the equation reduces to Eq. (17). In this equation, we recall that $C^{(i)}_{\rm w,2}$ are all independent, and $\lambda_{\rm w,2}$ and $\lambda_{\rm w,1\leftrightarrow 2}$ are defined to be a function of $\sum_{i=1}^{k} C^{(i)}_{\rm w,2}$. Sadly, however, this integral is not analytically tractable, therefore we numerically determine the integral using Mathematica. For large k, we suffer from the curse of dimensionality, therefore we use a cutoff: we ignore any k where $P_{S_1}(k) < \delta$, where $\delta = 10^{-4}$.

C. Comparison to Simulations

We compare the second-order cascade theory to simulations of DKs in order to check how well the assumptions approximate the DK mechanism. The first thing we notice is that, as expected, the agreement with simulations of DKs is not perfect, especially for large N (cf. Fig. 11b), because the theory is still a necessary but not sufficient condition for DKs (cf. Fig. 11a). This analysis can also give us a much better understanding of the relationship between $C_{\rm crit,DK}$, the critical value of $C_{\rm w,1}$ such that the probability of DKs is 1/2, and N. By understanding how $C_{\rm crit,DK}$ scales with N, we can determine whether DKs exist in the thermodynamic limit.

We can define $C_{\rm crit,2}$ to be the critical values of $C_{\rm w,1}$ when the probability of a second-order cascade is 1/2. We find that $C_{\rm crit,1}$, $C_{\rm crit,2}$, & $C_{\rm crit,DK}$ all appear to scale differently with system size. As Eq. (1) shows, in the first-order cascade theory, $C_{\rm crit,1} \sim N^b$ where $b = 0.5$, while we find that, for second-order cascade theory, $b = 0.55 \pm 0.01$, and for simulations of DKs, $b = 0.59 \pm 0.03$ (cf. Fig. 12). To see more explicitly how second-order and first-order cascade theory differ in their scaling, we can look at Fig. 6, where it is clear that the second-order cascade theory reaches probability 1/2 at larger and larger values of $C_{\rm w,1}$ as N increases, compared to first-order cascade theory.

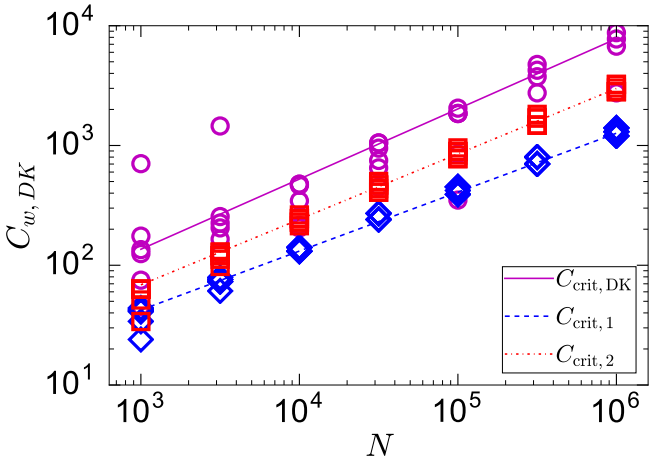

FIG. 12. (Color online) Scaling of $C_{\rm w,1}$ critical values versus N: $C_{\rm crit,DK}$ (magenta circles), $C_{\rm crit,1}$ (blue diamonds), and $C_{\rm crit,2}$ (red squares). We plot critical values for $3.2 \times 10^{-4} \leq \epsilon \leq 1.0 \times 10^{-1}$, and fit those values to the model $a \times N^b$, where $b = 0.59 \pm 0.03$ for simulations, $b = 0.49 \pm 0.01$ for the first-order cascade theory, and $b = 0.55 \pm 0.01$ for the second-order cascade theory.

Second-order cascades are a necessary, but not sufficient, condition for DKs, therefore $P({\rm SO}|C_{\rm w,1}) > P({\rm DK}|C_{\rm w,1})$, which implies that $C_{\rm crit,2} < C_{\rm crit,DK}$, where $C_{\rm crit,DK}$ is the critical size of $C_{\rm w,1}$ such that $P({\rm DK}|C_{\rm w,1}) = 1/2$. We find that these bounds agree with what we see numerically (cf. Fig. 2b), and $C_{\rm crit,2}$ is in much closer agreement than $C_{\rm crit,1}$ is to $C_{\rm crit,DK}$.

XI. APPENDIX B: ALTERNATIVE INITIAL CONDITIONS

We ask whether the steady state behavior of the IN and CC models is independent of the initial conditions. To check this, we calculate the equilibrium fraction of weak nodes in the network, $\langle p \rangle$, and the probability of a DK, across various initial conditions as we vary ϵ. In Fig. 13, we demonstrate that the fraction of weak nodes over time, $p(t)$, stabilizes to a value, $\langle p \rangle$, after a time, t, greater than $t_{\text{relax}} = 5 \times N$ timesteps across several different initial conditions in both models. To further demonstrate this, in Fig. 14 we plot the average fraction of weak nodes, $\langle p \rangle$, after $t_{\text{relax}} = 5 \times N$ for different reinforcement probabilities ϵ and find no statistically significant difference. Although we demonstrate that the number of weak nodes does not appear to be affected by the initial conditions, this does not guarantee that the distribution of cascade failures is unaffected. As a simple check for the CC model, Fig. 15 shows that p_{initial} has little effect on the probability of DK, $P(\text{DK})$. In fact, we find no statistically significant difference in $P(\text{DK})$ across initial conditions. Overall, it does not appear as though the dynamics are affected by initial conditions for $t > t_{\text{relax}}$.

XII. APPENDIX C: WEAK-NODE CLUSTER SIZE DISTRIBUTION

In this section, we show that the weak cluster size distribution is a power-law tail with an exponential cut-off. The power-law exponent is approximately 2.15 for the IN model, and 2.26 for the CC model (see Fig. 16), which differs considerably from what we would expect for SOC systems, where the exponent should be 2.5 for percolating clusters [39]. Future work, however, is necessary to determine whether the exponent is not equal to 2.5 in the thermodynamic limit.

We also notice that the exponential cut-off of the cluster distribution is approximately proportional to ϵ (Fig. 17). This agrees with our intuition that the weak-node cluster sizes should be power-law distributed as $\epsilon \to 0$. The reason it is directly proportional to ϵ, however, is in part because large clusters shrink in size when they fail. Nodes in the failed cluster are strengthened, which reduces the size of the weak-node cluster. Clusters of size C_{w} will not have any nodes strengthened with probability $(1-\epsilon)^{C_{\text{w}}} \approx \exp(-\epsilon C_{\text{w}})$. This suggests that the probability a cluster is of size C_{w} is $\sim \exp(-\epsilon C_{\text{w}})$, and therefore gives us intuition as to why $\lambda \sim \epsilon$, although it cannot explain why $\lambda \neq \epsilon$.

XIII. APPENDIX D: RECEIVER OPERATING CHARACTERISTICS AND CRITICAL VALUES

The ROC curve (Fig. 18) is created by plotting the true positive rate (TPR) against the false positive rate (FPR) as the discrimination threshold is varied [42]. Here, the TPR indicates the probability of DKs being correctly predicted. The FPR is the probability of DKs being wrongly predicted. We conclude that the second predictor, $C_{\text{w},1}$, is close to being an optimal predictor. Moreover, DKs are predicted by $C_{\text{w},1}$ with high accuracy, because DKs are unlikely to occur for $C_{\text{w},1} < C_{\text{crit,DK}}$, and due to the power-law distribution of $C_{\text{w},1}$ as shown in Fig. 19, most initial failures lead to small cascades. For example, when $\epsilon = 1.0 \times 10^{-3}$, the probability $P(C_{\text{w},1} < C_{\text{crit,DK}})$ is 99.986%.

Lastly, because nodes degrade at a fixed rate, the cost of upgrades is therefore also approximately the same. Our control strategy is able to upgrade nodes more intelligently, and can reduce the size of large cascades without the risk of upgrading more nodes than originally tolerated, as shown in Fig. 20.

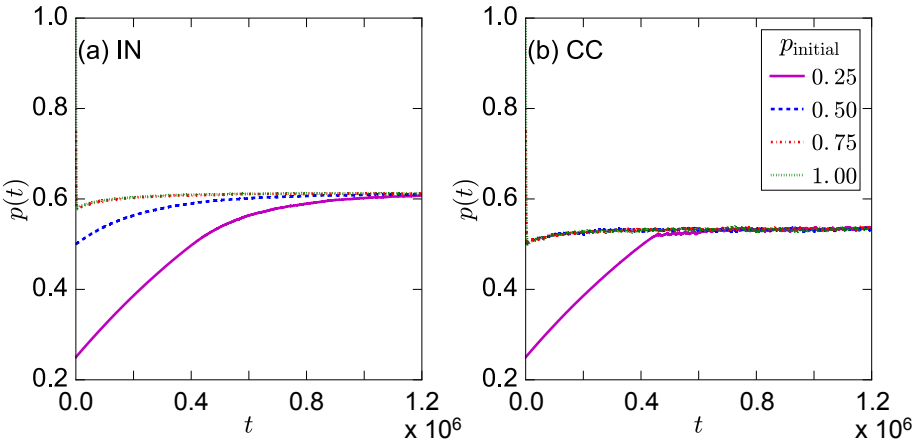

FIG. 13. (Color online) Fraction of weak nodes, $p(t)$, over time for various initial conditions. We vary the initial fraction of weak nodes, p_{initial}(0.25, solid line; 0.50, dashed line; 0.75, dash-dot line; 1.00, dotted line), from 0.25 to 1.0 for (a) the IN model and (b) the CC model. The equilibrium value, $\langle p \rangle$, is not significantly different for various initial conditions after a time, t, greater than $t_{\text{relax}} = 5 \times N$ timesteps. In this figure, we take one network realization with $N = 10^6$, $\epsilon = 3.2 \times 10^{-4}$.

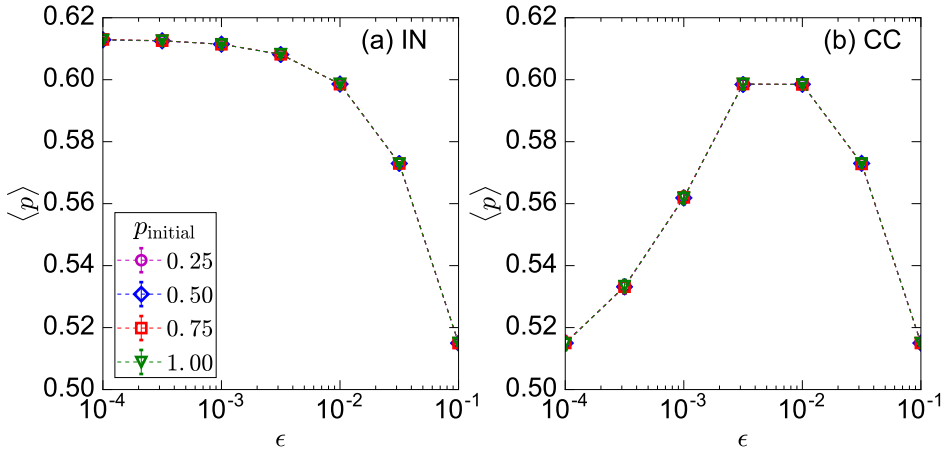

FIG. 14. (Color online) Average fraction of weak nodes after $t_{\text{relax}} = 5 \times N$ for different initial conditions (a) for the IN model and (b) for the CC model. The results are averaged for $N = 10^6$ over $15 \times N$ timesteps and 5 network realizations.

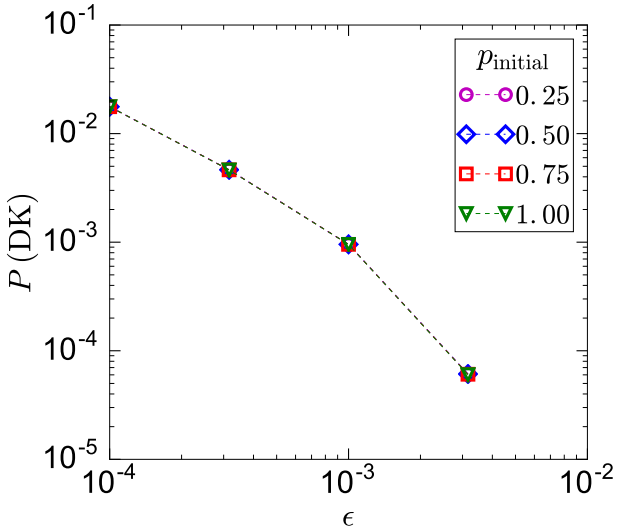

FIG. 15. (Color online) Probability of DK vs. ϵ for different p_{initial} in the CC model for network size with $N = 10^6$. Standard errors are smaller than marker sizes.

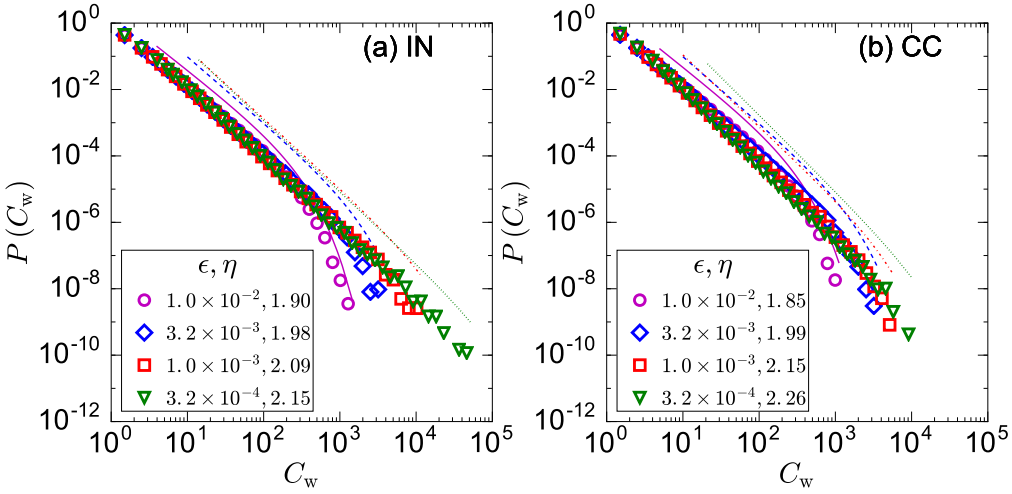

FIG. 16. (Color online) Weak-node cluster size distribution, $P(C_w)$, for $N = 10^6$ and 10 network realizations at $t = 9 \times N$. The parameter η is fitting exponent of the maximum likelihood power-law with exponential tail ($\epsilon = 1.0 \times 10^{-2}$, solid line; $\epsilon = 3.2 \times 10^{-3}$, dashed line; $\epsilon = 1.0 \times 10^{-3}$, dash-dot line; $\epsilon = 3.2 \times 10^{-4}$, dotted line).

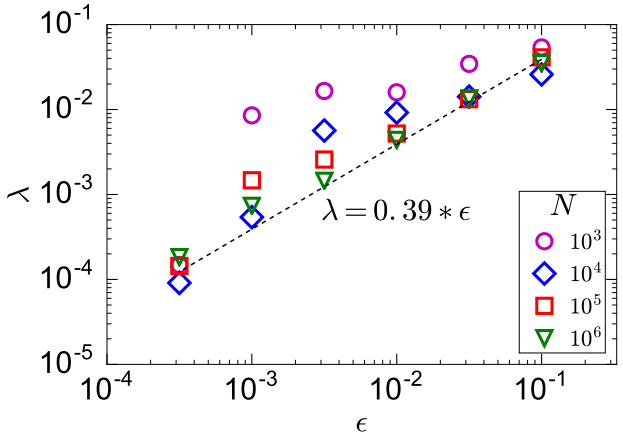

FIG. 17. The exponential cut-off parameter, λ, for the weak-node cluster distribution, $P(C_{\rm w})$, versus ϵ for the CC model (cf. Fig. 16 for $P(C_{\rm w})$). The best-fit relation between λ and ϵ is $\lambda = 0.39 \times \epsilon$ for $N = 10^5$ and 10^6.

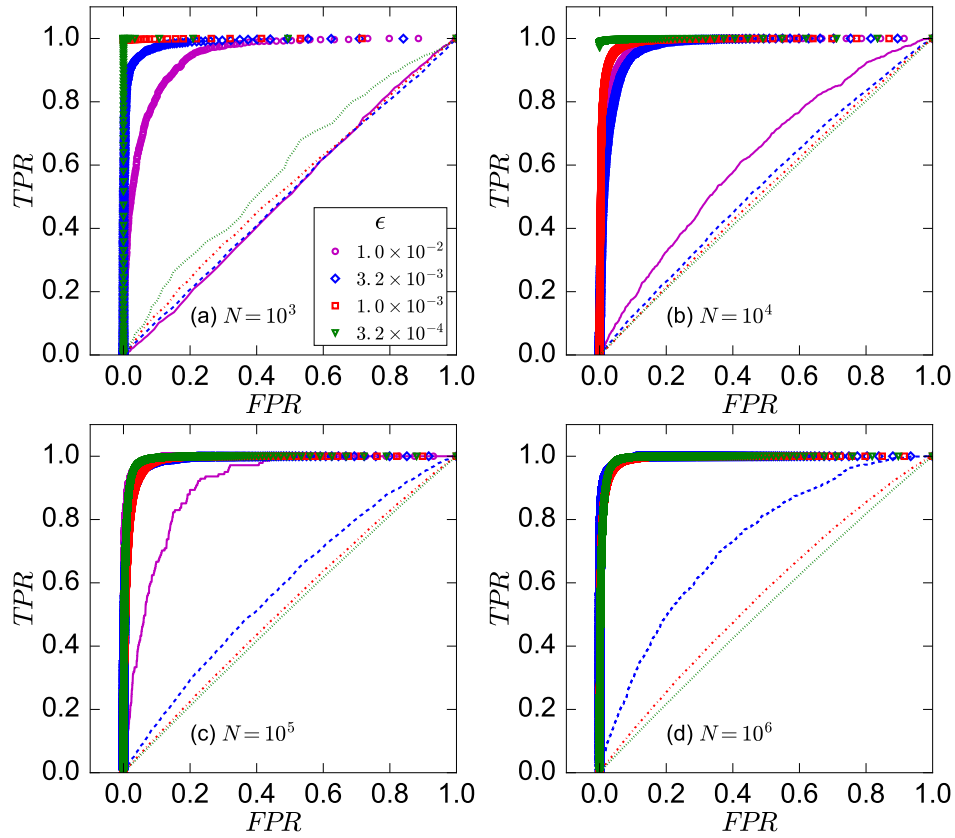

FIG. 18. (Color online) ROC curves for different N and different ϵ using two different predictors. Lines are the results predicted by the fraction of weak nodes ($\epsilon = 1.0\times10^{-2}$, solid line; $\epsilon = 3.2\times10^{-3}$, dashed line; $\epsilon = 1.0\times10^{-3}$, dash-dot line; $\epsilon = 3.2\times10^{-4}$, dotted line), and the open markers represent the results predicted by the size of the first weak-node cluster.

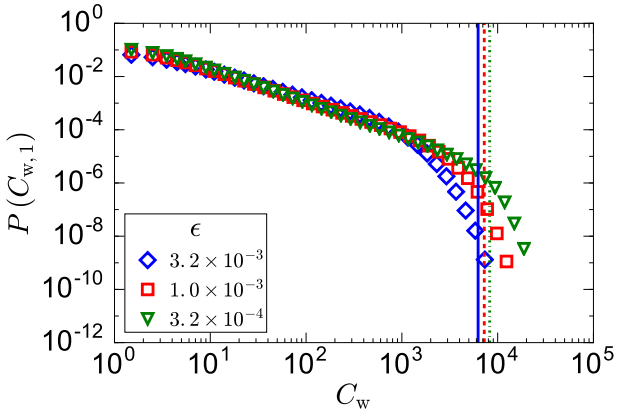

FIG. 19. (Color online) The probability the first weak-node cluster that fails is of size C_w, for $N = 10^6$ over $15 \times N$ timesteps and 1 network realization. The vertical lines show the value of $C_{crit,DK}$ ($\epsilon = 3.2\times10^{-3}$, solid line; $\epsilon = 1.0\times10^{-3}$, dashed line; $\epsilon = 3.2\times10^{-4}$, dash-dot line).

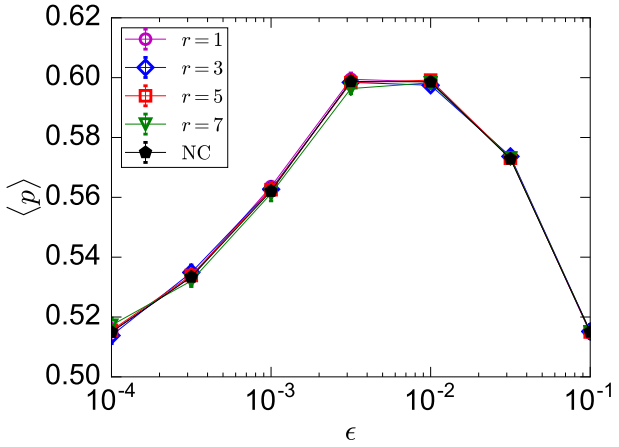

FIG. 20. (Color online) The fraction of weak nodes averaged over T timesteps with and without control for $N = 10^6$ and 10 network realizations. For the uncontrolled system, $T = 5\times N$, while for the controlled system $T = 10 \times N$.

[1] D. Sornette, International Journal of Terraspace Science and Engineering **2**, 1 (2009).
[2] D. Sornette and G. Ouillon, Eur. Phys. J. Special Topics **205**, 53 (2012).
[3] J. M. Carlson and J. Doyle, Phys. Rev. E **60**, 1412 (1999).
[4] P. Bak, *How Nature Works: the Science of Self-organized Criticality* (Copernicus, New York, 1996).
[5] P. Bak, C. Tang, and K. Wiesenfeld, Phys. Rev. Lett. **59**, 381 (1987).
[6] P. Bak, C. Tang, and K. Wiesenfeld, Phys. Rev. A **38**, 364 (1988).
[7] B. A. Carreras, V. E. Lynch, I. Dobson, and D. E. Newman, Chaos: An interdisciplinary journal of nonlinear science **12**, 985 (2002).
[8] H. Hoffmann and D. W. Payton, Chaos Soliton. Fract. **67**, 87 (2014).
[9] S. Wheatley, B. Sovacool, and D. Sornette, Risk analysis **37**, 99 (2017).
[10] J. Lorenz, S. Battiston, and F. Schweitzer, EPJ-B **71**, 441 (2009).
[11] C. J. Tessone, A. Garas, B. Guerra, and F. Schweitzer, Journal of Statistical Physics **151**, 765 (2013).
[12] R. M. D'Souza, Science **358**, 860 (2017).
[13] B. Podobnik, D. Horvatic, T. Lipic, M. Perc, J. M. Buld, and H. E. Stanley, J. R. Soc. Interface **12**, 20150770 (2015).
[14] X. Weng, Y. Hong, A. Xue, and S. Mei, Journal of Control Theory and Applications **4**, 235 (2006).
[15] M. Perc, J. R. Soc. Interface **11**, 20140378 (2014).
[16] A. Majdandzic, B. Podobnik, S. V. Buldyrev, D. Y. Kenett, S. Havlin, and H. E. Stanley, Nature Physics **10**, 34 (2014).
[17] L. Böttcher, M. Lukovic, J. Nagler, S. Havlin, and H. J. Herrmann, Sci. Rep. **7**, 41729 (2017).
[18] J. Gao, S. V. Buldyrev, S. Havlin, and H. E. Stanley, Phys. Rev. Lett. **107**, 195701 (2011).
[19] C. D. Brummitt, R. M. D'Souza, and E. A. Leicht, Proceedings of the National Academy of Sciences **109**, E680 (2012), http://www.pnas.org/content/109/12/E680.full.pdf.
[20] T. A. Short, *Electric Power Distribution Handbook* (CRC Press, Boca Raton, Fl, 2004).
[21] Y. Yang, T. Nishikawa, and A. E. Motter, Science **358**, eaan3184 (2017).
[22] W. O. Kermack and A. G. McKendrick, Proc. R. Soc. A: Math. Phys. Eng. Sci. **115**, 700 (1927).
[23] D. Centola and M. Macy, American journal of Sociology **113**, 702 (2007).
[24] G. J. Baxter, S. N. Dorogovtsev, A. V. Goltsev, and J. F. F. Mendes, Phys. Rev. E **82**, 011103 (2010).
[25] S. Janson, T. Luczak, T. Turova, and T. Vallier, The Annals of Applied Probability **22**, 1989 (2012).
[26] O. Ramos, E. Altshuler, and K. J. Måløy, Phys. Rev. Lett. **102**, 078701 (2009).
[27] D. E. Newman, B. A. Carreras, N. S. Degala, and I. Dobson, in *2012 45th Hawaii International Conference on System Sciences* (2012) pp. 2082–2090.
[28] R. J. Geller, D. D. Jackson, Y. Y. Kagan, and F. Mulargia, Science **275**, 1616 (1997).
[29] N. Taleb, *The Black Swan: The Impact of the Highly Improbable* (Random House, 2007).
[30] D. O. Cajueiro and R. F. S. Andrade, Phys. Rev. E **81**, 015102 (2010).
[31] P.-A. Noël, C. D. Brummitt, and R. M. D'Souza, Phys. Rev. Lett. **111**, 078701 (2013).
[32] H. L. d. S. Cavalcante, M. Oriá, D. Sornette, E. Ott, and D. J. Gauthier, Phys. Rev. Lett. **111**, 198701 (2013).
[33] A. O. Allen, *Probability, statistics, and queueing theory* (Academic Press, 2014).
[34] X. C. Zeng and D. W. Oxtoby, J. Chem. Phys. **94**, 4472 (1991).
[35] G. A. Pagini and M. Aiello, Physica A **392**, 2688 (2013).
[36] J. Mena-Lorca and H. Hethcote, J Math Biol. **30**, 693 (1992).
[37] R. Cohen, K. Erez, D. Ben-Avraham, and S. Havlin, Phys. Rev. Lett. **85**, 4626.
[38] P. Alstrøm, Phys. Rev. A **38**, 4905 (1988).
[39] K. Christensen, H. Flyvbjerg, and Z. Olami, Phys. Rev. Lett. **71**, 2737 (1993).
[40] D. J. Watts and S. Strogatz, Nature **393**, 440 (1998).
[41] S. E. Ahmed and R. J. McIntosh, Crux Math **26**, 151 (2000).
[42] C. D. Brown and H. T. Davis, Chemometrics and Intelligent Laboratory Systems **80**, 24 (2006).
[43] A. E. Motter, Physics **6**, 120 (2013).
[44] W. S. Lee, D. L. Grosh, F. A. Tillman, and C. H. Lie, IEEE Trans. Rel. **34**, 194 (1985).
[45] J. K. Vaurio, IEEE Trans. Rel. **47** (1998).
[46] H. Hong, J. D. Kubik, and J. C. Stein, J. Finance **59**, 137 (2004).
[47] W. Pan, Y. Altshuler, and A. Pentland, in *International Conference on Privacy, Security, Risk and Trust and 2012 International Conference on Social Computing* (IEEE, 2012) pp. 203–209.
[48] N. Johnson and B. Tivnan, Eur. Phys. J. Special Topics **205**, 65 (2012).
[49] N. Hodas and K. Lerman, Sci. Rep. **4** (2014), 10.1038/srep04343.
[50] D. Centola, Science **329**, 1194 (2010).
[51] D. J. Watts, Proceedings of the National Academy of Sciences **99**, 5766 (2002).
[52] M. Newman, SIAM Rev. **45**, 167 (2003).