

Explicit size distributions of failure cascades redefine systemic risk on finite networks

Rebekka Burkholz^{1,*}, Hans J. Herrmann^{2,3}, and Frank Schweitzer⁴

¹ETH Zurich, Computer Science, Zurich, 8092, Switzerland

²ETH Zurich, Computational Physics for Engineering Materials, IfB, Zurich, 8093, Switzerland

³Universidade Federal do Ceará, Física, Fortaleza Ceará, 60451-970, Brazil

⁴ETH Zurich, Management, Technology and Economics, Zurich, 8092, Switzerland

*rburkholz@ethz.ch

ABSTRACT

How big is the risk that a few initial failures of nodes in a network amplify to large cascades that span a substantial share of all nodes? Predicting the final cascade size is critical to ensure the functioning of a system as a whole. Yet, this task is hampered by uncertain or changing parameters and missing information. In infinitely large networks, the average cascade size can often be well estimated by established approaches building on local tree approximations and mean field approximations. Yet, as we demonstrate, in finite networks, this average does not even need to be a likely outcome. Instead, we find broad and even bimodal cascade size distributions. This phenomenon persists for system sizes up to 10^7 and different cascade models, i.e. it is relevant for most real systems. To show this, we derive explicit closed-form solutions for the full probability distribution of the final cascade size. We focus on two topological limit cases, the complete network representing a dense network with a very narrow degree distribution, and the star network representing a sparse network with a inhomogeneous degree distribution. Those topologies are of great interest, as they either minimize or maximize the average cascade size and are common motifs in many real world networks.

Introduction

Systemic risk is defined as the risk that a large fraction, $\rho \rightarrow 1$, of a system fails. This can happen because extreme events shock the system¹ or because the failures of a few system elements cause new failures. This propagation results in failure cascades that can eventually reach the size of the system. In the aftermath of the financial crisis in 2008/2009, systemic risks induced by such cascades received attention especially in the context of financial contagion of bank defaults²⁻⁴. Also the increasing complexity of economic value chains⁵ raised awareness for similar effects⁶, in particular concerning critical resources⁷ and food⁸. Many of these works are inspired by stability analyses of ecological food webs⁹⁻¹¹. But also generic models from physics and related areas have been utilized to study failure cascades. Examples include Ising models¹², models of fiber bundles that break under stress¹³, models of epidemic spreading^{14,15}, voter models¹⁶, and models of overload failures in power grids or other infrastructure¹⁷. Many of these cascade models have in common that they can be mapped to a generic threshold model for propagating failures on a graph G , or a complex network, where nodes represent the system elements and links between nodes their interactions.

Network model Formally $G = (V, E)$ is a constant undirected network consisting of $N = |V|$ nodes, which are elements of the node set V . The link set E contains tuples of the form (i, j) , each representing a pairwise connection between nodes i and j . Network topologies vary between Erdős-Rényi random networks, where the number of links per node can be described by an average degree, and scale-free networks, where the number of links vastly differs between core nodes, or hubs with very many links, and peripheral nodes with very few links. In this paper, we discuss limit cases of these two topologies, i.e. the complete network, where each node has the same number of links, $N - 1$, and the star network, where the center node has $N - 1$ links, but the peripheral nodes have only one link to the center. Both are recurring motifs of real world networks and an approximation of prevailing core periphery structures. In particular, financial and economic networks often have densely connected core nodes with peripheral nodes loosely connected to them^{18,19}. Most importantly, for many models, one of the two, the complete network or the star, minimizes while the other maximizes average systemic risk^{20,21}, when systemic risk is measured by the cascade size, i.e. the fraction of failed nodes in the network.

Cascade processes A failure of a node i at time t is indicated by its binary state variable $s_i(t) = 1$. Otherwise, it is healthy or functional, $s_i(t) = 0$, and can switch to failed in each of the discrete time steps $t = 0, \dots, T$. For simplicity, we exclude the possibility to recover, i.e. node states can only change from 0 to 1 but not vice versa. Each node state $s_i(t)$ is characterized by

two node variables, a threshold θ_i and the load $\lambda_i(t)$. θ_i summarizes the “robustness” of a node, its ability to withstand shocks. It is assumed to stay fixed over time. Individual thresholds are initially drawn independently at random from an arbitrary cumulative distribution function (cdf) $F(\theta)$, often also called the response function. A node i fails at time $t + 1$, if its load exceeds its threshold, i.e. $s_i(t + 1) = \Theta[\lambda_i(t) - \theta_i]$, where $\Theta[x]$ is the Heaviside function with $\Theta[x] = 1$ if $x \geq 0$ and $\Theta[x] = 0$ otherwise. The load $\lambda_i(t)$ can change because of interaction with neighboring nodes involving a load distribution process. This basically describes how a failure can propagate along a link if two nodes interact, because a load increase can result in $\lambda_i(t) \geq \theta_i$. Initially, all nodes with $\theta_i \leq \lambda_0$ fail, where $\lambda_i(0) = \lambda_0$ denotes the initial load of a node.

This allows us to define the size of a cascade at time t as the fraction of failed nodes in the network: $\rho(t) = 1/N \sum_{i=1}^N s_i(t)$. A cascade maximally lasts $T = N - 1$ time steps because at least one node needs to fail at each time step to keep the process going. The final size of the cascade, $\rho = \rho(T)$, i.e. the final fraction of failed nodes, is a common measure for systemic risk. While the cascade dynamics itself is deterministic, ρ is a random variable, as it depends on the thresholds that are drawn initially independently at random according to the distribution $F(\theta)$. This randomness can model uncertainty (e.g.g incomplete data or information) or changing exposures of nodes, for instance, due to fluctuating markets or aging system components in an engineering setting. Based on combinatorial arguments, in the materials and methods section and the SI, we derive a closed-form solution for the probability distribution of ρ on complete graphs, $p(\rho_\Delta)$, and stars, $p(\rho_\star)$, of finite arbitrary size N .

Three paradigmatic load distribution mechanisms We discuss three exemplary load distribution mechanisms well established in different fields. We refer to the first one²² as exposure diversification (ED). It means that a node is exposed equally to the risk that any of its neighbors fail. The more neighbors, the better for hedging against such exposures (in low systemic risk regimes)²³. A node carries simply the fraction of its failed neighbors as load. Thus, in complete networks each functional node carries the load $\lambda[k] = k/(N - 1)$ when k nodes have failed.

The ED mechanism has not only attracted theoretical interest^{24–28}, it is also used to explain opinion formation²⁹ and financial contagion^{2,30,31}, where random thresholds correspond to fluctuating exposures between banks and changing capital buffers. The randomness can also be interpreted as model uncertainty in such complex systems³².

The second mechanism, denoted as damage diversification (DD)^{23,33}, assumes that a failing node distributes its total load 1 equally among its network neighbors. Consequently, DD coincides with ED for complete networks.

The third mechanism defines the prominent fiber bundle model^{13,34} in material science. A force, applied to a bundle of fibers, is modeled by a load λ_0 carried by each node in a network, while a node’s threshold determines the strength of a fiber. When a fiber i breaks, it distributes its full load $\lambda_i(t)$ equally to its still functional network neighbors, which become connected afterwards. On complete networks, this mechanism is also known as global load sharing rule and translates into $\lambda[k] = \lambda_0 + k\lambda_0/(N - k)$ that each functional node carries when k nodes have failed. For the special and simplifying case of uniformly distributed thresholds, a formula for the cascade size distribution is known already³⁵. With our new derivations, now arbitrary threshold distributions and thus different fiber materials can be studied. For the star network, we focus on a local load sharing variant where no new edges are created^{36,37}

Results

The only random ingredient that we study is the cumulative threshold distribution $F(\theta)$. Our derivations apply to an arbitrary choice. Yet to be consistent and to allow for comparisons, we discuss only $F \sim \mathcal{N}(\mu + \lambda_0, \sigma^2)$, i.e. a normal distribution with mean $\mu + \lambda_0$ and standard deviation σ . For ED and DD, λ_0 does not influence the results, so we set the initial load to $\lambda_0 = 0$. For the fiber bundle model, we assume $\lambda_0 = 1$, i.e. a value large enough that we can observe both cascades that span the whole network ($\rho = 1$), and cascades that stop early. We are especially interested in parameter regions (μ, σ) that lead to medium average cascade sizes. This implies that such cascades would not encompass the whole system in case of an *infinitely* large network. We then investigate, for such “safe” parameter regions, the size of cascades for *finite* networks.

As a reference case, Fig. 1(a) shows the phase diagram for ρ_Δ for infinitely large systems, assuming a complete network and the ED/DD mechanism^{23,33}. ρ_Δ is taken as a measure of systemic risk. If the mean value of the threshold distribution, μ , is large, the system is less prone to large failure cascades. Taking a small, but fixed value of μ , we recall two important observations, (i) the sharp phase transition between completely safe systems (white, average cascade size zero) and completely failed systems at a critical value of (μ, σ) , (ii) the gradual decrease of systemic risk in areas with high σ . This means, the heterogeneity in thresholds plays a crucial role in amplifying, or preventing, systemic risks. Therefore, we first study the impact of increasing heterogeneity σ on failure cascades in *finite* networks. Fig. 1(b) shows the broad distribution of the final cascade size, spanning values between 0 and 1. Only for large values of σ we find distributions that match our expectations derived from the behavior of infinite systems, i.e. they are unimodal and symmetric. As σ decreases, however, the distributions become asymmetric and broader and even bi-modal. Only very close to a phase transition, broader distributions in finite systems are expected by theoretical physicists and, usually, explicit formulas for such distributions are unknown. But even for parameters

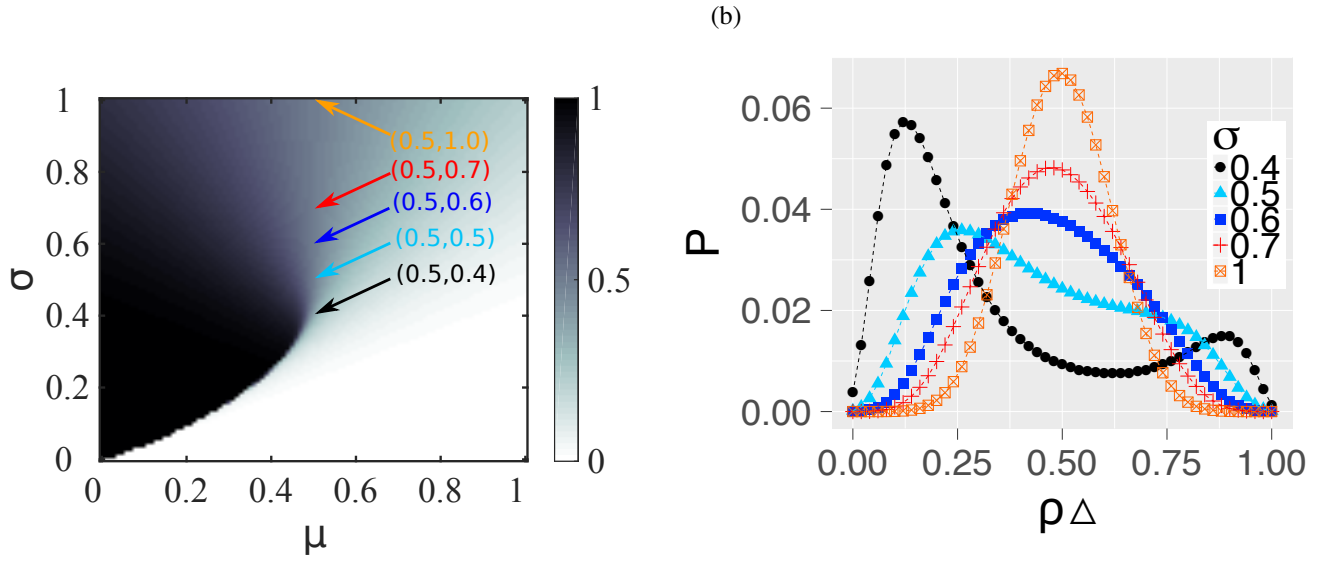


Figure 1. ED/DD mechanism on complete network and normally distributed thresholds, i.e $F \sim \mathcal{N}(\mu, \sigma^2)$, with mean μ and standard deviation σ . (a) (Average) final fraction of failed nodes ρ_Δ (color coded) on infinitely large network. The darker the color, the higher is ρ_Δ . (b) Distribution of ρ_Δ on finite network of size $N = 50$ with respect to different σ . $\mu = 0.5$ is fixed.

further away from the phase transition in Fig. 1(a), the average cascade size is meaningless to characterize the system. Instead, we can expect either small cascades or, as a new observation, very big ones that are likely to destroy the whole system and are relevant for systemic risk estimations. Counter-intuitively, the smaller the threshold heterogeneity, the the less predictive the system behavior becomes.

This strong bi-modality is a new observation that only applies to *finite* systems and is not captured by risk approximations for infinite systems. To further investigate how this behavior depends on the system size N , we simulate, for parameters $(\mu, \sigma) = (0.5, 0.4)$, the distribution of final cascade sizes. Monte Carlo simulations are necessary because, for systems up to $N = 10^8$, the products of large binomial terms and small probabilities that appear in our exact derivations become unfeasible to solve numerically. Fig. 2 shows the results. We see that the bi-modality persists even for networks of size $N = 10^7$. So, the question whether networks of that size can be reasonably described by approximations for infinite networks, is clearly answered by “no”. On the other hand, real economic networks which are prone to failure cascades are often of the order 10^4 - 10^6 , for example ownership networks or collaboration networks of firms⁵. Hence, it is of importance to correctly reflect finite-size effects in risk estimations of such systems.

Fig. 3 shows the corresponding cascade size distributions for the star-shaped networks. Here, we have to distinguish the ED and the DD mechanisms of load distribution, which coincide only for complete networks. Fig. 3(a) shows the results for the ED mechanism. Again, we observe bi-modality in the distribution of cascade sizes that vanishes only for large threshold heterogeneity σ . The bi-modality has a clear interpretation with respect to the central node. In smaller cascades, the center does not fail, which explains the peak at low values of ρ_Δ . But if the center fails, it triggers a significant number of further failures of leaf nodes that explains the right peak in the distribution. This bi-modality is so fundamental that it does not even vanish in the limit $N \rightarrow \infty$, as we show in the materials and methods section.

The results for the DD mechanism presented in Fig. 3(b) show that, in contrast to the ED mechanism, the distributions of the final cascade size are unimodal, even for smaller σ . Again, this can be explained with the role of the central node. Although the center fails with high probability, with the DD mechanism, its failure distributes only a small amount of load equally to the (large number) of leaf nodes. Hence, this does not cause substantial additional failures. As in infinitely large systems²³, in the presence of hubs, failure cascades involving the DD mechanism expose the system to a smaller failure risk than cascades involving the ED mechanism - at the expense that hubs usually fail.

Eventually, we also discuss the third load distribution mechanism used in the fiber bundle model. Because of the additional load λ_0 , more load is distributed after a cascade started. As a result, in complete networks large cascades amplify to a full system break-down shown in the very right peak at $\rho_\Delta = 1$ in Fig. 4(a). This is an extreme case of the bi-modal distribution, where also small cascades can occur with high probability. As Fig. 4(b) shows, in star-like networks the range of probable cascade sizes is usually bigger. Dependent on the (μ, σ) values of the threshold distribution, bi-modal cascade-size distributions can still occur. The peaks of the cascade-size distribution can be explained again by the survival or the failure of the central

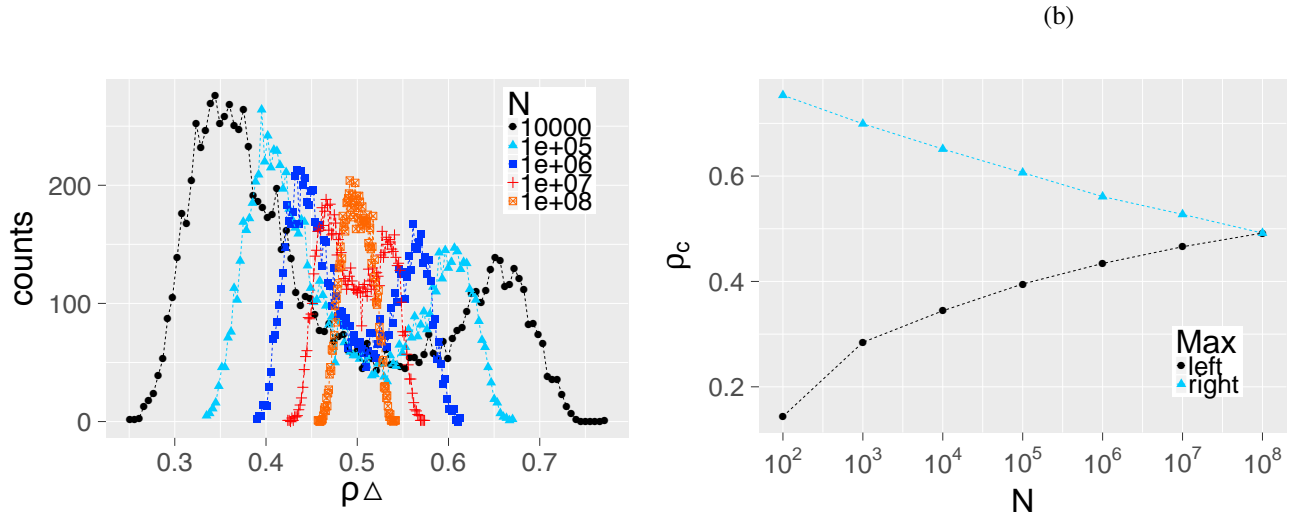


Figure 2. (a) Final cascade size distribution for complete networks of several sizes N obtained via 10^4 independent simulations of the ED mechanism, i.e. draws from $P(\rho_\Delta)$. The thresholds are independently normally distributed with mean $\mu = 0.5$ and standard deviation $\sigma = 0.4$ ($F \sim \mathcal{N}(0.5, 0.4^2)$). (b) Position of the two local maxima of the distributions in (a) with respect to the network size N .

node. In comparison with complete networks, these peaks are often close together. Further, the lower connectivity of the star-like networks reduces the risk of large cascades.

Discussion

Our driving question was whether systemic risk in finite systems can be reasonably estimated from approximations valid for infinite systems. We have to negate this question for two reasons. First, taking the final size of failure cascades as a risk measure, we could demonstrate, for the first time, that the distribution of this quantity changes fundamentally in finite systems. While a unimodal and narrow distribution, typical for infinite systems, would allow to use the average cascade size as a reasonable measure, we could show that for finite systems we often have a very broad and even bi-modal distribution. Hence, an estimated average does not quantify the real risk. Instead, small cascades but also extreme larger failure cascades become more likely. Even though a complete system break-down is unlikely as we show (which may be expected for infinite, but not for finite systems), large cascades in finite systems are much more frequent than expected so far. Secondly, we could also estimate the range of sizes for which such finite size effects play a role. We found that they are important for system sizes up to 10^7 , i.e. these systems are quite large, even though finite. This demonstrates that corrections are needed not just for small systems of a few to a few hundred nodes. Instead, they are already necessary for all real-world systems, up to a hundred million nodes.

In addition to these insights, we were able to derive novel explicit closed-form solutions to calculate the full distribution of the final cascade size, for two exemplary topologies. Our results are unique, as we derive exact solutions that apply to networks of arbitrary finite size N , a wide range of load distribution mechanisms, and arbitrary threshold distribution to characterize the robustness of the nodes.

In light of the combinatorial complexity of the problem, such results are rare and crucially depend on the high symmetry of the considered networks. Complete networks and stars are of particular interest, as they appear to be systemic risk minimizing or maximizing when only the average cascade size is considered^{20,21}. Intuitively, this is reasonable, as the complete graph represents the limit case of perfect risk diversification, where the risk is shared between all system components. The star network can be less prone to the amplification of failures in the course of a cascade, in particular when its central node is very robust. For that reason, central counter parties are often favored to reduce systemic risk. Yet, we have to reconsider our assessment of robust finite network topologies and take the probability of large system break-downs into account, as these are relevant for risk averse decision makers.

Many further real world networks, especially economic networks, show a core-periphery structure^{18,19} with densely connected core nodes and peripheral nodes loosely connected to the core. Thus, their topology can be understood as a superposition of the sample topologies discussed here, the complete network and the star network. Sparse topologies, with the star network as a paragon, occur frequently in real world networks³⁸ and play an important role in the spreading of failures³⁹. These systems usually have sizes that require to consider the full cascade size distribution for the estimation of systemic risk, as we propose with our solutions.

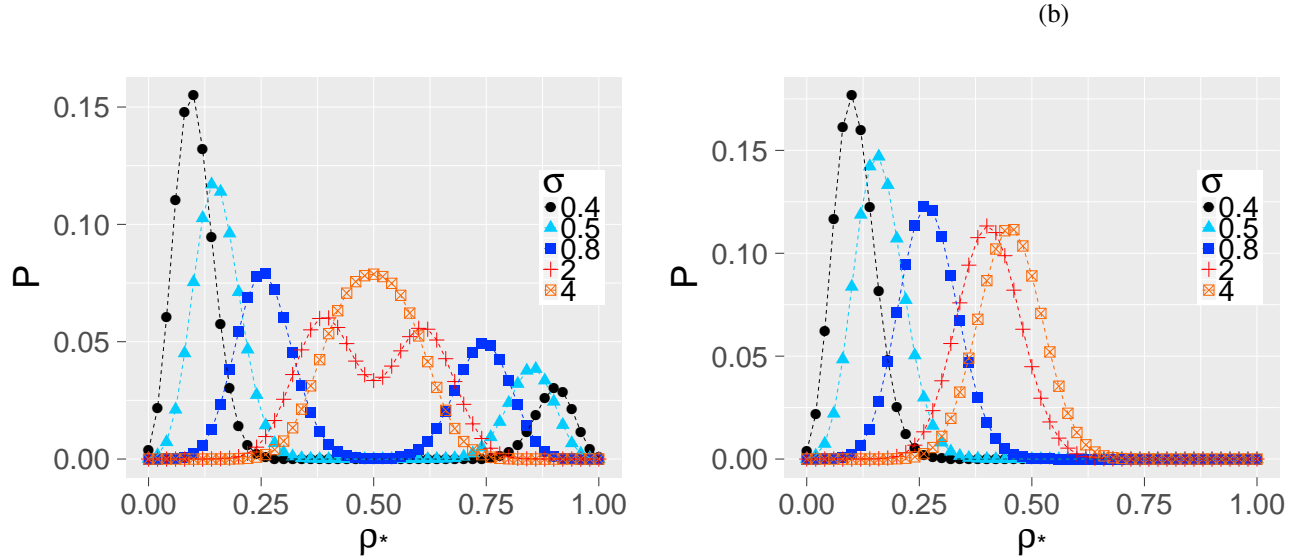


Figure 3. Probability distribution of the final fraction of failed nodes ρ_* on a star network with $N = 50$ nodes for the (a) ED mechanism and (b) DD mechanism. The thresholds are normally distributed with mean $\mu = 0.5$ and standard deviation σ (i.e. $F \sim \mathcal{N}(0.5, \sigma^2)$).

Beyond their elegance, such closed-form formulas have the advantage that they enable explicit optimization strategies for control of the involved parameters. For instance, we can derive derivatives with respect to those parameters that inform gradient based optimization methods. This allows to design robust systems, by controlling e.g. the broadness or the convergence of the cascade size distribution. It can be also applied in restoration strategies for failed systems⁴⁰.

Yet, we cannot expect to derive such analytic results for arbitrary network topologies. In this case, we have to rely on Monte Carlo simulations to estimate the broadness and shape of the cascade size distribution. We hope to inspire systemic risk analysts to base their insights and decisions on risk measures that can capture such broadness and multi-modality.

Methods

Derivation of cascade size distributions

In finite networks, the final cascade size ρ takes discrete values in the set $\{0, 1/N, 2/N, \dots, 1\}$. To determine the distributions $P(\rho_\Delta)$ and $P(\rho_*)$, we can calculate the probabilities for each event $\mathbb{P}(\rho_\Delta = k/N)$ and $\mathbb{P}(\rho_* = k/N)$, where k denotes the number of failed nodes at the end of the cascading process.

Complete networks

For complete networks, we can decompose the probability of k final failures into the product of the probability that $N - k$ out of N nodes do not fail and the probability p_k that exactly k nodes fail

$$P\left(\rho_\Delta = \frac{k}{N}\right) = \binom{N}{k} (1 - F(\lambda[k]))^{N-k} p_k. \quad (1)$$

Recall that $\lambda[k]$ denotes the load that a node with k failed neighbors carries so that $F(\lambda[k])$ is the probability that this load exceeds the node's threshold. With the help of the inclusion-exclusion principle⁴¹, p_k can be expressed as

$$p_k = \sum_{l=0}^{k-1} (-1)^{k+l+1} \binom{k}{l} F(\lambda[l])^{k-l} p_l. \quad (2)$$

The details of this derivation and proofs are provided in the supplementary information. Apparently, our derivations apply to an even larger class of models than threshold models. The failure probability of a node, here $F(\lambda[k])$, can be any function of the number of its failed neighbors.

Star networks

The central node with degree $d = N - 1$ has a prominent role, as it is the only node that can distribute accumulated load received by other nodes. We allow it to have a different threshold cdf F_c and load function λ_c than the remaining $N - 1$ leaf nodes, which

(b)

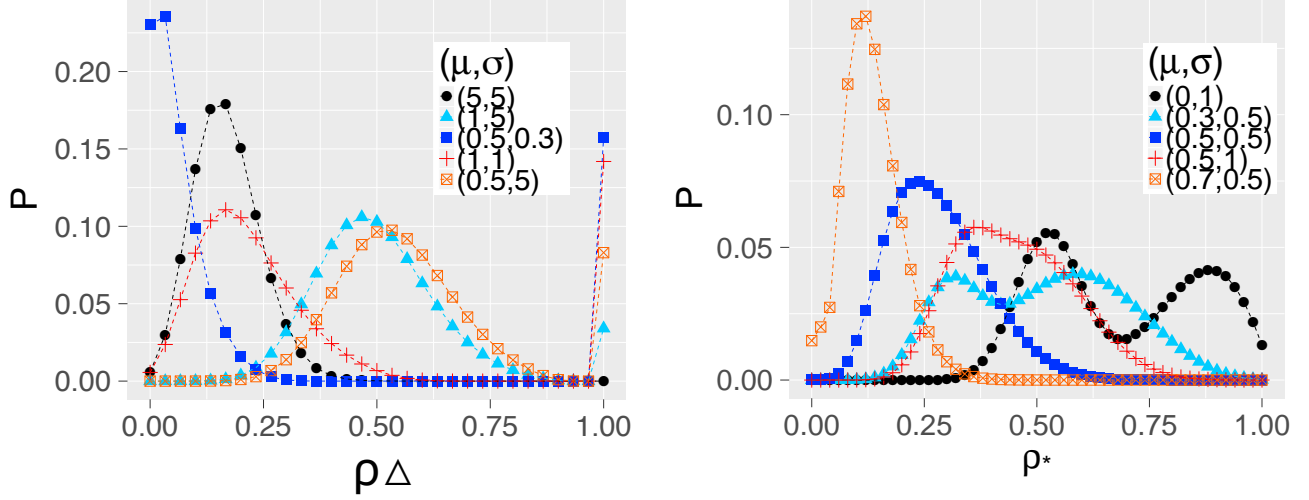


Figure 4. Fiber bundle model with initial load $\lambda_0 = 1$. Probability distribution of the final fraction of failed nodes with respect to normally distributed thresholds, i.e. $F \sim \mathcal{N}(\mu + \lambda_0, \sigma^2)$, on (a) complete network of size $N = 30$ and (b) star of size $N = 50$.

have degree 1, threshold cdf F and carry loads of the form λ_r . A leaf node can fail initially with probability $F(\lambda_r[0])$ (with 0 neighboring failures) or, if the center has failed before, it can carry the load $\lambda_r[j, l]$, which depends on two additional variables that are determined by the history of the center: j and l . j denotes the number of failed leaf nodes that have distributed load to the center, before the center has failed. l indicates the number of nodes among which the accumulated load of the center is shared (when the center fails). We simply add the probabilities for all possible cases and obtain:

$$\begin{aligned}
P\left(\rho_* = \frac{k}{N}\right) &= (1 - F_c(\lambda_c[k])) \binom{N-1}{k} F(\lambda_r[0])^k (1 - F(\lambda_r[0]))^{N-1-k} \\
&+ F_c(\lambda_c[0]) \binom{N-1}{k-1} \sum_{j=0}^{k-1} \binom{k-1}{j} F(\lambda_r[0])^j (F(\lambda_r[j, N-1-j]) - F(\lambda_r[0]))^{k-1-j} \\
&\times (1 - F(\lambda_r[j, N-1-j]))^{N-k} + \binom{N-1}{k-1} \sum_{j=1}^{k-1} \binom{k-1}{j} F(\lambda_r[0])^j (F_c(\lambda_c[j]) - F_c(\lambda_c[0])) \\
&\times (F(\lambda_r[j, N-1-j]) - F(\lambda_r[0]))^{k-1-j} (1 - F(\lambda_r[j, N-1-j]))^{N-k}.
\end{aligned} \tag{3}$$

The first summand considers the case when the center does not fail, while the second term adds the probability for the case when the center fails initially. Then, each of the other $k-1$ failures of leaves can either occur initially or because of a load distribution by the center. The size of this load might depend on the number of nodes l that fail initially together with the center, since these nodes cannot receive load after the failure of the center. The index j in the third term finally takes the events into account when j leaves have failed before the center.

Remarkably, this distribution converges for $N \rightarrow \infty$ to a discrete distribution with non-zero probability at three events, if $\lambda_c[k]$ and $\lambda_r[j, l]$ only depend on fractions of nodes ($\lambda_c[k] = \lambda_c[k/N]$, $\lambda_r[j, l] = \lambda_r[j/N, l/N]$). In an infinitely large star, exactly $F(\lambda_0)$ of the leaves fail initially. With probability $F_c(\lambda_0)$, the center node fails as well initially so that $\rho_* = F(\lambda_r[0, 1 - F(\lambda_0)])$ of the nodes fail in total. With probability $F_c(\lambda_c[F(\lambda_0)]) - F_c(\lambda_0)$, the center fails after the initial failures, which leads to $\rho_* = F(\lambda_r[F(\lambda_0), 1 - F(\lambda_0)])$. Or with probability $1 - F_c(\lambda_c[F(\lambda_0)])$, the center does not fail and we have $\rho_* = F(\lambda_0)$. In case of the ED model, two cases collapse to one because $\lambda_r[0, 1 - F(\lambda_0)] = \lambda_r[F(\lambda_0), 1 - F(\lambda_0)] = 1$. Thus, the observed bi-modality of the distributions in Fig. 3 (a) persists for increasing network size N .

Exemplary load distribution mechanisms

We have discussed three basic load distribution mechanisms corresponding to the exposure diversification (ED) model, the damage diversification (DD) model and a fiber bundle model. In the following, we specify how these models define the loads used in the derivations above.

ED load redistribution

In the complete graph, each load carries the load $\lambda[k] = k/(N-1)$. For the star network, $\lambda_c[k] = k/(N-1)$ holds for the center node and $\lambda_r[0] = 0$, $\lambda_r[j, l] = 1$ for the leaf nodes.

DD load redistribution

In the complete graph, each load carries the load $\lambda[k] = k/(N-1)$. On the star network, DD corresponds to the choice $\lambda_c[k] = k$, $\lambda_r[0] = 0$, $\lambda_r[j, l] = 1/(N-1)$.

Fiber bundle model load redistribution

On complete networks, the total load $\lambda_0 k$ has been redistributed so that each functional node carries $\lambda[k] = \lambda_0 + k\lambda_0/(N-k)$. For stars, we have $\lambda_c[k] = (k+1)\lambda_0$, $\lambda_r[0] = \lambda_0$, $\lambda_r[j, l] = (j+1)\lambda_0/l$.

For finite networks of arbitrary size N , we have obtained explicit closed forms of the full probability distribution of final cascade sizes for two prominent topologies. This is a major achievement that allows us to calculate the systemic risk in a highly efficient manner. More detailed information can be found in the supplementary information.

References

1. Tessone, C. J., Garas, A., Guerra, B. & Schweitzer, F. How Big Is Too Big? Critical Shocks for Systemic Failure Cascades. *J. Stat. Phys.* **151**, 765–783 (2013).
2. Gai, P. & Kapadia, S. Contagion in financial networks. *Proc. R. Soc. A* **466**, 2401–2423 (2010).
3. Haldane, A. G. & May, R. M. Systemic risk in banking ecosystems. *Nat.* **469**, 351–355 (2011).
4. Delli Gatti, D., Gallegati, M., Greenwald, B., Stiglitz, J. & Battiston, S. Liaisons dangereuses: Increasing connectivity, risk sharing and systemic risk. *J. economic dynamics control* **36**, 1121–1141 (2012).
5. Schweitzer, F. *et al.* Economic networks: the new challenges. *Sci.* **325**, 422 (2009).
6. Contreras, M. G. A. & Fagiolo, G. Propagation of economic shocks in input-output networks: A cross-country analysis. *Phys. Rev. E* **90**, 062812 (2014). URL <http://link.aps.org/doi/10.1103/PhysRevE.90.062812>. DOI 10.1103/PhysRevE.90.062812.
7. Klimek, P., Obersteiner, M. & Thurner, S. Systemic trade risk of critical resources. *Sci. Adv.* **1** (2015). URL <http://advances.sciencemag.org/content/1/10/e1500522>. DOI 10.1126/sciadv.1500522. <http://advances.sciencemag.org/content/1/10/e1500522.full.pdf>.
8. Burkholz, R. *Systemic Risk: From Generic Models to Food Trade Networks*. Ph.D. thesis, ETH Zurich (2016).
9. Battiston, S. *et al.* Complexity theory and financial regulation. *Sci.* **351**, 818–819 (2016). URL <http://science.sciencemag.org/content/351/6275/818>. DOI 10.1126/science.aad0299. <http://science.sciencemag.org/content/351/6275/818.full.pdf>.
10. Kondoh, M. Foraging adaptation and the relationship between food-web complexity and stability. *Sci.* **299**, 1388–1391 (2003). URL <http://science.sciencemag.org/content/299/5611/1388>. DOI 10.1126/science.1079154. <http://science.sciencemag.org/content/299/5611/1388.full.pdf>.
11. Briand, F., Cohen, J. E., Briand, F. & Newman, C. M. *Community food webs: data and theory* (Springer-Verlag Inc, Berlin; New York, 1990).
12. Sethna, J. P. *et al.* Hysteresis and hierarchies: Dynamics of disorder-driven first-order phase transformations. *Phys. Rev. Lett.* **70**, 3347–3350 (1993). URL <http://link.aps.org/doi/10.1103/PhysRevLett.70.3347>. DOI 10.1103/PhysRevLett.70.3347.
13. Pradhan, S., Hansen, A. & Chakrabarti, B. K. Failure processes in elastic fiber bundles. *Rev. Mod. Phys.* **82**, 499–555 (2010). URL <http://link.aps.org/doi/10.1103/RevModPhys.82.499>. DOI 10.1103/RevModPhys.82.499.
14. Pastor-Satorras, R., Castellano, C., Van Mieghem, P. & Vespignani, A. Epidemic processes in complex networks. *Rev. Mod. Phys.* **87**, 925–979 (2015). URL <http://link.aps.org/doi/10.1103/RevModPhys.87.925>. DOI 10.1103/RevModPhys.87.925.
15. Heesterbeek, H. *et al.* Modeling infectious disease dynamics in the complex landscape of global health. *Sci.* **347** (2015). URL <http://science.sciencemag.org/content/347/6227/aaa4339>. DOI 10.1126/science.aaa4339. <http://science.sciencemag.org/content/347/6227/aaa4339.full.pdf>.

16. Stark, H.-U., Tessone, C. J. & Schweitzer, F. Decelerating microdynamics can accelerate macrodynamics in the voter model. *Phys. Rev. Lett.* **101**, 018701 (2008). URL <http://link.aps.org/doi/10.1103/PhysRevLett.101.018701>. DOI 10.1103/PhysRevLett.101.018701.
17. Motter, A. E. & Lai, Y.-C. Cascade-based attacks on complex networks. *Phys. Rev. E* **66**, 065102 (2002). URL <http://link.aps.org/doi/10.1103/PhysRevE.66.065102>. DOI 10.1103/PhysRevE.66.065102.
18. Verma, T., Rusmann, F., Araujo, N. A. M., Nagler, J. & Herrmann, H. J. Emergence of core-peripheries in networks. *Nat. Commun.* **7**, 10441– (2016).
19. Glasserman, P. & Young, H. P. Contagion in financial networks. *J. Econ. Lit.* **54**, 779–831 (2016). URL <http://www.aeaweb.org/articles?id=10.1257/jel.20151228>. DOI 10.1257/jel.20151228.
20. Acemoglu, D., Ozdaglar, A. & Tahbaz-Salehi, A. Systemic risk and stability in financial networks. *Am. Econ. Rev.* **105**, 564–608 (2015). URL <http://www.aeaweb.org/articles?id=10.1257/aer.20130456>. DOI 10.1257/aer.20130456.
21. Acemoglu, D., Ozdaglar, A. & Tahbaz-Salehi, A. Networks, shocks, and systemic risk. Working Paper 20931, National Bureau of Economic Research (2015). URL <http://www.nber.org/papers/w20931>.
22. Watts, D. J. A Simple Model of global cascades on random networks. *Proc. Natl. Acad. Sci. U.S.A.* **99**, 5766 (2002).
23. Burkholz, R., Garas, A. & Schweitzer, F. How damage diversification can reduce systemic risk. *Phys. Rev. E* **93**, 042313 (2016). URL <http://link.aps.org/doi/10.1103/PhysRevE.93.042313>. DOI 10.1103/PhysRevE.93.042313.
24. Gleeson, J. P. & Cahalane, D. Seed size strongly affects cascades on random networks. *Phys. Rev. E* **75**, 1–4 (2007).
25. Payne, J., Dodds, P. & Eppstein, M. Information cascades on degree-correlated random networks. *Phys. Rev. E* **80**, 026125 (2009).
26. Hurd, T. & Gleeson, J. On Watts' cascade model with random link weights. *J. Complex Networks* **1**, 1–24 (2013).
27. Burkholz, R., Leduc, M. V., Garas, A. & Schweitzer, F. Systemic risk in multiplex networks with asymmetric coupling and threshold feedback. *Phys. D: Nonlinear Phenom.* **323–324**, 64 – 72 (2016). URL <http://www.sciencedirect.com/science/article/pii/S0167278915001943>. DOI <http://dx.doi.org/10.1016/j.physd.2015.10.004>. Non-linear Dynamics on Interconnected Networks.
28. Burkholz, R. & Schweitzer, F. Correlations between thresholds and degrees: An analytic approach to model attacks and failure cascades. *arXiv:1706.04451* (2017). [arXiv:1706.04451](https://arxiv.org/abs/1706.04451).
29. Watts, D. J. & Dodds, P. S. Influentials, networks, and public opinion formation. *J. consumer research* **34**, 441–458 (2007).
30. Battiston, S., Delli Gatti, D., Gallegati, M., Greenwald, B. C. N. & Stiglitz, J. E. Credit Default Cascades: When Does Risk Diversification Increase Stability? *J. Financial Stab.* **8**, 138–149 (2012).
31. Amini, H., Cont, R. & Minca, A. Resilience to contagion in financial networks. *Math. Finance* **26**, 329–365 (2016). URL <http://dx.doi.org/10.1111/mafi.12051>. DOI 10.1111/mafi.12051.
32. Battiston, S., Caldarelli, G., May, R. M., Roukny, T. & Stiglitz, J. E. The price of complexity in financial networks. *Proc. Natl. Acad. Sci.* **113**, 10031–10036 (2016). URL <http://www.pnas.org/content/113/36/10031.abstract>. DOI 10.1073/pnas.1521573113. <http://www.pnas.org/content/113/36/10031.full.pdf>.
33. Lorenz, J., Battiston, S. & Schweitzer, F. Systemic risk in a unifying framework for cascading processes on networks. *Eur. Phys. J. B* **71**, 441–460 (2009).
34. Daniels, H. E. The statistical theory of the strength of bundles of threads. i. *Proc. Royal Soc. Lond. Ser. A, Math. Phys. Sci. (1934-1990)* **183**, 405–435 (1945). URL <http://dx.doi.org/10.1098/rspa.1945.0011>.
35. Dobson, I., Carreras, B. A. & Newman, D. E. A loading-dependent model of probabilistic cascading failure. *Probab. Eng. Informational Sci.* **19**, 15–32 (2005). URL http://journals.cambridge.org/article_S0269964805050023. DOI 10.1017/S0269964805050023.
36. Moreno, Y., Gómez, J. B., Pacheco, A. F., Gomez, J. B. & Pacheco, A. F. Instability of scale-free networks under node-breaking avalanches. *Europhys. Lett. (EPL)* **58**, 630–636 (2002). URL <http://stacks.iop.org/0295-5075/58/i=4/a=630?key=crossref.806b054987dc0de948793132992f082d>. DOI 10.1209/epl/i2002-00442-2.
37. Burkholz, R. & Schweitzer, F. A framework for cascade size calculations on random networks. *arXiv preprint* (2017). URL <https://arxiv.org/abs/1701.06970>. 1701.06970.

38. Barabási, A.-L. & Albert, R. Emergence of scaling in random networks. *Sci.* **286**, 509–512 (1999). URL <http://science.sciencemag.org/content/286/5439/509>. DOI 10.1126/science.286.5439.509. <http://science.sciencemag.org/content/286/5439/509.full.pdf>.
39. Dorogovtsev, S. N. & Mendes, J. F. F. *Evolution of Networks: From Biological Nets to the Internet and WWW (Physics)* (Oxford University Press, Inc., New York, NY, USA, 2003).
40. Liu, C., Li, D., Zio, E. & Kang, R. A modeling framework for system restoration from cascading failures. *PLoS ONE* **9**, e112363 (2014). URL <http://dx.doi.org/10.1371/journal.pone.0112363>. DOI 10.1371/journal.pone.0112363.
41. Szpankowski, W. *Inclusion-Exclusion Principle*, 49–72 (John Wiley & Sons, Inc., 2001). URL <http://dx.doi.org/10.1002/9781118032770.ch3>.

Acknowledgements

RB acknowledges support by the ETH48 project. FS acknowledges support by the EU-FET project MULTIPLEX 317532.

Author contributions statement

R.B., F.S., H.J.H. designed the research. R.B. derived the formulas for the cascade size distribution, implemented them, and performed corresponding simulations. All authors discussed the results and wrote the paper.

Additional information

The authors declare no competing financial interests.