A simple model of fads and cascading failures

Duncan. J. Watts

Center for Nonlinear Earth Systems, 402 Fayerweather Hall, Mail Code 2571, Columbia University New York, NY 10027

Abstract

The origin of large but rare cascades that are triggered by small initial shocks is a problem that manifests itself in social and natural phenomena as diverse as cultural fads and innovations [1-3], social movements [4,5], and cascading failures in large infrastructure networks [6-8]. Here we present a possible explanation of such cascades in terms of a network of interacting agents whose decisions are determined by the actions of their neighbors according to a simple threshold rule. We identify conditions under which the network is susceptible to very rare, but very large cascades and explain why such cascades may be difficult to anticipate in practice.

How is it that small initial shocks can cascade to affect or disrupt large systems that have proven stable with respect to similar disturbances in the past? Why did a single, relatively inconspicuous, power line in Oregon trigger a massive cascading failure throughout the western US and Canada on 10 August 1996 [6], when similar failures in similar circumstances did not do so in the past? Why do some books, movies and albums emerge out of obscurity, and with small marketing budgets, to become popular hits [3], when many apparently indistinguishable efforts fail to rise above the noise? In this paper, we propose a possible explanation for this general phenomenon in terms of binary-state decision networks.

Our approach is motivated by considering a population of individuals each of whom must decide between two alternative actions, but whose decisions depend on those of their immediate acquaintances. In social and even economic systems, decision makers often pay attention to each other because they have limited knowledge of the problem, such as when deciding which movie to see or restaurant to visit, or limited ability to process the available information [2]. In other situations, the nature of the problem itself provides incentives for coordinated action, as is the case with social dilemmas [4] or competing technologies like personal computers or VCR's [1].

Although the detailed mechanisms involved in binary decision problems can vary widely across applications [1-5], the essence of many binary decision problems can be captured by the following threshold rule: An individual v adopts state 1 if at least a critical fraction ϕ_v of its k_v neighbors are in state 1, else it adopts state 0. In fact, when regarded more generally as a change of state, not just a decision, this rule is relevant to an even larger class of problems, including cascading failures in engineered networked systems such as power transmission networks [6,7] or the Internet [9-11]. Although motivated differently, the threshold rule is similar in nature to a family of models that are familiar to physicists, including random-field Ising models [12], models of spreading activation [13], self-organized criticality [7], percolation [14], and majority-vote cellular automata [15]. The model analyzed here, however, is distinguished from this literature in part because both the distribution of neighbors and also thresholds are heterogeneous and allowed to vary, and in part because the threshold rule is fractional. The latter feature implies that an individual's decision or change of state depends not only on its threshold and the states of its neighbors, but also upon its number of neighbors; hence it is the relationship between the two distributions--neighbors and thresholds--that is important for the dynamics.

Specifically, we consider a network (mathematically, a graph) of size *N*, in which each vertex (node) is connected to k other vertices with probability p_k and the average number of neighbors per vertex (degree) is denoted z. The population is initially all-off (state 0) and is perturbed at time $t = 0$ by a small fraction $\Phi_0 << 1$ of vertices that are switched on (state1). The population then evolves at successive time steps with all vertices updating their states according to the threshold rule above [16]. The thresholds ϕ_v are assigned at random according a distribution $f(\phi)$ defined on the unit interval and normalized such that $\int_0^1 f(\phi)d\phi = 1$. When $N \to \infty$ and each node is connected equally to all others $(z = N - 1)$, the dynamics of the model reduce to a one-dimensional map given by $\Phi_t = \int_0^{\Phi_{t-1}} f(\phi)d\phi$. From this map, it follows that for a smooth, unimodal distribution with $f(0) = f(1) = 0$, the system can have only three equilibrium states: all-on; all-off; and an unstable, intermediate equilibrium Φ_*. As long as $\Phi_0 < \Phi_*$, the initial perturbation will always die out, implying that arbitrarily small shocks can never cause large cascades in completely connected systems [17]. Many real networks, including those discussed above, are very sparsely connected in that the average degree is much less than the size of the system $(z << N)$ [18-20]. Here we address the relationship between the sparseness of a network and its vulnerability to cascades.

Our approach concentrates on two quantities: 1) the probability that a cascade will be triggered by a single node, or small seed of nodes; and 2) the expected size of a cascade once it is triggered. In the absence of any known geometry for the problem, a natural first choice of model for an interaction network is an undirected random graph [21], with N vertices and specified degree distribution p_k. In any sufficiently large random graph with $z << N$, no vertex neighboring the initial seed will be adjacent to more than one seed member. Hence the only way in which the seed can grow is if at least one of its immediate neighbors v has a threshold such that $\phi_v \le 1/k_v$, or equivalently has degree $k_v \leq K_v = \lfloor 1/\phi_v \rfloor$. We call vertices that are unstable in this one-step sense, *vulnerable*. Clearly, the more vulnerable vertices exist in the network, the more likely it is that an initial seed or shock will grow. But the extent of its growth, and hence the vulnerability of the network as a whole, depends largely on its global structure. Specifically, in an infinite network, cascades can be triggered with finite probability by a finite seed only when the network contains a percolating vulnerable cluster (consisting solely of vulnerable vertices).

The above condition, which we call the cascade condition, has the considerable advantage of reducing a complex dynamics problem to a static graph-theoretic problem that can be solved using the machinery of generating functions [22,23]. We proceed by defining the generating function of vertex degree:

$$G_0(x) = \sum_k \rho_k p_k x^k, \text{ where } \rho_k = \begin{cases} 1 & k = 0 \\ F(1/k) & k > 0 \end{cases}$$

and $F(\phi) = \int_0^{\phi} f(\varphi) d\varphi$. From this function, we can extract all the moments of the degree distribution of vulnerable vertices by evaluating its derivatives at $x = 1$; for example, the vulnerable fraction of the population is $P = G_0(1)$ and the average degree of vulnerable vertices is $G_0'(1)$. Next, we define the corresponding generating function $G_1(x)$ for vertices that we reach by following a randomly chosen edge. Because an edge arrives at such a vertex with probability proportional to its degree k, and we must discount the edge we came in along, the correctly normalized generating function is:

$$G_1(x) = \frac{\sum_k k \rho_k p_k x^{k-1}}{\sum_k k p_k} = \frac{G_0'(x)}{z}.$$

In addition, let $H_0(x)$ be the generating function for the vulnerable cluster size distribution; that is, $H_0(x)$ generates the probability that a randomly chosen vertex will belong to a vulnerable cluster of size n. Finally , let $H_1(x)$ generate the cluster size distribution for vertices that we arrive at by following a random edge. It follows that $H_1(x)$ satisfies the self-consistency equation $H_1(x) = 1 - G_1(1) + xG_1(H_1(x))$, from which $H_0(x)$ can be computed according to $H_0(x) = 1 - P + xG_0(H_1(x))$.

The quantity of primary importance to the problem of cascades is the average vulnerable cluster size $\langle n \rangle = H_0'(1)$ as this is the quantity that diverges at percolation. Substituting the expressions for $H_0(x)$ and $H_1(x)$ above, we find that $\langle n \rangle = P + (G_0'(1))^2 / (z - G_0''(1))$, which diverges at the critical point defined by:

$$G_0''(1) = \sum_k k(k-1)\rho_k p_k = z \qquad (1).$$

When $G_0''(1) < z$, all vulnerable clusters in the network are small, hence we do not expect to see any cascades. But when Equation 1 is satisfied, many vulnerable clusters become connected to form a percolating vulnerable cluster, from which it follows according to our claim above that cascades should be possible. We emphasize that, unlike standard percolation, ρ_k is a function both of the degree and threshold distributions; a feature that has important consequences for cascades.

We illustrate the cascade condition for the special case of a random graph, in which any pair of vertices is connected with probability $p = z/N$ [21]. Further, we assume initially that all vertices have the same threshold ϕ_c; that is, $f(\phi) = \delta(\phi - \phi_c)$. A characteristic of such uniform random graphs is that $p_k = e^{-z} z^k / k!$, the Poisson distribution, in which case our cascade condition (Equation 1) reduces to $zQ(K_c - 1, z) = 1$, where $K_c = \lfloor 1/\phi_c \rfloor$ and $Q(a, x)$ is the incomplete gamma function. Figure 1 expresses the cascade condition graphically as a boundary in the (ϕ_c, z) phase diagram (solid line) and compares it to the region (outlined by solid circles) in which cascades are observed in the full dynamical simulation. Because the simulated system is finite, the predicted and actual boundaries of the cascade window do not agree perfectly, but they are very similar; in particular, both display a lower and an upper boundary. That is, for each value of ϕ_c, the system has two critical points, as a function of the average degree z, at which the characteristic time scale of the dynamics diverges (see Figure 2a). The lower critical point is similar to the standard percolation transition for random graphs that occurs when the average connectivity exceeds $z = 1$ [21,22]. But the upper critical point is different: Here, the system has become so highly connected that most vertices are stable, and the percolating vulnerable cluster is ultimately lost.

To understand the nature of the upper transition, we solve exactly for the fractional size S_v of the vulnerable cluster inside the cascade window. From the definition of the generating functions, $S_v = 1 - H_0(1) = P - G_0(s)$, where s satisfies $s = 1 - G_1(1) + G_1(s)$. In the special case of a uniform random graph, with homogeneous thresholds, we obtain $S_v = Q(K_c + 1, z) - e^{z(s-1)} Q(K_c + 1, zs)$, in which s satisfies $s = 1 - Q(K_c, z) + e^{z(s-1)} Q(K_c, zs)$. We contrast this expression with that for size of the entire connected component of the graph, $S = 1 - e^{-zS}$ [22], which is equivalent to allowing $K_c \rightarrow \infty$. In Figure 2b we show the exact solutions for both S_v (dashed line) and S (solid line) for the case of $\phi_c = 0.2$ $(K_c = 5)$, along with the relevant order parameters of the dynamics: the probability and expected size of cascades. It is clear that while the probability of cascades (open circles) is approximated by S_v, the average size of the cascades that are actually triggered (solid circles) is given instead by S. The reason is that once a cascade has commenced, vertices can have multiple neighbors in the on state, and so even those vertices that were deemed stable with respect to the initial shock can be toppled, allowing the cascade to occupy the entire connected component of the graph.

The phase transition at the upper boundary of the cascade window thus exhibits a dual nature, depending on which order parameter one observes. The probability of a cascade increases continuously, in a second-order phase transition [24], as the critical point is approached from above, but the expected cascade size jumps discontinuously from zero to one in the manner of a first-order transition, once cascades become possible. In this region cascades, like fads [3], will occur only rarely and thus unpredictably, but when they do occur they will be extremely large . The unpredictability of cascades has a second, temporal, aspect: Instead of many potential innovations, only a few of which actually succeed, consider a single innovation that is introduced repeatedly until it cascades. Now there are two time scales to the problem: the introduction period T_i, given by the inverse of cascade frequency, and the adoption period T_a (from Figure 2a). From Figure 3 it is clear that near the critical points, the ratio T_i / T_a can vary rapidly over multiple orders of magnitude (for sufficiently large N), implying that even a potentially successful innovation could remain unnoticed for a long time before eventually cascading. Finally, when the connectivity of the network places it inside the cascade window but to the right of the peak in Figure 2b, then any increase in the connectivity has an ambiguous effect on the system's stability. On the one hand, cascades will become less frequent; but on the other hand, when they do occur, they will be larger. Hence the system becomes at once more robust, yet also more fragile; a feature thought to be endemic of complex, engineered systems [8].

These qualitative results are quite general within the class of random networks, applying to arbitrary distributions both of thresholds $f(\phi)$ and degree p_k . Variations in either distribution can affect the vulnerability of the system considerably, as is demonstrated in Figure 1. Outside the original cascade window (solid line) lie two windows corresponding to threshold distributions $f(\phi)$ that are normally distributed with mean ϕ_c and increasing standard deviation σ . In this case, increased heterogeneity of thresholds causes the system to be less stable. The dashed line inside the original window represents a homogeneous threshold distribution, but here p_k is distributed according to $p_k = Ck^{-\tau} e^{-k/\kappa} \quad (k > 0)$, where C,τ,κ are constants. This class of power-law random graphs has attracted much recent interest [9-11,19,20] as a model of many real networks, ranging from social networks to the Internet and the World Wide Web. The power law distribution has the effect that the mean connectivity z is dominated by the presence of a few highly connected nodes, while many nodes have only one edge. Because Equation 1 has the effect of excluding both these extremes from the vulnerable cluster, random graphs with power law degree distributions tend to be much less vulnerable to random shocks than uniform random graphs with the same z, a point observed elsewhere [9-11] with respect to the random deletion of nodes. Hence, unlike thresholds, increased heterogeneity of vertex degree appears to make the system more stable.

A significant theoretical challenge is to extend the results of this paper to include networks that exhibit local structure, such as clustering [18], which in general violates the assumption that vertices initially can have at most one neighbor in the on state. One possible extension is to assume that individuals are assigned to small groups, within which interactions are dense, and that the groups in turn interact randomly. This approximation has the effect of replacing z in the above analysis, with the density z_g of group interactions, leaving the model qualitatively unchanged, but making cascades generally more likely; a conclusion supported by simulations.

In conclusion, we have analyzed a simple model of networks of individuals making binary decisions as a function of the decisions of their neighbors. The results presented here are highly suggestive of phenomena observed in real-world examples of cultural fads and technological innovations: not only are the successes hard to separate apriori, from the failures, but some cascades happen almost instantaneously while some, like cellular pagers or the 1989 Leipzig parades [4], exhibit latency periods that are decades long. We hope that the introduction of this simple framework will stimulate theoretical and empirical efforts to analyze more realistic network models, and obtain comprehensive data on the frequency, size and time scales of cascades in real networked systems.

The author is grateful for the support of A. Lo, and acknowledges D. Callaway, M. Newman, and S. Strogatz for illuminating conversations.

1. W. B. Arthur and D. A. Lane. *Struct. Change and Econ. Dyn*. **4(1),** 81-103 (1993).

2. S. Bikhchandani, D. Hirshleifer, and I. Welch. *J. Pol. Econ*. **100(5),** 992-1026 (1992).

3. M. Gladwell. *The Tipping Point: How little things make can make a big difference*. (Little Brown, New York, 2000).

4. S. Lohmann. *World Politics* **47,** 42-101 (1994).

5. M. Granovetter. *Am. J. Soc.* **83(6)**, 1420-1443 (1978).

6. D. N. Kosterev, C. W. Taylor, and W. A. Mittelstadt. *IEEE Trans. on Power Systems*. **14(3),** 967-979 (1999).

7. M. L. Sachtjen, B. A. Carreras, and V. E. Lynch. *Phys. Rev. E.* **61(5),** 4877-4882 (2000).

8. J. M. Carlson and J. Doyle. *Phys. Rev. E*. **60(2),** 1412-1427 (1999).

9. R. Albert, H. Jeong, and A. L. Barabasi. *Nature* 406, 378-382 (2000).

10. R. Cohen, K. Diaz, D. ben-Avraham, and S. Havlin. http://xxx.lanl.gov/abs/cond-mat/0007048 (2000).

11. D. S. Callaway, M. E. J. Newman, S. H. Strogatz, and D. J. Watts. *Phys. Rev. Lett*. (submitted) Available at http://xxx.lanl.gov/abs/cond-mat/0007300 (2000).

12. J. P. Sethna et al. *Phys. Rev. Lett.* 70(21), 3347-3350 (1993).

13. J. Shrager, T. Hogg, and B. A. Huberman. *Science* 236, 1092-1094 (1987).

14. S. Solomon, G. Weisbuch, L. de Arcangelis, N. Jan, and D. Stauffer. *Physica A* 277, 239-247 (2000).

15. C. Moore. *J. Stat. Phys*. **88(3/4),** 795-805 (1997).

16. The analysis here assumes synchronous updating of agents, but the results are qualitatively unchanged if an asynchronous, random updating procedure is used instead.

17. Multimodal distributions $f(\phi)$ can exhibit stable equilibria that are intermediate between zero and one, but in this case, small shocks will always trigger cascades. In either case, the dynamics of the system is predictable.

18. D. J. Watts and S. H. Strogatz. *Nature* **393,** 440-442 (1998).

19. A. L. Barabasi and R. Albert. *Science* **286,** 509-512 (1999).

20. L. A. N. Amaral, A. Scala, M. Barthelemy, and H. E. Stanley. *Proc. Nat. Acad. Sci*. (in press). Also http://xxx.lanl.gov/abs/cond-mat/0001458 (2000).

21. B. Bollobas. *Random Graphs*. (Academic, London, 1985).

22. M. E. J. Newman, S. H. Strogatz, and D. J. Watts. http://xxx.lanl.gov/abs/cond-mat/0007235 (2000).

23. The method is exact for infinite N and in simulations is observed to work well when N is at least three orders of magnitude larger than the initial seed.

24. H. E. Stanley. *Introduction to Phase Transitions and Critical Phenomena* (Oxford University Press, Oxford, 1971).

FIGURES

Figure 1. Cascade windows for the threshold model. The solid line encloses the region of the (ϕ_c, z) plane in which the cascade condition (Equation 1 in text) is satisfied for a uniform random graph with a homogenous threshold distribution $f(\phi) = \delta(\phi - \phi_c)$. The solid circles outline the region in which global cascades occur for the same parameter settings in the full dynamical model for $N = 10,000$ (averaged over 100 random single node perturbations). The long dashes outline the equivalent cascade window for a random graph with a degree distribution that is a power law with exponent τ and exponential cut-off κ_0, as described in the text. The dash-dot line and short dashed line, represent cascade windows for uniform random graphs, but where the threshold distributions $g(\phi)$ are normally distributed with mean ϕ_c and standard deviation $\sigma = 0.05$ (dash-dot) and $\sigma = 0.1$ (short dashes) respectively.

Figure 2. Two cross sections of the solid cascade window from Figure 1, at $\phi_c = 0.2$. (a) The average time required for a cascade to terminate diverges at both the lower and upper boundaries of the cascade window. (b) The dashed line represents the fractional size of the largest vulnerable component, and the open circles represent the frequency of cascades that result from a single site being switched on at $t = 0$ in the full dynamical model (averaged over 1000 random perturbations for $N = 10,000$). The solid (upper) line is the average fractional size of the entire connected component of the random graph, and the solid circles correspond to the average size of global cascades in the dynamical model, when they are triggered. Hence rare but large cascades occur at the upper boundary.

Figure 3. Ratio of latency period T_l to adoption period T_a for a single innovation that is repeatedly introduced at regular intervals until it triggers a cascade. The ratio is small inside the cascade window, indicating rapid adoption, but diverges at the boundaries.

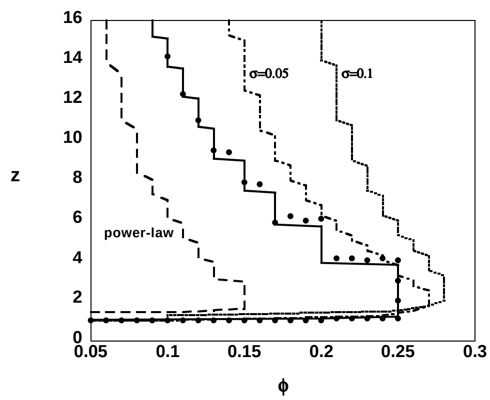

Figure 1

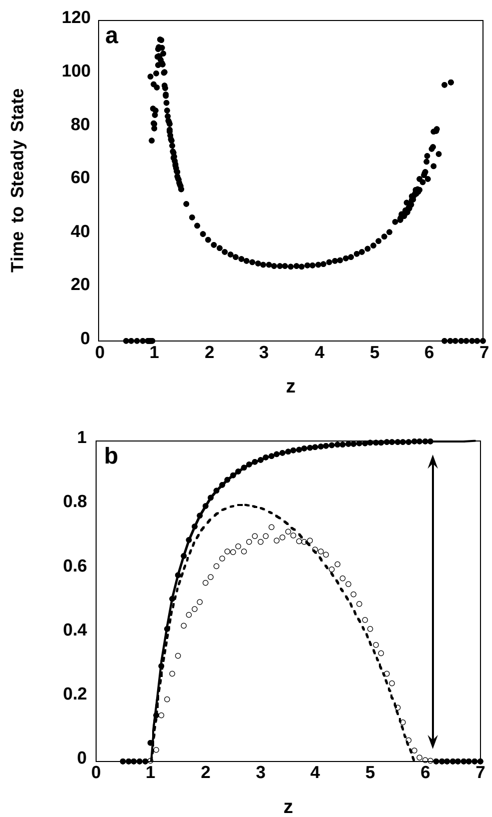
a
Time to Steady State
120
100
80
60
40
20
0
0 1 2 3 4 5 6 7
z
b
1
0.8
0.6
0.4
0.2
0
0 1 2 3 4 5 6 7
z

Figure 2

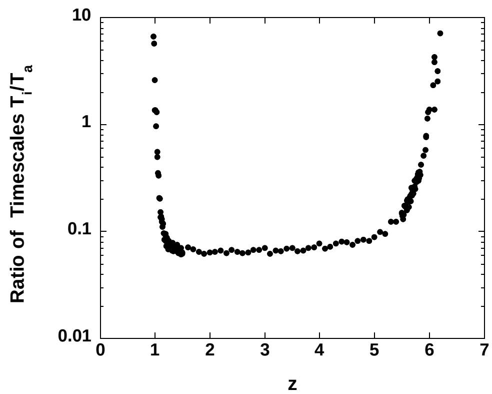

Figure 3