sociological science

Threshold Models of Collective Behavior II: The Predictability Paradox and Spontaneous Instigation

Michael W. Macy, Anna Evtushenko

Cornell University

Abstract: Collective behavior can be notoriously hard to predict. We revisited a possible explanation suggested by Granovetter's classic threshold model: collective behavior can unexpectedly fail, despite a group's strong interest in the outcome, because of the sensitivity of cascades to small random perturbations in group composition and the distribution of thresholds. Paradoxically, we found that a small amount of randomness in individual behavior can make collective behavior less sensitive to these perturbations and therefore more predictable. We also examined conditions in which collective behavior unexpectedly succeeds despite the group's weak interest in the outcome. In groups with an otherwise intractable start-up problem, individual randomness can lead to spontaneous instigation, making outcomes more sensitive to the strength of collective interests and therefore more predictable. These effects of chance behavior become much more pronounced as group size increases. Although randomness is often assumed to be a theoretically unimportant residual category, our findings point to the need to bring individual idiosyncrasy back into the study of collective behavior.

Keywords: collective behavior; thresholds; cascades; mobilization; unpredictability; randomness

Citation: Macy, Michael W., and Anna Evtushenko. 2020. "Threshold Models of Collective Behavior II: The Predictability Paradox and Spontaneous Instigation." Sociological Science 7: 628-648.

Received: October 18, 2020

Accepted: November 1, 2020

Published: December 8, 2020

Editor(s): Jesper Sørensen, Olav Sorenson

DOI: 10.15195/v7.a26

COLLECTIVE behavior can be notoriously hard to predict, even in groups whose members have a strong interest in the collective outcome. Proposed explanations include the free rider problem, path dependence, the Mathew effect, preference falsification, black swans, pluralistic ignorance, spiral of silence, unpopular norms, false enforcement, and dependence on a critical mass. The Mathew effect and path dependence in cascade dynamics make collective outcomes highly sensitive to the idiosyncratic choices of early movers, as demonstrated empirically by Salganik, Dodds, and Watts (2006) for the unpredictability of cultural "winners" and more recently by Macy et al. (2019) for the unpredictability of political alignments. Critical mass models show how a phase transition can make self-reinforcing collective behavior suddenly take off (Marwell and Oliver 2007). Preference falsification, pluralistic ignorance, and spiral of silence have been offered as explanations for the surprising and unanticipated revolution in Iran and the collapse of the former Soviet Union (Kuran 1991, 1995), whereas Arab Spring has been characterized as a black swan event (Ianchovichina 2018). Related models have also informed empirical studies of unanticipated events, including audience booing (Clayman 1993), crime sprees (LaFree 1999), unpopular norms (Willer, Kuwabara, and Macy 2009), revolutions (Kuran 1989, 1991, 1995; Kurzman 2004), school shootings (Larkin 2009), initial public offering (IPO) failures (Welch 1992), party defection (Hale and Colton 2017), the "Monday Demonstrations" in Leipzig, Germany (Braun 1995; Lohmann 1994), and the autism epidemic (Liu, King, and Bearman 2010). Most explanations of unpredictability incorporate the intuition that collective behavior is more likely to succeed when the shared interests of the members become strong enough to overcome social influence (Van de Rijt 2019), the temptation to free ride (Olson 1965), the fear of social disapproval (Bicchieri and Fukui 1999; Kuran 1991, 1995; Willer, Kuwabara, and Macy 2009), or the costs of contribution (Marwell and Oliver 2007).

In a theoretical article that has become a sociological classic with more than 5,000 citations (along with three follow-up articles: Granovetter and Soong [1983], [1986], [1988]), Granovetter (1978) shifted attention from the average *level* of interest in collective behavior to its *distribution*. The distribution matters because of threshold effects in complex contagion—when group members look around to see what others are doing before deciding to join in. The stronger an individual's interest in the outcome, the lower the number of others needed to trigger that member's participation. Although Granovetter originally assumed rational actors whose thresholds correspond to the costs and benefits of participation, the argument readily generalizes to emotional and normative cascades among individuals who vary in excitability (Granovetter and Soong 1983:167) or susceptibility to social pressure (Centola, Willer, and Macy 2005).

Granovetter showed how the equilibrium level of participation can be highly sensitive to small random variations in the distribution of thresholds, as might happen were "some individuals to enter or leave the situation" (1978:1429). He illustrated the dynamics with two groups, each with 100 members, one with a uniform threshold distribution $\{0,1,2,\ldots,99\}$ and an otherwise identical group except for the absence of anyone with threshold 1 $\{0,2,2,3,\ldots,99\}$ (1978:1425). The two groups have a nearly identical mean threshold, but the collective behavior of the two groups is radically different. In one, a cascade reaches all 100 members, whereas in the other, the cascade becomes trapped in a local equilibrium in which only the lone instigator participates.

More generally, for a cascade to reach all N members of a group, the critical condition is a cumulative threshold distribution such that, for every activation level L in the interval from zero to $N-1$, *the number of individuals with threshold less than or equal to L must be greater than L*. Cascades become trapped at the minimum of L for which this condition does not hold.

Depending on the location of perturbations in the distribution, collective behavior can be more likely in a group whose members have a weak interest in participation than in one where the interests are far stronger. To take an extreme example, imagine a group whose members' strong interest in collective behavior gives each a threshold of one, but they are missing an instigator with threshold zero. Despite their strong interest in acting together, they will wait in frustration for someone to throw the first rock. Now consider a group with uniformly distributed thresholds from 0 to 99, with a mean threshold of 49.5, indicating far weaker interest in the collective behavior. Nevertheless, a cascade will lead to mobilization of all 100 members, despite an average interest in participation that is far lower than that of the first group.

Granovetter also used a second model in which samples of size 100 are randomly drawn from an underlying population with thresholds that are normally distributed.

Sampling variation entails minor differences in the composition of the groups, and these small perturbations in the threshold distribution can cause cascades to stall at activation levels that vary randomly between groups whose average threshold is nearly identical, thereby making the outcomes highly unpredictable, even when the average preference of the underlying population is known.

In addition to sampling variations in the distribution of thresholds, Granovetter identified a second reason that cascades can unexpectedly succeed or fail: the distribution of social ties. “When threshold distributions have very stable equilibria it [the structure of social networks] may make very little difference; when these equilibria are unstable, however, the effects of social structure may overwhelm those of individual preferences” (1978:1430). The source of the instability is straightforward: those who decide to participate may not be connected to (and hence unable to influence) those whose thresholds happen to be strategically located in the threshold distribution. Simply put, the perturbations in the distribution can be created not only by sampling variation in group composition but also by the distribution of thresholds across clustered network locations such that the node next in line for activation at the global level only sees the local activation level, which may be insufficient to activate that node.

In sum, threshold models are “of particular value in understanding situations where the average level of preferences clearly runs strongly in favor of some action, but the action is not taken.” In these situations, “it is hazardous to infer individual dispositions from aggregate outcomes” (1978:1425).

The demonstration of this hazard is not unique to Granovetter’s model. Schelling (1971) famously used a threshold model to show how a population whose members will tolerate an ethnic out-group up to some critical proportion of their neighborhood can nevertheless end up with little or no ethnic diversity. Indeed, Granovetter “adapted the idea of behavioral thresholds from Schelling’s models of residential segregation,” and with “Schelling’s aim of predicting equilibrium outcomes from distributions of thresholds” (1978:1422, footnote 2). Both models provide compelling demonstrations of a foundational Durkheimian principle of sociological reasoning: collective behavior does not necessarily correspond to the aggregation of individual preferences among the members of a group.

Granovetter’s seminal theoretical research on threshold effects in cascade dynamics did not include empirical tests of the model. He used riot participation as an illustration and pointed to relevant applications, including “innovation adoption, rumor or disease spreading, strikes, voting, going to college, leaving social occasions, migration, or conformity” (1978:1425). In follow-up articles, Granovetter and Soong applied the threshold model to the adoption of new technology (1983), consumer behavior (1986), Chinese restaurants, and residential segregation (1988).

These studies, in turn, triggered a cascade of interest in social contagions and opinion dynamics that shows no sign of abating nearly five decades later. Threshold models have been incorporated in a large and growing theoretical and empirical literature on the spread of collective behavior, segregation of neighborhoods, and the mobilization of collective action (Acemoglu, Ozdaglar, and Yildiz 2011; Axelrod 1997; Braun 1995; Bruch and Mare 2006; Centola and Macy 2007; Chiang 2007; Chwe 1999; Fossett 2006, 2011; Hedström 1994; Macy 1991; Romero, Meeder, and

Kleinberg 2011; Valente 1996; Van de Rijt, Siegel, and Macy 2009; Watts 2002; Yin 1998; Zhang 2004).

From Deterministic to Stochastic Thresholds

Despite widespread interest in threshold effects, there is a paradoxical assumption in nearly all previous threshold models that has escaped attention: that human decision making is deterministic, a simplifying assumption inherited from Granovetter and Schelling. A deterministic threshold responds to a continuous distribution of inputs with a binary probability $p = \{0, 1\}$ based on a trigger point at which behavior is guaranteed to change. The deterministic specification assumes that individuals act with zero uncertainty; they know precisely the moment when their thresholds have been satisfied and will wait forever until that happens before they will act.

Deterministic thresholds are a curious assumption given Granovetter's emphasis on the importance of random sampling variation in the distribution of thresholds (1978:1428). The macrolevel distribution of thresholds is vulnerable to the luck of the draw in Granovetter's model, but the individual thresholds that comprise the distribution are deterministic—a stochastic distribution of deterministic thresholds. There is uncertainty in the composition of the group but no uncertainty in the decision making of the individuals that comprise it.

In a follow-up article, Granovetter and Soong (1983) defended the deterministic behavioral assumption as a useful simplification that facilitates analytical solutions in models that already build in stochasticity by allowing for heterogeneity between threshold distributions:

> . . . stochastic models of diffusion and epidemics are notoriously intractable . . . [hence] deterministic models ought not to be given up until absolutely necessary. Threshold models of collective behavior, then, may be viewed as one attempt to construct models of diffusion or contagion that explicitly builds in population heterogeneity thus avoiding the necessity of stochastic models (P. 166).

Modeling thresholds as stochastic would introduce a confound that obscures the effects of random perturbations in the distribution of thresholds.

Deterministic thresholds are not only analytically useful, they also appear to be a conservative assumption. Intuitively, the unpredictability of collective behavior can only increase with the introduction of noise at the individual level; the more idiosyncratic the decisions made by the individuals that comprise a group, the greater the aggregate unpredictability in the collective behavior. Adding random behavioral error on top of random sampling error can only amplify the unpredictability.[1]

Nevertheless, other studies have made important discoveries by introducing noise into deterministic models. For example, Shirado and Christakis (2017) showed empirically that stochastic behavior can improve a group's ability to solve coordination problems by jittering the group out of a local equilibrium. Also, Klemm et al. (2003) introduced noise into Axelrod's deterministic model of cultural segregation, showing that cultural boundaries eventually collapse if completely dissimilar agents retain a nonzero probability to interact. However, follow-up articles found that cultural boundaries are robust to noise in the presence of network-culture co-evolution (Centola et al. 2007) and social influence (Flache and Macy 2011).

Bruch and Mare (2006) introduced stochastic thresholds to Schelling's deterministic model of residential segregation, showing that local equilibrium in Schelling's classic model becomes unstable if thresholds for moving out of a changing neighborhood are continuous rather than binary. They modeled granularity as a logistic function that also introduced noise but emphasized granularity as the primary causal mechanism that allowed cascades to escape from a segregated equilibrium.

Although deterministic models of cultural and residential segregation have attracted critical attention, previous studies have not tested the robustness of deterministic assumptions in Granovetter's threshold model of collective behavior. Unpredictability poses a unique challenge for stochastic thresholds. Unlike cultural and residential boundaries that might collapse because of noise, unpredictability can be plausibly expected to increase when noise is introduced to a deterministic model.

We do not mean to suggest that deterministic assumptions should always be relaxed. On the contrary, theoretical models must rely on simplifying assumptions, and deterministic thresholds can be a useful simplification. Nor do we argue that collective behavior is predictable or that the unpredictability cannot be attributed in part to threshold effects. Indeed, we agree with Granovetter's central insight: that cascades can be trapped in a local equilibrium that need not correspond to aggregated individual preferences. That can happen whether thresholds are deterministic or stochastic.

Instead, our argument narrowly targets two new implications of threshold models that become apparent only with stochastic thresholds: the paradox of predictability and spontaneous instigation. To preview our results, when we introduced minor behavioral noise into Granovetter's threshold model, the unpredictability observed with deterministic thresholds largely vanished. The paradoxical effects of individual-level idiosyncrasy on collective unpredictability mirror the counterintuitive implications of threshold effects in the original Granovetter article. Just as Granovetter showed how collective behavior cannot always be predicted from individual preferences, we found that collective predictability need not correspond to the predictability of individual choices.

We also examined what happens just outside the critical region that Granovetter identified. Here outcomes are less sensitive to group composition because of a larger deficit of instigators than sampling error can overcome. When cascades can be expected to fail because of a seemingly intractable "start-up problem," introducing a very small amount of noise at the individual level can lead to an unexpected "spontaneous instigation" that may explain revolutionary surprises like Arab Spring and the collapse of the Soviet Union.

Models and Methods

In Granovetter's model, time proceeds in discrete intervals with synchronous updating in a population of 100 artificial agents. At time step t, each agent i

compares its activation threshold with the group's activation level for time $t-1$ and decides independently whether to activate, given the difference between the activation level L at $t-1$ and i's threshold T_i, where L and T are measured as the percentage of group members. Activation is governed by a step function with probability $p = 1$ if $L \geq T$ and $p = 0$ otherwise. No currently activated agents ever de-activate (an assumption that Granovetter and Soong relaxed in their 1983 follow-up article), hence the activation level at time t is the level at $t-1$ plus all additional members who chose to activate at time t. With infinite time, the outcome with stochastic thresholds becomes trivial if de-activation is precluded; we therefore permitted activated members to de-activate.

Following Granovetter, we used synchronous updating in which each member of the group decides whether to activate based on the group's level of activation at the end of the previous time step. (The results were nearly identical to those obtained with asynchronous updating in which each member decides whether to activate based on the level of activation at the moment of decision.) We made only two changes to the deterministic model: we allowed activated members to de-activate and we replaced the deterministic step function with a logistic function:

$$p_{i,t} = \frac{1}{1 + e^{m(T_i - L_{t-1})}}, \tag{1}$$

where $p_{i,t}$ is i's probability of activation at time t, T_i is the minimum percentage of the group that must be activated before i joins in, L_{t-1} is the percentage activated at the end of the previous time step, and m is a slope parameter ($m > 0$) that allows the function to approach a deterministic function as m goes to infinity. Conceptually, the logistic function generalizes Granovetter's discrete step function (with binary probabilities of 0 and 1) as the limiting case of a continuous sigmoidal distribution, as illustrated in Figure 1 with $m = 0.1$ and $m = 0.2$, the noise levels assumed for all subsequent figures; see Figure A2 in the online supplement for results with noise levels that range between nearly deterministic ($m = 1$) to the highest levels used in the current study ($m = 0.1$). Regardless of m, $p = 0.5$ when $L = T$. With $m = 0.2$, the probability of an error (i.e., deviation from deterministic behavior) is less than 0.001 across 64 percent of the distribution, with an average error rate of 0.03 across the entire distribution. This is sufficient error to test the robustness of deterministic thresholds but not so much as to make thresholds irrelevant.

With stochastic thresholds, errors can prevent activation even though $L_{t-1} \geq T_i$ or trigger activation even though $L_{t-1} < T_i$. A member i is highly unlikely to activate when $T_i << L$, but there remains a nonzero probability to activate in error; conversely, i is highly likely to activate when $L << T_i$, but there remains a nonzero probability for activation to fail in error as well (including the de-activation of those previously activated). Figure 1 shows that these two errors exactly offset each other, giving no aggregate change in i's probability of activation. The relative probability of activation errors and de-activation errors depends on the distribution of thresholds and the current level of activation. Even if every member of the group were activated, there remains a positive probability for an activated member to de-activate, which would then lower the group's activation level, making others more likely to de-activate, and so on. This avoids the trivial implication of stochastic

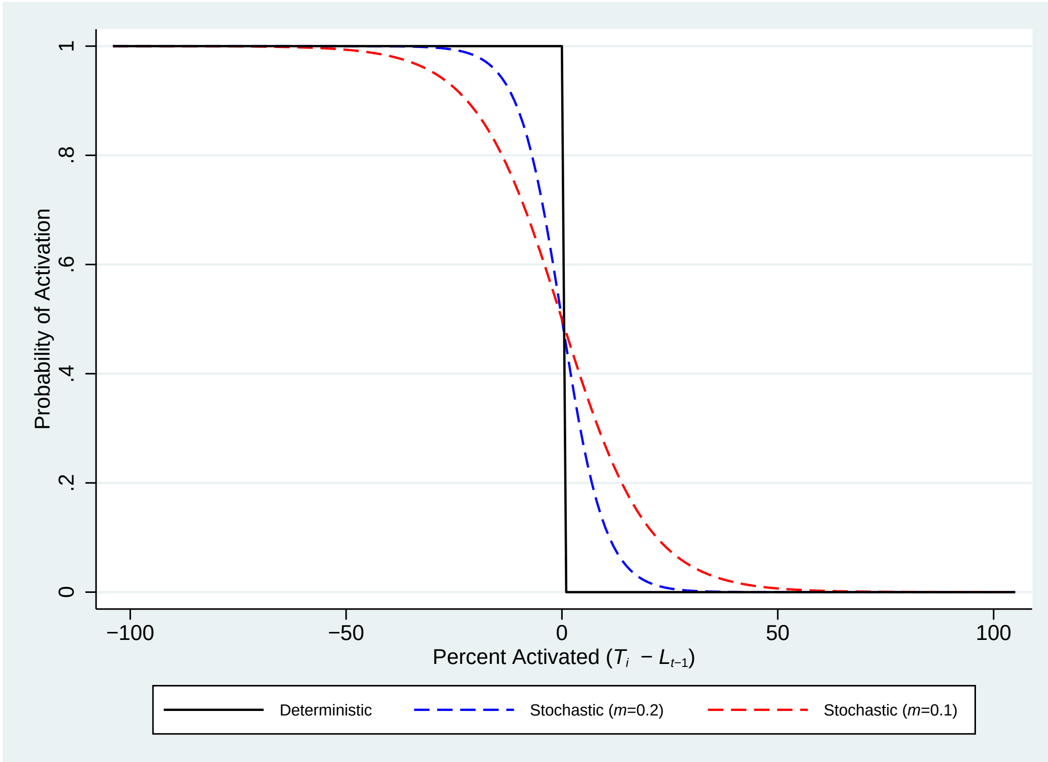

Figure 1: Deterministic and stochastic threshold functions. The solid line shows binary probabilities $p = \{0, 1\}$ as a deterministic function of the group's activation level L and individual threshold T, where L and T are the percentages of the group that are activated. The dashed lines show the continuous probability with stochastic thresholds with different slope parameters ($m = 0.1$ in red and $m = 0.2$ in blue). The functions are nearly identical if L is far from T except that activation is always possible even if $L >> T$ and activation can fail even if $T >> L$. These two errors offset each other, giving no increase in an individual's expected activation. As L approaches T, the probability of an error increases exponentially up to a maximum $p = 0.5$ when $L = T$.

thresholds that all members will eventually activate if time is infinite. Still, if the mean threshold is relatively low, de-activation errors are likely to quickly self-correct.

Uniformly Distributed Thresholds

We modeled microlevel unpredictability using both applications in Granovetter's original study: groups with uniform and normally distributed thresholds. Although empirically implausible, the uniform distribution affords a simple illustration of the equilibrium that is obtained when there is a randomly located perturbation in the distribution of thresholds. Following Granovetter, we modeled a perturbed uniform distribution by randomly removing one member of a group with N possible thresholds in the interval from 0 to $N - 1$, leaving one randomly located threshold without a host. For example, if the uniform distribution is perturbed at $T = 1$, the distribution becomes $T = \{0, 2, 2, 3, 4, 5, \ldots, N\}$. The perturbation at $T = 1$ is guaranteed to trap the cascade with incomplete participation when the threshold function is deterministic. The question is whether adding additional randomness in the form of activation errors at the individual level reduces the unpredictability of the outcomes in finite time.

Figure 2 reports two measures of cascade outcomes: the percentage activated (averaged across 100 samples that differ randomly in the location of the perturbed threshold) and the observed standard error[2] of the percentage activated (i.e., the standard deviation of the observed distribution of mean activations across 100 samples). The higher the standard error, the weaker the mean activation level as a predictor of the cascade outcome for a particular distribution of thresholds. The blue lines report the mean across 100 distributions, each with a randomly located perturbation. The red lines report the observed standard error of the mean as a measure of cascade unpredictability in a group with $N = 100$ members. The solid lines indicate deterministic thresholds and the dotted ($m = 0.1$) and dashed ($m = 0.2$) lines indicate increasing amounts of behavioral randomness.

When decision making is deterministic, the results confirm Granovetter's demonstration that random variation in the distribution of thresholds entails unpredictability of the equilibrium level of activation, with a standard error of 31 percent (in red) after 200 time steps. However, adding noise to the decision function does not make the outcomes even more unpredictable. As noise is added to individual behavior (going from the dashed lines to the dotted), activation errors overwhelm the effects of random perturbation in the distribution of thresholds. With deterministic thresholds, cascade outcomes are uniformly distributed around a mean that is expected to converge to 0.5 as the number of samples increases, given random perturbations. With stochastic thresholds, the outcomes converge toward a normal distribution with a similar mean level of activation, but the mean becomes an increasingly better predictor of particular cascade outcomes as more noise is introduced. This demonstrates the paradox that small microlevel perturbations can increase macrolevel predictability.

Normally Distributed Thresholds

Uniform threshold distributions are useful for illustrating cascade dynamics but are not empirically plausible. Granovetter (1978), therefore, also considered a model with normally distributed thresholds, as those may be

> . . . characteristic of populations where no strong tendencies of any kind exist to distort a distribution of preferences away from its regular variation about some central tendency. Yet, the results obtained are striking and counterintuitive, showing that paradoxical outcomes are not limited to special distributions such as the uniform (P. 1427).

Granovetter identified a critical region in which cascade outcomes are highly unpredictable. With too little heterogeneity, there can be a "start-up problem" (Marwell and Oliver 2007) due to the absence of instigators to trigger cascades. With too much heterogeneity, there can be a "follow-up problem" with too many "die

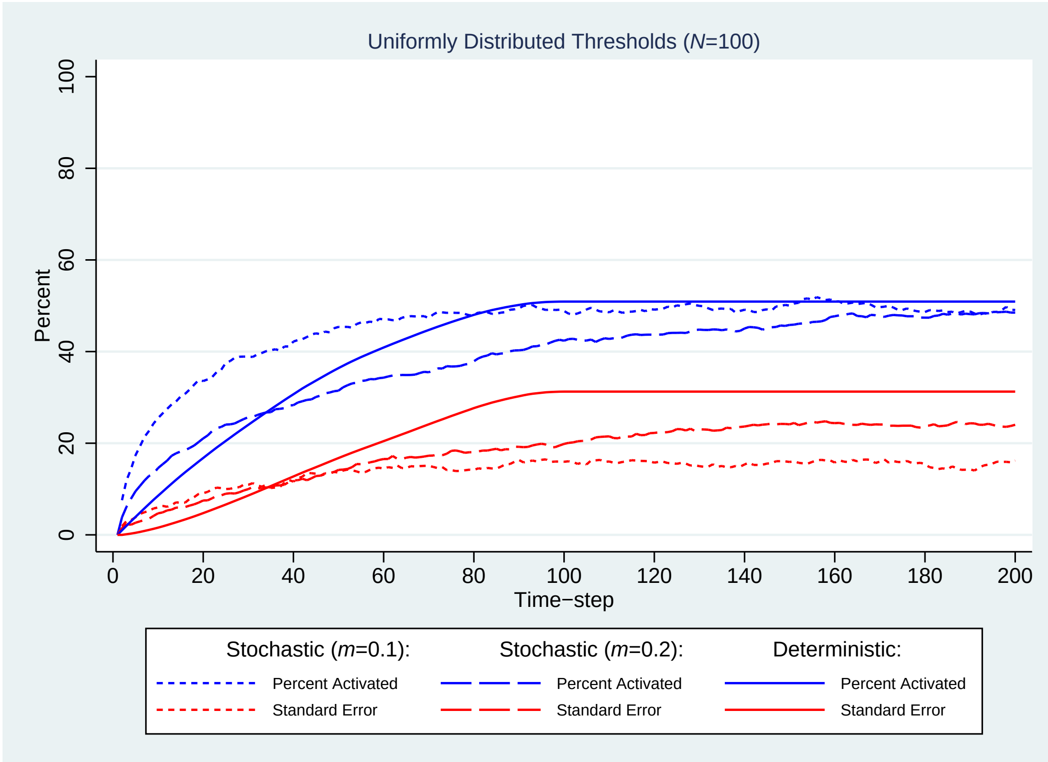

Figure 2: Unpredictability of activation level with a uniform threshold distribution and uniform random perturbations. The blue lines report the mean percentage activated in groups with $N = 100$ members. The red lines measure unpredictability as the observed standard error, that is, the standard deviation of the distribution of sample means over 100 observed distributions, each with a randomly located perturbation. Results show that as noise is added to individual behavior (going from the dashed lines with $m = 0.2$ to the dotted lines with $m = 0.1$), activation errors overwhelm the effects of random perturbation in the distribution of thresholds. With deterministic thresholds, cascade outcomes are uniformly distributed, reflecting the distribution of perturbations. With stochastic thresholds, the outcomes converge toward a normal distribution with a slight decline in the mean level of activation, but the mean becomes a modestly better predictor of particular cascades.

hards" who are resistant to influence and too few "lemmings" to sustain a cascade. In between, there is a narrow window of heterogeneity within which cascades are highly sensitive to random sampling error in the distribution of thresholds (see Figure A1 in the online supplement for Granovetter's original illustration of this window).

Figure 3 reports the outcomes of cascades within that critical window, with thresholds normally distributed around a mean $\mu = 25$ and $\sigma = 12.2$, the parameter combination that Granovetter identified as the critical point for a phase transition in the equilibrium activation level. The x axis reports time steps from 0 to 100 and the y axis reports the same outcome measures used in Figure 2: the mean percentage activated and the observed standard error of the mean across 100 samples of 100 members, drawn from an underlying population with normally distributed thresholds. With deterministic thresholds, 25 percent of the cascades reach nearly every

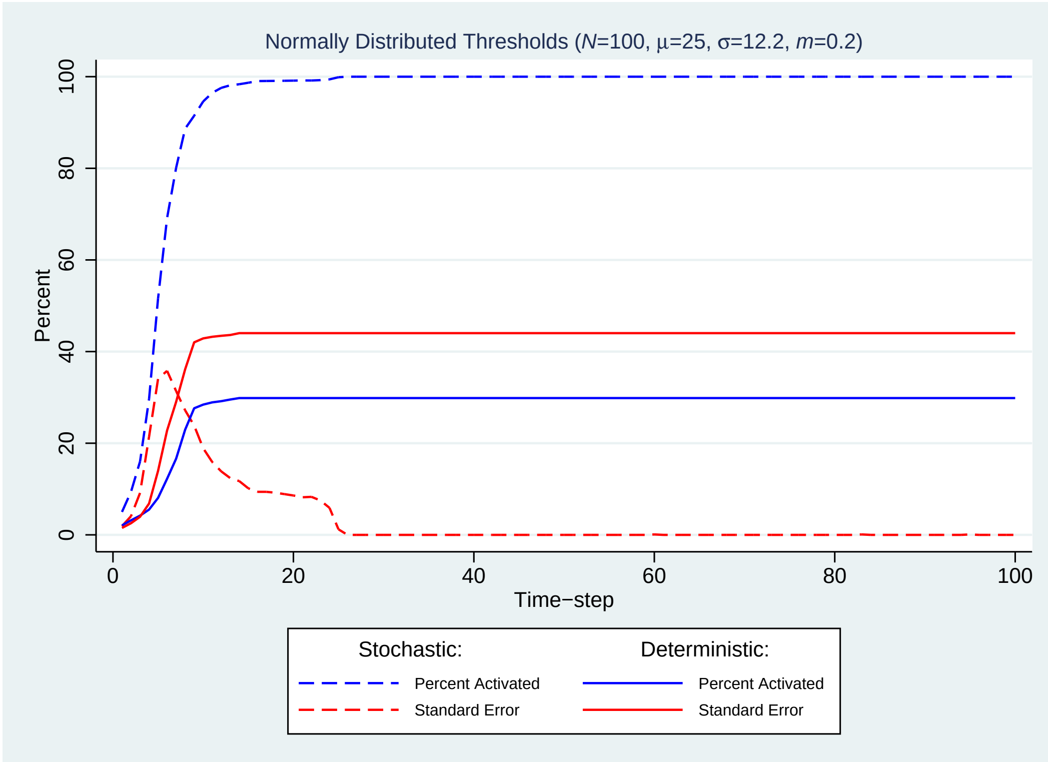

Figure 3: Unpredictability of activation level with a normal threshold distribution ($N = 100$, $\mu = 25$, $\sigma = 12.2$). Unpredictability is measured as the observed standard error of the mean percentage activated over 100 random samples from a normal distribution of thresholds. With deterministic thresholds, cascades reach an average of 25 percent activated with a standard error of 42 percent. With stochastic thresholds ($m = 0.2$), cascades are not trapped by perturbations in the distribution. The standard error first increases slightly as the percentage activated increases over time but then decreases to zero as the activation level increases to 100 percent. With greater noise (e.g., $m = 0.1$, not shown), the standard error reaches zero in four time steps, and with less noise (e.g., $m = 0.25$), the standard error remains slightly above zero.

member, whereas the others fail, with a standard error of 42 percent. With stochastic thresholds, however, outcomes converge over time to 100 percent activation. The standard error first increases slightly as the percentage activated increases from zero but then sharply decreases to near zero within 25 time steps. The results demonstrate the paradox that a small amount of individual randomness ($m = 0.2$) can make collective outcomes more predictable rather than less.

Is It Noise or Granularity?

The logistic function in Figure 1 differs from Granovetter's step function in two ways: not only is the function stochastic but it is also continuous rather than binary. Could the observed effects be falsely attributed to random activation errors when they are actually nothing more than the dependence of cascades on granularity (Bruch and Mare 2006)?

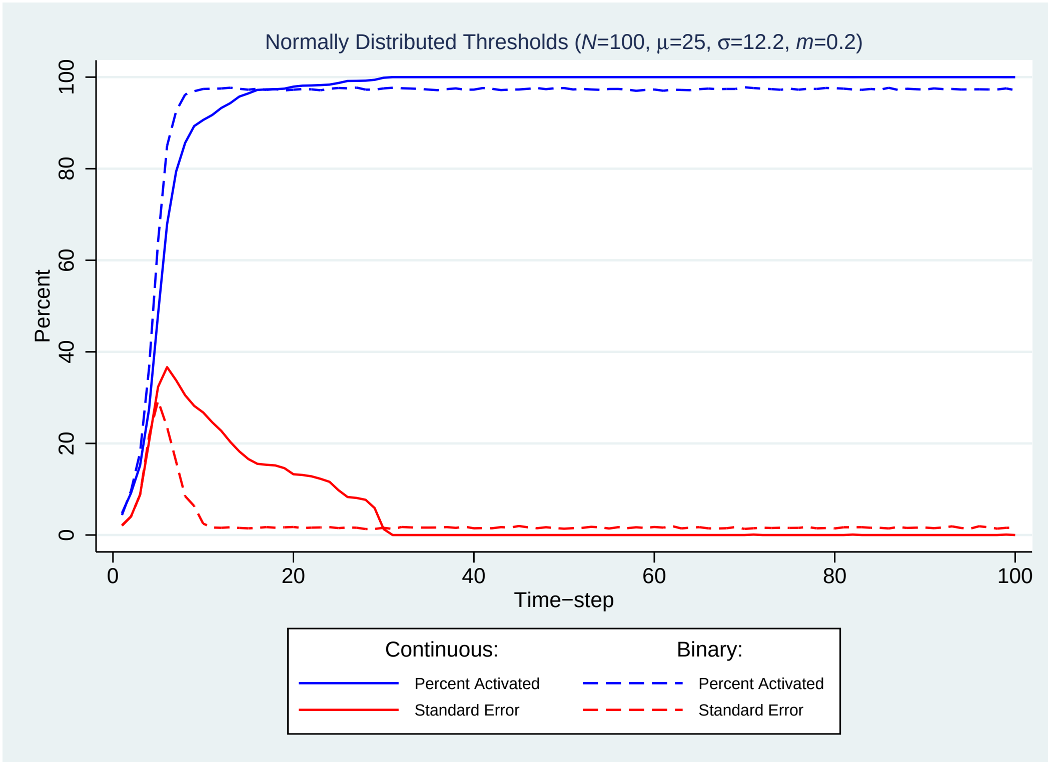

Figure 4: Unpredictability of activation level with continuous and binary stochastic thresholds. Figure 4 is identical to Figure 3 except that the deterministic step function with binary probabilities $p = \{0, 1\}$ is replaced by a stochastic step function with $p = \{0.03, 0.97\}$. Both functions are stochastic, with identical overall error rates of 0.03, but they differ in whether the probability distributions are continuous (solid line, with $m = 0.2$) or binary (dashed line). The results show that activation error, not granularity, allows cascades to overcome sampling perturbations. With greater noise (e.g., $m = 0.1$), the standard error reaches zero more quickly, and with less noise (e.g., $m = 0.25$), the standard error remains slightly above that for the binary function.

To find out, we tested what happens when Granovetter's step function is bounded away from the limits of probability by a small epsilon value ($\epsilon = 0.03$, equal to the average error rate with the logistic function) while retaining Granovetter's single trigger value at $T = L_{t-1}$. Figure 4 is identical to Figure 3 ($\mu = 25$, $\sigma = 12.2$, and $m = 0.2$) except that Granovetter's deterministic step function with binary probabilities $p = \{0, 1\}$ (shown in Figure 1) is replaced by a stochastic step function with $p = \{0.03, 0.97\}$. The two models in Figure 4 are both stochastic, but the probability distribution is continuous in one (solid line) and binary (dashed line) in the other. The results show that the explanatory mechanism in Figure 3 is activation error, not granularity. However, unlike the stochastic step function, the logistic function has the intuitive property that the probability of activation error approaches zero as the distance increases between T and L.

Strength in Numbers: The Effects of Group Size

Granovetter held group size constant at $N = 100$. As group size increases, sampling perturbations decline while stochastic decision making has more opportunities to accumulate random activation errors. Do stochastic thresholds then become the principal source of unpredictability?

Figure 5 shows the opposite. The parameters are identical to those in Figure 3 ($\mu = 25$, $\sigma = 12.2$, $m = 0.2$) except for N, which is an order of magnitude above and below 100. The results show that small groups ($N = 10$) are far less sensitive to activation error. After 100 time steps, the standard error with stochastic thresholds (21 percent) is only slightly below the standard error with deterministic thresholds (27 percent). In contrast, cascades quickly reach everyone in a group of 1,000 members with stochastic thresholds but not with thresholds that are deterministic (26 percent activation). Comparing the outcomes with $N = 1,000$ and $N = 100$ in Figure 3 shows that the time to reach all 100 members is reduced by more than half, from 25 time steps when $N = 100$ to 10 time steps when $N = 1,000$. However, we see a much more dramatic difference when comparing outcomes with $N = 10$ and $N = 100$. With stochastic thresholds, unpredictability remains very high in small groups but quickly disappears as group size increases.

The explanation is twofold. First, small groups are more susceptible to sampling errors that create larger perturbations in the distribution of thresholds. Second, small groups have fewer opportunities for the activation errors needed for cascades to bridge the wider ravines on a more rugged landscape. As a result, small groups can make outcomes highly unpredictable even with stochastic thresholds.

Figure 5 reveals a new mechanism supporting "the paradox of group size" (Marwell and Oliver 2007). Marwell and Oliver showed how large groups are more likely to contain a critical mass of instigators needed to jump-start collective action. We now see that large groups can also enjoy an additional advantage: they are more likely to self-generate a critical mass by leveraging behavioral idiosyncrasy.

This effect of group size has important implications for mobilizing strategies. Granovetter demonstrated the importance of the distribution of interest in the collective outcome, not just the mean, and even with stochastic thresholds, this conclusion still applies to small groups with short time horizons. Thus, leaders need to make sure everyone shows up, especially the instigators. With large groups, in contrast, leaders need to focus instead on mobilizing a high level of interest and enthusiasm and exercising patience while members are "milling about" deciding whether to participate. Don't worry as much about the distribution—if the level of interest is high enough, the need for instigators may eventually take care of itself.

Lastly, group size reveals the importance of the distinction between absolute and relative thresholds. Absolute thresholds refer to the number of activated members and model the assumption that nonactivated members of the group do not influence decision making. In contrast, relative thresholds refer to the percentage activated and correspond to the assumption that nonactivated members exert countervailing influence, such as an emergent norm that risks social disapproval from an orthodox majority, or the willingness of primary voters to support a candidate who is likely to win the general election, given the percentage of support. Granovetter models

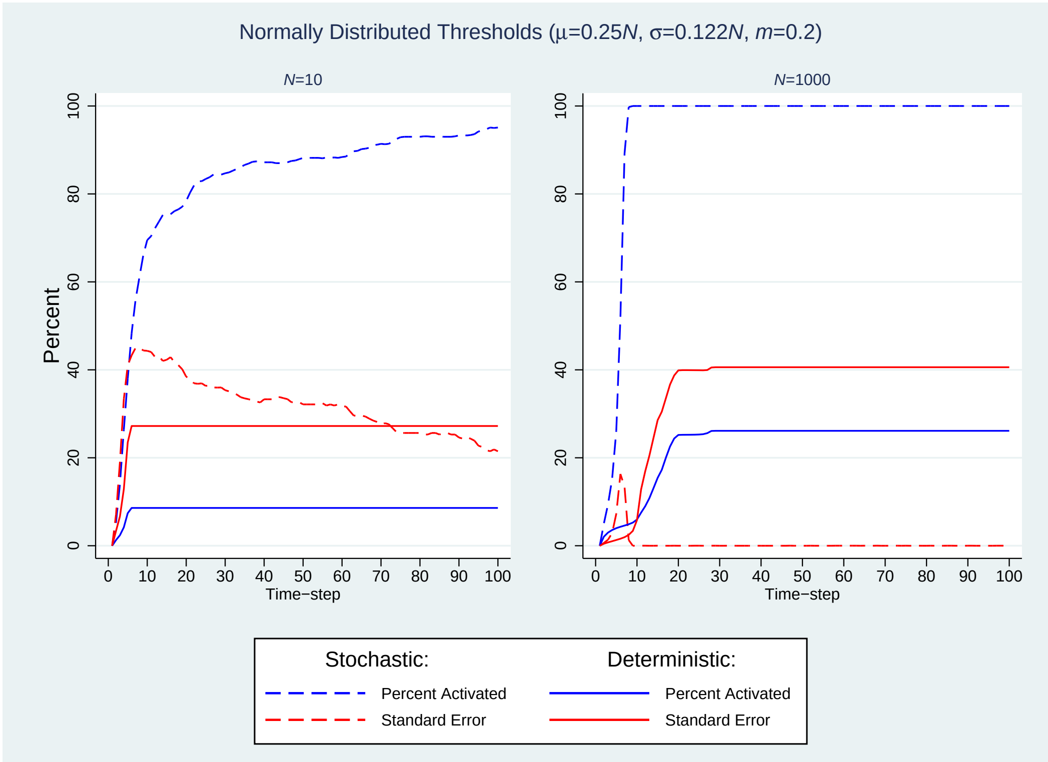

Figure 5: Unpredictability of activation level as N increases from 10 to 1,000. Group size is an order of magnitude above and below the size reported in Figure 3. Unpredictability is measured as the observed standard error of the mean activation level across 100 samples with $\mu = 0.25N$ and $\sigma = 0.122N$ randomly drawn from a normal threshold distribution. With deterministic thresholds, the equilibrium activation level increases from 8 percent with $N = 10$ to 26 percent with $N = 1,000$, with an increase in unpredictability from 28 percent with $N = 10$, to 41 percent with $N = 1,000$. In contrast, unpredictability *decreases* with group size if thresholds are stochastic ($m = 0.2$). With $N = 1,000$, the standard error decreases to zero within 10 time steps, whereas with $N = 10$, the standard error dropped below that for deterministic thresholds only after 75 time steps.

thresholds as the percentage of those who participate in a group of 100, making the absolute and relative proportions mathematically identical when group size is held constant (1978:1424, footnote 3). However, as group size increases, the same absolute threshold becomes a smaller percentage of the group. Had we held the mean threshold constant at $\mu = 25$ as we increased N to 1,000, the absence of unpredictability would be a trivial result.

Spontaneous Instigation of Collective Behavior

The preceding analyses have focused on the effects of activation errors in groups with a relatively strong interest in collective behavior ($\mu = 25$ and $\sigma = 12.2$) yet whose cascades unexpectedly fail because of perturbations in the distribution of deterministic thresholds. Figure 6 uses a heat map to extend the analysis to the full parameter space of possible distributions, including groups in which cascades might be expected to fail due to relatively weak interest in collective behavior. The heat map measures the mean and unpredictability of equilibrium outcomes after 100 time steps across 100 samples of size $N = 100$, randomly drawn from an underlying population with thresholds distributed around a mean ranging from $\mu = 0$ to $\mu = 99$ (from left to right within each cell) and standard deviation from $\sigma = 0$ to $\sigma = 99$ (from bottom to top in each cell). The small rectangular inset in each cell indicates $\mu = 25$ and $\sigma = 12$ as a reference point. Row 1 is the deterministic condition, included as a baseline for comparison with stochastic thresholds ($m = 0.2$) in row 2, with the difference in row 3. Row 4 reports the percentage activated and standard error attributable entirely to activation error by replacing the 100 sampled distributions with a single "prototypical" distribution (the average of the 100 ranked sample distributions, where the lowest prototypical threshold is the average of the 100 lowest sampled thresholds, and so on). Column A reports the percentage activated within 100 time steps, from indigo (0 percent activation) to red (100 percent). Green represents 25 percent activation, yellow 50 percent, and orange 75 percent. Column B reports unpredictability from indigo (0 percent) to red (50 percent), except for B3, which ranges from –50 to 50. (As in column A, green, yellow, and orange evenly divide the range, with yellow as the midpoint.) The heat map reveals three main findings:

1. The thin yellow boundary between the red and indigo regions in column A indicates a critical region of parameter combinations that are highly susceptible to random perturbations across samples. The critical region in which cascade outcomes are highly unpredictable with deterministic thresholds (B1) disappears within 100 time steps when thresholds are stochastic (B2). This is also evident in B3 as the dark region in which stochastic thresholds reduce the unpredictability observed with deterministic thresholds.

2. There is a second critical region (the yellow-red region in B3) of high unpredictability, even with stochastic thresholds. With deterministic thresholds, cascades in this region predictably fail because of insufficient numbers of low-threshold instigators. Stochastic thresholds can create spontaneous instigators who jump-start cascades, making collective behavior more unpredictable but also more likely to succeed. Removing sampling errors by fixing the threshold distribution (B4) reveals an inner core of unpredictability in this region that can be attributed entirely to activation errors, but most of the unpredictability is due to sampling error (B3).

3. Activation errors allow cascades to succeed (the bright yellow-red region in A3) that would have otherwise failed because of insufficient numbers of instigators. However, activation errors never cause a cascade to fail that would otherwise have succeeded. Simply put, when cascades unexpectedly succeed, the explanation is more likely to be spontaneous instigation, but when cascades unpredictably fail, the explanation is more likely to be group composition.

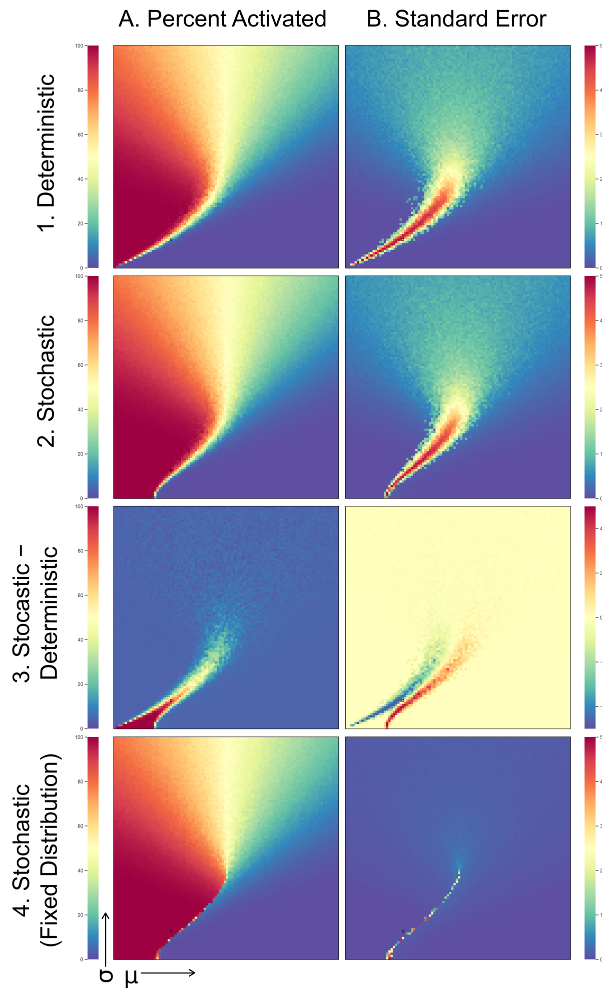

Figure 6: Unpredictability of activation level across an expanded parameter space. The heat map measures the level and unpredictability of cascade outcomes across 100 samples ($N = 100$) randomly drawn from an underlying population with thresholds normally distributed around a mean $\mu = \{0, \ldots, 100\}$ (from left to right on the x axis within each of the eight cells) and standard deviation $\sigma = \{0, \ldots, 100\}$ (from bottom to top on the y axis). Column A reports the percentage activated after 100 time steps, from indigo (0) to red (100). Column B reports the standard error from indigo (0) to red (50), except row 3 (–50 to 50). Row 3 reports the difference between deterministic thresholds (row 1) and stochastic thresholds (row 2, with $m = 0.2$). Row 4 reports the percentage activated and standard error attributable entirely to activation error when holding the threshold distribution constant across samples. The small rectangle in the lower left corner of each cell indicates $\mu = 25$ and $\sigma = 12.2$ as reference points for Figures 2 to 5. The heat map reveals three main findings: (1) The critical region in which the cascade outcomes are highly unpredictable with deterministic thresholds (B1) disappears within 100 time steps when thresholds are stochastic (B2 and B3). (2) There is a second critical region (the yellow-to-red region in B2) with high unpredictability even with stochastic thresholds. With deterministic thresholds, cascades in this region predictably fail because of insufficient numbers of low-threshold instigators, but activation errors can create spontaneous instigators who sometimes jump-start cascades, making collective behavior more unpredictable but also more likely to succeed. B4 shows that this region of cascade unpredictability is mainly due to sampling error, but there is an inner core of unpredictability that can be attributed entirely to activation errors. (3) Activation errors can make cascades more likely to succeed (the red region in A3), but they never cause a cascade to fail that might otherwise have succeeded.

Predictability and the Level of Collective Interest

Figure 6 measures unpredictability as the sensitivity of outcomes to chance variation in the distribution of thresholds. Unpredictability can also refer to the insensitivity of outcomes to the mean of the distribution. The mean threshold corresponds inversely to the group's level of interest in the collective behavior. Intuitively, we expect collective behavior to be more likely to succeed as that interest increases. If the likelihood of success does not increase, then the outcomes cannot be predicted by knowing the level of collective interest.

Figure 7 reports decreasing levels of activation after 100 time steps as the mean threshold increases. For comparability across two group sizes ($N = 100$ and $N = 1,000$), mean thresholds are expressed as the percentage of N and the standard deviation is normed by the mean ($\sigma = \mu/2$). The dotted vertical lines demarcate regions in which outcomes are relatively unpredictable from collective interests, either because thresholds are stochastic and collective interests are relatively strong (paradox of predictability) or thresholds are deterministic and collective interests are relatively weak (spontaneous instigation). With $N = 100$, there is also a midrange region in which outcomes are predictable from collective interests whether thresholds are deterministic or stochastic, but this region disappears as the size of the group increases.

The results show how the effects of individual activation errors on aggregate predictability varies with the level of collective interests. When collective interests are relatively strong (i.e., the mean threshold is low), stochastic thresholds reduce the sensitivity of cascades to chance variation in group composition, allowing cascades to consistently succeed regardless of the level of collective interest. In contrast, weak collective interests create start-up problems regardless of random perturbations in group composition, but idiosyncratic behavior ($m = 0.2$) can trigger successful cascades that are less sensitive to group composition and more sensitive to group interests.

Conclusion

Granovetter proposed an explanation for the unpredictability of collective behavior as a consequence of random sampling variation in the distribution of thresholds, but he modeled individual decision making as a deterministic function of the current level of activation relative to an individual's threshold. We tested the robustness of this explanation if the deterministic assumption is relaxed by the introduction of small amounts of behavioral noise. Results showed that stochastic thresholds made little difference in small groups, but as group size increased, the idiosyncrasies of behavior at the micro level can have a surprisingly stabilizing effect on the unpredictability that arises from sampling errors at the macro level—the predictability paradox.

Exploration of the full parameter space showed that this stabilizing effect is limited to a critical region in which deterministic thresholds are highly sensitive to perturbations in the distribution of thresholds. Outside this area, we discovered a second critical region in which activation errors can jump-start a cascade despite

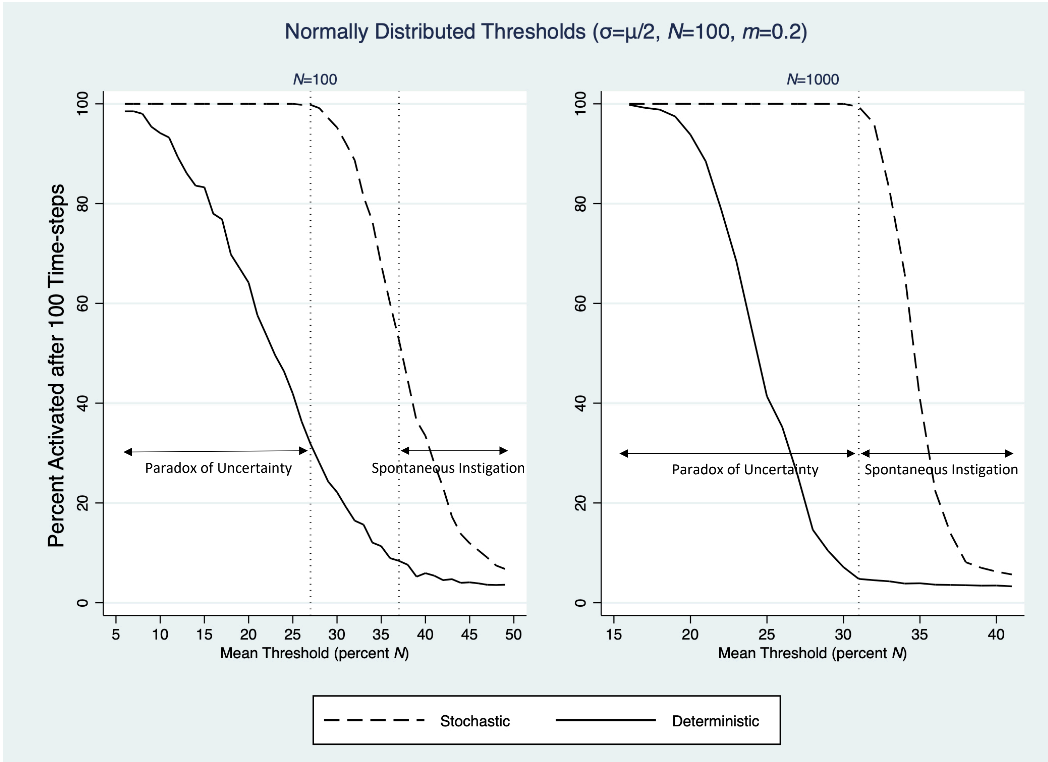

Figure 7: Predicting activation levels from collective interests as N increases from 100 to 1,000. Unpredictability is indicated by the change in percentage activated as the mean threshold increases (expressed as percent of N), with $\sigma = \mu/2$. The dotted vertical lines demarcate regions in which outcomes are comparatively insensitive to collective interests, either because thresholds are stochastic (paradox of predictability) or deterministic (spontaneous instigation). In the latter region, relatively weak collective interests create start-up problems regardless of random perturbations in group composition, but random errors in individual behavior ($m = 0.2$) can trigger successful cascades that are less sensitive to group composition and more sensitive to group interests.

the start-up problem caused by too few members with low thresholds, allowing cascades to succeed that would have otherwise failed. This spontaneous instigation makes cascades less sensitive to perturbations in the distribution of thresholds and more sensitive to the group's collective interests.

Conversely, de-activation errors do not cause cascades to fail that would otherwise succeed. This suggests the possibility that unexpected failures of mobilization may be more likely due to perturbations in group composition, whereas unexpected successes are more likely due to perturbations in individual behavior. Future empirical research is needed to disentangle compositional unpredictability (e.g., when instigators fail to show up for a protest) from behavioral idiosyncrasy (e.g., when someone steps forward who was not expected to lead). Although randomness is often assumed to be a theoretically unimportant residual category, our findings point to the need to bring individual randomness back into the study of collective behavior.

In addition to the theoretical implications, our study offers a methodological lesson. "By explaining paradoxical outcomes as the result of aggregation processes," Granovetter concludes, "threshold models take the 'strangeness' often associated with collective behavior out of the heads of actors and put it into the dynamics of situations" (1978:1442). Our analysis points to the need to bring at least a small bit of "strangeness" back into the heads of the actors. Humans are not automata, and the idiosyncrasies of human behavior can be assumed away only at the risk of misunderstanding "the dynamics of situations."

Notes

1 "Error" refers to random deviations from expected values and is not intended to imply "mistakes."

2 With observational data, the standard deviation of the distribution of sample means is typically unobserved and therefore must be estimated from the size and standard deviation of a single observed sample. Instead, we directly measure the distribution of the mean activations across 100 independent samples. We refer to the standard error over the distribution of mean activations to avoid confusion with the standard deviation of the distribution of individual activations in a given sample.

References

Acemoglu, Daron, Asuman Ozdaglar, and Ercan Yildiz. 2011. "Diffusion of Innovations in Social Networks." Pp. 2329–34 in *Proceedings of the 50th IEEE Conference on Decision and Control*. Piscataway, NJ: IEEE. https://doi.org/10.1109/CDC.2011.6160999.

Axelrod, Robert 1997. "The Dissemination of Culture: A Model with Local Convergence and Global Polarization." *Journal of Conflict Resolution* 41(2):203–26. https://doi.org/10.1177/0022002797041002001.

Bicchieri, Christina, and Yoshitaka Fukui. 1999. "The Great Illusion: Ignorance, Informational Cascades and the Persistence of Unpopular Norms." *Business Ethics Quarterly* 9(1):127–55. https://doi.org/10.2307/3857639.

Braun, Norman. 1995. "Individual Thresholds and Social Diffusion." *Rationality and Society* 7(2):167–82. https://doi.org/10.1177/1043463195007002005.

Bruch, Elizabeth E., and Robert D. Mare. 2006. "Neighborhood Choice and Neighborhood Change." *American Journal of Sociology* 112(3):667–709. https://doi.org/10.1086/507856.

Centola, Damon, Robb Willer, and Michael Macy. 2005. "The Emperor's Dilemma: a Computational Model of Self-Enforcing Norms." *American Journal of Sociology*, 110(4):1009–40. https://doi.org/10.1086/427321.

Centola, Damon, Juan Carlos González-Avella, Víctor M. Eguíluz, and Maxi San Miguel. 2007. "Homophily, Cultural Drift, and the Co-evolution of Cultural Groups." *Journal of Conflict Resolution* 51(6):905–29. https://doi.org/10.1177/0022002707307632.

Centola, Damon, and Michael Macy. 2007. "Complex Contagions and the Weakness of Long Ties." *American Journal of Sociology* 113(3):702–34. https://doi.org/10.1086/521848.

Chiang, Yen-Sheng. 2007. “Birds of Moderately Different Feathers: Bandwagon Dynamics and the Threshold Heterogeneity of Network Neighbors.” *Journal of Mathematical Sociology* 31(1):47–69. https://doi.org/10.1080/00222500601013536.

Chwe, Michael Suk-Young. 1999. “Structure and Strategy in Collective Action.” *American Journal of* Sociology 105(1):128–56. https://doi.org/10.1086/210269.

Clayman, Steven E. 1993. “Booing: The Anatomy of a Disaffiliative Response.” *American Sociological Review* 58(1):110–30. https://doi.org/10.2307/2096221.

Flache, Andreas, and Michael W. Macy. 2011. “Local Convergence and Global Diversity: from Interpersonal to Social Influence.” *Journal of Conflict Resolution* 55(6):970–95. https://doi.org/10.1177/0022002711414371.

Fossett, Mark. 2006. “Including Preference and Social Distance Dynamics in Multi-Factor Theories of Segregation.” *Journal of Mathematical Sociology* 30(3–4):289–98. https://doi.org/10.1080/00222500500544151.

Fossett, Mark. 2011. “Generative Models of Segregation: Investigating Model-Generated Patterns of Residential Segregation by Ethnicity and Socioeconomic Status.” *Journal of Mathematical Sociol*ogy 35(1–3):114–45. https://doi.org/10.1080/0022250X.2010.532367.

Granovetter, Mark S. 1973. “The Strength of Weak Ties.” *American Journal of Sociology* 78(6):1360–80. https://doi.org/10.1086/225469.

Granovetter, Mark. 1978. “Threshold Models of Collective Behavior.” *American Journal of Sociology* 83(6):1420–43. https://doi.org/10.1086/226707.

Granovetter, Mark, and Roland Soong. 1983. “Threshold Models of Diffusion and Collective Behavior.” *Journal of Mathematical Sociology* 9(3):165–79. https://doi.org/10.1080/0022250X.1983.9989941.

Granovetter, Mark, and Roland Soong. 1986. “Threshold Models of Interpersonal Effects in Consumer Demand.” *Journal of Economic Behavior and Organization* 7(1):83–99. https://doi.org/10.1016/0167-2681(86)90023-5.

Granovetter, Mark, and Roland Soong. 1988. “Threshold Models of Diversity: Chinese Restaurants, Residential Segregation, and the Spiral of Silence.” *Sociological Methodology* 18(1988):69–104. https://doi.org/10.2307/271045.

Hale, Henry E., and Timothy J. Colton. 2017. “Who Defects? Unpacking a Defection Cascade from Russia’s Dominant Party 2008–12.” *American Political Science Review* 111(2):322–37. https://doi.org/10.1017/S0003055416000642.

Hedström, Peter. 1994. “Contagious Collectivities: On the Spatial Diffusion of Swedish Trade Unions, 1890-1940.” *American Journal of Sociology* 99(5):1157–79. https://doi.org/10.1086/230408.

Ianchovichina, Elena. 2018. “Was the Arab Spring a Black Swan Event?” *Brookings Institution*, March 28. Retrieved November 13, 2020. https://www.brookings.edu/blog/future-development/2018/03/28/was-the-arab-spring-a-black-swan-event/.

Klemm, Konstantin, Victor M. Eguíluz, Raúl Toral, and Maxi San Miguel. 2003. “Global Culture: A Noise-Induced Transition in Finite Systems.” *Physical Review E* 67:045101. https://doi.org/10.1103/PhysRevE.67.045101.

Kuran, Timur. 1989. “Sparks and Prairie Fires: A Theory of Unanticipated Political Revolution.” *Public Choice* 61(1):41–74. https://doi.org/10.1007/BF00116762.

Kuran, Timur. 1991. “The East European Revolution of 1989: Is It Surprising That We Were Surprised? *American Economic Review* 81(2):121–25.

Kuran, Timur. 1995. "The Inevitability of Future Revolutionary Surprises." *American Journal of Sociology* 100(6):1528–51. https://doi.org/10.1086/230671.

Larkin, Ralph W. 2009. "The Columbine Legacy: Rampage Shootings as Political Acts." *American Behavioral Scientist* 52(9):1309–26. https://doi.org/10.1177/0002764209332548.

Liu, Ka-Yuet, Marissa King, and Peter S. Bearman. 2010. "Social Influence and the Autism Epidemic." *American Journal of Sociology* 115(5):1387–434. https://doi.org/10.1086/651448.

Lohmann, Susanne. 1994. "The Dynamics of Informational Cascades: The Monday Demonstrations in Leipzig, East Germany, 1989-91." *World Politics* 47(1):42–101. https://doi.org/10.2307/2950679.

Macy, Michael W. 1991. "Chains of Cooperation: Threshold Effects in Collective Action." *American Sociological Review* 56(6):730–47. https://doi.org/10.2307/2096252.

Macy, Michael, Sebastian Deri, Alexander Ruch, and Natalie Tong. 2019. "Opinion Cascades and the Unpredictability of Partisan Polarization." *Science Advances* 5(8):eaax0754. https://doi.org/10.1126/sciadv.aax0754.

Marwell, George, and Pam Oliver. 2007. *The Critical Mass in Collective Action*. Cambridge, UK: Cambridge University Press.

Olson, Mancur. 1965. *The Logic of Collective Action*. Cambridge, MA: Harvard University Press.

Romero, Daniel, Brendan Meeder, and Jon Kleinberg. 2011. "Differences in the Mechanics of Information Diffusion Across Topics: Idioms, Political Hashtags, and Complex Contagion on Twitter." Pp. 695–704 in *Proceedings of the 20th International World Wide Web Conference (WWW2011)*, New York, NY: Association for Computing Machinery. https://doi.org/10.1145/1963405.1963503.

Salganik, Mathew J., Peter Sheridan Dodds, and Duncan J. Watts. 2006. "Experimental Study of Inequality and Unpredictability in an Artificial Cultural Market." *Science* 311(5762):854–56. https://doi.org/10.1126/science.1121066.

Schelling, Thomas C. 1971. "Dynamic Models of Segregation." *Journal of Mathematical Sociology* 1(2):143–86. https://doi.org/10.1080/0022250X.1971.9989794.

Shirado, Hirokazu, and Nicholas A. Christakis. 2017. "Locally Noisy Autonomous Agents Improve Global Human Coordination in Network Experiments." *Nature* 545:370–4. https://doi.org/10.1038/nature22332.

Valente, Thomas W. 1996. "Social Network Thresholds in the Diffusion of Innovations." *Social Networks* 18(1):69–89. https://doi.org/10.1016/0378-8733(95)00256-1.

Van de Rijt, Arnout, David Siegel, and Michael Macy. 2009. "Neighborhood Chance and Neighborhood Change: A Comment on Bruch and Mare." *American Journal of Sociology* 114(4):1166–80. https://doi.org/10.1086/588795.

Van de Rijt, Arnout. 2019. "Self-Correcting Dynamics in Social Influence Processes." *American Journal of Sociology* 124(5):1468–95. https://doi.org/10.1086/702899.

Watts, Duncan J. 2002. "A Simple Model of Global Cascades on Random Networks." *Proceedings of the National Academy of Sciences* 99(9):5766–71. https://doi.org/10.1073/pnas.082090499.

Willer, Robb, Ko Kuwabara, and Michael W. Macy. 2009. "The False Enforcement of Unpopular Norms." *American Journal of Sociology* 115:451–90. https://doi.org/10.1086/599250.

Yin, Chien-Chung. 1998. "Equilibria of Collective Action in Different Distributions of Protest Thresholds." *Public Choice* 97(4):535–67. https://doi.org/10.1023/A:1004921725560.

Zhang, Junfu. 2004 “A Dynamic Model of Residential Segregation.” *Journal of Mathematical Sociology* 28:147–70. https://doi.org/10.1080/00222500490480202.

Acknowledgments: We thank Jon Kleinberg, Dana Warmsley, and Danielle Toupo for contributing ideas and technical suggestions. This research was supported by the National Science Foundation (SES 1756822, “Testing Unpredictability with Multiple Worlds”). Correspondence should be sent to Michael W. Macy, Department of Sociology, Cornell University, Ithaca, NY 14853.

Michael W. Macy: Department of Sociology and Department of Information Science, Cornell University. E-mail: mwm14@cornell.edu.
Anna Evtushenko: Department of Information Science, Cornell University. E-mail: ae392@cornell.edu.