NO. 4 · AUGUST 3, 2026 · THEHOLBROOKREPORT.COM

EXPLAINER

Thresholds: Why Two Nearly Identical Crowds Do Opposite Things

A walk through Granovetter's threshold model of collective behavior: what a threshold is, the arithmetic that turns a distribution of them into an outcome, the three mechanisms the paper's own author uses to dismantle its most famous result, and the fifty-year argument that followed about tipping points, complex contagion, and the 25 percent number.

ABOUT THIS REPORT
Written with Claude (Anthropic) and directed, fact-checked, and edited by Aaron Holbrook. Educational only — not financial, legal, or tax advice. Sources are listed at the end.

What this piece examines, and what it is: The subject is a formal model published in 1978 that has since become the standard apparatus for reasoning about riots, fads, bank runs, strikes, adoption curves, and revolutions. This report explains the model at full size, works its arithmetic, sets out the conditions under which its author says it is false, and traces the fifty years of testing and revision that followed. It is an educational explainer, not advice about how to start or stop anything. The report is AI-authored and directed by hand; every claim is labeled [FACT] (a sourced point of record), [CHARACTERIZATION] (a defensible label placed on those facts), or [PROJECTION] (a conditional "if X then Y," never a prediction).

Certain outcomes are routinely explained backward. A crowd riots, so the crowd must have been angry. A product spreads, so the product must have been good. A norm collapses, so opinion must have shifted. In each case the reasoning runs from the observed collective result to a claim about what the individuals in it wanted, and the strength of the result is taken as evidence of the strength of the wanting.

There is a class of situations where that inference is unreliable in a specific and demonstrable way: situations where each person's decision depends on how many other people have already decided. In those situations the outcome is governed by the *distribution* of individual tipping points across the group, and two groups whose distributions differ by a single member can arrive at opposite extremes. The average sentiment in the two groups is the same. The outcome is not.

The formal treatment is Mark Granovetter's "Threshold Models of Collective Behavior," published in the *American Journal of Sociology* in May 1978. [FACT] The sections below define the mechanism, work the arithmetic, and then spend as much space on where the model breaks as on where it holds, because the paper itself does.

The pills in this piece mark how load-bearing each point is to the model, not how dangerous anything is: CENTRAL MECHANISM for the parts the structure stands on, IMPORTANT CONTEXT for the supporting results, REFINEMENT for later corrections, and FAILURE CONDITION for the cases where the model does not apply.

The demonstration

One hundred people are milling around a square. Each has a threshold: the proportion of the group that would have to join a riot before that person joins it. Suppose those thresholds are perfectly uniform. One person has threshold 0, one has threshold 1, one has 2, and so on up to one with threshold 99. [FACT]

The person with threshold 0, the instigator, breaks a window. That satisfies the person with threshold 1, who joins. Now two people are rioting, which satisfies the person with threshold 2. The chain runs to the top. Equilibrium: 100 rioters.

Now alter exactly one person. Remove the individual with threshold 1 and replace them with a second individual with threshold 2. Every summary measure of the two crowds is essentially the same: same size, same mean threshold to within one unit, same range, same general disposition. The instigator breaks the window. Nobody has threshold 1. The cascade stops before it starts. Equilibrium: 1 rioter. [FACT]

---

1 The aggregation process, not the preferences, produced the difference CENTRAL MECHANISM

Granovetter's stated point is what an observer would conclude about each crowd. Of the first: "A crowd of radicals engaged in riotous behavior." Of the second: "A demented troublemaker broke a window while a group of solid citizens looked on." [FACT] Both descriptions are wrong about the people, and both are the natural reading of the events. The difference in outcome, in his words, "results only from the process of aggregation, and in particular from the gap in the frequency distribution in the second case." [FACT] [CHARACTERIZATION]

> **The operational consequence:** where behavior is interdependent, an aggregate outcome is not evidence about average sentiment. It is evidence about the joint distribution of sentiment and the sequence in which people were exposed to each other.

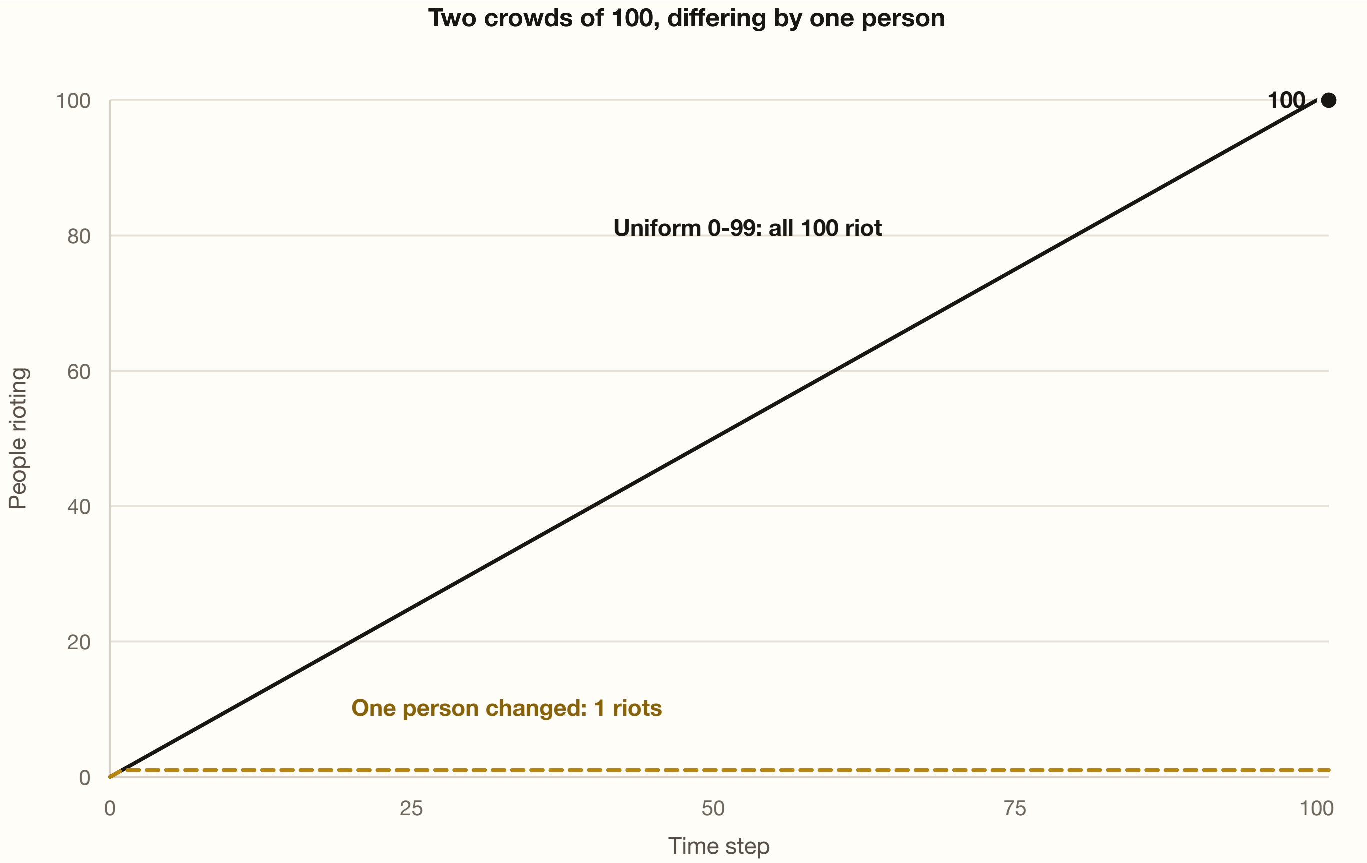

Forward recursion r(t+1) = F[r(t)] for both crowds. The only difference is that one person with threshold 1 is replaced by a second person with threshold 2. Computed from the distributions, not traced from the paper.

What a threshold is, and three things it is not

A threshold is the proportion of the group an individual must see act before acting, which is to say the point where perceived benefits first exceed perceived costs. [FACT] The concept is narrower than it looks, and most casual uses of the phrase "threshold model" get at least one of the following wrong.

Not an ideology score	Two people with identical thresholds may be politically opposite. A threshold nets many costs and benefits together, which is why Granovetter invokes the phrase "strange bedfellows." A radical has a low riot threshold because arrest is cheap to them and rioting is valuable; a person who hates being rained on has a low umbrella threshold for reasons that are not social at all. [FACT]
Not a norm or a belief	The sharpest distinction in the paper uses Matza's study of delinquent boys, most of whom did not think stealing cars was right and did not particularly want to do it, but who joined because refusing carried the cost of being called a sissy. [FACT] Granovetter: "The concept of threshold, then, is purely behavioral, connoting nothing about what the actor thinks is the 'right' thing to do." [FACT]
Not a trait the person carries	"An individual's riot threshold is not a number that he carries with him from one riot to another." [FACT] A threshold is a property of a person inside a particular configuration of costs and benefits, stable only as long as that configuration is.

Two endpoints define the scale. A threshold of 0 percent identifies an instigator, someone who acts with no prior participation at all. A threshold of 100 percent identifies someone who will not join under any circumstances. [FACT]

The arithmetic

Let `f(x)` be the frequency distribution of thresholds in the population and `F(x)` the cumulative distribution function, so `F(x)` is the proportion of people whose threshold is at or below `x`. Let `r(t)` be the proportion who have joined by time `t`.

If 60 percent have joined at time `t`, then by definition the people who will have joined at `t+1` are exactly those whose thresholds are at or below 60 percent. The entire process is one difference equation: [FACT]

```
r(t+1) = F[r(t)]
```

An equilibrium is a state where the next period equals the current one, so `r(t+1) = r(t)`, which is to say:

```
F(r) = r
```

Graphically, plot the cumulative distribution against the 45 degree line where `F(x) = x`. Starting from any participation level, move up to the curve, across to the diagonal, up to the curve again. The staircase converges, and the limit is the point where the cumulative distribution **first crosses the 45 degree line from above**. [FACT] In this version of the model, where nobody is permitted to quit once they have joined, oscillation is impossible and an equilibrium is always reached. [FACT]

This is the whole formal apparatus. Everything else in the literature is a question about what `F` looks like, who counts as “others,” and whether people obey the rule exactly.

Variance decides the outcome, and the mean does not

The model’s most useful result comes from holding the average constant and moving only the spread. Fix the mean threshold at 25 in a group of 100 and vary the standard deviation of a normal distribution of thresholds. [FACT]

STANDARD DEVIATION	EQUILIBRIUM PARTICIPANTS
Small	Near zero, rising gradually
Just below the critical value	About 6
Just above the critical value, roughly 12.2	**Nearly 100**
Large	Declines toward 50

Equilibrium participation against the spread of thresholds (mean fixed at 25, N = 100)

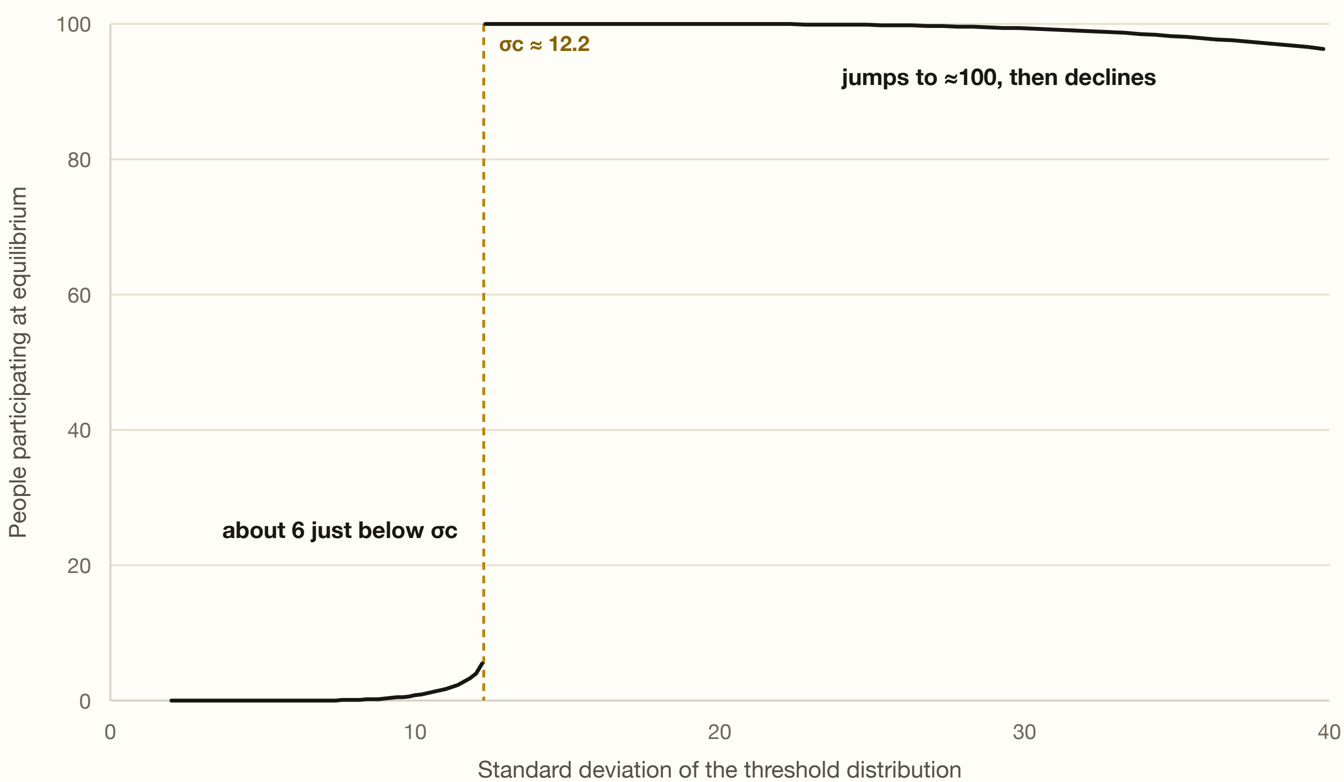

The mean never moves. Crossing a standard deviation of roughly 12.2 takes the outcome from about six participants to nearly the whole group. Computed by forward recursion on the normal CDF; the jump reproduces Granovetter's reported critical value.

---

2 A negligible change in variance flips the outcome discontinuously

CENTRAL MECHANISM

Crossing a standard deviation of approximately 12.2 moves the result from six participants to nearly the entire group. [FACT] The mean never moves. Granovetter is explicit that no sociological story attaches to this: "There is no obvious sociological way to explain why a slight perturbation of the normal distribution around the critical standard deviation should have a wholly discontinuous, striking qualitative effect. This perturbation might correspond to a minor fluctuation in the composition of a crowd ... a cause which would seem so insignificant in relation to its effect that causal attribution would never be made." [FACT] [CHARACTERIZATION]

> **Why this generalizes past riots:** the same structure governs any binary choice with interdependent payoffs. Whether to strike, adopt a technology, leave a boring lecture, migrate, believe a rumor, go to college, or hold a stock. In each case the interesting variable is the shape of the distribution near the current level of participation, not the average sentiment.

Three mechanisms that destroy the 100-person result

The uniform-distribution cascade is the part of the paper everyone quotes. It is also a knife-edge, and Granovetter spends the middle third of the paper taking it apart. [CHARACTERIZATION] The three mechanisms below matter more than the demonstration, because they are what determine whether a threshold distribution's equilibrium is worth believing.

3 Social structure: people do not weight everyone equally

IMPORTANT CONTEXT

The simple model assumes each person responds to the raw proportion of the group. Suppose instead that friends count double. Take an individual with a 50 percent threshold in a crowd where 48 have joined and 52 have not. Under anonymity they stay out. But if they know 20 people, 15 of whom have joined, they perceive (15 x 2) + (33 x 1) = 63 participants against (5 x 2) + (47 x 1) = 57 non-participants, or 63/120 = .525. Their threshold is passed and they join. [FACT]

Simulation over many randomly generated friendship structures found an asymmetry that matters more than the mechanism itself. The perturbed distribution's equilibrium of one participant is robust: where social structure changes it at all, the change rarely exceeds five to 10 people, and the modal result remains one. The uniform distribution's equilibrium of 100 is unstable against nearly any social structure at all. [FACT] The largest effects occur at a moderate density of acquaintance, where people know on average about a quarter of the group, a result Granovetter reports and declines to explain. [FACT]

> **The rule that follows:** "When threshold distributions have very stable equilibria it may make very little difference; when these equilibria are unstable, however, the effects of social structure may overwhelm those of individual preferences." [FACT] Network structure is not always the important variable. It is the important variable precisely when the preference-based prediction is fragile.

4 Sampling: the same population produces different crowds

IMPORTANT CONTEXT

Suppose a city's population carries the uniform threshold distribution, 1 percent of people at each value from 0 to 99, and that crowds of 100 form by random draw from it. The deterministic answer of 100 participants is now only one possible outcome among many. The probability that a drawn crowd contains no instigator is (1 - .01)^100 = .37, which yields zero participants. The probability of exactly one instigator and no one-percenter is .14, which yields one participant. Over half the time, .51, the result is zero or one. [FACT]

This converts a modeling caveat into an explanation. Two cities with identical underlying threshold distributions can produce a riot and a quiet evening for reasons internal to neither. Granovetter uses it to reinterpret Spilerman's finding that the only robust predictor of 1960s racial disorders was the absolute size of a city's Black population. [FACT] If each gathering is an independent trial with a fixed probability of tipping, more gatherings produce more riots with no difference in disposition required: at a per-incident probability of .10, a city with one incident has a 90 percent chance of no large riot, while a city with 10 incidents faces (.90)^10 = .35. [FACT] [CHARACTERIZATION]

5 Spatial dispersion, and the reverse bandwagon

IMPORTANT CONTEXT

The simple model assumes complete connectedness, meaning every person observes every other. Real events are not like this. Stark and colleagues found that the 1965 Watts riot was not a single incident but more than 1,850 separate cases of riot action over five days across a wide area. [FACT] Modeling clusters with limited movement between them produces a result that runs against intuition: too much movement out of a cluster that has already tipped deactivates some participants, whose departure deactivates others, and the group settles at a new and possibly much lower equilibrium. Granovetter calls this a reverse bandwagon. [FACT] For some distributions, therefore, small movements between clusters have larger effects than big ones. [FACT]

The conditions under which the model is false

Granovetter states his own falsification conditions directly, which is unusual enough to be worth reproducing.

6 Where behavior is not contingent, the model has nothing to say

FAILURE CONDITION

The model "may be said to have been, in effect, falsified" where all or most thresholds sit at 0 or 100 percent, which is to say where people's behavior is not in fact contingent on others'. [FACT] The example is Weber's crowd of pedestrians opening umbrellas at the start of a shower: not action oriented to each other, but many people reacting the same way to the same rain. [FACT] Granovetter amends even this, noting that some people surely do wait to see whether others open theirs, but the principle stands. Assigning someone a threshold is not a claim that social influence dominates their behavior; it is only the claim that a proportion exists at which they would switch. [FACT]

Three distinct measurement problems are set out separately, and they behave differently: [FACT]

Systematic misperception	If everyone misjudges participation by the same amount, the model is unchanged. If radicals overestimate and conservatives underestimate, the model is adjustable. If misperception is random, the stability of the underlying equilibrium becomes the critical question.
Pluralistic ignorance	People who have decided may not make it public, so many thresholds can be passed with nobody aware of it and nobody acting. Granovetter notes this is where the model earns its keep for policy, because it names the specific obstacle to full adoption rather than attributing the shortfall to resistance.
Granularity	Nobody can reliably distinguish 15 percent from 17 percent, a limit Granovetter sources to Miller's work on human information processing. [FACT] The operational threshold is the lowest value the individual can actually perceive.

One structural requirement is easy to miss. The model requires each individual’s net-benefit curve to cross zero **only once**. The paper works an example of an individual with a 38 percent threshold whose net benefit later turns negative again: cautious participants who join at 50 percent but leave when participation passes 90 percent for fear of official reprisal. [FACT] The analogy offered is an unfamiliar restaurant, where an empty room is a bad sign and a full one means an unbearable wait. [FACT] Curves that cross more than once require different apparatus. The observable signature is aggregate behavior oscillating between levels while the underlying cost structure appears stable. [PROJECTION]

Measurement, and the one empirical test in the paper

Thresholds are dispositions, so they are difficult to observe before the behavior occurs. Granovetter's defense is an analogy to microeconomics, where supply and demand schedules are also almost never measured directly and the discipline has learned to infer them from characteristics and aggregate outcomes. [FACT]

The proposed procedure: index an individual's threshold by the exact number of others who had already adopted at the moment that individual adopted, regress those thresholds on background characteristics, and use the fitted equations to predict threshold *distributions* in populations not yet exposed. [FACT]

Dozier tested this against family-planning adoption in 23 rural Korean villages over 1964 to 1973, fitting on 18 villages and predicting five. [FACT] The result was conditional. Where the regressions explained less than 50 percent of variance in thresholds, predictions were worse than those from standard methods. Above 50 percent, "the accuracy of the recursion model improves dramatically," with many predictions approaching full accuracy. [FACT] Granovetter flags the caution himself: some of the regressors were measured after the period the data covered and correlate with thresholds in ways that strip them of explanatory power. [FACT] [CHARACTERIZATION]

The fifty-year argument

The model did not arrive from nowhere and it did not stay as published.

The predecessor. Thomas Schelling's segregation models, published 1971 to 1972, gave individuals a tolerance for the proportion of unlike neighbors and showed that mild individual preferences aggregate into near-total separation. [FACT] Granovetter credits them directly and generalizes past the specific substance to any binary choice, adding the insistence that exact distributions rather than averages determine outcomes. [FACT]

Concealment. Timur Kuran's *Sparks and Prairie Fires* (1989) and *Now Out of Never* (1991) apply threshold logic to political revolution by adding preference falsification. [FACT] If people conceal opposition while opposition appears weak, private thresholds can fall for years with no visible change, accumulating what Kuran calls a latent bandwagon that a minor event then ignites. [FACT] It remains the standard account of why essentially nobody forecast the events of 1989, and structurally it is Granovetter plus a theory of why the distribution is unobservable. [CHARACTERIZATION]

Networks. Duncan Watts placed threshold agents on sparse random graphs in "A Simple Model of Global Cascades on Random Networks" (2002), identifying two regimes in which very large cascades occur very rarely, with a power-law distribution of cascade sizes where network connectivity is the binding constraint. [FACT] This is the direct ancestor of most contemporary modeling of virality and of cascading infrastructure failure. [CHARACTERIZATION]

Complex contagion. Centola and Macy drew the consequence that matters most in “Complex Contagions and the Weakness of Long Ties” (2007). [FACT] Information transmits on a single contact, so long-range weak ties accelerate it. Behavior with a threshold above one requires reinforcement from several sources, and long ties then *impede* diffusion rather than aiding it. What determines whether a behavior crosses between groups is not the length of a bridge but its width. [FACT] [CHARACTERIZATION]

7 The 25 percent number is one experimental result, not a constant

REFINEMENT

Centola and colleagues ran online groups converging on an arbitrary naming convention, then introduced committed confederates advocating a different name. Below roughly 25 percent of the group the minority failed; at about 25 percent the majority flipped. The design used 10 groups of 20 participants. [FACT] The number has since circulated widely as a general law of social change. It should not be read that way. The authors state that the critical fraction moves with parameters, notably the memory length that stands in for how entrenched a convention is, and independent real-world replication remains thin. [FACT] A separate line of work by Wiedermann and colleagues derives the broad threshold distribution from local network interaction rather than assuming it, recovering saddle-node bifurcation and hysteresis with critical minority sizes near 20 percent. [FACT] [CHARACTERIZATION]

The live disagreement: is the knife-edge real?

Granovetter’s central claim is that cascades are hypersensitive to small perturbations in group composition. That claim is the source of the model’s explanatory power and also its most direct challenge.

8 A small amount of individual randomness makes collective behavior more predictable, not less

REFINEMENT

Macy and Evtushenko revisited the model in "Threshold Models of Collective Behavior II" (2020) with a minimal change: allow individuals to behave randomly some small fraction of the time rather than obeying their thresholds exactly. [FACT] The results run against expectation. A little noise makes collective outcomes *less* sensitive to small perturbations in composition and therefore more predictable. In groups facing an otherwise intractable start-up problem, meaning no instigator, randomness produces what they term spontaneous instigation, so that outcomes come to track the strength of the group's collective interest rather than the accident of who happened to be present. Both effects grow more pronounced as group size increases. [FACT] Their stated conclusion is that "randomness is often assumed to be a theoretically unimportant residual category" and that individual idiosyncrasy needs to be brought back into the study of collective behavior. [FACT]

> **What this does to the original:** [CHARACTERIZATION]

How the model gets misapplied

Threshold language has been absorbed into marketing and social-media strategy in a form the source material does not support, usually reduced to the instruction to locate and recruit the people with threshold zero.

Watts spent much of the following decade arguing that this reverses the causal structure. [FACT] In a threshold cascade the outcome is determined by the distribution across the whole population; who happens to occupy the instigator role is largely incidental, because instigators are necessary but not sufficient and are not scarce. [CHARACTERIZATION] [PROJECTION]

A second misapplication is treating a threshold model as a forecasting instrument. The exact-distribution sensitivity that makes the 100-person demonstration memorable is the same property that makes forecasting impractical, since an answer that turns on a one-person gap requires an input measured to a precision nobody has. [CHARACTERIZATION]

What the model licenses

The question it replaces	Not "what does this group believe," which is often unanswerable and frequently not the operative variable, but "at what level of participation would each person switch."
The diagnostic	Look at the shape of the threshold distribution immediately above the current participation level. If there is nobody who would join at the level already reached, the process stops there regardless of how favorable average sentiment is.
The inference it forbids	Reading individual dispositions off a collective outcome. A large result is not evidence of strong preference, and a null result is not evidence of weak preference, wherever behavior is interdependent.
The stability question	Before treating any predicted equilibrium as real, ask whether it survives a one-person change. If it does not, network structure and sampling accident will dominate preferences, and the prediction is not usable.
Where it does not apply	Situations where behavior is not contingent on others', where thresholds cluster at 0 or 100 percent, or where individuals' net benefit from participating rises and then falls again.

The durable contribution is the negative one. Granovetter supplied a mechanism and a worked counterexample for an inference people make constantly, in political commentary, in post-mortems, in market narratives, and in retrospectives of every kind: that a group which did something must have wanted to. Under interdependence, sometimes it did not, and the arithmetic shows how.

Defined terms

bandwagon

A self-reinforcing process in which each additional participant lowers the barrier for the next; the mechanism by which a threshold distribution converts a small initial action into a large outcome.

cumulative distribution function

A function giving the proportion of a population whose value falls at or below a given point. Here: the proportion of people whose threshold is at or below x.

hysteresis

A system property where the path back is not the path forward: the level of support needed to establish a new state is higher than the level needed to maintain it, so reversing a tip requires more than undoing the trigger.

instigator

A person whose threshold is zero: someone who will act with no prior participation by anyone else. Necessary to start a cascade in the deterministic version of the model.

Pluralistic ignorance

Merton's term for a situation in which many members of a group privately hold a position but each believes they are alone in holding it, because nobody has made their position public.

preference falsification

Kuran's term for the act of misrepresenting one's private preferences under perceived social or political pressure, producing a public opinion distribution that diverges systematically from the private one.

threshold

The proportion of a group that must already have taken an action before a given individual will take it; the point where that individual's perceived benefits first exceed perceived costs.

width

Centola and Macy's term for the number of independent ties connecting two clusters. A wide bridge can deliver the multiple reinforcements a high-threshold behavior requires; a single long tie cannot.

Sources

- Mark Granovetter, "Threshold Models of Collective Behavior," *American Journal of Sociology* 83(6), 1978, pp. 1420-1443 - the source of the model, the 100-person demonstration, the difference equation, the normal-distribution result, the social-structure and sampling analyses, the falsification conditions, and the Dozier test. Full text PDF.
- Thomas Schelling, "Dynamic Models of Segregation," *Journal of Mathematical Sociology* 1, 1971, pp. 143-186 - the predecessor model Granovetter credits, applying tolerance thresholds to residential choice.

- Mark Granovetter and Roland Soong, "Threshold Models of Diffusion and Collective Behavior," *Journal of Mathematical Sociology* 9, 1983, pp. 165-179 - the follow-up work extending the model to diffusion, later applied to consumer behavior and residential segregation.
- Timur Kuran, "Now Out of Never: The Element of Surprise in the East European Revolution of 1989," *World Politics* 44(1), 1991, pp. 7-48 - preference falsification, the latent bandwagon, and the unforecastability of revolution.
- Timur Kuran, *Private Truths, Public Lies: The Social Consequences of Preference Falsification* (Harvard University Press, 1995) - the book-length treatment.
- Duncan J. Watts, "A Simple Model of Global Cascades on Random Networks," *PNAS* 99(9), 2002, pp. 5766-5771 - threshold agents on sparse random graphs; the two cascade regimes and the power-law size distribution.
- Damon Centola and Michael Macy, "Complex Contagions and the Weakness of Long Ties," *American Journal of Sociology* 113(3), 2007, pp. 702-734 - the distinction between simple and complex contagion and the argument that bridge width, not length, governs the spread of high-threshold behavior.
- Damon Centola, Joshua Becker, Devon Brackbill, and Andrea Baronchelli, "Experimental Evidence for Tipping Points in Social Convention," *Science* 360, 2018, pp. 1116-1119 - the committed-minority experiment and the roughly 25 percent critical fraction, with the authors' own caveats about its dependence on parameters. Open-access version.
- Marc Wiedermann, E. Keith Smith, Jobst Heitzig, and Jonathan F. Donges, "A Network-Based Microfoundation of Granovetter's Threshold Model for Social Tipping," *Scientific Reports* 10, 2020 - deriving the threshold distribution from local network interaction; saddle-node bifurcation, hysteresis, and critical minorities near 20 percent.
- Michael W. Macy and Anna Evtushenko, "Threshold Models of Collective Behavior II: The Predictability Paradox and Spontaneous Instigation," *Sociological Science* 7, 2020, pp. 628-648 - the effect of individual randomness on cascade sensitivity, and the case for returning idiosyncrasy to the model.
- David Matza, *Delinquency and Drift* (Wiley, 1964) - the study of delinquent boys Granovetter uses to separate thresholds from norms.
- George A. Miller, "The Magical Number Seven, Plus or Minus Two," *Psychological Review* 63, 1956, pp. 81-97 - the limit on perceptual distinctions underlying the granularity problem.