When Adoption Is Sudden, Not Gradual

Jon Smirl

2026-03-01

The S-Curve Lie

Every technology forecaster draws the same picture: adoption starts slow, accelerates through a middle phase, and tapers off as the market saturates. The S-curve. It is elegant, intuitive, and deeply misleading for an important class of technologies.

The S-curve assumes adoption is a *continuous* process — that each new user adds a little momentum, which attracts another user, which adds a little more. Ten percent adoption leads to twenty, which leads to forty. Smooth and predictable. If you can measure early adoption rates, you can forecast the timeline.

But some technologies do not work this way. For network technologies — where the value of joining depends on who else has joined — adoption is not a smooth ascent. It is a regime shift. Below a critical mass N^*, adoption is essentially zero no matter how good the technology is. Above N^*, adoption is rapid, discontinuous, and practically irreversible. The technology goes from "nobody uses this" to "everybody uses this" with very little time in between.

This distinction matters enormously for prediction, for investment, and for policy. If you are watching for gradual acceleration and the actual mechanism is a sudden jump, you will be blindsided every time.

Why Critical Mass Exists

Consider a simple decision. You are choosing between a centralized service (reliable, established, works today) and a mesh alternative (potentially better, but only if enough other people join). Your expected payoff from joining the mesh depends on how many others have joined:

$$\pi_{\text{mesh}}(N) = b \cdot N^{\gamma} - c$$

where b is the per-connection benefit, N is the number of current participants, $\gamma > 0$ captures how network effects scale, and c is the switching cost. When N is small, $\pi_{\text{mesh}} < 0$ — joining is a losing proposition. You pay the switching cost c and get almost nothing back because there is nobody to connect with.

The critical mass N^* is the value of N at which π_{mesh} first crosses zero. Below N^*, every rational agent stays with the centralized option. The mesh sits empty despite being technologically superior. Above N^*, joining becomes individually rational, and each new joiner makes it *more* rational for the next person, triggering a cascade.

Definition (Critical Mass Threshold).

$$N^* = \left(\frac{c}{b}\right)^{1/\gamma}$$

The critical mass is the number of participants at which the network benefit first exceeds the switching cost. It falls when switching costs c decrease or per-connection benefits b increase.

This is why the transition is discontinuous. There is no stable equilibrium between "nobody joins" and "nearly everyone joins." Once N crosses N^*, the cascade runs to completion.

The Epidemiological Analogy

Epidemiologists have a concept that captures this perfectly: the basic reproduction number R_0. For a disease, R_0 is the average number of new infections caused by a single infected individual. When $R_0 < 1$, each case generates fewer than one new case, and the outbreak dies. When $R_0 > 1$, each case generates more than one new case, and the epidemic grows exponentially.

The same logic applies to mesh adoption. Define a mesh reproduction number: the average number of new adopters recruited by a single existing mesh participant. Below the threshold — the mesh equivalent of $R_0 = 1$ — each adopter fails to recruit enough new participants to sustain growth. The technology sputters and stalls. Above the threshold, each adopter recruits more than one new participant, and adoption becomes self-sustaining.

Theorem (Activation Threshold).

The mesh network admits a nontrivial equilibrium (positive adoption) if and only if its reproduction number exceeds 1. The transition from zero adoption to positive adoption is a first-order (discontinuous) regime shift, not a gradual increase. See *activation_threshold_iff_product*.

The mathematics here borrow from the Potts model in materials science, repurposed for economics. In the original setting, the Potts model describes how atoms in a material choose orientations — and how, at a critical temperature, the material snaps from disordered to ordered in a discontinuous jump. The economic version replaces atoms with agents, orientations with technology choices, and temperature with uncertainty. The result is the same: a first-order transition with a discontinuous jump at the critical point, unlike a gradual (second-order) transition where the change is smooth.

This distinction between first-order and second-order is crucial. A second-order transition *can* be predicted by watching early trends — the change is continuous, and small movements foreshadow larger ones. A first-order transition *cannot*. Early adoption rates near zero carry no information about how close the system is to N^*. You can be one unit below the threshold with zero visible adoption, then one unit above with explosive growth.

Three Examples

Email (1990s). Email existed from the early 1970s, used almost exclusively by academics and military researchers. For twenty years, adoption was negligible — not because the technology was bad, but because there was nobody to email. Critical mass arrived in the mid-1990s with consumer internet access. Within roughly three years (1995–1998), email went from a curiosity to a necessity. The transition was not gradual. It was a regime shift triggered by crossing N^*.

WhatsApp (2010s). A messaging app is useless until your contacts use it. WhatsApp launched in 2009 and grew slowly. Then, country by country, it hit critical mass. In each market the pattern was the same: near-zero relevance, then a sharp jump to near-universal adoption within 12–18 months once enough of a user's contact list had joined. The aggregate global growth curve looks like an S-curve, but that is an artifact of averaging across countries that each experienced a discontinuous jump at different times.

China's Bitcoin Mining Ban (2021). In May 2021, China banned cryptocurrency mining. The Bitcoin network's hashrate — a measure of total computational power — collapsed by 46 percentage points almost overnight. This looked like a death blow. But within months, mining operations had reformed in the United States, Kazakhstan, and elsewhere. The mesh reformed because the conditions for $R_0 > 1$ still held: mining was profitable, hardware was portable, and electricity was available. The ban removed participants, temporarily pushing the network below N^* in China, but the global mesh crossed the threshold again quickly. This episode demonstrates both the discontinuity of the transition and the resilience of mesh structures once formed.

Why S-Curves Mislead

The standard technology adoption S-curve assumes a **diffusion** process: awareness spreads gradually, each adopter tells two friends, and penetration climbs smoothly toward saturation. This model works well for technologies with individual value — a better toaster, a faster car. You do not need other people to own the same toaster for yours to work.

But for network technologies, diffusion is the wrong model. The correct model is a *knock-out_triggered_regime_shift* — a collective shift that happens all at once when conditions are right. The practical consequence is severe: if you are forecasting mesh adoption by fitting an S-curve to early data, you will dramatically overestimate the time to majority adoption *and* dramatically underestimate the speed of the final transition. The early data points are near-zero and contain almost no information about when the jump will occur.

What *does* predict the timing? The distance to the threshold. How close is N to N^*? This depends on switching costs, per-connection benefits, and the existing installed base — not on the current adoption rate. A technology can sit at 2% adoption for a decade, then jump to 80% in a year, if something reduces the switching cost c (cheaper hardware, better software, regulatory approval) or increases the per-connection benefit b (a new use case, interoperability with existing systems).

Implications

For investors: watching early adoption rates is nearly useless for network technologies. Instead, track the variables that determine N^* — switching costs, interoperability, and the density of potential participants. A technology at 1% adoption with falling switching costs is closer to breakthrough than a technology at 10% adoption with stable switching costs.

For policymakers: regulation that raises switching costs (licensing requirements, compliance burdens, interoperability restrictions) can keep adoption below N^* indefinitely, even for a clearly superior technology. Conversely, regulation that lowers switching costs (open standards, portability mandates) can trigger a cascade that no subsequent regulation can reverse.

For researchers: the distinction between diffusion and regime shift is not semantic. It changes which mathematical models are appropriate, which data is informative, and which predictions are possible. The *mesh-formation* framework and the *r0-crossing* criterion provide the right tools for analyzing technologies where value depends on participation. As (Perez2002) documented across two centuries of technological revolutions, the pattern of sudden adoption following long dormancy is not an anomaly — it is the norm for network technologies. The CES framework explains why.

References