

Hype Cycles, Moral Panic, and the Epistemics of Consequential Technology

Noah Podolefsky and Claude Sonnet 4.6 | March 2026

Abstract

The Gartner Hype Cycle is widely used as a descriptive model for how public and market enthusiasm around emerging technologies evolves over time. It assumes a broadly homogeneous audience moving through sequential phases: a peak of inflated expectations, followed by a trough of disillusionment, and eventually a plateau of productivity. This paper argues that the model is structurally incomplete for technologies with significant social, political, or ethical consequences. In those cases, a second, oppositional response curve emerges simultaneously with the hype peak, driven by anticipated harm rather than disappointed expectations. Drawing on Stanley Cohen's moral panic framework and Sheila Jasanoff's concept of sociotechnical imaginaries, we propose a dual-curve model in which techno-optimists (TOs) and concerned skeptics (CSs) occupy opposite poles of a polarized response landscape. A third population, cautious optimists (COs), does not ride a third curve but instead serves as an analytical lens on the space between the two: defined by a distinct epistemology rather than a directional response trajectory. The fundamental fault line runs between narrative-first and evidence-first reasoning. That distinction has implications for how technology governance and public discourse might be more productively structured.

1. Introduction

When Gartner introduced the Hype Cycle in the 1990s, it offered a useful if simplified account of how technologies move through public consciousness. The model described a recognizable pattern: an initial trigger gives rise to rapid enthusiasm, expectations overshoot reality, disillusionment follows, and the technology eventually finds its appropriate level of adoption. The model has proven durable precisely because it captures something real about how markets and media process novelty.

Its implicit assumption, however, is that the relevant audience is relatively homogeneous in its concerns. The audience in the Gartner model is primarily composed of potential adopters, investors, and commentators who care mainly about whether the technology will deliver on its promises. Disillusionment arrives when it does not. The trough is a consequence of unmet expectations.

This model fits reasonably well for technologies whose primary significance is economic or functional: faster processors, better batteries, more efficient logistics software. It fits poorly for technologies that implicate social structure, political power, labor markets, civil liberties, or deeply held values. For those technologies, a second population emerges immediately and in parallel, reacting to anticipated consequences rather than waiting for expectations to be disappointed. The backlash does not wait for the trough; it arrives at the peak.

This paper develops a framework for understanding that collision. Consequential technologies generate two response curves running in parallel and colliding at the peak: techno-optimists (TOs) riding the familiar Gartner trajectory and concerned skeptics (CSs) following an opposing trajectory driven by anticipated harm. A third population, cautious optimists (COs), occupies a

qualitatively different position. Rather than constituting a third curve, COs serve as an analytical lens on the space between the two: their evidence-first reasoning stands in contrast to the narrative-first reasoning shared by both TOs and CSs, and that contrast illuminates what is at stake in the collision.

2. The Gartner Model and Its Limits

The Gartner Hype Cycle identifies five phases in the public life of a technology: the Technology Trigger, the Peak of Inflated Expectations, the Trough of Disillusionment, and the Plateau of Productivity. The model is descriptive rather than predictive, and it has been criticized on various methodological grounds, including the absence of a robust empirical basis for its phase boundaries and the circularity of classifying technologies after the fact. Nevertheless, it captures a genuine pattern in how enthusiasm and skepticism play out over time in market and media discourse.

The model's key limitation for our purposes is its temporal linearity. It assumes that skepticism follows enthusiasm, that disillusionment arrives downstream when inflated expectations collide with reality. This structure holds when the skepticism is functional: when critics are asking whether the technology works as advertised. It breaks down when skepticism is consequential: when critics are asking what the technology will do to society, regardless of whether it works as advertised.

A technology like generative artificial intelligence, for instance, generates simultaneous responses of intense enthusiasm from TOs (developers, investors, and early adopters) and intense concern from CSs (educators, labor economists, civil society organizations, and affected workers), with both responses occurring before the technology has had time to either deliver or disappoint at scale. The Gartner model has no vocabulary for this simultaneity.

3. Moral Panic as a Parallel Curve

Stanley Cohen's concept of moral panic, developed in his 1972 study of the Mods and Rockers phenomenon in Britain, describes a process by which a society identifies a threat to its existing values and structures, amplifies that threat through media and institutional response, and mobilizes against it. The threat is represented by a "folk devil," a figure or phenomenon onto which anxieties are projected. The moral panic draws on real social tensions and legitimate concerns, even when the specific threat is exaggerated or mischaracterized.

The moral panic framework has been applied to a range of technology-related anxieties, from video games to social media to genetic engineering. In each case, the pattern is recognizable: a new phenomenon triggers rapid, disproportionate concern among CSs; media coverage amplifies the threat; institutional responses are mobilized; and the concern eventually subsides, often without the feared outcomes materializing at the anticipated scale.

A meaningful distinction separates classical moral panics from the CS response to consequential technologies. Classical moral panics tend to be conservative in the sociological sense: defensive reactions that seek to protect an existing social order against perceived transgression. CSs responding to technologies like artificial intelligence or social media surveillance may have different motivations entirely; many are responding to genuinely novel risks with no historical precedent. The threats they anticipate, including algorithmic discrimination, mass labor displacement, and erosion of epistemic autonomy, are plausible extrapolations from observable trends rather than imagined dangers.

We propose the term "grounded moral panic" to describe this variant. A grounded moral panic shares the structural features of Cohen's original formulation (rapid mobilization, emotional amplification, narrative simplification, disproportionate response) but the object of concern carries a real, if uncertain, probability of materializing. The scenario is imaginable because it is logically derivable from known properties of the technology, not merely because it feels alarming. What distinguishes CS epistemology from legitimate consequence-anticipation is the sequence: the conclusion arrives before the evidence, and evidence is subsequently recruited to support a position already reached on other grounds.

4. Sociotechnical Imaginaries and Competing Visions

Sheila Jasanoff's concept of sociotechnical imaginaries provides a complementary analytical lens. Imaginaries are collectively held, institutionally stabilized visions of desirable futures, enacted through science and technology. Different national, political, and cultural communities construct different imaginaries around the same technology, each projecting its own values, fears, and aspirations onto the technological object.

The imaginary framework helps explain why the two response curves do not converge. TOs and CSs are constructing fundamentally different visions of what the technology is and means, rather than simply offering different assessments of the same evidence. The TO imaginary emphasizes capability, efficiency, democratization of access, and economic growth. The CS imaginary emphasizes disruption, concentration of power, erosion of human agency, and exacerbation of existing inequalities. These framings differ not merely in their predictions about the same future, but in what counts as a relevant outcome.

This framing divergence is important because evidence alone cannot settle the dispute. A deployment of AI that improves diagnostic accuracy in radiology can be simultaneously evidence for the TO imaginary (technology augmenting human capability) and the CS imaginary (technology displacing radiologists and concentrating value in platform owners). Both readings are defensible from the same data. The imaginaries shift only, if at all, under the accumulated weight of evidence over long time horizons, and even then imperfectly.

5. A Visual Model of the Collision

Figure 1 presents the dual-curve model alongside the CO epistemic band. The TO trajectory (solid curve, above zero) follows the familiar Gartner shape: rising from the Technology Trigger to a Peak of Inflated Expectations, falling to a Trough of Disillusionment, and settling at a Plateau of Productivity. The CS trajectory (dashed curve, below zero) mirrors this structure in opposing valence: rising in alarm from the same trigger to a Peak of Moral Alarm, moderating to a Trough of Partial Acceptance as the most extreme fears go unrealized, and settling at a Plateau of Grudging Recognition that remains permanently subzero. The CS curve never crosses the zero line; concerns moderate but do not resolve into net positive assessment.

The CO population is represented as a band straddling zero throughout. The band is narrow at the trigger (little evidence, few settled questions), widens through the middle as the space of serious empirical and normative questions expands, and narrows again as evidence accumulates toward emerging consensus. The CO centerline hugs zero without committing to a direction, reflecting genuine sustained uncertainty rather than centrist fence-sitting. The width of the band at any point represents the range of live questions that careful, evidence-based reasoning must hold open.

The vertical shading marking the collision zone indicates the period when the TO peak and CS peak of moral alarm coincide. This simultaneity is the diagnostic signal: when a technology generates enthusiasm and alarm at the same moment rather than in sequence, it indicates that the technology touches something beyond functional capability, implicating social structure, power distribution, or deeply held values.

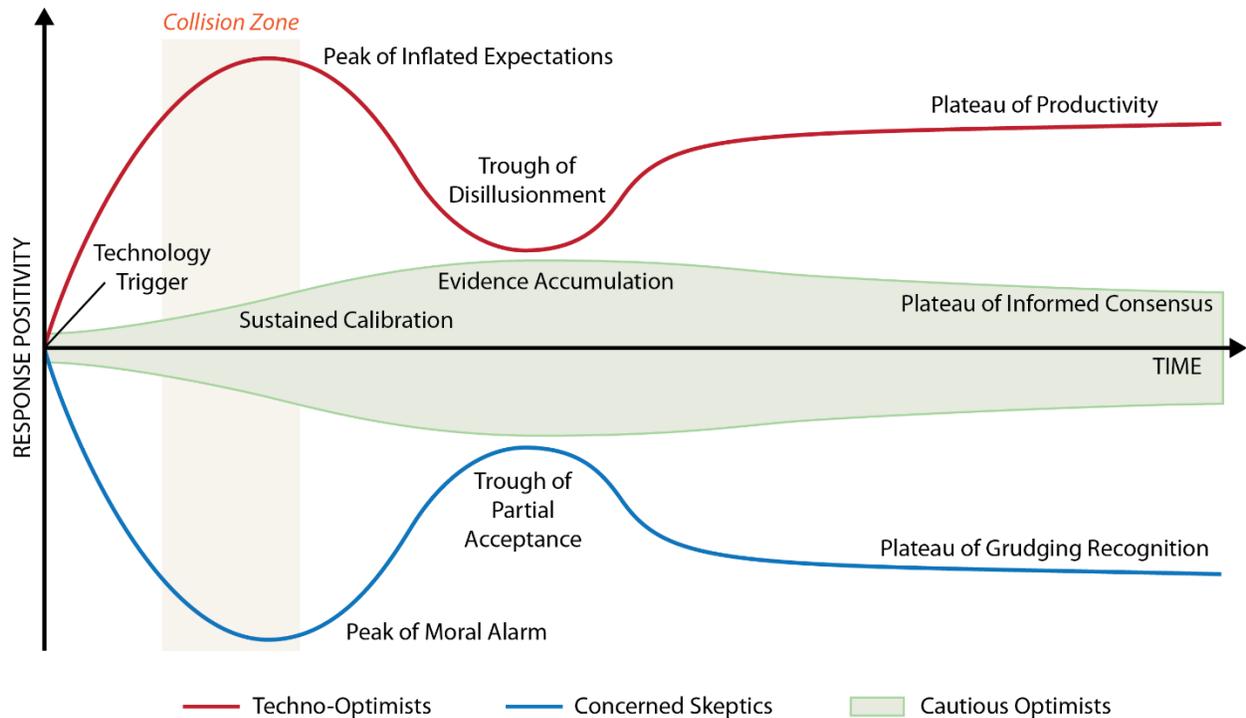


Figure 1. Response trajectories for consequential technologies. Solid curve (above zero): techno-optimist (TO) Gartner trajectory. Dashed curve (below zero): concerned skeptic (CS) grounded moral panic trajectory. Shaded band straddling zero: cautious optimist (CO) epistemic band, width proportional to range of live questions.

6. Three Populations, Two Epistemologies

The dual-curve model identifies three distinct populations in the public response to consequential technologies. Understanding them requires attending to how they reason, not just what positions they hold.

6.1 Techno-Optimists (TOs)

TOs are the primary population described by the Gartner model: investors, developers, early adopters, and media commentators who project high expectations onto an emerging technology. Their enthusiasm is typically genuine and often partially justified; new technologies do frequently deliver significant benefits. The characteristic error of TO epistemology is a tendency to discount negative externalities, assume that problems will be solved by further technical development, and treat the imagined endpoint of the technology as its likely trajectory.

TOs tend toward narrative-first reasoning. A compelling story about what a technology could become is constructed early, and subsequent information is filtered through that story.

Confirming evidence is weighted heavily; disconfirming evidence is treated as a temporary obstacle or a problem for future versions. When expectations are high and evidence is thin, emotional intuition reliably fills the gap, though is generally in good faith.

6.2 Concerned Skeptics (CSs)

CSs are the population that generates the second curve. They arrive at the peak rather than after disillusionment, mobilizing against a technology whose consequences they anticipate or are already observing in early form. Their concerns may originate from various motivations: ideological opposition to the interests or actors associated with the technology; genuine technical or social analysis that arrives early; or an emotional response to a vivid narrative of harm.

CSs and TOs share the same underlying epistemic structure. Both populations start from a prior belief and then look for evidence to support it, rather than building their view from the evidence up. Both are susceptible to narrative capture, social reinforcement, and motivated cognition. Their priors point in opposite directions, but the cognitive machinery is the same. A TO and a CS looking at the same deployment of facial recognition technology will each find abundant support for their pre-existing view.

This shared epistemic structure has a significant implication: TOs and CSs, despite appearing to be opposites, are more similar to each other than either is to the CO population. Shifting someone from one camp to the other requires only reversing their prior; the underlying habits of mind stay the same. Moving either population toward genuinely evidence-calibrated reasoning requires a much deeper change.

6.3 Cautious Optimists (COs) as Analytical Lens

Rather than constituting a third mass response, COs occupy a qualitatively different position from TOs and CSs. The CO disposition serves as an analytical lens: by showing what evidence-first reasoning looks like in practice, it clarifies what is absent from the TO and CS positions. COs hold that the technology may deliver significant benefits while also believing those benefits are not inevitable, and that the path to them runs through confronting real risks and making deliberate governance choices. They maintain genuine uncertainty about both the upside and the consequences simultaneously.

Domain expertise is common among COs but does not define them. They may bring technical knowledge, historical perspective, media literacy, or sociological analysis. What they share is an epistemology: grounded in evidence, comfortable with sustained uncertainty, resistant to the pull of a compelling narrative, and capable of holding multiple framings simultaneously without forcing a premature resolution. Psychologists call this actively open-minded thinking: the disposition to seek out evidence that might overturn your current view, revise beliefs when new information warrants it, and treat conclusions as provisional. Research by Jonathan Baron, Keith Stanovich, and others consistently finds this to be a minority disposition, even among highly educated populations.

The real fault line runs between narrative-first and evidence-first reasoning. The TO/CS axis is downstream of that more fundamental divide.

7. Structural Disadvantage and the Information Environment

COs are doing the most careful and productive reasoning in the public response to consequential technologies, but they are doing it under conditions of structural disadvantage. Contemporary

information environments, particularly social media platforms, systematically reward certainty, simplicity, and emotional intensity. Nuanced positions that resist easy summary and acknowledge genuine uncertainty do not travel well. Outrage and enthusiasm are both high-engagement; measured uncertainty generates neither.

There is a further irony: the TO and CS populations amplify each other through their conflict. Extravagant TO claims generate material for CSs to mobilize against; that CS mobilization generates defensive responses from TOs; and the resulting conflict produces high-engagement content that further marginalizes the CO voice. The collision of the two curves, rather than generating synthesis, produces a feedback loop that drives both poles further from evidence-calibrated positions.

The speed of contemporary information propagation also alters the timing structure of the collision. In earlier eras, the CS response would have developed more slowly, allowing more time for evidence to accumulate before alarm peaked. Now the CS curve can peak almost simultaneously with the TO hype curve, before the technology has had time to either deliver or disappoint. The backlash is anticipatory rather than reactive, and both populations are operating primarily on narrative and prior because the evidence simply has not arrived yet.

8. Conclusion

The Gartner Hype Cycle describes a real phenomenon, but only part of the public response to consequential technologies. When a technology implicates social structure, political power, or widely held values, it generates two simultaneous curves running in opposite directions and colliding at the peak: TOs riding the familiar hype trajectory, CSs following a grounded moral panic trajectory.

The collision is not a dialogue. TOs and CSs talk past each other, reinforced by the dynamics of contemporary information environments and by the fundamental incompatibility of their visions of what the technology means. Providing better arguments or more evidence is unlikely to shift either population, since both are reasoning from prior to evidence, and more evidence simply gives both sides more material to work with.

The CO disposition, understood as an analytical lens rather than a third trajectory, clarifies what is at stake in this collision. COs are the population best positioned to do productive work at the interface of technology development and governance, but their actual influence on outcomes is often limited. They matter most in specific institutional contexts, such as technical standards bodies, scientific advisory panels, certain regulatory processes, and least in the places with the most power: industry and investment, which are dominated by TOs, and political advocacy and regulation, which are largely driven by CSs. The dual-curve collision, when it occurs, is itself a signal worth attending to: it indicates that a technology is consequential in ways that purely functional analysis will miss. Recognizing that signal early, before positions have hardened, may be the most useful contribution this framework can offer.

References

- Baron, J. (1995). Myside bias in thinking about abortion. *Thinking and Reasoning*, 1(3), 221-235.
- Cohen, S. (1972). *Folk Devils and Moral Panics*. London: MacGibbon and Kee.
- Collingridge, D. (1980). *The Social Control of Technology*. New York: St. Martin's Press.

- Fenn, J., & Raskino, M. (2008). *Mastering the Hype Cycle*. Harvard Business Press.
- Fenn, J. (1995). When to leap on the hype cycle. Gartner Group Research Note SPA-401-068.
- Jasanoff, S. (2015). *Future Imperfect: Science, Technology, and the Imaginations of Modernity*. In *Dreamscapes of Modernity*. University of Chicago Press.
- Stanovich, K. E., West, R. F., & Toplak, M. E. (2016). *The Rationality Quotient*. MIT Press.
- Stirling, A. (2008). Opening up and closing down: Power, participation, and pluralism in the social appraisal of technology. *Science, Technology, & Human Values*, 33(2), 262-294.