



© Banstanks/AdobeStock

LIGHTS - Technology

# To Tell the Truth : Five voices. Five competing realities. And the question no algorithm can answer

ESCP Impact Paper No.2026-45-ENG

Kenneth KRUSHEL  
NYU Stern School of Business

# **[ To Tell the Truth : Five voices. Five competing realities. And the question no algorithm can answer**

Kenneth Krushel\*

ESCP Business School

## **Abstract**

We have more information and assessment of artificial intelligence than any previous generation. We agree about it less. The most consequential decisions about AI, who builds it, who governs it, who profits from it, and who pays its hidden costs, are being shaped not by evidence but by competing narratives, each coherent, each partial, each backed by interests that rarely announce themselves as such. The top AI labs did not reach trillion-dollar valuations by underselling themselves, and what sounds like an existential warning often turns out, on closer inspection, to be a very effective sales pitch. This essay reads five of those narratives, traces their intellectual origins, and exposes the gap between what each voice illuminates and what it protects. What matters most right now is not building models or writing code. It is learning to hear what a story is selling before you decide whether to buy it.

Keywords: narrative, epistemology, surveillance capitalism, governance, critical thinking

\*Professor, NYU Stern School of Business

**To Tell the Truth:**  
**Five voices, five competing realities.**  
**And the question no algorithm can answer**

At 3:45 in the morning on April 10, 2026, a man named Daniel Alejandro Moreno-Gama walked up to 855 Chestnut Street in San Francisco's Russian Hill neighborhood and threw a bottle containing a flaming rag at the metal gate of the house. Security guards extinguished the fire. The suspect fled on foot. Hours later, the owner of the house sat down and wrote a blog post. He included a photograph of his husband and their child. And then Sam Altman, CEO of OpenAI and arguably the most consequential figure in the artificial intelligence industry, wrote something that, read carefully, is more honest than almost anything else he has said publicly about AI.

« *I have underestimated the power of words and narratives.* » Sam Altman, April 10, 2026<sup>1</sup>

*This essay is about that underestimation. Not Altman's specifically. Everyone's.* The confession is small; the gap it reveals is not.

Altman was referring to a long profile in *The New Yorker* (April 6, 2026) that had appeared days earlier, questioning whether he could be trusted. But the sentence reaches further than he perhaps intended. We are in a moment when the most powerful people in technology, finance, and government are competing to define what AI is, what it will do, and who it will serve. Each is telling a story, and each story is a bid for authority. And most of us, including the people telling the stories, have underestimated what those narratives are doing to us.

In 1956, a television show called *To Tell the Truth* premiered in America. Three people stood before a panel, all claiming the same identity. Two were coached to lie; only one was bound by the rules to be honest. Germany had *Sag die Wahrheit*, a cultural fixture on SWR for decades. France watched *Qui dit vrai?* on TF1 through the seventies and eighties. The British made *Tell the Truth* a staple on ITV. The show was not really about deception. It was about discernment: the specific human skill of separating performance from reality when both arrive dressed in the same suit. We are back in that studio. Except now, as Altman's morning made clear, the stakes are not television embarrassment.

The argument this essay makes is not that any of these voices is wrong. It is that each is incomplete in a specific, determined way, shaped by the position from which it speaks, and that understanding that shape is more useful than choosing a side.

## **The end of fact**

Before we can ask who is telling the truth, we need to understand why the question has become so hard. Harvard historian Jill Lepore, writing in *The New Yorker* in 2016, observed that facts are being taken over by data, producing what she called epistemological mayhem, in part because weighing facts requires investigation, discernment, and judgment, which are irreducibly human activities, while the collection and analysis of data are outsourced to machines.<sup>2</sup> She noted, with characteristic dryness, the dream attributed to Google's Larry Page: that if you think about a fact, a machine will simply tell you the answer, because we will all have implants. This is not a fantasy: as of March 2026, China approved the world's first commercially cleared invasive brain-computer interface.<sup>3</sup>

What has arrived ahead of them is harder to contain: infinite information, and the collapse of any shared agreement about how to trust it. Two governing philosophies have emerged in direct tension. The European Union's AI Act, in force since 2024, establishes a risk-tiered framework that classifies AI systems by the harm they can cause, imposes transparency requirements on general-purpose models, and prohibits outright the uses deemed incompatible with fundamental rights.<sup>4</sup> The United States, notwithstanding a June 2026 Executive Order establishing a voluntary pre-release review process, has moved under its current administration, with the sustained pressure of the Big Tech lobby, to eliminate most federal oversight, treating the technology as infrastructure to be built at speed. These are not policy disagreements. They are disagreements about what kind of thing AI is.

As philosopher Michael Lynch argued in 2012, without shared standards for what counts as reliable evidence, we will not be able to agree on the facts, let alone on values.<sup>5</sup> We have not merely lost that common ground. We have lost the habit of noticing its absence.

What follows reads five positions already contending for authority over this question, not to adjudicate between them, but to see them clearly enough to stop being managed by them.

## **The sovereign accelerationists**

In June 2023, Marc Andreessen published a manifesto titled *Why AI Will Save the World*. The word *save* is doing significant moral labor in that title. Salvation implies a prior condition of damnation, an urgency that places objection on the wrong side of history. Andreessen was explicit: "*AI is not just another technology*," he wrote. "*It is the fundamental technology, the one that will solve all of mankind's greatest problems.*"<sup>6</sup> The word *all* is doing as much work as *save*.

The intellectual architecture draws from Ayn Rand's *The Fountainhead*, with its moral primacy of the productive creator, and from Friedrich Hayek's *The Road to Serfdom*, which argued that central planning, however well intentioned, tends toward tyranny. For Andreessen and his political ally David Sacks, who served as the White House AI and crypto czar until March 2026, AI safety regulation is not a reasonable precaution. It is a road to serfdom: entrenching incumbents, empowering bureaucrats, surrendering the future to China. Andreessen Horowitz, the venture capital firm he cofounded, has also placed some of the largest AI bets in the industry. The financial exposure does not invalidate the argument. But it clarifies what the argument is protecting.

## **The containment insiders**

Dario Amodei left OpenAI to found Anthropic because he believed the technology he was building was dangerous. Mustafa Suleyman cofounded DeepMind, sold it to Google, and then moved to Microsoft AI. Both now use the language of containment to describe technology they built and sold. The word containment was borrowed from Cold War nuclear strategy, which tells you something about the scale of what they believe they are managing. Suleyman, to his credit, names the contradiction. "*The coming wave*," Suleyman writes, "*will be the defining political challenge of the twenty-first century and one of the greatest tests democratic societies have ever faced.*"<sup>7</sup> He argues for radical new governance architecture while acknowledging that containment may already be impossible.

Amodèi's interpretability research, in his essay «*The Urgency of Interpretability*», represents a theory of trust through transparency. The EU AI Act's transparency requirements for general-purpose models formalize exactly this demand: explain how the system works, or accept that it cannot be trusted in high-stakes decisions. Whether Amodèi's research meets that standard in practice, or whether it serves mainly to preempt the regulators, is a question neither his essay nor the Act has yet answered. The reader must decide whether that is wisdom or marketing, particularly as Anthropic moves toward commercial scale.

## **The structural power critics**

In 1914, the lawyer Louis Brandeis, who would become a US Supreme Court Justice, published *Other People's Money and How the Bankers Use It*, arguing that concentrated financial power was as dangerous to democracy as any political tyranny. A century later, Columbia law professor Tim Wu showed in *The Age of Extraction* (2025) that the same logic now governs AI: tech platforms have become instruments of extraction, seizing money, data, and attention from hundreds of millions of people who have no practical alternative. The convenience is real. So is the capture.

For Wu, and for the social theorist Shoshana Zuboff, the truth of AI is not in the code. It is in who controls the infrastructure through which AI reaches everyone else: the cloud platforms, the data pipelines, the server networks that now function as the invisible architecture of the digital economy. Whoever owns the pipes controls the water.<sup>8</sup>

Zuboff's term surveillance capitalism, the conversion of private human behavior into predictive data assets sold to whoever will pay, has migrated from academic argument to policy vocabulary in less than a decade. "Surveillance capitalism," she writes, "unilaterally claims human experience as free raw material for translation into behavioral data."<sup>9</sup> The speed of that migration is its own evidence: the phrase named something people already sensed but could not say. Altman's confession applies here too, in reverse. Zuboff estimated exactly how much power a precise phrase could carry.

## **The materialist skeptics**

In 2023, *TIME* magazine reported that OpenAI had contracted workers in Kenya, paid between \$1.32 and \$2 per hour, to label traumatic content: graphic violence, sexual abuse, descriptions of torture. The purpose was to filter that material out of ChatGPT before the public encountered it.<sup>10</sup> The people who made the product feel clean were themselves exposed to industrial quantities of human suffering, for less than what a San Francisco software engineer spends on coffee before 9am. By 2024, nearly a hundred of those workers had signed an open letter to the US President describing their conditions as modern day slavery.<sup>11</sup> A 2025 survey across Kenya, Colombia, and Ghana documented sixty independent incidents of psychological harm among data workers, including PTSD and substance dependence.<sup>12</sup>

This is the territory Kate Crawford maps in *Atlas of AI* (2021), anchored in Karl Polanyi's argument in *The Great Transformation* (1944) that market systems are always embedded in physical and social reality, regardless of how weightless they appear to their architects.<sup>13</sup> Crawford follows the supply chain the industry's language obscures: lithium extracted from salt flats in Chile, water drawn from rivers in drought regions to cool the servers, human labor sitting beneath the automation. "AI is neither artificial nor intelligent," she has written. "It is made from vast amounts of human labor, data, and resources."<sup>14</sup>

Microsoft's own sustainability report showed a 34% increase in water consumption between 2021 and 2023, driven largely by AI cooling demands.<sup>15</sup> By 2025, North American data centers were consuming nearly one trillion liters of water annually; a single Google facility in Iowa used enough in 2024 to supply all of the state's residential water for five days.<sup>16</sup>

There is no cloud. There is a building, a turbine, a river, and a person in Nairobi reading something they cannot unsee. The materialist skeptic's narrative insists on this: that every claim about AI must be held accountable to the bodies, landscapes, and labor it actually consumes.

### **The marginalist pragmatists**

Tyler Cowen is the least dramatic figure in this essay, issuing no manifestos, fearing no extinction, celebrating no salvation. He is an economist at George Mason University who publishes a widely read blog called *Marginal Revolution*, and his intellectual ancestor is Joseph Schumpeter, who argued in *Capitalism, Socialism and Democracy* (1942) that technological disruption is not a crisis to be managed but a permanent condition to be navigated, with gains and losses distributed unevenly across time and populations.<sup>17</sup> Cowen's contribution is a single, inconvenient question: is the productivity data real?

The evidence at the task level says yes. A 2025 study across three major technology companies found that AI coding tools increased programmer output by 26%, measured by tasks completed per hour and lines of reviewed code per day, with the largest gains among the least experienced developers.<sup>18</sup> The macroeconomic picture is less settled. A March 2026 survey of more than 700 corporate executives by the Federal Reserve Bank of Atlanta found companies reporting a 1.8% productivity gain from AI in 2024. But productivity gains reported in executive surveys measure perception, not output. When the Atlanta Fed researchers cross-checked those self-reports against actual business metrics, the numbers were substantially smaller. The authors called it a productivity paradox, invoking Robert Solow's famous 1987 observation that computers were visible everywhere except in the productivity statistics.<sup>19</sup>

A coder works faster. Whether that makes the company more productive, or the economy more prosperous, is a different and still-open question. Cowen's question is not pessimism. It is the most important empirical question in the room: is this working, at scale, for anyone other than the people selling it? What his framework does not examine is whether the communities that bear the transitional costs of creative destruction have the institutional resilience to survive until the economy catches up.

### **What the humanists saw**

Cowen's skepticism operates at the level of data. There is a deeper skepticism, and it operates at the level of knowledge itself.

When an organization adopts AI as a cost reduction tool and stops asking what it is doing to the people it no longer employs, or to the customers whose behavior it now predicts, something has happened that is more than a management failure. Martin Heidegger, in *The Question Concerning Technology* (1954), had a name for it: 'Gestell' (enframing): the tendency of technology to turn everything it touches, including human beings, into resources to be assessed and optimized. To see a forest as standing timber rather than as a living system is *Gestell* in action; to see a worker as a productivity metric rather than as a person is the same movement of mind.<sup>20</sup> Where Heidegger diagnosed how technology

shapes perception, his contemporary Günther Anders identified the moral consequence. Writing in the aftermath of Hiroshima, Anders named the Promethean gap: the growing distance between what we are capable of making and what we are capable of comprehending or taking responsibility for. We can build the bomb; we cannot fully hold in mind what it means to have built it. The gap has not closed. It has industrialized.<sup>21</sup> The danger is not any particular machine. It is the disappearance of the capacity to question what the machines are doing to our picture of the world.

The third piece of this diagnosis comes from Walter Benjamin. Where Anders measures the distance between what we make and what we can take moral responsibility for, Benjamin measures the distance between what is true and what merely circulates as true. Writing in 1935, he observed that when images and ideas are reproduced infinitely, they lose what he called their aura: the quality of being anchored in a specific time, place, and history. A painting carries the singular presence of the moment it was made; a postcard of the same painting is a copy with no origin. Anders showed us the moral gap between making and reckoning; Benjamin shows us what happens to truth in a world where every claim can be replicated, packaged, and distributed at scale until the original and the imitation are indistinguishable.<sup>22</sup> When every position in the AI debate arrives pre-equipped with its own experts, its own evidence, its own archive of outrage and vindication, the distance between what is true and what is merely persuasive collapses.

The question of what AI can truthfully claim about itself is not a technical question. It requires us to know what truth is, how knowledge is justified, what it means for a claim to be reliable. These questions run through the humanistic tradition as surely as the scientific one. To exclude them from this inquiry would be to accept precisely the narrowing Heidegger warned against: treating every question as an optimization problem, including the question of what is worth optimizing for.

Heidegger, Anders, and Benjamin each diagnosed the problem. The poet John Keats, writing nearly a century before any of them, named the capacity we need in order to resist it. In December 1817, Keats wrote a letter to his brothers George and Tom in which he described, in a flash of private recognition, what he called negative capability: the capacity “of being in uncertainties, mysteries, doubts, without any irritable reaching after fact and reason.” He was 22, writing not a treatise but a thought he could not let go. It has not dated.<sup>23</sup> This is not antirational. It is the willingness to sit with what you do not yet understand, rather than reach for the nearest explanation simply because the uncertainty is uncomfortable. Every voice in this essay reaches after certainty. Andreessen’s is triumphant; Zuboff’s is diagnostic; Cowen’s is empirical. None practices negative capability. That is, in part, what makes each of them persuasive, and each of them insufficient. It is what the current moment requires, and what no AI system has yet demonstrated the capacity to practice.

## **The underestimation we share**

You have just spent time with people who have devoted careers, capital, and considerable intelligence to this question, and they cannot agree on what is happening. The disagreement is the situation.

Consider it plainly. You are an executive whose board wants to know whether to restructure operations around AI tools whose return on investment remains genuinely unresolved. Or you are a student deciding whether to stake the next decade of your career on skills whose value may be transformed at any point by the very technology you are studying. Waiting is itself a decision.

Put the five voices side by side and something becomes clear: each holds a genuine piece of the truth. None holds all of it, and none will tell you so. The game show panel did not simply listen to arguments. It asked questions designed to produce involuntary specificity, the kind of detail only someone who has lived a life would know. The scuff marks on the shoes, not the script.

The question to ask is simple: whose interests does this argument protect? Andreessen's investment portfolio is not exposed to labor displacement, which is why his manifesto does not dwell on it. Amodei's interpretability research is genuine and also positions Anthropic as the responsible adult in the room. Zuboff's phrase became a policy argument and a brand. Surveillance capitalism has made her one of the most cited scholars in technology policy, which is a reminder that naming a system does not place you outside it. Crawford follows the supply chain and is sometimes less attentive to what the technology does for the people using it. Cowen's empirical patience reflects a longstanding faith in markets to absorb shocks that not every community can afford to wait out.

Knowing the ancestry, the financial exposures, the temperamental inclinations, does not make any of these arguments false. It makes them readable in a way they were not before. You are not moving through five random opinions. You are moving through five structured ways of seeing, each with its own internal logic and its own blind spot. Altman wrote, shaken at 4am with a scorched gate outside his front door, that he had underestimated the power of words and narratives. So have most of the figures in this essay, in different ways, for different reasons.

So have the institutions supposed to prepare people to lead through this kind of disruption. The five voices gathered here are not a menu from which you choose one. They are a field, a set of forces whose combined pressure shapes the decisions you make whether you name them or not. The question is whether you understand the field well enough to act within it deliberately, rather than be swept along by it.

When responsibility for a technology is distributed across its builders, its investors, its regulators, and its users, accountability has not been shared. It has been dissolved. Andreessen cannot see displacement, Suleyman cannot see complicity, Zuboff cannot see utility, Crawford cannot see scale, and Cowen cannot see urgency. None of them is wrong. None of them is sufficient. This essay does not resolve the question of what AI is or what should be done about it. That is not evasion. Resolving it prematurely, before the evidence is in and the interests are visible, is precisely what each of these voices has done. What this demands, practically, is not a methodology but a habit: before adopting a technology, a policy, or an argument, ask who built it, who benefits from its spread, and whose costs are not in the prospectus. That is not skepticism. It is the minimum condition for judgment. The panel on the old game show always asked, at the end, for the real person to please stand up. In our version, no one stands up. Everyone remains seated. The work of discernment belongs to the person watching. That is you.

## References

1. Sam Altman, personal blog post, April 10, 2026. <https://blog.samaltman.com>
2. Jill Lepore, "After the Fact," The New Yorker, March 21, 2016. <https://www.newyorker.com/magazine/2016/03/21/the-internet-of-us-and-the-end-of-facts>
3. You Xiaoying, "China has approved the world's first invasive brain-computer chip – here's what's next," MIT Technology Review, June 2026.

4. European Parliament, Regulation (EU) 2024/1689 on Artificial Intelligence (the AI Act), in force August 1, 2024. <https://artificialintelligenceact.eu>
5. Michael Lynch, "In Praise of Reason: Why Rationality Matters for Democracy", MIT Press, 2012, pp. 4–7. <https://mitpress.mit.edu/9780262518420/in-praise-of-reason/>
6. Marc Andreessen, "Why AI Will Save the World," Andreessen Horowitz, June 6, 2023. <https://a16z.com/ai-will-save-the-world/>
7. Mustafa Suleyman with Michael Bhaskar, "The Coming Wave: Technology, Power, and the Twenty-first Century's Greatest Dilemma", Crown, 2023, p. 11. <https://www.the-coming-wave.com>
8. Louis D. Brandeis, *Other People's Money and How the Bankers Use It* (Frederick A. Stokes, 1914), Chapter 1; Tim Wu, *The Age of Extraction: How Tech Platforms Conquered the Economy and Threaten Our Future Prosperity*, Knopf, 2025. <https://www.penguinrandomhouse.com/books/691177/the-age-of-extraction-by-tim-wu/>
9. Shoshana Zuboff, "The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power, PublicAffairs, 2019, p. 8. <https://shoshanazuboff.com/book/about/>
10. Billy Perrigo, "OpenAI Used Kenyan Workers on Less Than \$2 Per Hour to Make ChatGPT Less Toxic," TIME, January 18, 2023. <https://time.com/6247678/openai-chatgpt-kenya-workers/>
11. Open letter from Kenyan data workers to President Biden, 2024. Reported in The Final Call and multiple outlets, May 2024. <https://www.finalcalldigital.com/articles/exploitation-and-rising-concerns-facing-kenyan-ai-data-labeling-workers>
12. Equidem, survey of 76 data workers across Kenya, Colombia, and Ghana, 2025, cited in Brookings Institution, "Reimagining the future of data and AI labor in the Global South," October 2025. <https://www.brookings.edu/articles/reimagining-the-future-of-data-and-ai-labor-in-the-global-south/>
13. Karl Polanyi, "The Great Transformation: The Political and Economic Origins of Our Time", Farrar and Rinehart, 1944, pp. 68–76.
14. Kate Crawford, "Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence", Yale University Press, 2021, p. 8. <https://yalebooks.yale.edu/book/9780300264630/atlas-of-ai/>
15. Microsoft 2023 Environmental Sustainability Report. <https://query.prod.cms.rt.microsoft.com/cms/api/am/binary/RW15mgm>
16. Mordor Intelligence market data, cited in multiple sources, 2025; Google 2024 Environmental Report, Council Bluffs data center. <https://www.tomshardware.com/tech-industry/investors-push-amazon-microsoft-and-google-to-disclose-data-center-water-and-power-consumption>
17. Joseph Schumpeter, "Capitalism, Socialism and Democracy", Harper and Brothers, 1942, Chapter 7, "The Process of Creative Destruction," pp. 81–86.

18. Mert Demirer et al., "The Effects of Generative AI on High Skilled Work," MIT Sloan, 2025. <https://mitsloan.mit.edu/ideas-made-to-matter/how-generative-ai-affects-highly-skilled-workers>
19. Federal Reserve Bank of Atlanta, "How Might AI Change the Workplace? Evidence from Corporate Executives," March 25, 2026. <https://www.atlantafed.org/research-and-data/publications/policy-hub-macroblog/2026/03/25/how-might-ai-change-the-workplace-evidence-from-corporate-executives>
20. Martin Heidegger, "The Question Concerning Technology," in "The Question Concerning Technology and Other Essays", trans. William Lovitt, Harper and Row, 1977, pp. 3–35.
21. Günther Anders, "Die Antiquiertheit des Menschen", Vol. 1, C.H. Beck, 1956; Vol. 2, C.H. Beck, 1980. The concept of the Promethean gap (Prometheisches Gefälle) appears in Vol. 1, pp. 15–95. Anders argued that modern technology creates a structural disparity between human productive capacity and human moral comprehension—we can build what we cannot fully reckon with.
22. Walter Benjamin, "The Work of Art in the Age of Mechanical Reproduction" (1935), in *Illuminations*, ed. Hannah Arendt, trans. Harry Zohn, Schocken Books, 1969, pp. 217–252.
23. John Keats, letter to George and Tom Keats, December 21, 1817, in "The Letters of John Keats, 1814–1821", ed. Hyder Edward Rollins, Vol. 1, Harvard University Press, 1958, p. 193.