

The “Big Tech” Governance Dilemma: Lessons from OpenAI’s Struggle to Balance Innovation and Accountability

Ayanna Varma

“Sam is out.” At noon on November 17, 2023, a single WhatsApp text to Silicon Valley executives triggered a crisis at OpenAI, the company at the forefront of artificial intelligence.¹ The days that followed produced widespread confusion and chaos: the company’s board fired CEO Sam Altman, cycled through interim CEOs, faced overwhelming employee revolt, and ultimately reinstated him only after the board itself was completely recomposed.² The speed and opacity of this conflict left observers confused, struggling to answer a basic question: how could a company this influential unravel so quickly, and what exactly caused the breakdown?³

The crisis mattered not only because of its internal dysfunction, but because OpenAI occupies an unusually consequential position in today’s technological and economic landscape, where the scale of AI integration into daily life is increasingly difficult to ignore:

Goldman Sachs projects that AI could contribute seven trillion dollars to global GDP over the next decade⁴; Forbes estimates that more than 378 million people will use AI tools in 2025; and in the health sector alone, roughly 90 percent of US hospitals already rely on AI for diagnosis and monitoring.⁵

At the same time, analysts have warned of “potential competition bottlenecks in AI industries,” noting that the financial resources and computing power required to build advanced models place meaningful constraints on who can participate in the ecosystem.⁶ Commentators also note that “the gap between

¹ Lawrence J. Trautman and Larry D. Foster II, “Sam Altman, OpenAI, and the Importance of Corporate Governance,” *Journal of Law, Technology and the Internet* 16, no. 1 (2025).

² Kyle Wiggers, “A timeline of Sam Altman’s firing from OpenAI — and the fallout,” *TechCrunch*, 5 Jan. 2024, <https://techcrunch.com/2024/01/05/a-timeline-of-sam-altmans-firing-from-openai-and-the-fallout/>.

³ David Goldman, “How OpenAI So Royally Screwed Up the Sam Altman Firing,” *CNN* (20 Nov. 2023).

⁴ Dylan Walsh, “A New Look at the Economics of AI,” *MIT Sloan School of Management*, 21 Jan. 2025, <https://mitsloan.mit.edu/ideas-made-to-matter/a-new-look-economics-ai>.

⁵ Bernard Marr, “Mind-Blowing AI Statistics Everyone Must Know About Now in 2025,” *Forbes*, 3 June 2025, 06:37 a.m. EDT, <https://www.forbes.com/sites/bernardmarr/2025/06/03/mind-blowing-ai-statistics-everyone-must-know-about-now-in-2025/>.

⁶ Bertin Martens, “Why Artificial Intelligence Is Creating Fundamental Challenges for Competition Policy,” *Bruegel Policy Brief* (Brussels: Bruegel, July 18 2024).

technological advancements and regulatory frameworks is widening,” leaving companies themselves responsible for decisions that have broad societal consequences.⁷ This framework creates a structural tension in which corporate actors focused on growth and deployment now play a central role in shaping AI’s societal impact, raising difficult questions about who guides this technology and according to what values.

OpenAI provides an interesting case study of this dilemma. Founded with the mission to ensure that Artificial General Intelligence (AGI - AI systems that are generally smarter than humans) benefits all of humanity, the company under CEO Sam Altman has struggled to balance ambitious growth with public-interest goals.⁸ The company started as a nonprofit but later created a for-profit subsidiary (OpenAI Global, LLC) to secure the talent, compute, and capital needed to compete, placing a safety-oriented nonprofit board above a commercially driven research lab developing increasingly powerful systems.⁹ The 2023 leadership crisis, when Altman was abruptly fired and reinstated within days, revealed breakdowns in the company’s governance and decision-making.¹⁰ Studying Altman’s, and by extension OpenAI’s, decisions around leadership, partnerships, and priorities provides a lens for understanding how companies can advance AI responsibly while mitigating its risks to society.

This paper therefore asks: **How do Sam Altman’s leadership decisions at OpenAI illuminate the strengths and limitations of existing governance frameworks in AI and, more broadly, how might these insights inform the development of more effective regulatory approaches for AI/emerging technologies as a whole?**

In order to examine the governance of AI and emerging technologies, this paper focuses on OpenAI and Sam Altman’s decision-making from 2020–2023, as well as subsequent organizational developments.

⁷ Gary Drenik, “The Pitfalls of AI Self-Regulation,” *Forbes*, 22 Oct. 2024, <https://www.forbes.com/sites/garydrenik/2024/10/22/the-pitfalls-of-ai-self-regulation/>.

⁸ “Sam Altman Discusses OpenAI’s Growth Challenges and Ambitious Future in AI Innovation,” *LUNARTECH Blog*, 3 Sept. 2025, <https://www.lunartech.ai/blog/sam-altman-discusses-openai-s-growth-challenges-and-ambitious-future-in-ai-innovation>.

⁹ The Midas Project and Tech Oversight Project, “The OpenAI Files,” June 18, 2025, <https://www.openaifiles.org/>.

¹⁰ David Goldman, “How OpenAI So Royally Screwed Up the Sam Altman Firing,” *CNN* (20 Nov. 2023).

This focus is guided by both the practical limits of the query and its conceptual goals. OpenAI's hybrid non-profit/for-profit structure, market influence, and high visibility make it a unique case to observe how regulatory and governance challenges take shape within a high-stakes technological institution.

Additionally, the abrupt firing and reinstatement of Altman provides a bounded event that highlights structural tensions relevant to the research question, allowing this paper to analyze governance not as an abstraction but as a lived institutional crisis.

In conducting my analysis, I approach the OpenAI crisis as a window into the governance dilemmas that arise when frontier technologies evolve faster than the institutions meant to oversee them. Rather than reconstructing the episode chronologically, I am interested in examining *why* governance failed: to what extent did the breakdown reflect shortcomings in leadership, signaled by the board's claim that Altman was "not consistently candid," and to what extent did it stem from structural pressures embedded in OpenAI's hybrid nonprofit/for-profit design? To interpret these dynamics, I analyze how formal and informal power shifted among the CEO, board, employees, and strategic partners, noting that more than 700 employees threatened to resign and Microsoft publicly supported Altman's return, revealing the limits of the board's authority.¹¹

I then consider whether these dynamics can be extrapolated by situating the crisis within a broader historical pattern of tensions that have emerged during earlier technological transitions, analyzing whether the mechanisms that undermined oversight at OpenAI resemble those that have historically surfaced when institutional design lags behind technical and commercial momentum. Existing research on digital regulation shows that rapid technological innovation often widens the information asymmetries that place regulators at a structural disadvantage.¹² As Lancieri, Edelson, and Bechtold note, "regulatory capture is a recurring theme for digital regulations," driven less by intentional design than by the persistent imbalance

¹¹ "The Employee Letter to OpenAI's Board," *The New York Times*, 20 Nov. 2023, <https://www.nytimes.com/interactive/2023/11/20/technology/letter-to-the-open-ai-board.html>.

¹² Andrea Saltelli, Dorothy J. Dankel, Monica Di Fiore, Nina Holland, and Martin Pigeon, "Science, the endless frontier of regulatory capture," *Futures* 135 (2022): 102860, <https://doi.org/10.1016/j.futures.2021.102860>.

between institutional capacity and technological complexity.¹³ Finally, because high-stakes emerging technologies lack unified empirical frameworks for assessing governance risk, I treat the OpenAI crisis as an analytical case through which to clarify these tensions and identify principles that may inform more effective regulatory approaches.

In this first part of my analysis, I examine what the OpenAI crisis reveals about the strengths and limitations of contemporary governance frameworks for frontier AI. I consider two competing explanations for the breakdown: one locates responsibility in leadership, the other in structural failures embedded in organizational design.

One interpretation is that the crisis was precipitated by Altman himself. The board publicly stated that he “was not consistently candid in his communications with the board, hindering its ability to exercise its responsibilities,” language that signaled a breakdown of trust between the CEO and the body charged with mission oversight.¹⁴ Former board members later suggested that the decision reflected broader concerns about Altman’s transparency and alignment with the organization’s safety mandate.¹⁵ Public commentary also captured this tension; *The Economist* described Altman as “a visionary with a trustworthiness problem,” encapsulating doubts about his suitability to lead the development of increasingly powerful AI systems.¹⁶

In my view, Altman’s authority did not resemble traditional political charisma but something closer to network-based visionary influence, rooted in his long-standing position within Silicon Valley’s entrepreneurial elite. As president of Y Combinator, he was, in Paul Graham’s words, “one of those rare people who manage to be both fearsomely effective and yet fundamentally benevolent,” and his portfolio

¹³ Filippo Lancieri, Laura Edelson, and Stefan Bechtold, “AI Regulation: The Politics of Fragmentation and Regulatory Capture,” *Oxford Business Law Blog*, June 19, 2025, <https://blogs.law.ox.ac.uk/oblb/blog-post/2025/06/ai-regulation-politics-fragmentation-and-regulatory-capture>.

¹⁴ “OpenAI announces leadership transition,” OpenAI, November 17, 2023, <https://openai.com/index/openai-announces-leadership-transition/>.

¹⁵ Kit Eaton, “2 OpenAI Ex-Board Members Rip Altman’s Leadership,” *Inc.com*, May 28, 2024, <https://www.inc.com/kit-eaton/two-openai-ex-board-members-rip-altmans-leadership.html>.

¹⁶ “Sam Altman Is a Visionary with a Trustworthiness Problem,” *The Economist*, May 20, 2025, <https://www.economist.com/culture/2025/05/20/sam-altman-is-a-visionary-with-a-trustworthiness-problem>.

included incredibly successful companies such as Stripe, Reddit, and Neuralink.¹⁷ Public profiles routinely highlighted his blend of “technical acumen” and sensitivity to the “broader societal implications of AI,” portraying him as a leader whose legitimacy derived from vision, expertise, and an ability to mobilize talent and resources.¹⁸ His public rhetoric functioned as a strategic extension of this influence. In a White House appearance, he described himself as “a child of the internet revolution” and imagined that “some kid... staying up late playing with ChatGPT” might found the next great American company, intelligently positioning AI as both a patriotic project and a generational opportunity.¹⁹

This networked legitimacy translated into extraordinary organizational influence. Within forty-eight hours of Altman’s firing, more than 730 of OpenAI’s roughly 770 employees signed a letter declaring the board “incapable of overseeing OpenAI,” threatening collective resignation and moving to join a Microsoft-backed venture unless he was reinstated. The letter noted that “Microsoft has assured us that there are positions for all OpenAI employees,” demonstrating how employee loyalty and partner confidence had consolidated around Altman himself.²⁰

By contrast, the board, oriented toward caution and oversight, lacked the emotional and symbolic resonance Altman commanded. Several members were aligned with effective-altruist perspectives that emphasized caution and long-term safety, but the board as a whole had limited corporate governance experience and, following the transition to a capped-profit structure, “was entirely devoid of investor representation or anyone who could give voice to economic interests.”²¹ This left the board well positioned to evaluate AGI risks but poorly equipped to anticipate how its decision would reverberate

¹⁷ Paul Graham, “Sam Altman for President,” Y Combinator, February 21, 2014, <https://www.ycombinator.com/blog/sam-altman-for-president/>; Sissi Cao, “Sam Altman’s Startup Portfolio: 12 Companies Owned By the OpenAI CEO,” *Observer*, May 25, 2023, <https://observer.com/2023/05/openai-sam-altman-startup-portfolio/>.

¹⁸ North American Executive, “Tracing the Rise of CEO Sam Altman,” North American Executive, n.d., <https://www.northamericanexec.com/news/tracing-the-rise-of-ceo-sam-altman/>.

¹⁹ U.S. Congress. Senate. Committee on Commerce, Science, and Transportation. *OpenAI CEO Sam Altman testifies at Senate hearing on artificial intelligence | full video*. YouTube video, Posted by CBS News, May 8, 2025. <https://www.youtube.com/watch?v=pxGE41V04fs>.

²⁰ “The Employee Letter to OpenAI’s Board,” *The New York Times*.

²¹ Louis Lehot, Lyman Thai, and Patrick Daugherty, “Navigating Board Turbulence: Lessons From OpenAI,” Foley & Lardner LLP, January 9, 2024, <https://www.foley.com/insights/publications/2024/01/navigating-board-turbulence-lessons-openai/>.

across employees, investors, and key partners. Altman, by contrast, had arranged a tender offer that valued the company at an impressive “\$86 billion and provided significant liquidity for employees’ otherwise untradeable private company shares.”²² As one governance commentary observed, “positions don’t equate to power,” a dynamic clearly visible at OpenAI, where the board held constitutional authority but lacked the practical leverage needed to enforce its decision.²³ Ultimately, it is difficult for any regulatory body, formal or internal, to compete with a figure who embodies opportunity, progress, and shared purpose. From this perspective, Altman’s accumulation of informal authority may have weakened the viability of a governance model that relied on a safety-oriented board being able to assert independent judgment. A system cannot function if its formal oversight mechanisms lack the legitimacy or internal support required to intervene.

The alternative explanation, and ultimately the more compelling one, is that the crisis reflects structural tensions within OpenAI’s governance model rather than an individual failure. The organization’s nonprofit board and its for-profit subsidiary responded to fundamentally different incentives. As the OpenAI Charter makes clear, the company’s overarching purpose is to ensure that AGI “benefits all of humanity,” yet its for-profit arm operated under competitive, capital-intensive pressures and a capped-profit structure that explicitly promised “returns for our first round of investors... capped at 100x their investment.”²⁴ Such features were intended to align commercial incentives with public benefit, but governance analysts observed that the structure instead created objectives that were “not only confusing but ultimately combustible” at precisely the moment the organization scaled.²⁵

²² Lehot, Thai, and Daugherty, “Navigating Board Turbulence.”

²³ Lehot, Thai, and Daugherty, “Navigating Board Turbulence.”

²⁴ Greg Brockman, Ilya Sutskever, and OpenAI, “OpenAI LP,” OpenAI, March 11, 2019, <https://openai.com/index/openai-lp/>.

²⁵ Andrea Bonime-Blanc, “The OpenAI Governance Crisis: Early Tech Company Lessons,” *Directorship Magazine* (National Association of Corporate Directors), January 25, 2024, <https://www.nacdonline.org/all-governance/governance-resources/directorship-magazine/online-exclusives/2024/January2024/OpenAI-governance-crisis-early-tech-company-lessons/>.




Figure 1. OpenAI's structural model, illustrating the hybrid non-profit/capped-profit arrangement. (Adapted from Crouse 2025)²⁶

In this context, Altman's influence was less an aberration and more a function of the organization's fundamental reliance on him as both its experienced corporate leader and its defining technological visionary. The governance model, though imaginative in its attempt to combine public-interest oversight with startup dynamism, lacked mechanisms to mediate conflicts between these fundamentally incompatible roles. The board's action was legitimate in theory, but unenforceable in practice.

Together, these interpretations suggest that while Altman's leadership shaped the crisis's trajectory, the deeper cause of the breakdown was structural. OpenAI's governance model created a divergence between formal authority, informal influence, and organizational incentives, making effective oversight difficult

²⁶ Megan Crouse, "How OpenAI's Corporate Structure Works and Why Changing It Matters," *TechRepublic*, September 19, 2025, <https://www.techrepublic.com/article/news-openai-structure-explained/>.

even for well-intentioned actors. The board's attempt at self-regulation illustrates this tension: it acted on safety concerns, but the rapid collapse of its decision exposed the limits of safety-oriented governance when institutional capacity and incentive alignment are lacking. This insight frames the remainder of the paper, which examines whether the dynamics observed at OpenAI reflect a recurring historical pattern and what these patterns imply for the design of more effective governance systems going forward.

In the next part of my analysis, I take a comparative approach to ask whether the governance failure observed at OpenAI reflects a broader pattern in the history of transformative technologies or whether it is a singular episode shaped by the idiosyncrasies of this particular organization. This question matters because if the dynamics at OpenAI mirror patterns that have appeared in earlier technological transitions, the crisis can be understood not as an isolated anomaly but as evidence of recurring structural weaknesses that emerge when mission-oriented governance coexists with powerful commercial pressures. Conversely, if frontier AI differs in meaningful ways from past technologies, then the implications of the crisis point towards governance challenges that cannot be derived from historical analogy alone, instead requiring innovative and adaptive alternative oversight structures.

There is considerable evidence that the OpenAI crisis fits a familiar historical pattern. Mission-driven institutions often face mounting strain as they scale, confronting both the “complex governance challenges” of hybrid structures and the practical reality that “purpose-minded leaders are torn between the need to scale up profits and the zeal to make a difference.”²⁷ This tension has surfaced repeatedly in transformative technologies.

Mozilla Corporation offers a clear example. The company was created in 2005 as a “wholly owned taxable subsidiary” of the Mozilla Foundation to serve the nonprofit mission of sustaining an “open,

²⁷ Steven Rathgeb Smith, “Hybridization and nonprofit organizations: The governance challenge,” *Policy and Society* 29, no. 3 (August 2010): 223–24, <https://doi.org/10.1016/j.polsoc.2010.06.003>; Divyesh Patel, “Prioritizing Purpose Over Profit Is Key to Consistent Growth and Sustainable Profit — Here's Why,” *Entrepreneur*, May 20, 2024, <https://www.entrepreneur.com/growing-a-business/prioritizing-purpose-over-profit-is-key-to-consistent/472953>.

global internet” and promoting “shared humanity.”²⁸ Similar to OpenAI’s hybrid structure, this was intended to secure resources necessary for long-term innovation, yet the commercial pressure immediately complicated this design; within a year, 85 percent of the Mozilla Corporation’s 2006 revenue came from Google, leaving the organization financially dependent on the very platform whose dominance it aimed to counter, ultimately opposing its core mission.²⁹

The commercialization of the internet offers another point of comparison to OpenAI, revealing how technologies born from public missions and national competition can shift toward private control as they scale. ARPANET began in 1958 as a publicly funded DARPA (Defense Advanced Research Projects Agency) project, created to “help American military technology stay ahead of its enemies” and was initially used as a tool for “academic engineers and computer scientists” linking major research universities.³⁰ Its early architecture relied on “a universal language” enabling any two networked computers to communicate across technical boundaries.³¹ Yet once the network was privatized in the 1990s, governance authority shifted from public agencies to private firms, with platforms and ISPs (Internet Service Providers) determining the commercial landscape of the web.³² Public oversight lagged behind this transition, reflecting what governance scholars describe as a persistent “pacing problem,” in which new innovations “emerge unexpectedly... and forge ahead at a rapid pace compared to the slower pace at which the regulatory system progresses.”³³ The result was a form of “kludgeocracy,” a patchwork of “quick, inelegant, or makeshift solutions” adopted as regulators scrambled to keep up.³⁴ This trajectory

²⁸ Mozilla, “Mozilla Foundation: About the Mozilla Corporation,” accessed December 2, 2025, <https://www.mozilla.org/en-US/foundation/moco/>; Mozilla, “The Mozilla Manifesto Addendum: Pledge for a Healthy Internet,” accessed December 2, 2025, <https://www.mozilla.org/en-US/about/manifesto/>.

²⁹ Noam Cohen, “Will Success, or All That Money From Google, Spoil Firefox?,” *The New York Times*, November 12, 2007, <https://www.nytimes.com/2007/11/12/technology/12link.html>.

³⁰ “From ARPANET to the Internet,” *Science Museum*, 2 Nov. 2018, <https://www.sciencemuseum.org.uk/objects-and-stories/arpamet-internet>.

³¹ “From ARPANET to the Internet,” *Science Museum*.

³² Peter H. Lewis, “U.S. Begins Privatizing Internet’s Operations,” *The New York Times*, 24 Oct. 1994, <https://www.nytimes.com/1994/10/24/business/us-begins-privatizing-internet-s-operations.html>.

³³ Srivathsan Karanai Margan, “The Pacing Problem Unplugged Part 1,” *Contingencies*, 1 Jan. 2025, <https://www.actuary.org/article/the-pacing-problem-unplugged-part-1/>.

³⁴ “The Pacing Problem Unplugged Part 1,” *Contingencies*.

parallels frontier AI: a technology increasingly framed as part of a geopolitical race, yet driven in practice by private actors whose priorities shape its development long before public institutions can respond.

Biotechnology provides yet another instructive precedent. Defined as the use of living organisms “to modify or improve human health and human environment,” the field initially developed without established biosafety frameworks, prompting scientists to warn that recombinant DNA research carried “both beneficial and destructive potentials.”³⁵ As with frontier AI, the earliest risk signals came from researchers rather than policymakers.

At the 1975 Asilomar Conference, scientists concluded that recombinant DNA experiments posed potential biohazards and agreed that work should continue only under “stringent restrictions,” a voluntary system of self-regulation that “proved remarkably effective” in stabilizing early risks, much like today’s AI safety workshops and technical standards discussions.³⁶ Formal oversight followed only afterward, when the United States issued federal recombinant DNA guidelines in 1976, and more comprehensive regulatory structures developed gradually as scientific and commercial pressures grew.³⁷

Each of these cases suggest that governance failures are not unique to OpenAI but are characteristic of technologies in which innovation outpaces formal mechanisms of accountability.

At the same time, it would be reductive to treat OpenAI as simply another iteration of this pattern.

Artificial intelligence differs from earlier technologies in ways that meaningfully elevate the stakes of governance failure. As Desmarais puts it, “Unlike the Internet 1.0 revolution... ChatGPT hit the world stage all at once, growing in most global regions simultaneously.”³⁸ AI capabilities advance rapidly, driven by what analysts describe as “a staggering leap... of architectural and materials innovation that is

³⁵ Ashish Swarup Verma, Shishir Agrahari, Shruti Rastogi & Anchal Singh, “Biotechnology in the Realm of History,” *Journal of Pharmacy & Bioallied Sciences* 3, no. 3 (2011): 321–323, <https://doi.org/10.4103/0975-7406.84430>.

³⁶ Paul Berg, “Asilomar 1975: DNA Modification Secured,” *Nature* 455 (17 Sept. 2008): 290–291, <https://doi.org/10.1038/455290a>.

³⁷ Berg, “Asilomar 1975,” 290.

³⁸ Anna Desmarais, “The AI Revolution Is ‘Unprecedented’ in the Scale of Human History, New Report Finds,” *Euronews*, 6 Nov. 2025, <https://www.euronews.com/next/2025/06/11/the-ai-revolution-is-unprecedented-in-the-scale-of-human-history-new-report-finds>.

reshaping what's possible at the hardware level.”³⁹ The stakes are correspondingly higher: AI researchers warn that the technology “poses unprecedented threats,” which require immediate action. Forecasts of AGI arriving as early as 2026 underscore how little margin exists for reactive governance.⁴⁰ Furthermore, the very public nature of the OpenAI crisis, unfolding over days and under intense scrutiny, provides a uniquely early glimpse into tensions that historically only become visible after a decade or more of commercial presence.⁴¹ In this sense, AI is not exceptional because it escapes historical precedent but because it accelerates and amplifies the familiar dynamics through which governance tends to fail.

Taken together, these similarities and differences indicate that the OpenAI crisis can indeed be extrapolated, but with caution. The underlying governance failure aligns with a recurring pattern in technological history: when innovation outpaces regulation, when commercial imperatives overshadow mission driven oversight, and when institutional designs rely on self regulation without external scaffolding, structural breakdown becomes likely. Yet AI's unprecedented reach and velocity mean that societies cannot rely on reactive governance or assume that institutional corrections will arrive in time to prevent harm.

In the final part of my analysis, I consider what a more effective governance architecture for frontier AI might look like and whether existing institutions are positioned to support it. If the OpenAI crisis revealed the limits of self regulation and if historical precedent suggests that governance failures often follow predictable patterns, the question then becomes how to move beyond episodic, reactive models of oversight? For analytic clarity, I frame the debate in terms of two broad approaches. The first centers on preserving innovation and emphasizes improved internal governance, on the premise that technical

³⁹ Keumars Afifi-Sabet, “AI Is Entering an ‘Unprecedented Regime.’ Should We Stop It — and Can We — Before It Destroys Us?,” *LiveScience*, 1 Aug. 2025, <https://www.livescience.com/technology/artificial-intelligence/ai-is-entering-an-unprecedented-regime-should-we-stop-it-and-can-we-before-it-destroys-us>.

⁴⁰ “AI poses unprecedented threats. Congress must act now,” *The Guardian*, 2 Dec. 2025, <https://www.theguardian.com/commentisfree/2025/dec/02/artificial-intelligence-threats-congress>; Keumars Afifi-Sabet, “AI Is Entering an ‘Unprecedented Regime.’ Should We Stop It — and Can We — Before It Destroys Us?,” *LiveScience*, 1 Aug. 2025, <https://www.livescience.com/technology/artificial-intelligence/ai-is-entering-an-unprecedented-regime-should-we-stop-it-and-can-we-before-it-destroys-us>.

⁴¹ OpenAI, “OpenAI Progress,” <https://progress.openai.com/>.

proximity and market pressures will correct failures more effectively than early regulatory constraints. The second argues that AI, like finance or nuclear energy, poses system level risks that require stronger institutional involvement. My view is that neither model is sufficient on its own. Effective oversight will require a layered approach that integrates internal governance, sector-specific regulation, and system-level monitoring, coordinated in a way that remains adaptive as the technology evolves.

One view holds that firms developing advanced AI should regulate themselves because they possess the technical expertise and agility regulators lack, and because, as one analysis warns, “one-size-fits-all regulation will over-regulate in some instances and under-regulate in others.”⁴² According to this perspective, slowing innovation risks weakening national competitiveness in what observers describe as “a race... happening without a referee,” where rapid capability development is seen as strategically essential.⁴³ Proponents also argue that government intervention can threaten innovation and expressive freedom, since “broad regulations that limit innovation” or restrict specific AI tools risk preventing new entrants and undermining AI’s potential benefits.⁴⁴ The patchwork of state-level proposals reinforces this concern; as one policy analysis notes, an AI “pause” or fragmented state rules would create a “legal and compliance nightmare” that hampers development.⁴⁵

The second view argues that leaving frontier AI governance entirely to firms is incompatible with the technology’s scale and societal reach, since AI is “opaque to outsiders” and produces “emergent behaviour” whose outcomes “cannot be known *a priori*.”⁴⁶ In the United States, regulators have begun issuing AI guidance, but the prevailing model remains “use case and sector specific,” requiring agencies like the FDA, SEC, and NHTSA to “mold or expand the way they think about regulation” even as nobody

⁴² Tom Wheeler, “The three challenges of AI regulation,” *Brookings Institution*, 15 June 2023, <https://www.brookings.edu/articles/the-three-challenges-of-ai-regulation/>.

⁴³ Tom Wheeler, “The three challenges of AI regulation.”

⁴⁴ David Inserra, “Artificial Intelligence Regulation Threatens Free Expression,” *Cato Institute*, 16 July 2024, <https://www.cato.org/briefing-paper/artificial-intelligence-regulation-threatens-free-expression>.

⁴⁵ Nicole Alvarez & Adam Conner, “4 Reasons the Senate’s AI Pause Should Be Opposed,” *Center for American Progress*, 27 June 2025, <https://www.americanprogress.org/article/4-reasons-the-senates-ai-pause-should-be-opposed/>.

⁴⁶ Rahul Tongia, “Why Regulating AI Can Be Surprisingly Straightforward, When Teamed with Eternal Vigilance,” *World Economic Forum*, 28 May 2024, <https://www.weforum.org/stories/2024/05/why-regulating-ai-can-be-surprisingly-straightforward-providing-you-have-eternal-vigilance/>.

oversees frontier model development itself.⁴⁷ Policymakers and scholars warn that without stronger institutional involvement, AI risks replicating the policy failures of earlier tech waves; as FTC Chair Lina Khan writes, “the trajectory of the Web 2.0 era was not inevitable” and public officials “have a responsibility to ensure this hard-learned history doesn’t repeat itself.”⁴⁸ Khan also cautions that expanding AI adoption “risks further locking in the market dominance of large incumbent technology firms,” since a handful of companies control the data, compute, and cloud infrastructure on which the entire ecosystem depends.⁴⁹

Given the scope of this paper, it is not possible to design a full institutional blueprint for frontier AI governance, but it is possible to identify the structural features a credible framework would require. Research on emerging technologies emphasizes that regulators face a profound pacing problem in which technological innovations “forge ahead at a rapid pace compared to the slower pace at which the regulatory system progresses,” leaving regulators with “only two options: control the pace of technological innovations or reinvigorate regulatory entities and experiment with different rulemaking approaches,” and the former is “nearly impossible in the current technological landscape.”⁵⁰ A growing body of work therefore recommends governance models that begin early and evolve iteratively. Anticipatory governance stresses early monitoring and cross-sector coordination, while planned adaptive regulation requires a “prior commitment... to periodic re-evaluation and potential revision” and institutional mechanisms “to monitor and synthesize new information” as technologies develop.⁵¹ Soft-law tools such as voluntary codes, standards, certifications, and collaborative frameworks, can

⁴⁷ The Wall Street Journal, *SPECIAL WSJ Tech Live: The Man Leading Trump’s AI Charge Against China*, interview with Michael Kratsios, The Journal Podcast, 28 Nov. 2025, <https://www.wsj.com/video/series/wsj-the-future-of-everything/special-wsj-tech-live-the-man-leading-trumps-ai-charge-against-china-the-journal-podcast/7056AED7-E68D-49F3-9AD9-3BB2B27C0B2D>.

⁴⁸ Lina M. Khan, “We Must Regulate A.I. Here’s How,” *The New York Times*, 3 May 2023, <https://www.nytimes.com/2023/05/03/opinion/ai-lina-khan-ftc-technology.html>.

⁴⁹ Khan, “We Must Regulate A.I. Here’s How.”

⁵⁰ Srivathsan Karanai Margan, “The Pacing Problem Unplugged Part 2,” *Contingencies*, 18 Apr. 2025, <https://www.actuary.org/article/the-pacing-problem-unplugged-part-2/>.

⁵¹ Jeremy A. Carp, “Autonomous Vehicles: Problems and Principles for Future Regulation,” *University of Pennsylvania Journal of Law & Public Affairs* 4, no. 1 (Nov. 2018), <https://scholarship.law.upenn.edu/cgi/viewcontent.cgi?article=1048&context=jlpa>.

complement these approaches by allowing norms to adjust faster than statutory regimes and by distributing oversight across multiple stakeholders.⁵²

To address risks that cut across sectors and exceed the capacity of any single firm or regulator, a governance architecture for frontier AI must also include system-level oversight. Scholars describe this landscape as a “regime complex,” a network of overlapping and sometimes conflicting regulatory regimes that require coordination rather than replacement.⁵³ Models from other sectors suggest that this role could be housed within a specialized unit in an existing institution or an independent review body modeled on transportation safety boards or financial stability councils.

It is important to acknowledge that no governance structure will eliminate uncertainty or guarantee safe outcomes. There is no neutral framing in AI governance. Decisions about what counts as risk and what values should guide oversight inevitably reflect political and ethical judgment. At the same time, the OpenAI crisis in conjunction with historical precedent shows that continuing to rely on structures that distribute authority, power, and incentives in conflicting directions is unsustainable. Governance must be iterative and adaptive rather than assuming that any first attempt will be definitive. The aim is not to halt innovation or impose uniform controls, but to avoid repeating known governance failures at a scale where consequences may be far more difficult to reverse.

This paper has shown that Sam Altman’s leadership decisions illuminate both the possibilities and the limits of self regulation in frontier AI. The OpenAI crisis revealed that while internal governance can raise concerns, it cannot enforce them when authority, influence, and incentives diverge. Historical precedent suggests that these tensions are not unique, but AI heightens their stakes and compresses the time available for correction. The analysis therefore supports the need for a more robust, layered governance architecture that combines internal reform, sector specific regulation, and system level oversight. The crisis at OpenAI serves as an early reminder that effective governance in emerging

⁵² Margan, “The Pacing Problem Unplugged Part 2.”

⁵³ Carp, “Autonomous Vehicles.”

technologies will depend not on existing structures correcting themselves, but on the future of intentional institutional design. What remains is the harder question: how societies can build institutions capable not only of responding to today's failures, but of anticipating tomorrow's risks before they materialize.

Works Cited

- "A Brief History of the Internet." Stanford Computer Science.
<https://cs.stanford.edu/people/eroberts/courses/soco/projects/distributed-computing/html/history.html>.
- Afifi-Sabet, Keumars. "AI Is Entering an 'Unprecedented Regime.' Should We Stop It - and Can We Before It Destroys Us?" *LiveScience*. 1 Aug. 2025.
<https://www.livescience.com/technology/artificial-intelligence/ai-is-entering-an-unprecedented-regime-should-we-stop-it-and-can-we-before-it-destroys-us>.
- AI.gov. "Homepage." Accessed December 2, 2025. <https://www.ai.gov/>.
- "AI poses unprecedented threats. Congress must act now." *The Guardian*. 2 Dec. 2025.
<https://www.theguardian.com/commentisfree/2025/dec/02/artificial-intelligence-threats-congress>.
- Altman, Samuel. Testimony Before the Senate Judiciary Subcommittee on Privacy, Technology and the Law. U.S. Senate, May 16, 2023.
<https://www.judiciary.senate.gov/imo/media/doc/2023-05-16%20-%20Bio%20%26%20Testimony%20-%20Altman.pdf>.
- Alvarez, Nicole, and Adam Conner. "4 Reasons the Senate's AI Pause Should Be Opposed." Center for American Progress. 27 June 2025.
<https://www.americanprogress.org/article/4-reasons-the-senates-ai-pause-should-be-opposed/>.
- "The AI Revolution Is 'Unprecedented' in the Scale of Human History, New Report Finds." *Euronews*. 6 Nov. 2025.
<https://www.euronews.com/next/2025/06/11/the-ai-revolution-is-unprecedented-in-the-scale-of-human-history-new-report-finds>.
- Benkler, Yochai. "The Role of Technology in Political Economy: Part 1." Law and Political Economy Project, July 25, 2018.
<https://lpeproject.org/blog/the-role-of-technology-in-political-economy-part-1/>.
- Berg, Paul. "Asilomar 1975: DNA Modification Secured." *Nature* 455 (17 Sept. 2008): 290-291.
<https://doi.org/10.1038/455290a>.
- Bonime-Blanc, Andrea. "The OpenAI Governance Crisis: Early Tech Company Lessons." *Directorship Magazine* (National Association of Corporate Directors), January 25, 2024.
<https://www.nacdonline.org/all-governance/governance-resources/directorship-magazine/online-exclusives/2024/January2024/OpenAI-governance-crisis-early-tech-company-lessons/>.
- Bouttell, Laura. "What Leadership Style Is Sam Altman: Transformational Visionary." Quarterdeck, 1 June 2025. <https://quarterdeck.co.uk/articles/what-leadership-style-is-sam-altman>.
- Bredbenner, Melissa. The Regulatory Review. "Sunday Spotlight: Risk Identification in AI Regulation." October 8, 2023.
<https://www.theregreview.org/2023/10/08/sunday-spotlight-risk-identification-in-ai-regulation/>.
- Brockman, Greg, Ilya Sutskever, and OpenAI. "OpenAI LP." OpenAI, March 11, 2019.
<https://openai.com/index/openai-lp/>.

- Buchholz, Katharina. "The Extreme Cost of Training AI Models." *Forbes*, 23 Aug. 2024, <https://www.forbes.com/sites/katharinabuchholz/2024/08/23/the-extreme-cost-of-training-ai-models/>.
- Bytebridge. "History of OpenAI Founders, Early Contributors, and Investors." Medium, 2024, <https://bytebridge.medium.com/history-of-openai-founders-early-contributors-and-investors-6845e3bc2be4>.
- Cao, Sissi. "Sam Altman's Startup Portfolio: 12 Companies Owned By the OpenAI CEO." *Observer*, May 25, 2023. <https://observer.com/2023/05/openai-sam-altman-startup-portfolio/>.
- Carp, Jeremy A. "Autonomous Vehicles: Problems and Principles for Future Regulation." *University of Pennsylvania Journal of Law & Public Affairs* 4, no. 1 (Nov. 2018). <https://scholarship.law.upenn.edu/cgi/viewcontent.cgi?article=1048&context=jlpa>.
- Chamath. "A Short History of OpenAI." Substack, 2024, <https://chamath.substack.com/p/a-short-history-of-openai>.
- Cohen, Noam. "Will Success, or All That Money From Google, Spoil Firefox?" *The New York Times*, November 12, 2007. <https://www.nytimes.com/2007/11/12/technology/12link.html>.
- Cohen, Noam. "Why Has Google Spent a Half-Billion Dollars on Firefox?" *Bloomberg*, 5 May 2023. <https://www.bloomberg.com/news/newsletters/2023-05-05/why-google-keeps-paying-mozilla-s-firefox-even-as-chrome-dominates>.
- Cowhey, Peter. *The International Political Economy of the Regulation of Digital Technology*. UC Institute on Global Conflict and Cooperation, 2 June 2025, <https://ucigcc.org/publication/the-international-political-economy-of-the-regulation-of-digital-technology/>.
- Crouse, Megan. "How OpenAI's Corporate Structure Works and Why Changing It Matters." TechRepublic, September 19, 2025. <https://www.techrepublic.com/article/news-openai-structure-explained/>.
- Drenik, Gary. "The Pitfalls of AI Self-Regulation." *Forbes*, 22 Oct. 2024, <https://www.forbes.com/sites/garydrenik/2024/10/22/the-pitfalls-of-ai-self-regulation/>.
- Eaton, Kit. "2 OpenAI Ex-Board Members Rip Altman's Leadership." *Inc.com*. May 28, 2024. <https://www.inc.com/kit-eaton/two-openai-ex-board-members-rip-altmans-leadership.html>.
- Field, Hayden. "Former OpenAI Board Member Explains Why CEO Sam Altman Was Fired." *CNBC*, 29 May 2024, <https://www.cnbc.com/2024/05/29/former-openai-board-member-explains-why-ceo-sam-altman-was-fired.html>.
- Financial Wise. "Elon Musk Explains Why Sam Altman Was Fired." YouTube video, 2024. <https://www.youtube.com/watch?v=-AB7b-XGaCU>.
- "From ARPANET to the Internet." Science Museum. 2 Nov. 2018. <https://www.sciencemuseum.org.uk/objects-and-stories/arpamet-internet>.

- Goldman, David. "How OpenAI So Royally Screwed Up the Sam Altman Firing." CNN, 20 Nov. 2023, www.cnn.com/2023/11/19/tech/sam-altman-open-ai-firing-board. Accessed 26 Oct. 2025.
- Gracias, Sofia. *Comparing the EU AI Act to Proposed AI-Related Legislation in the US*. The University of Chicago Business Law Review Online Edition, 2024, <https://businesslawreview.uchicago.edu/online-archive/comparing-eu-ai-act-proposed-ai-related-legislation-us>.
- Graham, Paul. "Sam Altman for President." Y Combinator, February 21, 2014. <https://www.ycombinator.com/blog/sam-altman-for-president/>.
- Guha, Neel, Christie M. Lawrence, Lindsey A. Gailmard, Kit T. Rodolfa, Faiz Surani, Rishi Bommasani, Inioluwa Deborah Raji, Mariano-Florentino Cuéllar, Colleen Honigsberg, Percy Liang & Daniel E. Ho. "AI Regulation Has Its Own Alignment Problem: The Technical and Institutional Feasibility of Disclosure, Registration, Licensing, and Auditing." *George Washington Law Review* 92, no. 6 (Dec. 2024): 1473-1557. <https://dho.stanford.edu/wp-content/uploads/AI%20Regulation.pdf>.
- "History of Biotechnology." Lone Star College, <https://www.lonestar.edu/history-of-biotechnology.htm>.
- Huang, Han, and Shumin Wu. "The Impact of Transformational Leadership on Organizational Performance: Evidence from Service Organizations." *Journal of Service Science and Management*, vol. 13, no. 3, 18 May 2020, <https://www.scirp.org/reference/referencespapers?referenceid=2737955>.
- Huang, Lilin, Ashley Grice, Wanying Zhu, and Whitaker Rapp. "The Purpose-Driven Digital Transformation." BCG Publications, 1 Apr. 2024. <https://www.bcg.com/publications/2024/reaching-purpose-driven-digital-transformation-success>.
- Hughes, Susan A. "What Does a Tech Oligarchy Look Like?" Harvard Kennedy School, 17 June 2025, <https://www.hks.harvard.edu/faculty-research/policy-topics/science-technology-data/what-does-tech-oligarchy-look>.
- Inserra, David. "Artificial Intelligence Regulation Threatens Free Expression." Cato Institute. 16 July 2024. <https://www.cato.org/briefing-paper/artificial-intelligence-regulation-threatens-free-expression>.
- ISSARICE. "Timeline of OpenAI." Timelines, 2024, https://timelines.issarice.com/wiki/Timeline_of_OpenAI.
- Janeway, William H. "In Search of the AI Bubble's Economic Fundamentals." Project Syndicate. 7 Nov. 2025. <https://www.billjaneway.com/in-search-of-the-ai-bubbles-economic-fundamentals>.
- Janeway, William H. "The Political Economy of Technology." Project Syndicate, 25 Aug. 2023, <https://www.project-syndicate.org/onpoint/technological-progress-by-william-h-janeway-2023-08>.
- Kambur, Hüseyin, and Ayhan Dolunay. "A Research on Copyright Issues Impacting Artists' Emotional States in the Framework of Artificial Intelligence." *Frontiers in Psychology*, vol. 15, 2024, 1409646, <https://doi.org/10.3389/fpsyg.2024.1409646>.

- Kang, Cecilia, Tripp Mickle, Ryan Mac, David Yaffe-Bellany, and Theodore Schleifer. "Silicon Valley's Man in the White House Is Benefiting Himself and His Friends." *The New York Times*. 30 Nov. 2025. <https://www.nytimes.com/2025/11/30/technology/david-sacks-white-house-profits.html>.
- Khan, Lina M. "We Must Regulate A.I. Here's How." *The New York Times*. 3 May 2023. <https://www.nytimes.com/2023/05/03/opinion/ai-lina-khan-ftc-technology.html>.
- Krishna, Abhinav C. S. "The Evolution of the Internet, The World Wide Web, and Search From ARPANET to AI Agents." SEO Central, 7 Sept. 2025. <https://www.theseocentral.com/the-evolution-of-the-internet-the-world-wide-web-and-search-from-arpnet-to-ai-agents>.
- Lancieri, Filippo, Laura Edelson, and Stefan Bechtold. "AI Regulation: The Politics of Fragmentation and Regulatory Capture." Oxford Business Law Blog, June 19, 2025. <https://blogs.law.ox.ac.uk/oblb/blog-post/2025/06/ai-regulation-politics-fragmentation-and-regulatory-capture>.
- Lehot, Louis, Lyman Thai, and Patrick Daugherty. "Navigating Board Turbulence: Lessons From OpenAI." Foley & Lardner LLP, January 9, 2024. <https://www.foley.com/insights/publications/2024/01/navigating-board-turbulence-lessons-openai/>.
- Lewis, Peter H. "U.S. Begins Privatizing Internet's Operations." *The New York Times*. 24 Oct. 1994. <https://www.nytimes.com/1994/10/24/business/us-begins-privatizing-internet-s-operations.html>.
- Margan, Srivathsan Karanai. "The Pacing Problem Unplugged Part 1." *Contingencies*. 1 Jan. 2025. <https://www.actuary.org/article/the-pacing-problem-unplugged-part-1/>.
- Marr, Bernard. "Mind-Blowing AI Statistics Everyone Must Know About Now in 2025." *Forbes*, 3 June 2025, <https://www.forbes.com/sites/bernardmarr/2025/06/03/mind-blowing-ai-statistics-everyone-must-know-about-now-in-2025/>.
- Martens, Bertin. "Why Artificial Intelligence Is Creating Fundamental Challenges for Competition Policy." Bruegel Policy Brief (Brussels: Bruegel, July 18 2024). <https://www.bruegel.org/policy-brief/why-artificial-intelligence-creating-fundamental-challenges-competition-policy>.
- The Midas Project and Tech Oversight Project. "The OpenAI Files." June 18, 2025. <https://www.openaifiles.org/>.
- Mozilla. "Mozilla Foundation: About the Mozilla Corporation." Accessed December 2, 2025. <https://www.mozilla.org/en-US/foundation/moco/>.
- Mozilla. "The Mozilla Manifesto Addendum: Pledge for a Healthy Internet." Accessed December 2, 2025. <https://www.mozilla.org/en-US/about/manifesto/>.
- Nigam, Gaurav. "Leadership Lessons from Sam Altman's OpenAI Saga." Medium, 30 Nov. 2023. <https://medium.com/@gauravnigam/leadership-lessons-from-sam-altmans-openai-saga-80f98509e1b0>.

NHLOCAL. AI Timeline. GitHub, 2025, <https://nhlocal.github.io/AiTimeline/>.

North American Executive. "Tracing the Rise of CEO Sam Altman." North American Executive, <https://www.northamericanexec.com/news/tracing-the-rise-of-ceo-sam-altman/>.

OECD. Tackling Policy Challenges Through Public Sector Innovation: A Strategic Portfolio Approach. OECD Public Governance Reviews. Paris: OECD Publishing, 2022. https://www.oecd.org/content/dam/oecd/en/publications/reports/2022/11/tackling-policy-challenges-through-public-sector-innovation_49144212/052b06b7-en.pdf.

OpenAI. "Elon Musk Wanted an OpenAI for-Profit." 13 Dec. 2024, <https://openai.com/index/elon-musk-wanted-an-openai-for-profit/>.

"OpenAI announces leadership transition." OpenAI, November 17, 2023. <https://openai.com/index/openai-announces-leadership-transition/>.

Pearson, Catherine. "A Smartphone Before Age 12 Could Carry Health Risks, Study Says." The New York Times. 1 Dec. 2025. <https://www.nytimes.com/2025/12/01/well/family/early-smartphone-ownership-study.html>.

Patel, Divyesh. "Prioritizing Purpose Over Profit Is Key to Consistent Growth and Sustainable Profit Here's Why." *Entrepreneur*, May 20, 2024. <https://www.entrepreneur.com/growing-a-business/prioritizing-purpose-over-profit-is-key-to-consistent/472953>.

ProQuest. "Analysis of AI Industry Funding Trends." Blogs, Podcasts, & Websites, 2024, <https://www.proquest.com/docview/2893405418/F3C36CA0AF564E62PQ/17accountid=15054&sourcetype=Blogs,%20Podcasts,%20%20Websites>.

ProQuest. "OpenAI Leadership and Organizational Decisions." Newspapers, 2024, <https://www.proquest.com/docview/2892378652?pq-origsite=primo&accountid=15054&sourcetype=Newspapers>.

Saltelli, Andrea, Dorothy J. Dankel, Monica Di Fiore, Nina Holland, and Martin Pigeon. "Science, the endless frontier of regulatory capture." *Futures* 135 (2022): 102860. <https://doi.org/10.1016/j.futures.2021.102860>.

"Sam Altman Discusses OpenAI's Growth Challenges and Ambitious Future in AI Innovation." LUNARTECH Blog, 3 Sept. 2025, <https://www.lunartech.ai/blog/sam-altman-discusses-openai-s-growth-challenges-and-ambitious-future-in-ai-innovation>.

"Sam Altman Is a Visionary with a Trustworthiness Problem." *The Economist*, May 20, 2025. <https://www.economist.com/culture/2025/05/20/sam-altman-is-a-visionary-with-a-trustworthiness-problem>.

Smith, Steven Rathgeb. "Hybridization and nonprofit organizations: The governance challenge." *Policy and Society* 29, no. 3 (August 2010): 219-29. <https://doi.org/10.1016/j.polsoc.2010.06.003>.

- Stilgoe, Jack, Richard Owen, and Phil Macnaghten. "Developing a Framework for Responsible Innovation." *Research Policy*, vol. 42, no. 9, 2013, pp. 1568-1580, <https://doi.org/10.1016/j.respol.2013.05.008>.
- Stegman, Chris. "What Plato Has Taught Me About Artificial Intelligence." Medium, 2024, <https://medium.com/@cstegman/what-plato-has-taught-me-about-artificial-intelligence-3a17f31a5805>.
- "The 2025 AI Index Report." Stanford Institute for Human-Centered Artificial Intelligence, 2025, <https://hai.stanford.edu/ai-index/2025-ai-index-report>.
- "The Employee Letter to OpenAI's Board," *The New York Times*, 20 Nov. 2023, <https://www.nytimes.com/interactive/2023/11/20/technology/letter-to-the-open-ai-board.html>.
- The Wall Street Journal*. SPECIAL WSJ Tech Live: The Man Leading Trump's AI Charge Against China. Interview with Michael Kratsios. *The Journal Podcast*. 28 Nov. 2025. <https://www.wsj.com/video/series/wsj-the-future-of-everything/special-wsj-tech-live-the-man-leading-trumps-ai-charge-against-china-the-journal-podcast/7056AED7-E68D-49F3-9AD9-3BB2B27C0B2D>.
- Tongia, Rahul. "Why *Regulating AI Can Be Surprisingly Straightforward*, When Teamed with Eternal Vigilance." World Economic Forum. 28 May 2024. <https://www.weforum.org/stories/2024/05/why-regulating-ai-can-be-surprisingly-straightforward-providing-you-have-eternal-vigilance/>.
- Trautman, Lawrence J., and Larry D. Foster II. "Sam Altman, OpenAI, and the Importance of Corporate Governance." *Journal of Law, Technology and the Internet*, vol. 16, no. 1, annual 2025, pp. 133+. Gale Academic OneFile, e354aa. Accessed 26 Oct. 2025.
- Ulanoff, Lance. "Trump's Genesis Mission AI Plan Smacks of ARPANET but Smells Like Skynet." Yahoo News, 25 Nov. 2025. <https://www.yahoo.com/news/articles/trumps-genesis-mission-ai-plan-175625145.html>.
- U.S. Congress, Senate, Committee on Commerce, Science, and Transportation, OpenAI CEO Sam Altman testifies at Senate hearing on artificial intelligence | full video, YouTube video, posted by CBS News, May 8, 2025, <https://www.youtube.com/watch?v=pxGE41V04fs>.
- U.S. Environmental Protection Agency (EPA). "Regulation of Biotechnology under TSCA and FIFRA." Last modified November 10, 2025. <https://www.epa.gov/regulation-biotechnology-under-tsca-and-fifra>.
- Tushman, Michael L., and Charles A. O'Reilly III. "Ambidextrous Organizations: Managing Evolutionary and Revolutionary Change." *California Management Review*, vol. 38, no. 4, Summer 1996, pp. 8-30.
- Valz, Duane. "The Corporate Governance Failure at the Heart of Sam Altman's Ouster from Open AI." Medium (blog), November 20, 2023. <https://medium.com/@duanevalz/the-corporate-governance-failure-at-the-heart-of-sam-altmans-ouster-from-openai-cbdede5c7d21>.

- Verma, Ashish Swarup, Shishir Agrahari, Shruti Rastogi & Anchal Singh. "Biotechnology in the Realm of History." *Journal of Pharmacy & Bioallied Sciences* 3, no. 3 (2011): 321-323.
<https://doi.org/10.4103/0975-7406.84430>.
- Walsh, Dylan. "A New Look at the Economics of AI." MIT Sloan School of Management, 21 Jan. 2025,
<https://mitsloan.mit.edu/ideas-made-to-matter/a-new-look-economics-ai>.
- Wheeler, Kitty. AI Magazine. "OpenAI vs Anthropic: The Results of the AI Safety Test." September 1, 2025. <https://aimagazine.com/news/openai-vs-anthropic-the-results-of-the-ai-safety-test>.
- Wheeler, Tom. "The three challenges of AI regulation." Brookings Institution. 15 June 2023.
<https://www.brookings.edu/articles/the-three-challenges-of-ai-regulation/>.
- Wiggers, Kyle. "A timeline of Sam Altman's firing from OpenAI - and the fallout." TechCrunch. 5 Jan. 2024.
<https://techcrunch.com/2024/01/05/a-timeline-of-sam-altmans-firing-from-openai-and-the-fallout/>.
- Yale SOM Insights. "This Is How the AI Bubble Bursts." Yale School of Management, 2024,
<https://insights.som.yale.edu/insights/this-is-how-the-ai-bubble-bursts>.
- YouTube. "Sam Altman Interview - Capabilities and Industry Perspective." YouTube, uploaded by TED, 2024, https://www.youtube.com/watch?v=5MWT_doo68k.