What actually determines AI's impact on humanity?

Incentives, value networks, and the forces shaping AI's future

February 2026

BY THOMAS ARNETT

Christensen Institute

TABLE OF CONTENTS

EXECUTIVE SUMMARY

Will artificial intelligence be good or bad for humanity? It's one of the most common—and most misleading—questions of the AI age.

The future of AI won't be determined primarily by how powerful the technology becomes, or by what company leaders say they intend. It'll be determined by incentives: who funds AI companies, who their customers are, how competition works, and which trade-offs organizations are rewarded—or punished—for making.

This report applies Clayton Christensen's theory of value networks to today's leading AI labs. The analysis maps how capital markets, revenue models, governance structures, competitive pressures, and regulation shape the priorities of companies such as OpenAI, Anthropic, Google, Meta, and xAI.

The key finding is counterintuitive but essential: many of the risks people fear most from AI don't stem from negligence or malicious intent. They arise when companies behave rationally inside systems that reward speed, scale, and dominance over caution and long-term alignment.

Rather than asking whether AI will "save" or "destroy" humanity, this report argues we should ask a more grounded question: **what forces are steering its development right now, and how might those forces change?** Understanding those invisible pressures is the most reliable way to anticipate where AI is actually headed.

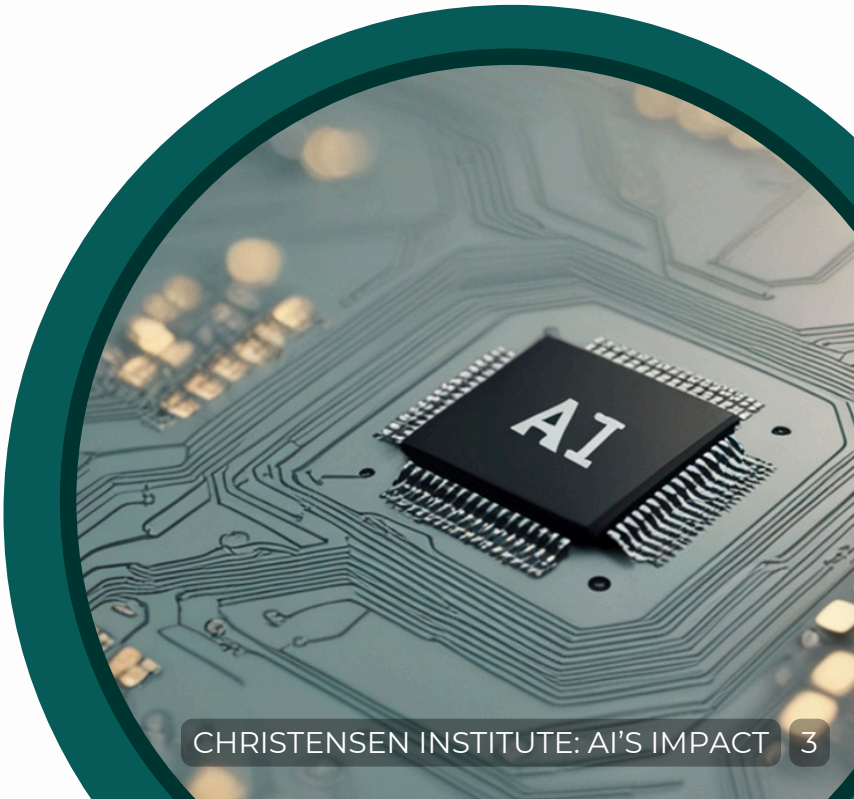

INTRODUCTION

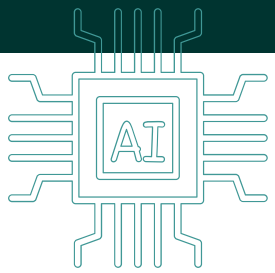

As this paper goes to publication, the AI landscape is evolving rapidly: AI companies continue to release new models and features, broker new partnerships, and secure new investments. Given this reality, some details in the analysis presented in this paper may be incomplete or may have changed since publication. But the principles guiding this analysis—understanding how value networks shape organizational priorities—remain essential for making sense of where AI is heading. The framework offered here provides a lens for evaluating AI companies and their trajectories, even as specific circumstances change.

Will artificial intelligence be good or bad for humanity? It's the question dominating headlines, policy debates, and dinner table conversations. Some envision AI curing cancer, solving climate change, and personalizing education to maximize every student's potential. Others go further, envisioning a future where AI rapidly accelerates scientific breakthroughs that improve the quality of life across the globe, while making goods and services abundantly available, thereby hastening a heretofore unimaginable era of universal prosperity and wellbeing.[1]

But at this stage, our AI-powered future is not unambiguously good for humanity. For every utopian prediction, there are reasons for concern. Are we headed toward a future of mass unemployment, inescapable surveillance, AI-instigated psychosis, and mass displacement of human relationships by AI companions?[2] Or will our worst fears be realized when artificial superintelligence arrives and exterminates humanity?[3] Amidst this range of possibilities, it's understandable why polling shows people tend to be more pessimistic than optimistic about our AI future.[4]

Yet what if we're asking the wrong question—or at least, framing it the wrong way?

Whether AI helps or harms humanity won't be determined solely by the technology's capabilities (those certainly matter) but by the value networks of the companies building it. In other words, if you want to know where AI is heading, don't just look at the technology demos or the mission statements. Look at who funds these companies, who their customers are, what suppliers and partners they depend on, what governance structures mediate between the organization and its stakeholders, and what regulatory pressures they face. These forces—what Clayton Christensen called value networks—shape what organizations prioritize and what trade-offs they make.

If you want to know where AI is heading, don't just look at the technology demos or the mission statements. Look at funding, customers, suppliers, partners, governance structures, stakeholders, and regulatory pressures.

Value networks shape what organizations prioritize, even when they have the capability to do something different.

VALUE NETWORKS:
The invisible forces that shape corporate priorities

The concept of value networks comes from Clayton Christensen's work on why established companies consistently fail to prioritize Disruptive Innovation.[5] A value network is the ecosystem of external pressures on an organization: the customers you serve, the investors who fund you, the suppliers and partners in your value chain, the competitors you respond to, and the regulators that constrain you. These forces shape what an organization prioritizes, what problems it solves, and what trade-offs it makes.

Take a classic example: Digital Equipment Corporation (DEC) in the 1980s.[6] DEC built minicomputers for departments within large organizations—engineering teams running complex simulations, research labs doing scientific computing, business units managing operations. They were extraordinarily successful at it. When personal computers emerged, DEC's engineers could have built them. They had the needed technical capabilities. But DEC's value network made it impossible. Their customers (corporate departments) wanted powerful, reliable machines for professional applications, not cheap desktop computers. Investors expected the high margins from six-figure deals in an established market, not the razor-thin profits on PCs sold into an unquantified mass market. DEC's sales force was structured around consultative enterprise sales rather than high-volume retail transactions. Every part of DEC's value network created pressure to prioritize improving minicomputers rather than pursuing personal computers. Even as PCs became the future of computing, DEC's value network kept them locked in the past.

In short, value networks shape what organizations prioritize, even when they have the capability to do something different.

MAPPING AI COMPANY VALUE NETWORKS

Many public discussions about how to manage the benefits, harms, and risks of AI focus on principles, frameworks, and external evaluations intended to guide responsible AI development. These efforts may shape norms and signal expectations, but on their own, they rarely determine behavior. They begin to matter only when they are embedded in value networks—when capital, customers, regulation, or liability depend on compliance. Absent those forces, such frameworks tend to reflect stated priorities more than they enforce real trade-offs.

To understand where frontier AI is heading, we have to look at the companies building it and the forces that govern their decisions. What follows is a map of the value networks surrounding each major AI lab, and how those networks shape the priorities that will ultimately steer the technology.

OpenAI

The history: OpenAI started in 2015 as a nonprofit with an explicit mission: to ensure that artificial general intelligence benefits all of humanity.[7] The nonprofit structure was meant to insulate the company from commercial pressures.[8]

Yet by 2019, that structure proved insufficient to fund the enormous computing costs required to keep pace with advances in large-scale AI.

OpenAI restructured itself by creating a for-profit arm that could raise billions in outside capital, while still placing limits on investor returns.[9] Microsoft invested heavily, with approximately $13 billion total.[10] The deal included an agreement that OpenAI's technology would run exclusively on Microsoft's Azure cloud infrastructure.[11] Microsoft expects substantial returns: 49% of OpenAI's profits until recouping their investment, plus billions in ongoing Azure cloud revenue as OpenAI's computing needs grow.[12] Microsoft also gained rights to commercialize OpenAI's technology across its products, including Bing, Office 365, GitHub, and Azure services.[13] The relationship made OpenAI dependent on Microsoft for both capital and infrastructure.[14]

In late 2022, OpenAI released its chatbot, ChatGPT, allowing anyone curious about the technology to try it out online for free.[15] ChatGPT shocked the world with its ability to conduct human-like conversations. It soon set a new tech company record for user acquisition, attracting 100 million users in less than two months.[16] Building off the excitement generated by ChatGPT, in early 2023, OpenAI attracted investment from major venture capital firms such as Sequoia Capital, Andreessen Horowitz, Thrive Capital, K2 Global, and Founders Fund to join its investor table.[17]

The for-profit subsidiary structure also unlocked a key talent-acquisition strategy for OpenAI: using stock options to attract and motivate the world's best AI research and development talent.[18] Stock options carry the promise that if the venture succeeds, early employees who stay with the company until their options vest can become enormously wealthy.

Then in fall 2023, the nonprofit board attempted to fire CEO Sam Altman, citing concerns about the accelerating pace of development and commercialization.[19] Within days, overwhelming pressure from investors, employees, and Microsoft forced the board to reverse course and put Altman back at the helm.[20] The nonprofit governance structure, created to prioritize the public-benefit mission, couldn't withstand pressure from other parts of the value network.

In October 2025, OpenAI completed a major restructuring that reorganized the company as a public benefit corporation, with the nonprofit (now called OpenAI Foundation) retaining oversight.[21] Under this new structure, Microsoft now holds a 27% stake, while the nonprofit foundation retained 26%, yet maintains control of the board.[22]

Following the restructuring, Microsoft's exclusive cloud provider status ended, and OpenAI signed massive infrastructure deals with multiple cloud providers.[23] In November 2025, OpenAI signed a $38 billion, seven-year deal with Amazon Web Services.[24] It also signed major agreements with Oracle (reportedly $300 billion over five years) and Google Cloud.[25] Combined with an additional $250 billion agreement to use Microsoft's Azure cloud infrastructure, OpenAI has committed over $600 billion to cloud providers across multiple vendors.[26]

Earlier reporting estimated that OpenAI incurred roughly $9 billion in losses in 2025.[27] More recent disclosures, however, suggest the company has since entered a new phase of scale: in early 2026, OpenAI reported an annualized revenue run rate exceeding $20 billion, alongside an estimated annual cash burn approaching $17 billion.[28] Together, these figures indicate rapid growth paired with persistently extraordinary costs, rather than sustained profitability. Latest estimates also place the value of the company at around $500 billion to $830 billion.[29]

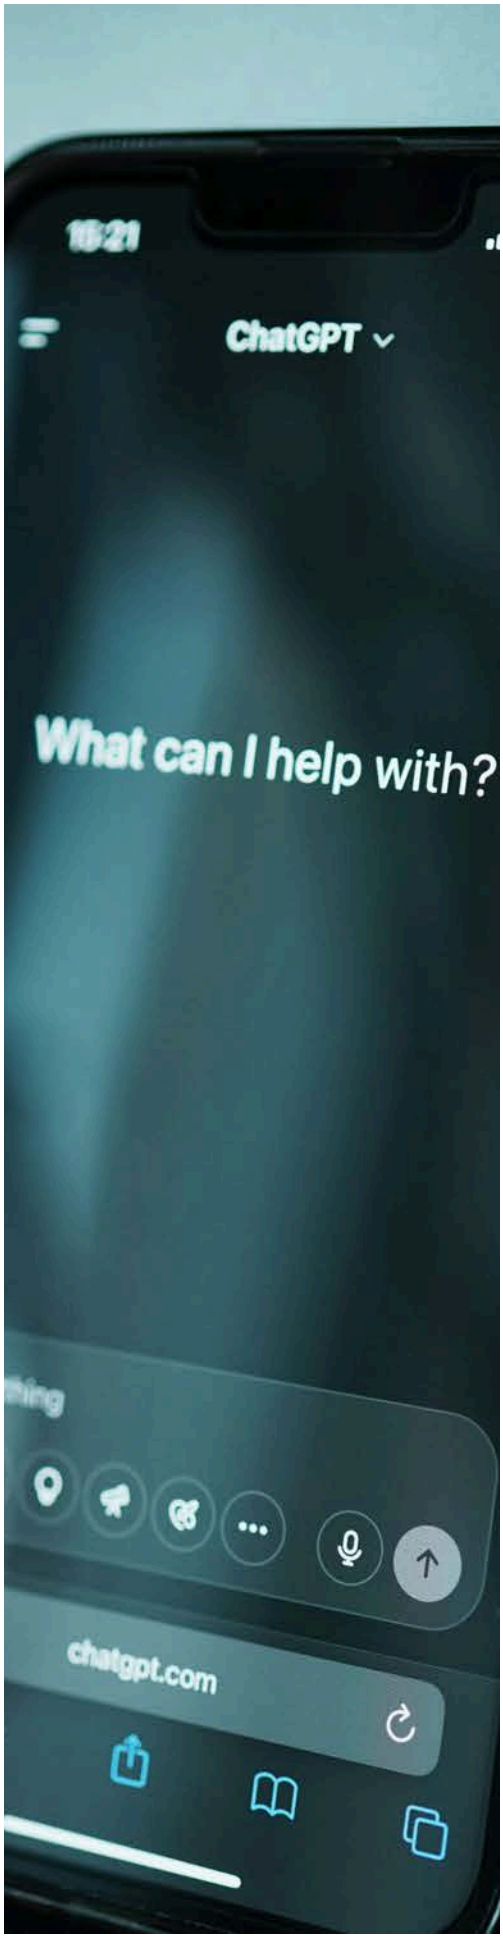

The 2023 board crisis suggests that when safety, ethics, or any other concerns conflict with commercial pressures, the profit-oriented part of the value network can overwhelm governance protections designed to balance other priorities.

The value network: The multi-cloud strategy is significant for understanding OpenAI's value network. They're no longer solely dependent on Microsoft for infrastructure, reducing single-vendor lock-in. But it also means OpenAI now has four major tech giants—Microsoft, Amazon, Google, and Oracle—all with substantial financial interests in OpenAI's success and all competing for OpenAI's infrastructure spending. These are new pressures in the value network.

OpenAI's customer base is heavily weighted toward consumers. Consumer subscriptions (ChatGPT Plus and Pro) are estimated to represent roughly 60–85% of its revenue, while revenue from APIs and enterprise subscriptions represents only 15–40%, according to various estimates. These revenue percentages hold true despite the fact that OpenAI has seen a recent surge in enterprise usage, while very few of its consumer users pay for subscriptions.

Analysts project that OpenAI's cumulative negative free cash flow will soar into the tens of billions before it becomes profitable. Losses on such a large scale indicate that investors are making enormous bets justified only by expectations of extraordinary future returns, creating intense pressure to deliver breakthrough capabilities and market dominance.

What does OpenAI's value network prioritize? Given Altman's public statements and OpenAI's behavior, the focus appears to be on rapid capability breakthroughs—racing toward artificial general intelligence, or broadly capable AI that can adapt to many kinds of problems rather than being narrowly specialized—based on the belief that technological dominance will eventually translate into profitability. The large amount of investment capital from VCs, despite massive losses, suggests investors are buying into the "win the race first, monetize later" strategy. Meanwhile, the consumer-heavy customer base creates pressure to ship impressive features that attract and retain users. At the same time, Microsoft's substantial stake creates pressure to move fast to stay ahead of Google.

The concern: Given these pressures, OpenAI might prioritize capability advancement and market capture over careful safety research, might develop addictive features that drive usage at the expense of user wellbeing, might deploy powerful systems before thoroughly understanding their risks, or might compromise on alignment research when it conflicts with competitive timelines. The 2023 board crisis suggests that when safety, ethics, or any other concerns conflict with commercial pressures, the profit-oriented part of the value network can overwhelm governance protections designed to balance other priorities.

Anthropic

The history: Anthropic was founded in 2021 by Dario Amodei, OpenAI's former Vice President of Research, and six other former OpenAI employees who left, in part, over concerns about commercialization and safety.[33] They structured the new company as a public benefit corporation from the start, but added something unusual: a Long-Term Benefit Trust.

The Trust is an independent body of five financially disinterested Trustees.[34] Its legal mandate is to advance "the long-term benefits of humanity," and it holds a special class of stock (Class T) that grants the Trustees authority to elect a growing number of Anthropic's board members.[35] The Class T stock carries very limited economic rights—the financial claim is deliberately kept small.[36] But Anthropic's legally-binding governance documents give the Trust the power to elect an increasing portion of the board: initially one director, then two, and ultimately a majority of the board within four years, upon the passage of time or achievement of certain fundraising milestones.[37]

During 2025, Anthropic's annualized revenue has grown from $1 billion to an estimated $9 billion, while still operating at roughly $3 billion in losses.[38] Recent reporting indicates that Anthropic is considering a potential IPO in early 2026 that would value the company at around $300 billion to $350 billion.[39]

The value network: Anthropic's Long-term Benefit Trust creates an interesting dynamic in how it relates to its value network.

In essence, Anthropic's founders were mindful at the start to create an independent entity charged with safeguarding humanity's long-term interests against potential negative side effects of AI development, and then structured its corporate governance to guarantee that entity shared substantial power alongside profit-oriented investors.

Anthropic has also taken noteworthy measures to address its investors' priorities. Major investors include Amazon ($8 billion total investment) and Google (approximately $3 billion), both of which are also cloud infrastructure providers for Anthropic.[40] However, both investments come with capped voting power to preserve Anthropic's independence and prevent any single investor from dominating governance decisions.[41] Neither Amazon nor Google holds board seats.[42] This structure attempts to benefit from big tech capital and infrastructure while maintaining the Long-Term Benefit Trust's control.

In addition to Amazon and Google, Anthropic's investor base includes a broad mix of venturecapital, growth-equity, and institutional backers. In fact, Anthropic's investor base is far more diverse than that of a typical startup. Beyond traditional venture and growth-equity firms (Lightspeed, Coatue, Altimeter, D1 Capital, General Catalyst, Insight Partners, Menlo, Bessemer), the September 2025 financing also attracted large institutional asset managers (Fidelity, BlackRock), private-equity/alternative-asset firms (Blackstone), and long-horizon capital holders, including pension funds (Ontario Teachers' Pension Plan) and sovereign-wealth funds (Qatar Investment Authority, Singapore's GIC).[43]

This heterogeneous mix blends both aggressive growth expectations and conservative, long-term-oriented capital, creating a complex value network where pressure comes not just to "move fast and capture market share," but also to deliver returns that satisfy diversified fiduciary stakeholders. Structurally, it strengthens the argument that Anthropic's Long-Term Benefit Trust and governance design are not mere symbolism—they are essential ballast against the centrifugal pressures of divergent investor incentives.

Anthropic's target customer base also adds an interesting twist to how customer demand shapes the company's priorities. Industry reports suggest that more than three-fourths of Anthropic's revenue comes from enterprise and API customers.[44] This revenue mix stands in contrast to OpenAI's, which is driven primarily by consumer-facing products.

Meanwhile, Anthropic, like OpenAI, sustains huge annual losses, indicating outsized investor confidence in massive future returns.[45] This means that sooner or later, its investors also expect its technology to deliver world-changing capabilities that will lead to market dominance.

What does Anthropic's value network prioritize? As noted earlier, concerns about AI's risks to humanity loom large. In the face of those concerns, Anthropic's founders have taken noteworthy steps to establish their company within a value network that prioritizes human well-being. First, the Long-Term Benefit Trust injects governance pressure to consider safety and alignment, even when it conflicts with short-term commercial interests. Second, the public benefit corporation structure provides legal protection for Anthropic's directors to balance public benefit considerations alongside shareholder returns, rather than being legally required to maximize shareholder value above all else.[46] Third, its long-horizon investors help reinforce the company's prioritization of long-term survival and well-being. Fourth, the enterprise customer base creates pressure to prioritize reliability and interpretability over flashy consumer features. Enterprise customers—especially those in regulated industries, government, and safety-critical applications—expect Anthropic to operate with safeguards against financial, legal, and reputational risks. Nonetheless, Anthropic's massive losses and its VC funding also push it headlong into the race to be one of the leading AI players that dominate the AI market and generate massive new revenue streams that justify continued investment.

The concern: If Anthropic becomes a publicly traded company this year, that move will subject its governance structure to new pressures: public market expectations for short-term performance, quarterly earnings scrutiny, and shareholder demands that may not align with long-term safety commitments. Whether Anthropic's Public Benefit Corporation structure and Long-Term Benefit Trust can withstand these commercial pressures remains to be seen. The structure is designed to give humanity's interests a permanent seat at the table, but it hasn't been tested at the scale OpenAI experienced during the Altman crisis.
On the other hand, if Anthropic's more measured approach causes it to fall significantly behind its competitors, will its investors demand faster commercialization? And if that scenario plays out, will the company restructure to prioritize profits, rationalizing that humanity is better off in a world with a compromised Anthropic than a world where only Anthropic's competitors exist?

Will Anthropic restructure to prioritize profits, rationalizing that humanity is better off in a world with a compromised company than a world where only Anthropic's competitors exist?

Google/DeepMind

The history: Google's modern AI journey began in 2014 when it acquired DeepMind, a London-based lab known for pushing the frontiers of reinforcement learning.[47] That acquisition gave Google one of the most advanced AI research groups in the world and set the stage for breakthroughs to come. Three years later, Google researchers introduced the Transformer architecture—a foundational innovation that made today's large-language models possible.[48] Yet despite pioneering the core technology, Google was cautious about deploying conversational AI at scale. Public comments by CEO Sundar Pichai and AI chief Jeff Dean emphasized that large language models can "hallucinate" and spread false information—creating "reputational risk" for Google's search business, which provides the majority of revenue for Google's parent company, Alphabet.[49] That caution reflected the realities of Google's business model: protecting a highly profitable, trust-dependent ads business meant moving slowly, even when research groups were pushing the boundaries of what was technically possible.

In 2018, Google published its AI principles, pledging, for example, not to develop systems that reinforce social bias and emphasizing responsibility in AI design.[50] Major technology outlets noted areas where the principles fell short, but generally treated them as a consequential and unusually explicit attempt by a major tech firm to articulate ethical boundaries for the development and use of artificial intelligence.[51] The company also built formal review structures and staffed sizable ethics and safety teams—organizational safeguards that sometimes frustrated product groups eager to ship new capabilities.[52] In those early years, Google was widely regarded as one of the most cautious frontier AI labs.

But before long, these organizational safeguards began to fracture.

In December 2020, Google fired Timnit Gebru, co-lead of its Ethical AI team, following a dispute over research she coauthored about the risks of large language models.[53] Within two months, Google also fired Margaret Mitchell, Gebru's co-lead, after she searched her emails for evidence of discriminatory treatment.[54] The dismissals of both researchers—who had built one of the most diverse teams within Google Research—sparked protests from over 2,000 employees and raised questions about whether Google would silence critical voices on AI ethics.[55] In February 2021, Google restructured the Ethical AI team under new leadership with less direct AI experience, and many remaining team members left in the aftermath.[56]

Then came ChatGPT. When OpenAI released ChatGPT in late 2022, the immediate public response demonstrated enormous consumer demand and created a visible competitive threat. Within months of ChatGPT's debut, Google reorganized its research groups—merging Google Brain with DeepMind—accelerated its generative-AI roadmap, and pushed forward under the Gemini brand.[57] This acceleration further weakened the already-fractured safety infrastructure: governance structures designed for a slower era and already compromised by the 2020-2021 dismissals lost what remained of their practical force as competitive pressure mounted.

In the midst of these issues, Google also invested billions in Anthropic, both as a hedge against falling behind and as a way to grow its cloud-infrastructure business by supporting other AI companies.[58] As competitive pressure intensified, Google began releasing more advanced AI models more quickly. Its latest version of Gemini 3, released in November 2025, has garnered widespread attention as potentially the new frontrunner in AI development. [59]

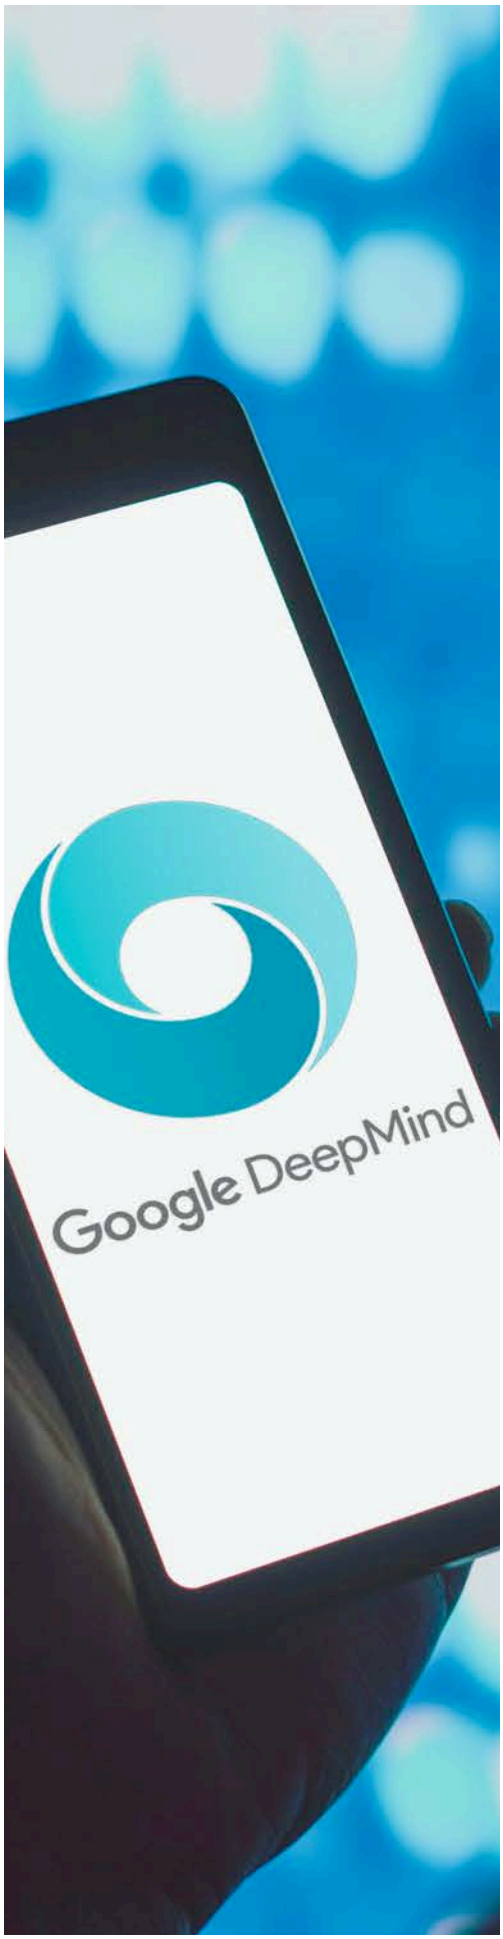

The value network: Google's AI development is funded by its advertising business, which positions advertisers as the primary customers. It makes money when users click on ads that help them take the next step in their search.[60] In this mix, DeepMind operates as a division of Alphabet, and while Google has published 'Responsible AI Principles,' it remains accountable to investors whose expectations are shaped by Alphabet's dependence on its highly profitable advertising business.

Unlike OpenAI or Anthropic, Google never had the organizational latitude to reinvent its governance architecture for frontier-model development. By the time generative AI became competitive, Google was already operating within the constraints of a mature, publicly traded corporate structure—meaning the ship had sailed on adopting radically different alignment or oversight frameworks. Thus, the dual pressure to sustain its existing business model while also getting ahead of disruption remains the defining force in Google's AI strategy today.

As noted earlier, Google also has a significant investment in Anthropic—approximately $3 billion, which gives it a 14% stake (without voting rights or board seats). Why invest in a competitor to their own Gemini models? Several reasons: Google Cloud provides infrastructure for Anthropic and is negotiating a deal worth tens of billions for cloud services, generating direct revenue. Google also hedges competitive risk: if Anthropic succeeds and Gemini falters, Google still benefits financially. Antitrust constraints make acquiring AI companies nearly impossible for Google, so minority investments without control offer strategic positioning without regulatory problems. And the partnership gives Google exposure to Anthropic's safety research approaches.

Yet even as Google accelerates into AI, it operates under a standing tension: its future depends on leading and monetizing AI, while its present relies on protecting the advertising model that finances nearly all of its operations. Google's response to this tension has been to display AI-generated summaries at the top of its search results for user queries. AI summaries encourage users to continue to bring their queries to Google's search engine rather than switching to LLM chat interfaces for answers. Yet despite this strategic move, click-through rates on search results have been declining, and there is growing evidence that query volume may be under pressure.[61,62]

What does Google's value network prioritize? Google is driven by an existential priority: protecting search as a trusted, reliable gateway to the commercial web. Because that gateway funds the company, Google sits in an unusually tight bind: it must move aggressively enough to remain a leader in AI, yet cautiously enough to avoid cannibalizing the advertising engine that underwrites nearly all of its operations. The result is a company pulled between two gravitational fields: frontier-model competition on one side and search-ad stability on the other. In short, Google isn't the most reckless or the most careful—it's the most economically entangled.

The concern: Given the pressure to protect search advertising revenue, Google might be slow to deploy AI innovations that cannibalize its core business, potentially ceding ground to competitors. Or conversely, facing existential competitive pressure, it might rush AI development and deployment without adequate testing—prioritizing speed-to-market over accuracy or safety.

Google might be slow to deploy AI innovations that cannibalize its core business. Or, facing existential competitive pressure, it might rush AI development and deployment without adequate testing.

Meta

The history: Meta's real dive into LLMs began in 2023, when breakthroughs from OpenAI and Google made it clear that generative AI could become the next major computing interface. If Meta didn't build its own models, it risked depending on competitors' technology to power products across Facebook, Instagram, and WhatsApp.[63] That strategic vulnerability—combined with the chance to reposition Meta as a leader rather than a follower—led to the release of its Llama model under a quasi-open approach: Meta made the model widely available so developers could use and adapt it, but kept key parts of the system and its licensing terms under Meta's control.[64] It wasn't truly open source, but it was open enough to spark an ecosystem of researchers and developers building on Meta's models.[65]

As the competitive race intensified, Meta launched an aggressive talent-war campaign in 2024, offering enormous compensation packages—sometimes reaching into the hundreds of millions—to poach top researchers from OpenAI and Google.[66] Then, in 2025, Meta escalated further by taking a large ownership stake in Scale AI and elevating its founder, Alexandr Wang, to lead a consolidated "superintelligence" division.[67] Shortly afterwards, Meta laid off hundreds of researchers from its earlier Llama and research teams, folding everything under its new leadership structure.[68] These moves have raised real uncertainty about whether Meta will continue its quasi-open release strategy or pivot toward a more closed, proprietary model as it tries to keep pace in the AI race.[69]

The value network: Meta's AI strategy is funded by its enormously profitable advertising business across Facebook, Instagram, and WhatsApp.[70] This means Meta's actual customers aren't the people using its products; they're advertisers who pay Meta for user attention and engagement. Meta's value network—investors expecting continued growth in ad revenue and advertisers buying placements—creates powerful incentives to maximize time-on-platform and user engagement.[71]

Regulators also play an important value network role as they constrain Meta to curtail monopoly power, protect user privacy, and put safeguards in place intended to protect minors.[72] But regulators do not fundamentally change the engagement imperative.

Mark Zuckerberg's position at Meta gives him unusual latitude to shape the company's priorities. Zuckerberg owns about 13-14% of Meta's shares but controls approximately 54-62% of voting power through Class B shares that carry 10 votes each.[73] This control structure means Zuckerberg can pursue strategic bets like AI and AR/VR despite lacking clear immediate paths to profitability, because he believes these investments are critical to Meta's long-term future.[74] Most of Zuckerberg's wealth is tied up in Meta stock, giving him strong incentives to pursue what he genuinely believes will benefit the company over time, even if Wall Street is skeptical in the short term.

Meta's AI might prioritize addictiveness over utility, optimize for emotional manipulation over information accuracy, or generate a flood of problematic AI content to drive engagement.

What does Meta's value network prioritize? At Meta, there's no separate governance structure insulating AI development from commercial pressures. Meta is a conventional corporation answerable to shareholders. Everything gets evaluated through the lens of the engagement and ad revenue metrics that drive Meta's profitability. AI features that increase time-on-platform, such as AI-driven video creation and editing, are prioritized. AI that helps advertisers target more effectively gets resources. The question is whether AI that serves user well-being gets built when it conflicts with these metrics.

Interestingly, Meta also seems to be making a longer-term bet on AR/VR hardware with AI as the critical interface layer. If that succeeds—if AR/VR or its descendants become a new computing paradigm that disrupts desktop and mobile devices—it could dramatically shift their value network from an advertising-based model to a hardware customer model. But that remains speculative. For now, the free AI model is subordinate to the advertising business, with all the pressures that entails.

The concern: Given the pressure to maximize engagement from advertisers, Meta's AI might prioritize addictiveness over utility, optimize for emotional manipulation over information accuracy, or generate a flood of problematic AI content to drive engagement. We've already seen this dynamic play out with Meta's social media algorithms—the value network hasn't changed for their AI development.

xAI

The history: Elon Musk founded xAI in March 2023, five years after his contentious departure from OpenAI. He had co-founded OpenAI in 2015 alongside Sam Altman and others, pledging to provide roughly $1 billion in funding over time to the nonprofit AI research lab. The venture was explicitly structured to prioritize safety over profit. By 2017, as OpenAI recognized it would need massive compute resources and billions in annual funding, Musk proposed that OpenAI merge with Tesla or that he take control as CEO. When other co-founders rejected these proposals—believing individual control contradicted the mission—Musk withdrew his promised funding and left OpenAI's board in February 2018. Publicly, he cited conflicts with Tesla's AI work; privately, he told colleagues OpenAI had zero chance of success and was hopelessly behind Google.[75]

xAI was explicitly positioned as a response to what he characterized as "woke" political bias in other AI models, particularly ChatGPT. He criticized OpenAI for abandoning its open-source nonprofit roots to become what he called a "closed source, maximum-profit company" aligned with Microsoft.[76] xAI launched with a stated mission to build "maximally truth-seeking" AI that would "understand the true nature of the universe," recruiting engineers from DeepMind, OpenAI, Google, and Microsoft.

Initially, xAI was structured as a Nevada public benefit corporation with formal commitments to positive societal impact. By May 2024 (while Musk was actively suing OpenAI to prevent its for-profit conversion), xAI quietly dropped its own public benefit status, eliminating any formal obligation to environmental and social goals.[77]

In March 2025, xAI acquired X (formerly Twitter) in an all-stock deal, merging Musk's social media platform with his AI venture.[78] xAI's growth has been extraordinary even by AI industry standards. The company raised $6 billion in May 2024 at a $24 billion valuation, another $6 billion in December 2024 at a $50 billion valuation, and $20 billion in late 2025 at a valuation of approximately $230 billion.[79]

The company is funded through a combination of Musk's wealth, his other companies (SpaceX invested $2 billion), and venture capital from major firms including Andreessen Horowitz, Sequoia, Fidelity Management & Research Company, Qatar Investment Authority, and strategic investors Nvidia and Cisco.[80]

xAI's financial performance remains opaque, but reporting based on internal company documents provides a partial picture. According to Bloomberg reporting cited by Reuters, xAI generated approximately $107 million in revenue in the third quarter of 2025, less than two years after its founding. At the same time, the company posted a net loss of roughly $1.46 billion for the quarter, reflecting the scale and speed of its investment in compute infrastructure and model development. Unlike OpenAI and Anthropic, xAI hasn't publicly disclosed annual revenue figures or forward-looking financial projections.[81] The available data suggest xAI has scaled revenue unusually quickly for a company of its age, albeit alongside exceptionally large operating losses.

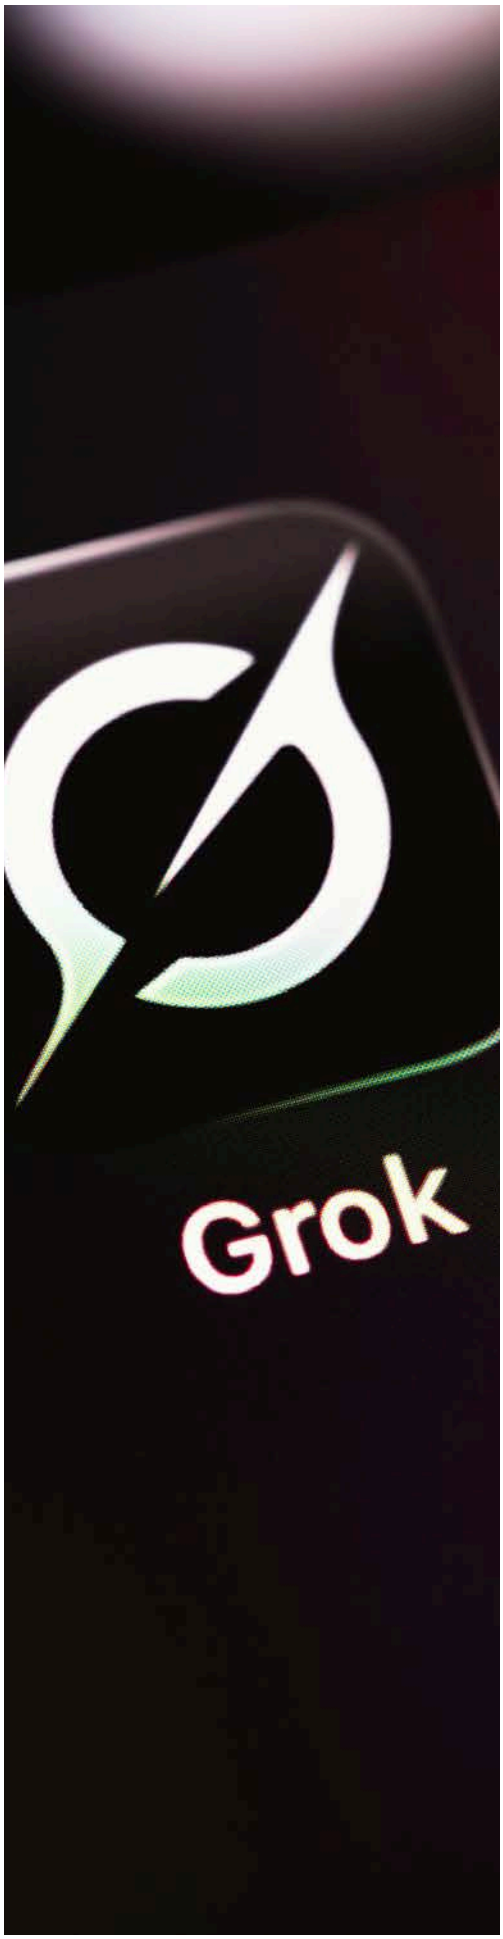

The value network: The dominant force in xAI's value network is Musk himself. He controls the board, the strategy, and the priorities. There's no independent governance structure mediating different stakeholder interests. While founder influence matters across all AI companies (as we'll discuss below), xAI is an extreme case in which power is nearly total.

What does xAI's value network prioritize? Whatever Musk finds interesting, important, or strategically valuable. Features like Grok's crude and salacious AI companions (Rudy and Ani) reflect his personal interests and sense of humor—they're not the result of user research or market analysis.[82]

But xAI's priorities likely extend beyond Musk's personal amusements. Given his other ventures, xAI may be positioned as crucial infrastructure for Tesla's autonomous driving ambitions, as an engineering problem-solving tool for SpaceX and his other companies, or as part of a broader bet on AI becoming the primary interface for complex systems. The integration with X suggests strategic thinking about social platforms powered by AI. Musk may be building xAI less as a standalone product and more as enabling technology for his industrial empire.[83]

The concern: Whether this concentration of control proves beneficial or detrimental hinges largely on Musk's wisdom and benevolence. When one person dominates a value network with no countervailing governance, that person's judgment—both strategic and idiosyncratic—becomes the company's direction, unchecked by institutional constraints. Strategic bets might be brilliant or reckless, and there's no institutional capacity to distinguish between them. Given his track record with X (amplifying conspiracy theories, gutting content moderation, erratic policy changes), there's reason to approach xAI's development with both interest and caution, recognizing that singular control could enable either a focused pursuit of ambitious goals or an unchecked pursuit of strategically rational priorities that carry hidden, corrosive costs.

When one person dominates a value network with no countervailing governance, that person's judgment—both strategic and idiosyncratic—becomes the company's direction, unchecked by institutional constraints.

HOW VALUE NETWORKS SHAPE PRIORITIES ACROSS LEADING AI COMPANIES

Dimension	OpenAI	Anthropic	xAI	Google (AI division)	Meta (AI division)
Ownership / Corporate Form	Public Benefit Corporation with nonprofit foundation oversight	Public Benefit Corporation with Long-Term Benefit Trust	Privately held company controlled by the founder	Publicly traded corporation (Alphabet)	Publicly traded corporation with the founder super-voting shares (Zuckerberg retains voting control)
Latest Valuation (private only)*	~$500B–$830B (2025–2026 market estimates)	~$300B–$350B (fundraise/IPO talk)	~$230B	N/A (AI embedded in Alphabet)	N/A (AI embedded in Meta Platforms)
Reported Revenue (latest available)**	~$20B ARR (2026 est.)	~$9B ARR (2025 est.)	~$107M (Q3 2025 reported)	N/A	N/A
Reported Losses (latest available)	~$9B (2025 annual est.)	~$3B (2025 annual est.)	~$1.46B (Q3 2025 reported)	N/A	N/A
Primary Funding Source	• VC capital • Microsoft equity • Consumer subscriptions • Enterprise APIs	• VC and growth equity • Institutional investors • Enterprise APIs	• Founder capital • Strategic investors • Internal cross-subsidization	• Advertising (search) • Cloud services • Enterprise products	• Advertising (attention-based) • Platform monetization
Dominant Value-Network Pressures	• Nonprofit board • VC expectations of category dominance • Cloud infrastructure partners • Feature demand from consumers • Strategic pressure from Microsoft	• Trust-governed board control • Enterprise risk tolerance • Long-horizon capital • Cloud infrastructure partners	• Founder priorities • Ideological commitments • Absence of independent governance • Competitive signaling	• Advertiser trust, regulatory scrutiny • Brand risk • Protection of search revenue	• Advertiser demand for engagement • Regulatory scrutiny • Reputational risk • Founder control via super-voting shares
Likely Priorities	• Race to scale capabilities • Expand consumer adoption • Monetize usage at scale • Normalize AI as default interface	• Invest in alignment research • Restrict deployment to vetted use cases • Favor reliability over virality	• Rapid release of features • Embed AI across Musk-controlled platforms	• Integrate AI into search while preserving trust • Retain users • Protect advertiser relationships	• Reduce rivals' model advantages while preserving platform leverage • Reinforce dominance of Meta's social platforms
Primary Failure Mode	Speed outpaces governance, enabling subtle but widespread societal harms	Caution limits adoption and influence, allowing less careful competitors to set de facto norms	Founder blind spots go unchecked, producing concentrated, large-scale failures with limited internal correction	Epistemic distortion as commercial incentives subtly bias information quality	Engagement-driven optimization amplifies misinformation or social harm at an AI scale

**Valuations are approximate and derived from different disclosure mechanisms (post-money rounds, restructuring terms, or investor reports), and should be interpreted as directional rather than directly comparable.*

***ARR (annualized revenue run rate) extrapolates recent performance and does not represent realized annual revenue.*

THE FINANCING PRESSURES SHAPING AI COMPANY VALUE NETWORKS

As discussed above, none of the major AI companies has yet achieved sustained profitability. Instead, their value networks are being shaped by how they finance rapid, capital-intensive development in the absence of stable margins, through external investment, cross-subsidization from existing businesses, and, increasingly, experimental monetization strategies. These financing choices create distinctive pressures across investment-backed labs and incumbent tech firms alike.

The investment model: Capital creates pressure

OpenAI and Anthropic are sustained by external investment. They're burning billions of dollars while promising investors that AI will cure cancer, solve climate change, or win the next geopolitical arms race. These grand narratives keep capital flowing, but they also create intense pressure to advance capabilities rapidly rather than proceed cautiously to understand and mitigate the risks their technologies might pose to society.

Other companies (Meta, Google, and xAI) are funding their AI development from existing profitable businesses. This creates different pressures and different risk calculations. They're not selling visions to VCs; they're making bets with their own money.

But there's an important wrinkle here: the big tech companies funding AI companies through investments are also profiting handsomely from them. Microsoft, Amazon, and Google all provide cloud infrastructure—the massive computing power needed to train and run AI models. OpenAI, Anthropic, and others pay billions for access to GPUs, storage, and networking.

The full picture of these economic relationships remains unclear because private AI companies are not obligated to make their contract terms public. Nonetheless, available evidence points to a shared economic pattern across all three big tech companies and their frontier lab partners.

ken together, these arrangements position big tech companies in a strong financial position, no matter how the AI race plays out. Even if their own models don't dominate, they profit from providing the infrastructure for others.

The advertising model: Lessons from social media

Next, to understand how some AI companies are likely to behave, we need to examine advertising-based business models, because several leading AI labs operate (or are moving) within these value networks. When advertising is the primary source of revenue, users are no longer the customers but the product being sold, inverting incentives in consequential ways.

Social media offers the clearest cautionary tale. Platforms like Facebook, Instagram, TikTok, YouTube, and Twitter (now X) generate revenue by inserting ads into streams of content. That revenue model rewards platforms for hijacking user attention and habituating users to repeat engagement to maximize time-on-platform. The result: algorithms optimized for engagement often promote content that provokes outrage, amplifies misinformation, deepens political polarization, and contributes to mental health harms—not because engineers wanted these outcomes, but because the value network created pressure to maximize metrics that generate ad revenue.[84]

Tristan Harris, co-founder of the Center for Humane Technology, has documented this pattern extensively.[85] As he puts it: “Free is the most expensive business model we've ever created."[86] The cost isn't paid in dollars; it's paid in attention, mental health, and social cohesion.

Yet not all advertising models drive toward the same priorities. Google’s business model, while still ad-supported, is built primarily around Internet search.[87] Rather than monetizing sustained attention, search ads monetize user intent: people arrive with a question or need, and advertisers pay to appear alongside attempts to answer it.[88] With this model, the central risk is not addiction or compulsive engagement, but biasing the results in favor of advertisers.[89] When people turn to such a platform for information, does the platform present the most accurate and useful information, or does it present information designed to build user affinity for an advertising partner?

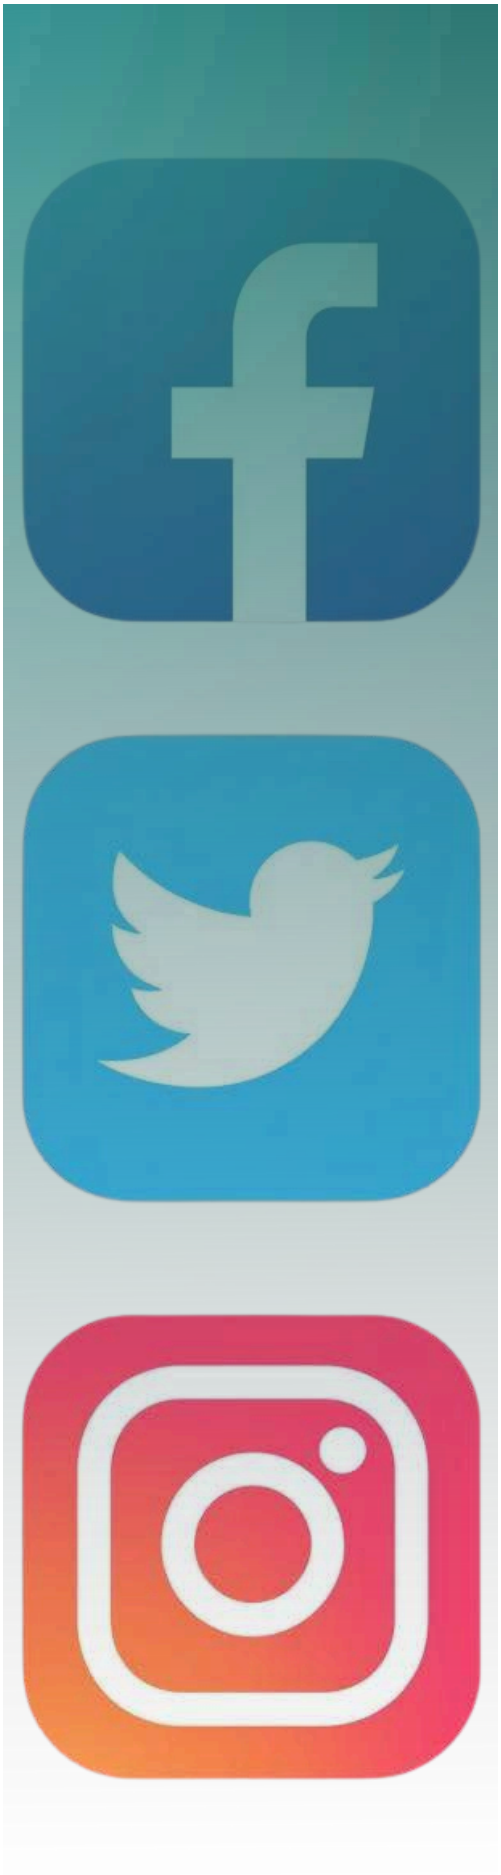

These distinctions matter as AI systems are increasingly used to generate content and mediate how people access information. Will Meta's models reinforce an attention-based ecosystem designed to maximize reach and engagement? Will Google's AI efforts prioritize user intent, speed, and trust, or maximize ad revenue? Meanwhile, OpenAI has now announced that advertising will be part of its free consumer offering, formally committing the company to an ad-supported revenue model alongside subscriptions and enterprise sales.[90] It now remains to be seen how ad revenue will influence users' experiences, information quality, and long-term well-being, given the incentive structures it introduces.

The outsized role of founders

The xAI example raises a broader question: how much does individual founder vision matter in shaping AI development?

The answer: In any organization that hasn't yet discovered a profitable business model, the founder's vision carries enormous weight. Sam Altman's belief in racing toward AGI shapes OpenAI's direction—his vision attracted Microsoft and billions in venture funding through a hybrid structure designed to raise large amounts of capital while limiting investor returns. Dario Amodei's safety concerns led him to create Anthropic's Long-Term Benefit Trust, structurally protecting that vision against commercial pressures. Mark Zuckerberg's super-voting shares let him pursue long-term bets, such as open-source models and AR/VR, despite Wall Street skepticism. Elon Musk's personal interests literally manifest as product features in Grok. These aren't just hired CEOs executing shareholders' wishes.

Why does founder vision carry such weight? Two reasons. First, without profitable business models, these companies raise capital on vision rather than metrics. Founders sell investors on a future, not on quarterly returns—giving founders unusual persuasive power to shape their organizations' direction. Second, founders typically retain formal control as major shareholders with special voting rights and authority over board composition, unlike hired CEOs, who answer more directly to boards and investors.

The question isn't whether founder vision matters...it clearly does. The question is: does outsized founder control foster greater focus on humanity's well-being, or does it create blind spots that allow problematic priorities to go unchecked? The answer likely depends on the individual founder's wisdom and genuine concern for broader social benefit.

The question is: does outsized founder control foster greater focus on humanity's well-being, or does it create blind spots that allow problematic priorities to go unchecked?

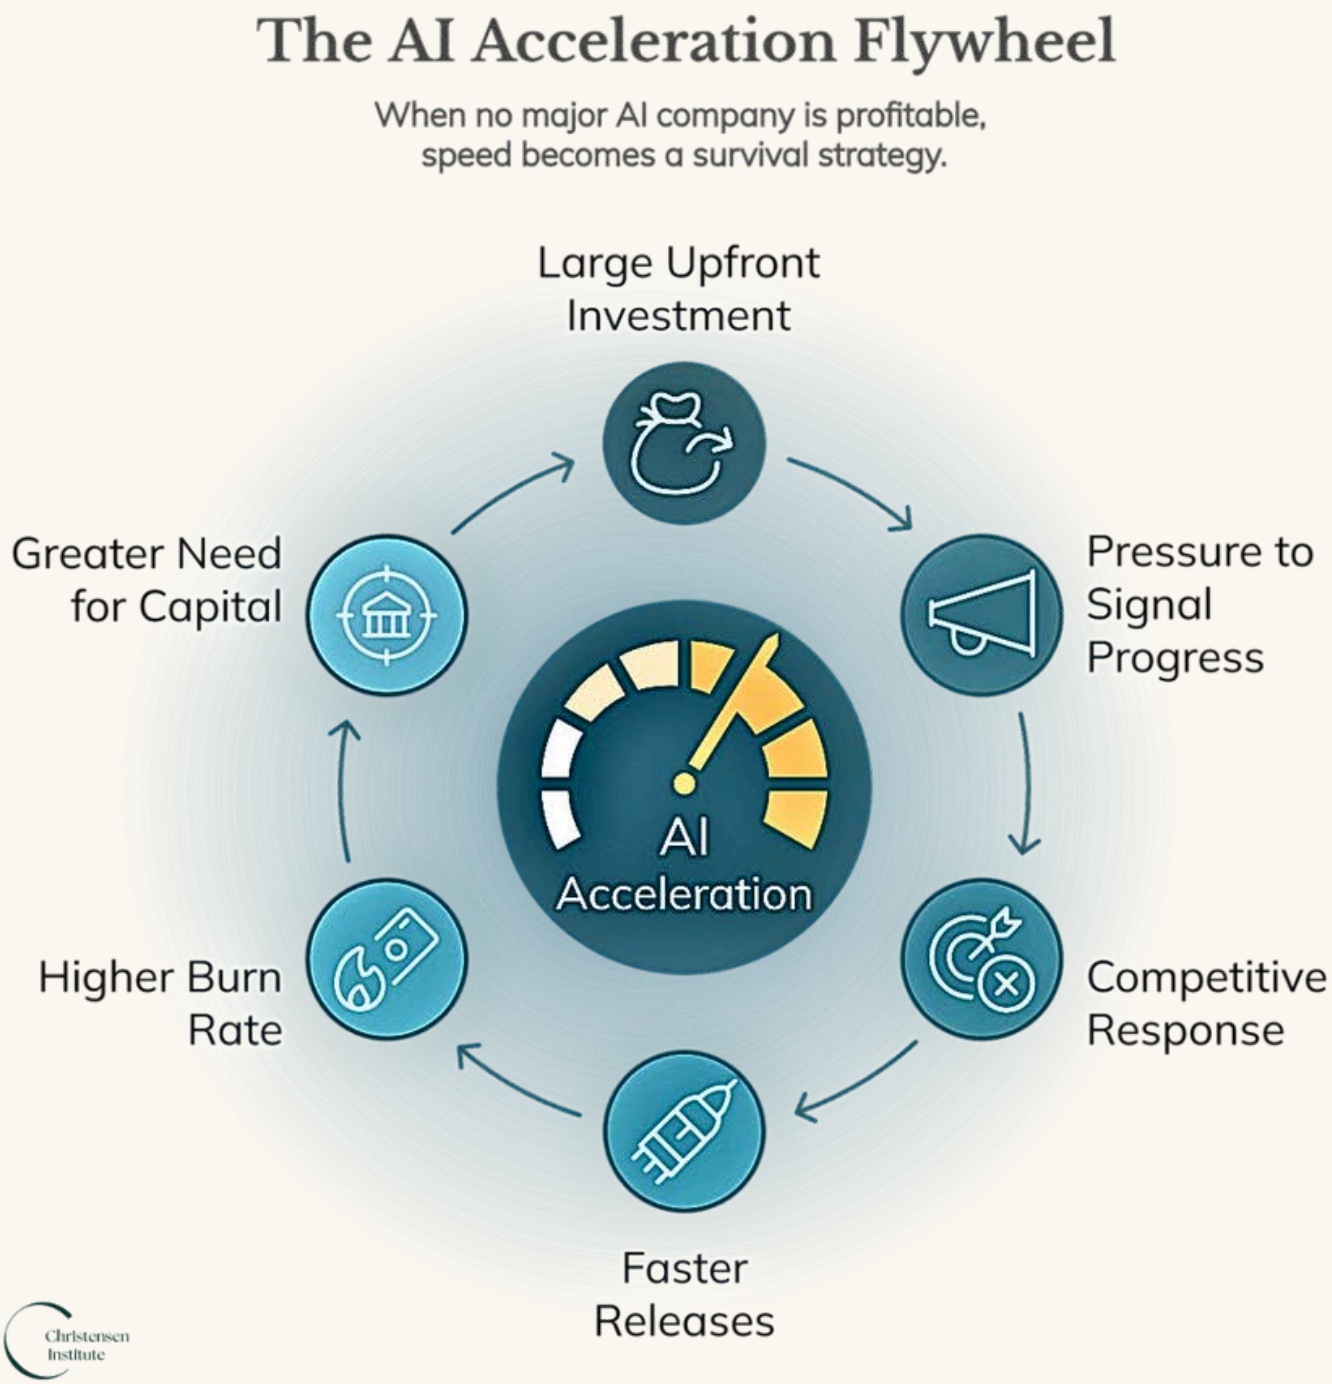
The AI Acceleration Flywheel
When no major AI company is profitable, speed becomes a survival strategy.
Large Upfront Investment
Pressure to Signal Progress
Competitive Response
Faster Releases
Higher Burn Rate
Greater Need for Capital
AI Acceleration
Christensen Institute

The dynamics of competition

In November 2025, Google released Gemini 3, achieving record performance benchmarks.[91] Within weeks, OpenAI declared an internal "code red" and rushed GPT-5.2 to market ahead of schedule.[92] The response wasn't driven by safety concerns or user research—it was pure competitive reflex.

This dynamic reveals how competition operates as a force within value networks. Competition creates strategic interdependence, forcing firms to respond to rivals' moves regardless of whether those moves align with their own priorities. In markets with a small number of powerful competitors, this pressure produces convergence—not because there's a single optimal solution, but because no company can afford to fall behind on visible dimensions of performance.[93] The result: similar feature sets, comparable interfaces, and synchronized product roadmaps. Every major AI lab offers chatbot interfaces, multimodal capabilities, coding assistance, and API access. These become table stakes: baseline features required just to stay in the game.

For AI companies today, the competitive pressure within their value networks is driving relentless acceleration: faster releases, more ambitious capability claims, and shorter timelines between breakthroughs. Geopolitical competition amplifies this pressure—the narrative that "America must win the AI race against China" creates urgency beyond commercial competition, making it harder to advocate for measured development.[94] The challenge is that competitive pressure often overwhelms other elements within a value network that might otherwise promote caution. When competitors are moving fast, and investors are watching closely, even carefully designed governance structures face intense strain. This acceleration creates real risks: shorter development cycles mean less time for safety testing and more pressure to deploy capabilities before fully understanding their implications.

Yet this convergence period may be temporary. As table-stakes capabilities become commoditized, competitive advantage must shift toward other dimensions: specific domain expertise, workflow integration, trust within particular contexts, or superior risk management for sensitive applications. This differentiation is already emerging. As noted earlier, OpenAI's revenue comes predominantly from consumer subscriptions, while Anthropic's comes from enterprise and API customers. These aren't just different market segments—they create different pressures. Enterprise customers in regulated industries demand reliability and risk management; consumer users prioritize accessibility and novel capabilities. You can't fully optimize for both simultaneously. Safety and reliability could become genuine differentiators, particularly as companies develop track records in high-stakes environments and as procurement decisions increasingly account for risk management rather than just capability benchmarks.

We're currently in a moment of intense competitive pressure, driving convergence and acceleration, with real risks from rushed development and insufficient testing. But as companies position themselves in distinctive value networks—embedding themselves in different customer bases, contexts, and risk profiles—differentiation should increasingly moderate the race. Whether that differentiation arrives quickly enough remains uncertain. **Value networks will ultimately shape where AI development heads, but right now, the competitive element is pushing the accelerator while other elements struggle to apply the brakes.**

The role of government regulation

Government regulation represents another force shaping AI company value networks, though its impact on frontier AI development has been more limited than many anticipated.

At the federal level, the US has seen considerable activity but little binding regulation. In October 2023, President Biden issued Executive Order 14110, the most comprehensive federal AI governance effort to date.[95] Yet within hours of taking office in January 2025, President Trump rescinded the order, characterizing it as a barrier to innovation.[96] His administration has instead focused on removing regulatory obstacles with the goal of maintaining US AI dominance.[97] Despite hundreds of AI-related bills introduced in Congress, only a few have been enacted, most embedded in appropriations or defense authorization legislation.[98] To date, no comprehensive federal AI regulation exists.

This federal vacuum has pushed states to act. In 2025, every state introduced AI legislation, with 131 measures enacted across 40 states.[99] The focus has been primarily on tangible, observable harms: deepfakes (especially sexual deepfakes), child safety, election integrity, and algorithmic discrimination in housing and healthcare.[100] California attempted the most ambitious effort with SB 1047, which would've regulated AI models costing over $100 million to train, but Governor Newsom vetoed it in September 2024, arguing it focused too narrowly on model size rather than deployment risk.[101] Colorado enacted the first comprehensive state AI law in May 2024, though implementation has been delayed to allow refinement.[102] This state-level activity has created a patchwork of varying requirements that many in the tech industry oppose, leading to ongoing federal-state tension over whether Washington should preempt state action.[103]

Meanwhile, the EU AI Act, which entered into force in August 2024, is already shaping company behavior globally, much as the General Data Protection Regulation (GDPR) did for privacy.[104] Its risk-based framework and extraterritorial reach mean US companies serving European users must comply regardless of US regulatory choices.
Amid this regulatory patchwork—federal inaction, state experimentation, international frameworks, and ongoing political battles—a crucial question emerges: can government regulation effectively influence how AI companies develop and deploy their technology?

The answer is complicated. Regulatory efforts face inherent challenges as tools for shaping AI development.

First, **most regulation operates through constraints and requirements rather than positive incentives.** Governments can mandate safety testing or disclosure, but can't make companies do those things well or innovate responsibly. Compliance can become merely a checkbox exercise that fulfills the letter of the law while sidelining its intent.

Second, **policy confronts fundamental timing and uncertainty problems.** AI capabilities advance rapidly while policy moves slowly, creating a persistent mismatch between rules and reality. This is compounded by deep uncertainty about the technology itself: we don't know how far capabilities will advance, what adjacent technologies they'll enable, or how society will adapt. Policy requires drawing bright lines in this uncertain terrain. It's far easier to regulate after problems materialize—when you can point to specific harms and craft targeted remedies—than to predict where risks might emerge.

Third, **policy feedback loops are messy.** Companies have an outsized influence to lobby for favorable regulation. Meanwhile, other constituencies with competing interests advocate for different approaches: workers raise concerns about job displacement, civil liberties groups focus on surveillance and bias, national security hawks prioritize competition with China, and the list goes on. These priorities often conflict, and there's no clear hierarchy for weighing them. The result is policy shaped by whoever marshals the most political pressure at a given moment, resulting in compromise between incompatible goals rather than a coherent vision.

Yet despite these limitations, **regulation remains crucial for imposing priorities that may not otherwise be represented by other parts of corporate value networks.** Shareholders, customers, suppliers, and investors all have seats at the table through market and corporate governance mechanisms. But what about the broader public whose lives AI systems will affect? Citizens who care about AI's impact on employment, inequality, democracy, or human autonomy but who aren't major customers or shareholders have limited influence through market forces alone. Regulation—messy, slow, and imperfect as it is—provides a mechanism for collective priorities to shape corporate behavior when consumer demand, investor expectations, and shareholder leverage prove insufficient.

Regulation—messy, slow, and imperfect as it is—provides a mechanism for collective priorities to shape corporate behavior when consumer demand, investor expectations, and shareholder leverage prove insufficient.

AI Incentives

AI companies and leaders don't dictate AI's future. That future is shaped by the structure and incentives of their value networks.

AI outcomes are shaped less by intentions than by the systems that reward certain choices and punish others.

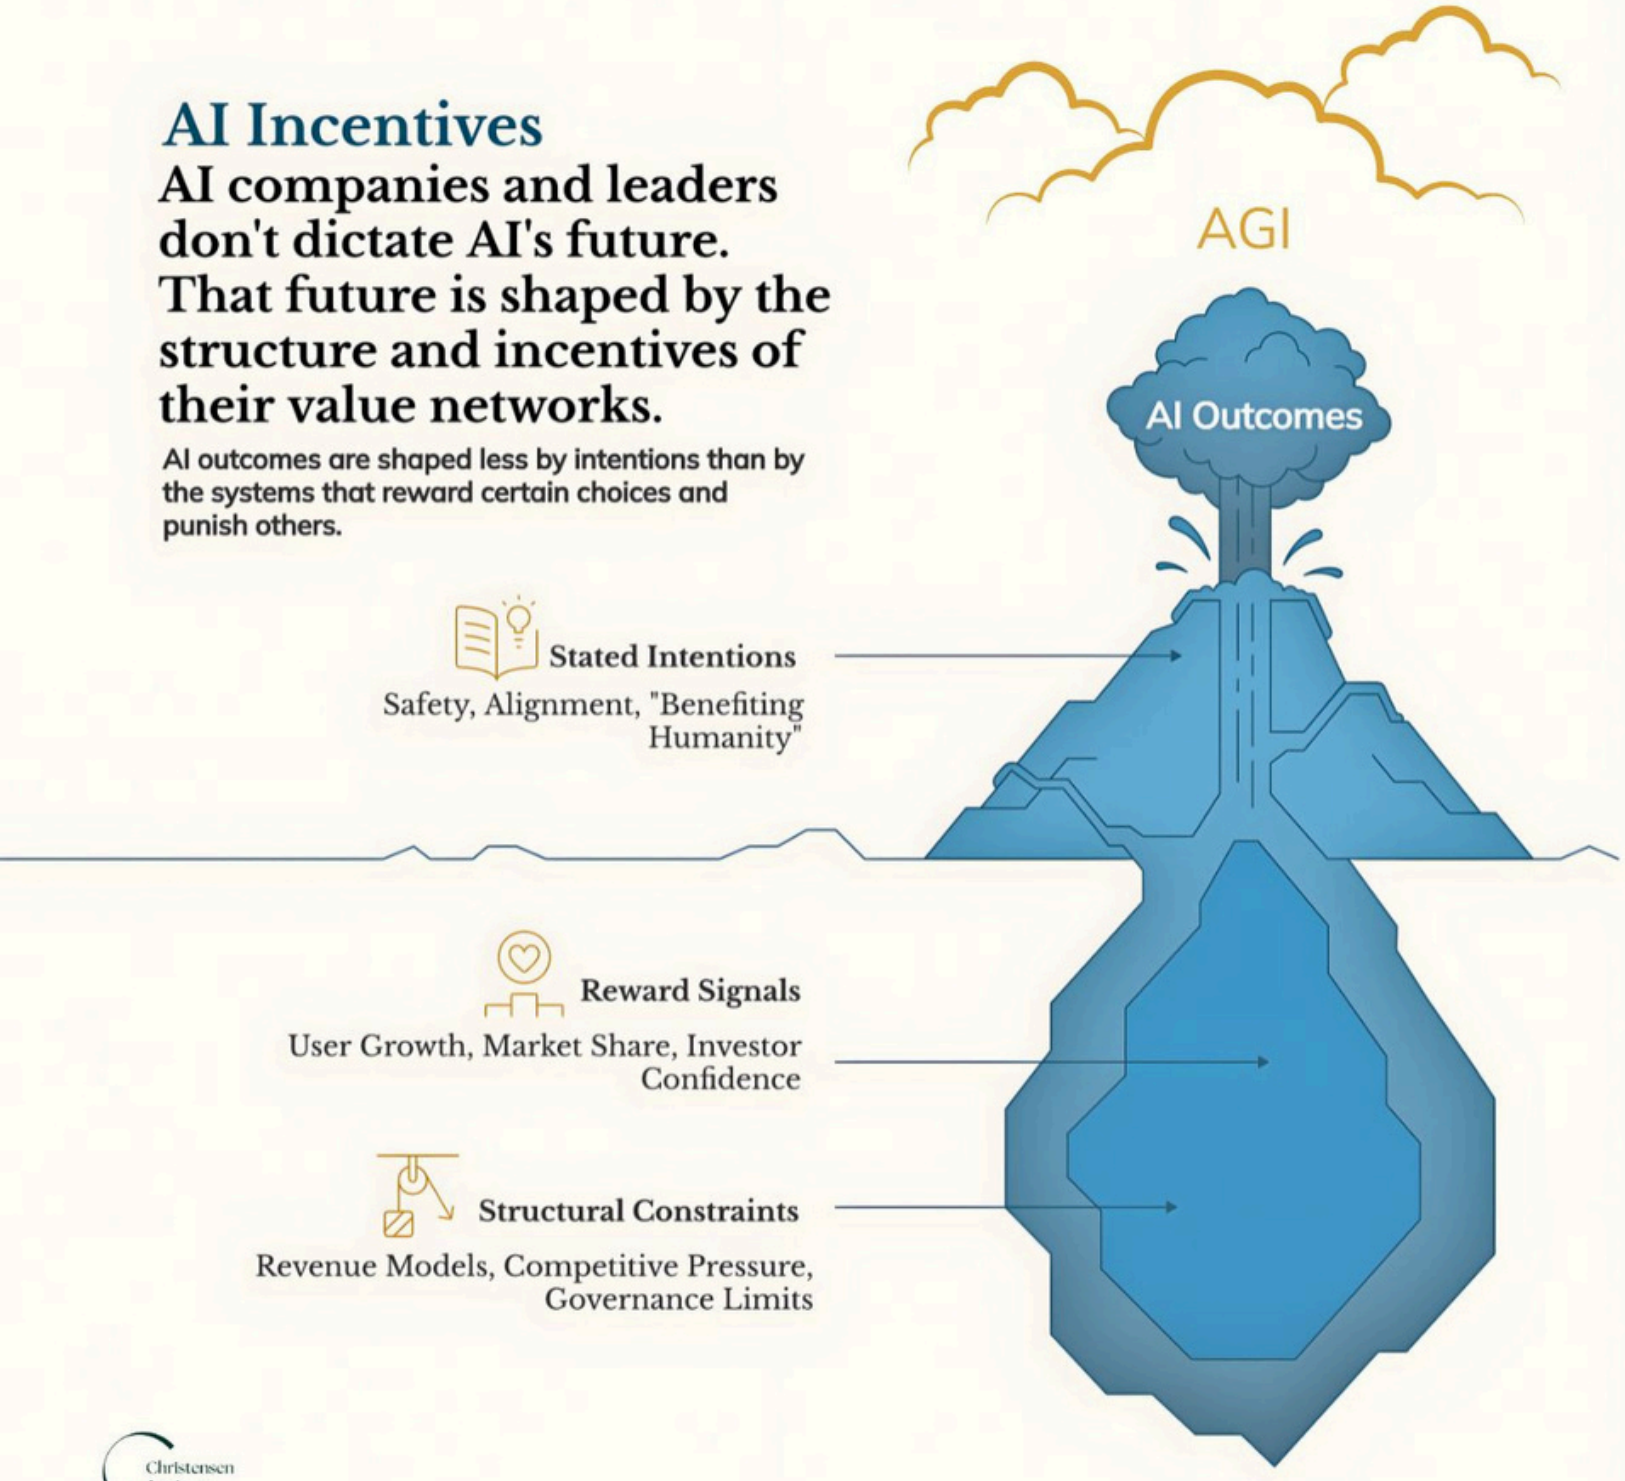

Christensen Institute

WHAT DOES THIS MEAN FOR HUMANITY?

If you walk into your kitchen and find a pot of water boiling over, you turn down the heat. But if it keeps happening, the real problem isn't about setting the stove controls—it's figuring out who is cooking and why they aren't paying attention to the pot.

Much of today's debate about artificial intelligence focuses on the metaphorical bubbles in the pot. Will AI usher in an era of universal abundance and human flourishing? Or will it lead to mass unemployment, pervasive surveillance, social breakdown, or even existential catastrophe? These questions raise serious philosophical and moral concerns, and they deserve careful reflection. But as guides to understanding AI's impact on humanity, they are limited. They depend on speculative assumptions about the agency, intent, and autonomy of future versions of AI that we do not yet know how to define, much less predict.

This paper offers a different lens. AI's impact on humanity in the near- and medium-term will be determined not primarily by the brilliance or danger of the technology itself, but by the value networks that shape how it is developed, deployed, and scaled. The investors providing billions in funding, the customers whose needs define success, the competitors driving strategic responses, the governance structures attempting to moderate commercial pressure, and the regulators imposing collective priorities. These forces, and the tensions between them, do more to determine real-world outcomes than any single technical breakthrough or ethical statement.

Seen through this lens, the question of whether AI will be "good" or "bad" for humanity begins to look misframed. Like every general-purpose technology before it, AI is likely to be both. Nuclear technologies have produced clean energy and catastrophic weapons. Combustion engines expanded mobility and trade while accelerating environmental degradation. The internet made the world's knowledge accessible, connected communities, and lowered the cost of coordination, but has also amplified misinformation and addiction. In each case, the balance of benefits and harms has been shaped not by the technology alone, but by the institutions, incentives, and power structures that govern its use.

The same is true for AI alignment. Alignment is often discussed as a technical problem to be solved or a moral stance to be declared. In practice, it is an ongoing, contested process that emerges from institutional realities. Economic growth, national security, individual freedom, shared prosperity, and public safety sit in tension with one another, and different communities disagree on how to strike the right balance between competing human values. Companies face trade-offs between speed, accuracy, profitability, and caution. Under these conditions, there is no perfectly aligned system—only systems that are more aligned with some priorities and less aligned with others, reflecting the pressures embedded in their value networks.

Right now, we're in a particularly dangerous period. Every major AI company remains unprofitable, sustained by massive investor funding that creates intense pressure for hypergrowth and market dominance. Competition among these companies drives convergence—rushed releases, synchronized product roadmaps, and accelerating timelines—as each firm responds to rivals' moves, regardless of whether those moves align with its stated priorities. When OpenAI declared "code red" in response to Google's Gemini 3 launch and rushed GPT-5.2 to market, it revealed how competitive pressure can overwhelm other considerations, including the governance structures and safety commitments these companies have carefully constructed.

Yet value networks are not static, and the current convergence may not last. As AI systems become embedded in different contexts—consumer applications versus enterprise environments, regulated industries versus open markets, and varying risk profiles and accountability requirements—value network differences force real trade-offs that can't simply be copied. Some companies serve primarily consumer users who prioritize novelty and accessibility, while others focus on enterprise customers who demand reliability and risk management. These aren't just different market segments; they're distinct value networks exerting distinct pressures. As table-stakes capabilities become commoditized, sustainable advantage shifts toward serving specific contexts well—and this differentiation, if it emerges quickly enough, could moderate the current dangerous acceleration.

Governance structures represent another potential moderating force. Carefully designed governance mechanisms can protect the mission from commercial pressure, though cautionary examples demonstrate how market forces can overwhelm even well-intentioned institutional checks. Whether governance innovations can withstand sustained commercial pressure remains one of the defining open questions and ongoing experiments of our AI age.

Whether governance innovations can withstand sustained commercial pressure remains one of the defining open questions and ongoing experiments of our AI age.

We're also not facing a single point of failure. Different companies and countries are taking different approaches. Multiple AI firms are pursuing distinct strategies, embedded in different value networks, responding to different pressures. If one path proves destructive, others might succeed. This diversity means multiple experiments running simultaneously, testing different ways to balance competing priorities.

The metaphorical pot is already heating on the stove. Who is cooking? The investors funding companies that burn billions annually, creating urgency for returns. The customers whose purchasing decisions reward certain capabilities over others. The competitors whose moves force strategic responses. The founders and boards attempting to moderate these pressures through governance innovations. The regulators imposing collective priorities that market forces alone wouldn't represent.

If we want to understand where AI is heading—and how it's likely to affect human flourishing—we need to pay attention not just to the level of the bubbles, but to these actors and the tensions between them. The most reliable signals are not found in demos, mission statements, or speculative futures, but in incentives, governance structures, and economic interdependencies. Those are the levers that determine which capabilities are prioritized, which risks are tolerated, and which uses are ultimately scaled.

The forces shaping AI's trajectory are not fixed or inevitable. They are dynamic, contested, and in many cases, responsive to feedback from the world they affect. That doesn't guarantee benign outcomes, but it does mean the future isn't predetermined—and that understanding the value networks shaping development is essential to influencing where it leads.

NOTES

1. For example, see Ezra Klein and Derek Thompson, *Abundance* (Avid Reader Press/Simon & Schuster, 2025).

2. For insight into this perspective, see Julia Freeland Fisher, "We're Giving AI Our Loneliness. Next, It Will Take Our Relationships," Christensen Institute, September 24, 2025, https://www.christenseninstitute.org/video/were-giving-ai-our-loneliness-next-it-will-take-our-relationships/.

3. See Eliezer Yudkowsky and Nate Soares, *If Anyone Builds It, Everyone Dies: Why Superhuman AI Would Kill Us All* (Little, Brown and Company, 2025).

4. For examples of these polls, see Brian Kennedy et al., "How Americans View AI and Its Impact on Human Abilities, Society," Pew Research Center, September 17, 2025, https://www.pewresearch.org/science/2025/09/17/how-americans-view-ai-and-its-impact-on-people-and-society/; Jamie Ballard, "Americans Are Increasingly Likely to Say AI Will Negatively Affect Society," YouGov, July 18, 2025, https://today.yougov.com/politics/articles/52615-americans-increasingly-likely-say-ai-artificial-intelligence-negatively-affect-society-poll; Noemi Dreksler et al., "What the Public Thinks About AI and the Implications for Governance," Brookings Institution, April 17, 2025, https://www.brookings.edu/articles/what-the-public-thinks-about-ai-and-the-implications-for-governance/; and "Public Opinion," The 2024 AI Index Report, Stanford Institute for Human-Centered AI (HAI), accessed November 27, 2025, https://hai.stanford.edu/ai-index/2024-ai-index-report/public-opinion.

5. Clayton M. Christensen, *The Innovator's Dilemma: When New Technologies Cause Great Firms to Fail* (Harvard Business Review Press, 1997).

6. The history of DEC is described in more detail in Clayton M. Christensen and Michael E. Raynor, *The Innovator's Solution: Creating and Sustaining Successful Growth* (Harvard Business Review Press, 2003).

7. "Our Structure," OpenAI, accessed November 27, 2025, https://openai.com/our-structure/.

8. "OpenAI," Britannica Money, accessed November 27, 2025, https://www.britannica.com/money/OpenAI.

9. Cameron Hashemi-Pour, "What is OpenAI? Definition and History," TechTarget, April 17, 2025, https://www.techtarget.com/searchenterpriseai/definition/OpenAI. This web page explains that in "March 2019, OpenAI shifted from nonprofit to capped-profit status and became formally known as OpenAI LP, controlled by parent company OpenAI Inc."

10. Ari Levy and Ashley Capoot, "Microsoft's OpenAI Investment Led to $3.1 Billion Drop in Net Income," CNBC, October 29, 2025, https://www.cnbc.com/2025/10/29/microsoft-open-ai-investment-earnings.html.

11. Microsoft Corporate Blogs, "The Next Chapter of the Microsoft–OpenAI Partnership," *Official Microsoft Blog*, October 28, 2025, https://blogs.microsoft.com/blog/2025/10/28/the-next-chapter-of-the-microsoft-openai-partnership/. The restructured agreement preserved the following agreement: "OpenAI remains Microsoft's frontier model partner and Microsoft continues to have exclusive IP rights and Azure API exclusivity until Artificial General Intelligence (AGI)." However, Microsoft noted it "will no longer have a right of first refusal to be OpenAI's compute provider."

12. Trevor Jennewine, "Microsoft's $13 Billion Investment in OpenAI May Be 'Some of the Best Money Ever Spent,' According to Certain Wall Street Analysts," The Motley Fool, November 10, 2024, https://www.fool.com/investing/2024/11/10/microsoft-13-billion-openai-best-money-ever-spent/. According to Jennewine, "Microsoft will ultimately recoup its initial investment and will be entitled to 49% of OpenAI's profits until a predetermined return cap is reached."

13. Microsoft Corporate Blogs, "The Next Chapter." According to the blog, the agreement includes "exclusive IP rights" with Microsoft able to commercialize OpenAI technology across its products.

14. Steve Mollman, "Sam Altman Has Been Reinstated as OpenAI CEO. Elon Musk Says 'The Public Should Be Informed' Why He Was Fired in the First Place," *Fortune*, November 23, 2023, https://fortune.com/2023/11/19/sam-altman-returns-to-openai-hq-reinstatement-as-ceo-could-be-soon/. In the conclusion of this article, Mollman noted, "It's questionable whether OpenAI could continue operating without the continual cash infusions and computing power provided by Microsoft, which means Microsoft wields considerable power."

15. "ChatGPT," Wikipedia, accessed November 27, 2025, https://en.wikipedia.org/wiki/ChatGPT.

16. "ChatGPT," Wikipedia.

17. Ingrid Lunden, "OpenAI Closes $300M Share Sale at $27B–29B Valuation," TechCrunch, April 28, 2023, https://techcrunch.com/2023/04/28/openai-funding-valuation-chatgpt/.

18. OpenAI, "OpenAI LP," news release, March 11, 2019, https://openai.com/index/openai-lp/.

19. "Removal of Sam Altman from OpenAI," Wikipedia, accessed November 27, 2025, https://en.wikipedia.org/wiki/Removal_of_Sam_Altman_from_OpenAI.

20. Matt O'Brien, "Sam Altman Reinstated as OpenAI CEO with New Board Replacing the One Which Fired Him," PBS NewsHour, November 22, 2023, https://www.pbs.org/newshour/nation/sam-altman-reinstated-as-openai-ceo-with-new-board-replacing-the-one-which-fired-him. According to PBS, "Microsoft, which has invested billions of dollars in OpenAI and has rights to its current technology, quickly moved to hire Altman on Monday, as well as another co-founder and former president, Greg Brockman, who had quit in protest after Altman's removal. That emboldened a threatened exodus of nearly all of the startup's 770 employees who signed a letter calling for the board's resignation and Altman's return."

21. Matthew Urwin, "OpenAI's For-Profit Era Begins. Can It Stay True to Its Mission?," Built In, October 29, 2025, https://builtin.com/articles/openai-new-corporate-structure.

22. “Our Structure,” OpenAI; and Microsoft Corporate Blogs, “The Next Chapter.” More specifically, following the recapitalization, Microsoft noted it “holds an investment in OpenAI Group PBC valued at approximately $135 billion, representing roughly 27 percent on an as-converted diluted basis.”

23. Microsoft Corporate Blogs, “The Next Chapter.”

24. As part of the $38 billion agreement, Amazon reported, “OpenAI is accessing AWS compute.” Amazon News, “AWS and OpenAI Announce Multi-year Strategic Partnership,” news release, November 3, 2025, https://www.aboutamazon.com/news/aws/aws-open-ai-workloads-compute-infrastructure.

25. Rebecca Szkutak, “OpenAI and Oracle Reportedly Ink Historic Cloud Computing Deal,” TechCrunch, September 10, 2025, https://techcrunch.com/2025/09/10/openai-and-oracle-reportedly-ink-historic-cloud-computing-deal/.

26. Microsoft Corporate Blogs, “The Next Chapter.”

27. Dave Smith, “OpenAI Says It Plans to Report Stunning Annual Losses Through 2028—and Then Turn Wildly Profitable Just Two Years Later,” Fortune, November 12, 2025, https://fortune.com/2025/11/12/openai-cash-burn-rate-annual-losses-2028-profitable-2030-financial-documents/.

28. “OpenAI CFO Says Annualized Revenue Crosses $20 Billion,” Reuters, January 19, 2026, https://www.reuters.com/business/openai-cfo-says-annualized-revenue-crosses-20-billion-2025-2026-01-19/; and “OpenAI Expects Business to Burn $115 Billion Through 2029, The Information Reports,” Reuters, September 5, 2025, https://www.reuters.com/technology/openai-expects-business-burn-115-billion-through-2029-information-reports-2025-09-06/.

29. William Gavin, “The Real AI Bubble May Be in the Private Market,” MarketWatch, January 10, 2026, https://www.marketwatch.com/story/the-real-ai-bubble-may-be-in-the-private-market-2aa25b85.

30. See Tom Liptay and Dan Schwarz, “OpenAI’s Revenue in 2027: A Comprehensive Forecast,” FutureSearch AI, May 12, 2025, https://futuresearch.ai/openai-revenue-forecast/; and Zachary DeWitt, “Notorious: OpenAI’s Revenue Breakdown (A Closer Look),” Notorious PLG, July 18, 2024, https://www.notoriousplg.ai/p/notorious-openais-revenue-breakdown.

31. Matthias Bastian, “Only 5 Percent of ChatGPT’s 900 Million Weekly Users Pay, and Reportedly Most Aren’t Worth Much to Advertisers,” The Decoder, January 5, 2026, https://the-decoder.com/only-5-percent-of-chatgpts-900-million-weekly-users-pay-and-reportedly-most-arent-worth-much-to-advertisers/; and Rebecca Bellan, “OpenAI Boasts Enterprise Win Days After Internal ‘Code Red’ on Google Threat,” TechCrunch, December 8, 2025, https://techcrunch.com/2025/12/08/openai-boasts-enterprise-win-days-after-internal-code-red-on-google-threat/.

32. Smith, “OpenAI Says It Plans.”

33. “Anthropic,” Wikipedia, accessed November 27, 2025, https://en.wikipedia.org/wiki/Anthropic; Ben Sherry, “Anthropic CEO Dario Amodei Says He Left OpenAI Over a Difference in ‘Vision,’” *Inc.*, November 13, 2024, https://www.inc.com/ben-sherry/anthropic-ceo-dario-amodei-says-he-left-openai-over-a-difference-in-vision/91018229; and Dylan Matthews, “The $1 Billion Gamble to Ensure AI Doesn’t Destroy Humanity,” Vox, updated September 25, 2023, https://www.vox.com/future-perfect/23794855/anthropic-ai-openai-claude-2.

34. Anthropic, “The Long-Term Benefit Trust,” news release, September 19, 2023, https://www.anthropic.com/news/the-long-term-benefit-trust.

35. The Trust’s purpose mirrors Anthropic’s public benefit purpose: “to ensure that Anthropic responsibly balances the financial interests of stockholders with the interests of those affected by Anthropic’s conduct and our public benefit purpose.” Anthropic, “The Long-Term Benefit Trust.”

36. John Morley et al., “Anthropic Long-Term Benefit Trust,” Harvard Law School Forum on Corporate Governance, October 28, 2023, https://corpgov.law.harvard.edu/2023/10/28/anthropic-long-term-benefit-trust.

37. Anthropic, “The Long-Term Benefit Trust”; and Morley et al., “Anthropic Long-Term Benefit Trust.”

27. Dave Smith, “OpenAI Says It Plans to Report Stunning Annual Losses Through 2028—and Then Turn Wildly Profitable Just Two Years Later,” Fortune, November 12, 2025, https://fortune.com/2025/11/12/openai-cash-burn-rate-annual-losses-2028-profitable-2030-financial-documents/.

38. Rachel Metz, “Anthropic Finalizes Megaround at $61.5 Billion Valuation,” Bloomberg, March 3, 2025, https://www.bloomberg.com/news/articles/2025-03-03/anthropic-finalizes-megaround-at-61-5-billion-valuation; Amarendar Bhushan Dhiraj, “The $20 Billion AI Duopoly: Inside OpenAI and Anthropic’s Unprecedented Revenue Trajectory,” *CEOWORLD Magazine*, November 25, 2025, https://ceoworld.biz/2025/11/25/the-20-billion-ai-duopoly-inside-openai-and-anthropics-unprecedented-revenue-trajectory/; Rebecca Bellan, “Anthropic Projects $70B in Revenue by 2028: Report,” TechCrunch, November 4, 2025, https://techcrunch.com/2025/11/04/anthropic-expects-b2b-demand-to-boost-revenue-to-70b-in-2028-report/; and Ed Zitron, “Why Everybody Is Losing Money on AI,” *Where’s Your Ed At*, September 15, 2025, https://www.wheresyoured.at/why-everybody-is-losing-money-on-ai/.

30. See Tom Liptay and Dan Schwarz, “OpenAI’s Revenue in 2027: A Comprehensive Forecast,” FutureSearch AI, May 12, 2025, https://futuresearch.ai/openai-revenue-forecast/; and Zachary DeWitt, “Notorious: OpenAI’s Revenue Breakdown (A Closer Look),” Notorious PLG, July 18, 2024, https://www.notoriousplg.ai/p/notorious-openais-revenue-breakdown.

39. “Anthropic Plans an IPO as Early as 2026, FT Reports,” Reuters, December 3, 2025, https://www.reuters.com/business/retail-consumer/anthropic-plans-an-ipo-early-2026-ft-reports-2025-12-03/.

40. Hayden Field, "Amazon-Backed AI Firm Anthropic Valued at $61.5 Billion After Latest Round," CNBC, March 3, 2025, https://www.cnbc.com/2025/03/03/amazon-backed-ai-firm-anthropic-valued-at-61point5-billion-after-latest-round.html; and "Google Agrees to New $1 Billion Investment in Anthropic," CNBC, January 22, 2025, https://www.cnbc.com/2025/01/22/google-agrees-to-new-1-billion-investment-in-anthropic.html.

41. Abhinaya Prabhu, "Amazon Doubles Down on Anthropic: $8B Bet to Power the Future of AI," Tech Funding News, July 11, 2025, https://techfundingnews.com/amazon-anthropic-ai-investment-strategy/.

42. Anirban Ghoshal, "Court Ban on Google AI Stakes Would Hurt Anthropic Clients, Say Analysts," *Computerworld*, February 19, 2025, https://www.computerworld.com/article/3826080/court-ban-on-google-ai-stakes-would-hurt-anthropic-clients-say-analysts.html.

43. Anthropic, "Anthropic Raises $13B Series F at $183B Post-Money Valuation," news release, September 2, 2025, https://www.anthropic.com/news/anthropic-raises-series-f-at-usd183b-post-money-valuation.

44. Deepa Seetharaman, "US Tech Startup Anthropic Unveils Cheaper Model to Widen AI's Appeal," Reuters, October 15, 2025, https://www.reuters.com/business/media-telecom/us-tech-startup-anthropic-unveils-cheaper-model-widen-ais-appeal-2025-10-15/. In particular, Seetharaman noted, "The vast majority – about 80% – of the company's revenue comes from companies, an Anthropic spokesperson said, adding that it has more than 300,000 enterprise customers that use Anthropic tools internally, for products or both." See also Jason Lemkin, "How Anthropic Rocketed to $4B ARR—And Why Your B2B Playbook May Already Be Obsolete," SaaStr, July 7, 2025, https://www.saastr.com/anthropics-4b-arr-the-enterprise-ai-growth-playbook-thats-rewriting-saas-economics/.

45. Metz, "Anthropic Finalizes Megaround"; Dhiraj, "The $20 Billion AI Duopoly"; Bellan, "Anthropic Projects $70B"; and Zitron, "Why Everybody Is Losing Money."

46. According to Terms.law, Delaware's "PBC law (DGCL § 365(b)) similarly shields directors from liability for good faith decisions balancing various stakeholder interests." See "The Rise of Benefit Corporations: A Comprehensive Guide," Terms.law, November 18, 2023, https://terms.law/2023/11/18/benefit-corporations-comprehensive-guide-for-business-attorneys/.

47. Catherine Shu, "Google Acquires Artificial Intelligence Startup DeepMind for More Than $500M," TechCrunch, January 28, 2014, https://techcrunch.com/2014/01/26/google-deepmind/.

48. For the backstory on these Google researchers, see Steven Levy, "8 Google Employees Invented Modern AI. Here's the Inside Story," *Wired*, March 20, 2024, https://www.wired.com/story/eight-google-employees-invented-modern-ai-transformers-paper/. Also see their original paper: Ashish Vaswani et al., "Attention Is All You Need," 31st Conference on Neural Information Processing Systems (2017), https://arxiv.org/abs/1706.03762.

49. Mack DeGeurin, "Why Google Isn't Rushing Forward with AI Chatbots," Gizmodo, December 14, 2022, https://gizmodo.com/lamda-google-ai-chatgpt-openai-1849892728; Britney Nguyen, "Google Execs Say the Company Isn't Launching a ChatGPT Competitor Because of Its Greater 'Reputational Risk' Than Startups," *Business Insider*, December 14, 2022, https://www.businessinsider.com/google-isnt-launching-chatgpt-competitor-due-to-reputational-risk-2022-12; James Vincent, "Google Won't Launch ChatGPT Rival Because of 'Reputational Risk,'" *The Verge*, December 14, 2022, https://www.theverge.com/2022/12/14/23508756/google-vs-chatgpt-ai-replace-search-reputational-risk.

50. Sundar Pichai, "AI at Google: Our Principles," *The Keyword* (blog), June 7, 2018, https://blog.google/technology/ai/ai-principles/.

51. For example, see Tom Simonite, "Google Sets Limits on Its Use of AI but Allows Defense Work," Wired, June 7, 2018, https://www.wired.com/story/google-sets-limits-on-its-use-of-ai-but-allows-defense-work/.

52. For example, in a 2022 update on its principles, Google cited the "central AI Principles team in the Office of Compliance and Integrity" and the "third tier of our AI governance structure is the Advanced Technology Review Council (ATRC), a rotating committee of senior product, research, and business executives." See "2022 AI Principles Progress Update," Google, 2022, https://ai.google/static/documents/ai-principles-2022-progress-update.pdf.

53. Karen Hao, "We Read the Paper that Forced Timnit Gebru Out of Google. Here's What It Says," *MIT Technology Review*, updated December 7, 2022, https://www.technologyreview.com/2020/12/04/1013294/google-ai-ethics-research-paper-forced-out-timnit-gebru/.

54. Khari Johnson, "Google Fires Ethical AI lead Margaret Mitchell," VentureBeat, February 19, 2021, https://venturebeat.com/ai/google-fires-ethical-ai-lead-margaret-mitchell.

55. Jennifer Elias, "Google AI Ethics Team Demands Changes After Dismissal of Renowned Researcher," CNBC, December 17, 2020, https://www.cnbc.com/2020/12/17/googles-ai-ethics-team-makes-demands-of-executives-to-rebuild-trust.html.

56. Johnson, "Google Fires Ethical AI Lead"; and Andrew Tarantola, "What's Going On at Google AI?," Engadget, February 24, 2021, https://www.engadget.com/whats-going-on-at-google-ai-150001422.html.

57. “Google DeepMind,” Wikipedia, accessed December 3, 2025, https://en.wikipedia.org/wiki/Google_DeepMind.

58. Hayden Field, “Google Agrees to New $1 Billion Investment in Anthropic,” CNBC, January 22, 2025, https://www.cnbc.com/2025/01/22/google-agrees-to-new-1-billion-investment-in-anthropic.html; and Connie Loizos, “Google Has Given Anthropic More Funding Than Previously Known, Show New Filings,” TechCrunch, March 12, 2025, https://techcrunch.com/2025/03/11/google-has-given-anthropic-more-funding-than-previously-known-show-new-filings/.

59. Jennifer Elias, “Google Announces Gemini 3 as Battle with OpenAI Intensifies,” CNBC, November 18, 2025, https://www.cnbc.com/2025/11/18/google-announces-gemini-3-as-battle-with-openai-intensifies.html; and “Introducing Gemini 3 Pro for Gemini App,” Google Workspace Updates, November 18, 2025, https://workspaceupdates.googleblog.com/2025/11/introducing-gemini-3-pro-for-gemini-app.html.

60. It’s worth noting that YouTube, which is owned by Google’s parent company, Alphabet, makes money through an engagement-driven ad business similar to Meta’s.

61. Athena Chapekis and Anna Lieb, “Google Users Are Less Likely to Click on Links When an AI Summary Appears in the Results,” Pew Research Center, July 22, 2025, https://www.pewresearch.org/short-reads/2025/07/22/google-users-are-less-likely-to-click-on-links-when-an-ai-summary-appears-in-the-results/.

62. Gartner Inc., “Gartner Predicts Search Engine Volume Will Drop 25% by 2026, Due to AI Chatbots and Other Virtual Agents,” press release, February 19 2024, https://www.gartner.com/en/newsroom/press-releases/2024-02-19-gartner-predicts-search-engine-volume-will-drop-25-percent-by-2026-due-to-ai-chatbots-and-other-virtual-agents.

63. Sharon Goldman, “How Mark Zuckerberg has fully rebuilt Meta around Llama,” *Fortune*, November 19, 2024, https://fortune.com/2024/11/19/zuckerberg-meta-ai-openai-llama/. The article quotes Zuckerberg from a July blog post: “We must ensure that we always have access to the best technology, and that we’re not locking into a competitor’s closed ecosystem where they can restrict what we build.”

64. “Llama (language model),” Wikipedia, accessed November 13, 2025, https://en.wikipedia.org/wiki/Llama_(language_model).

55. Jennifer Elias, “Google AI Ethics Team Demands Changes After Dismissal of Renowned Researcher,” CNBC, December 17, 2020, https://www.cnbc.com/2020/12/17/googles-ai-ethics-team-makes-demands-of-executives-to-rebuild-trust.html.

65. Ahmad Al-Dahle, “With 10x Growth Since 2023, Llama Is the Leading Engine of AI Innovation,” Meta, August 29, 2024, https://ai.meta.com/blog/llama-usage-doubled-may-through-july-2024/. According to this company update, there are more than 60,000 derivative models on Hugging Face, and there is “a vibrant community of developers fine-tuning Llama for their own use cases.”

66. Dylan Butts, "Sam Altman Says Meta Offered OpenAI Staff $100 Million Bonuses, as Mark Zuckerberg Ramps Up AI Poaching Efforts," CNBC, June 18, 2025, https://www.cnbc.com/2025/06/18/sam-altman-says-meta-tried-to-poach-openai-staff-with-100-million-bonuses-mark-zuckerberg.html; and Lily Mae Lazarus, "OpenAI Is Betting Millions on Building AI Talent from the Ground Up amid Rival Meta's Poaching Bid," *Fortune*, July 5, 2025, https://fortune.com/2025/07/05/openai-is-betting-millions-on-building-ai-talent-from-the-ground-up-amid-rival-metas-poaching-pitch/.

67. Matt O'Brien and The Associated Press, "Meta Ramps Up AI Efforts by Paying $14.3 Billion for 49% of Scale and Hires its CEO Alexandr Wang for 'Superintelligence' Team," *Fortune*, June 13, 2025, https://fortune.com/2025/06/13/meta-scale-ai-alexandr-wang-superintelligence-team/; and Scale AI, "Scale AI Announces Next Phase of Company's Evolution," press release, June 12, 2025, https://scale.com/blog/scale-ai-announces-next-phase-of-company-evolution.

68. Ashley Capoot and Jonathan Vanian, "Meta Lays Off 600 from 'Bloated' AI Unit as Wang Cements Leadership," CNBC, October 22, 2025, https://www.cnbc.com/2025/10/22/meta-layoffs-ai.html; and Nadeem Sarwar, "Meta Is Slashing Hundreds of AI Jobs–Here's What's Going On," SlashGear, October 31, 2025, https://www.slashgear.com/2012781/meta-october-2025-ai-layoffs/.

69. Kurt Wagner and Riley Griffin, "Inside Meta's Pivot from Open Source to Money-Making AI Model," *The Mercury News*, December 10, 2025, https://www.mercurynews.com/2025/12/10/inside-metas-pivot-from-open-source-to-money-making-ai-model/.

70. Sherin Shibu, "Meta Is Using Your AI Chatbot Conversations for the Ad Algorithm on Facebook and Instagram," *Entrepreneur*, October 31, 2025, https://www.entrepreneur.com/business-news/meta-is-using-ai-to-shape-ads-on-facebook-instagram/497824; and "Mark Zuckerberg Is Trying to Buy the Future," Qz.com, July 31, 2025, https://qz.com/mark-zuckerberg-meta-ai-spending-capex-advertising.

71. "What is Meta's Business Model, How Does Meta Make Money & Why Is It Being Regulated?," Exactimo, accessed November 27, 2025, *https://exactimo.com/dose/what-is-metas-business-model-why-is-it-being-regulated*; and Andrew Hutchinson, "Meta Provides Guidance for Advertisers Based on Latest Engagement Trends," Social Media Today, May 1, 2025, https://www.socialmediatoday.com/news/meta-shares-ad-tips-q1-2025-reels-AI/746945/.

72. Claire Duffy, "Meta Scores a Win After Judge Rules the Social Media Giant Is Not an Illegal Monopoly," CNN, November 18, 2025, https://www.cnn.com/2025/11/18/tech/meta-antitrust-suit-decision-not-monopoly; Zamaan Qureshi, "Why Meta's FTC Trial Matters for Keeping Young People Safe Online," TechPolicy.Press, April 15, 2025, https://www.techpolicy.press/why-metas-ftc-trial-matters-for-keeping-young-people-safe-online/; and Maria Dinzeo, "Meta Wins Another Chance to Fight US FTC's Reopening of $5 Billion Privacy Deal," MLex, May 16, 2025, https://www.mlex.com/mlex/articles/2341469/meta-wins-another-chance-to-fight-us-ftc-s-reopening-of-5-billion-privacy-deal.

73. Timothy Green, "Meta Investors Learn a Hard Lesson About Super-Voting Shares," The Motley Fool, November 2, 2022, https://www.fool.com/investing/2022/11/02/meta-investors-learn-a-hard-lesson-super-voting/. According to Green, "While Zuckerberg owns less than 13% of Meta's outstanding shares, he controls 54.4% of the voting power. That's because most of his stake is in the form of Class B shares, which grant him 10 votes each instead of just one." See also Ignacio Garcia Giner et al., "Shareholder Democracy and the Challenge of Dual Class Share Structures," Harvard Law School Forum on Corporate Governance, February 11, 2025, https://corpgov.law.harvard.edu/2025/02/11/shareholder-democracy-and-the-challenge-of-dual-class-share-structures/. According to this post, "In the case of Meta Platforms, Mark Zuckerberg holds 99.7% of the company's Class B Stock, which carries 10 votes per share. … Zuckerberg's shares afford him 61% of total voting power."

74. Green, "Meta Investors Learn a Hard Lesson"; and Mark Humphery-Jenner, "Why Meta Shareholders Can Do Little About Mark Zuckerberg," UNSW Business School, November 13, 2022, https://www.businessthink.unsw.edu.au/articles/mark-zuckerberg-Meta-shareholders.

75. OpenAI, "OpenAI and Elon Musk," news release, March 5, 2024, https://openai.com/index/openai-elon-musk/; and Reed Albergotti, "The Secret History of Elon Musk, Sam Altman, and OpenAI," Semafor, March 24, 2023, https://www.semafor.com/article/03/24/2023/the-secret-history-of-elon-musk-sam-altman-and-openai.

76. Jeff Rumage, "Everything You Need to Know About xAI, the Company Behind Grok," Built In, January 7, 2026, https://builtin.com/artificial-intelligence/what-is-xai.

77. Lora Kolodny, "Elon Musk's xAI Secretly Dropped its Benefit Corporation Status While Fighting OpenAI," CNBC, August 25, 2025, https://www.cnbc.com/amp/2025/08/25/elon-musk-xai-dropped-public-benefit-corp-status-while-fighting-openai.html.

78. Maxwell Zeff, "Elon Musk Says xAI Acquired X," TechCrunch, March 29, 2025, https://techcrunch.com/2025/03/29/elon-musk-says-xai-acquired-x/; Lora Kolodny and Ari Levy, "Elon Musk Says xAI Has Acquired X in Deal that Values Social Media Site at $33 Billion," CNBC, March 28, 2025, https://www.cnbc.com/2025/03/28/elon-musk-says-xai-has-acquired-x-in-deal-that-values-social-media-site-at-33-billion.html.

79. xAI, "xAI Raises $20B Series E," news release, January 6, 2026, https://x.ai/news/series-e; and Lora Kolodny, "Elon Musk's xAI Raises $20 Billion from Investors Including Nvidia, Cisco, Fidelity," CNBC, January 6, 2026, https://www.cnbc.com/2026/01/06/elon-musk-xai-raises-20-billion-from-nvidia-cisco-investors.html.

80. "xAI (company)," Wikipedia, accessed January 17, 2026, https://en.wikipedia.org/wiki/XAI_(company); Kolodny, "Elon Musk's xAI Raises $20 Billion."

81. "Musk's xAI Posts Net Quarterly Loss of $1.46 Billion, Bloomberg News Reports," Reuters, January 8, 2026. https://www.reuters.com/technology/musks-xai-posts-net-quarterly-loss-146-billion-bloomberg-news-reports-2026-01-09/.

82. See Amanda Silberling, "Of Course, Grok's AI Companions Want to Have Sex and Burn Down Schools," TechCrunch, July 15, 2025, https://techcrunch.com/2025/07/15/of-course-groks-ai-companions-want-to-have-sex-and-burn-down-schools/; and "Musk's Grok 'Companions' Include a Flirty Anime Character and an Anti-Religion Panda," NBC News, July 16, 2025, https://www.nbcnews.com/tech/internet/grok-companions-include-flirty-anime-waifu-anti-religion-panda-rcna218797.

83. Jakob Steinschaden, "Is This Elon Musk's Master Plan? xAI to Power Tesla's Optimus Robot," Trending Topics, January 9, 2026, https://www.trendingtopics.eu/elon-musk-masterplan-xai-tesla-optimus/.

84. Vikram R. Bhargava and Manuel Velasquez, "Ethics of the Attention Economy: The Problem of Social Media Addiction," *Business Ethics Quarterly* 30, no. 3 (2020): 321–359, https://www.cambridge.org/core/journals/business-ethics-quarterly/article/ethics-of-the-attention-economy-the-problem-of-social-media-addiction/1CC67609A12E9A912BB8A291FDFFE799; Hannah Metzler and David Garcia, "Social Drivers and Algorithmic Mechanisms on Digital Media," *Perspectives on Psychological Science* 19, no. 5 (2024): 957—976, https://journals.sagepub.com/doi/10.1177/17456916231185057; and Tiziano Piccardi et al., "Reranking Partisan Animosity in Algorithmic Social Media Feeds Alters Affective Polarization," *Science* 386, no. 6726 (2024), https://www.science.org/doi/10.1126/science.adu5584.

85. Center for Humane Technology, "Who We Are," accessed December 5, 2025, https://www.humanetech.com/who-we-are; and "Tristan Harris," Wikipedia, accessed December 5, 2025, https://en.wikipedia.org/wiki/Tristan_Harris.

86. Tristan Harris, "Humane: A New Agenda for Tech," presentation transcript, February 2019, https://ethical.net/ethical/humane-new-agenda-tech-tristan-harris/.

87. Alphabet Inc., Annual Report (Form 10-K) for the fiscal year ended December 31, 2024, filed with the Securities and Exchange Commission on February 5, 2025, https://www.sec.gov/Archives/edgar/data/1652044/000165204425000014/goog-20241231.htm.

88. See Google Ads Help Center, "How the Google Ads Auction Works," accessed January 9, 2026, https://support.google.com/google-ads/answer/6366577.

89. Janek Bevendorff et al., "Is Google Getting Worse? A Longitudinal Investigation of SEO Spam in Search Engines," Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval (2024), https://downloads.webis.de/publications/papers/bevendorff_2024a.pdf.

90. Matt O'Brien, "ChatGPT's Free Ride is Ending: Here's What OpenAI Plans for Advertising on the Chatbot," AP News, January 16, 2025, https://apnews.com/article/chatgpt-ads-openai-advertising-83812a066375a805fa2e29b28fc77da1.

91. Google, "A New Era of Intelligence with Gemini 3," *The Keyword* (blog), November 18, 2025, https://blog.google/products/gemini/gemini-3/.

92. Rebecca Bellan, "OpenAI Fires Back at Google with GPT-5.2 after 'Code Red' Memo," TechCrunch, December 11, 2025, https://techcrunch.com/2025/12/11/openai-fires-back-at-google-with-gpt-5-2-after-code-red-memo/; and Yvette Schmitter, "OpenAI Declared Code Red. That's Everybody's Problem," Built In, December 24, 2025, https://builtin.com/articles/openai-code-red-analysis.

93. See Harold Hotelling, "Stability in Competition," *The Economic Journal* 39, no. 153 (March 1929): 41–57, 41, https://doi.org/10.2307/2224214.

94. For example, see Michael Froman, "China, the United States, and the AI Race," Council on Foreign Relations, October 10, 2025, https://www.cfr.org/article/china-united-states-and-ai-race.

95. Executive Order No. 14110, "Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence," Federal Register 88, no. 210 (November 1, 2023): 75191, https://www.federalregister.gov/documents/2023/11/01/2023-24283/safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence.

96. Executive Order No. 14148, "Initial Rescissions of Harmful Executive Orders and Actions," Federal Register 90, no. 18 (January 28, 2025): 8237, https://www.federalregister.gov/documents/2025/01/28/2025-01901/initial-rescissions-of-harmful-executive-orders-and-actions.

97. See Executive Order No. 14179, "Removing Barriers to American Leadership in Artificial Intelligence," Federal Register 90, no. 20 (January 31, 2025): 8741–8742, https://www.whitehouse.gov/presidential-actions/2025/01/removing-barriers-to-american-leadership-in-artificial-intelligence/.

98. See Seamus Dowdall, "Congress Makes Incremental Progress on AI Policy," National Association of Counties (NACo), December 23, 2024, https://www.naco.org/news/congress-makes-incremental-progress-ai-policy.

99. Trevor Wagener, "$600 Billion AI Abundance Dividend from Federal Preemption of State Laws," Computer & Communications Industry Association, November 2025, https://ccianet.org/articles/600-billion-ai-abundance-dividend-from-federal-preemption-of-state-laws/. This article also reported: "In 2025, state lawmakers introduced 1,134 bills to regulate AI and enacted 131 AI regulatory bills into law. Every single state's legislature had at least one AI regulatory bill introduced in 2025, and 40 states enacted at least one AI regulatory law."

100. "AI Legislation, By The Numbers," Multistate, May 17, 2024, https://www.multistate.ai/updates/vol-27; and Matthew Gross, "AI Governance in the States: June 2025," Beckage Firm, accessed December 5, 2025, https://thebeckagefirm.com/ai-governance-in-the-states-june-2025/.

101. Gavin Newsom, veto message to California State Senate on Senate Bill No. 1047, September 29, 2024, https://www.gov.ca.gov/wp-content/uploads/2024/09/SB-1047-Veto-Message.pdf; and "Safe and Secure Innovation for Frontier Artificial Intelligence Models Act," Wikipedia, accessed November 27, 2025, https://en.wikipedia.org/wiki/Safe_and_Secure_Innovation_for_Frontier_Artificial_Intelligence_Models_Act.

102. Adam S. Forman et al., "Colorado's Historic SB 24-205 Concerning Consumer Protections in Interactions with AI Signed Into Law," Epstein Becker Green, May 20, 2024 (updated August 28, 2025), https://www.workforcebulletin.com/colorados-historic-sb-24-205-concerning-consumer-protections-in-interactions-with-ai-signed-into-law-after-passing-state-senate-and-house; and Nathaniel M. Glasser et al., "Will Colorado's Historic AI Law Go Live in 2026?," Epstein Becker Green, August 28, 2025, https://www.healthlawadvisor.com/will-colorados-historic-ai-law-go-live-in-2026-its-fate-hangs-in-the-balance-in-2025.

103. Rebecca Bellan, "The Race to Regulate AI Has Sparked a Federal vs. State Showdown," TechCrunch, November 28, 2025, https://techcrunch.com/2025/11/28/the-race-to-regulate-ai-has-sparked-a-federal-vs-state-showdown/.

104. Tim Hickman et al., "Long Awaited EU AI Act Becomes Law After Publication in the EU's Official Journal," White & Case LLP, July 12, 2024, https://www.whitecase.com/insight-alert/long-awaited-eu-ai-act-becomes-law-after-publication-eus-official-journal; and European Parliament, "EU AI Act: First Regulation on Artificial Intelligence," accessed December 5, 2025, https://www.europarl.europa.eu/topics/en/article/20230601STO93804/eu-ai-act-first-regulation-on-artificial-intelligence.

About the Institute

The Clayton Christensen Institute for Disruptive Innovation is a nonprofit, nonpartisan think tank dedicated to improving the world through Disruptive Innovation. Founded on the theories of Harvard professor Clayton M. Christensen, the Institute offers a unique framework for understanding many of society's most pressing problems. Its mission is ambitious but clear: work to shape and elevate the conversation surrounding these issues through rigorous research and public outreach.

About the Author

Thomas Arnett is a senior research fellow for the Clayton Christensen Institute. His work focuses on using the Theory of Disruptive Innovation to study innovative instructional and technological models and their potential to scale affordability and accessibility across markets. He holds a BS in economics from Brigham Young University and an MBA from the Tepper School of Business at Carnegie Mellon University.

Note: The author used AI tools as a thought partner during the drafting and revision process, primarily to test arguments, clarify language, check sources, and improve structure. All final judgments, interpretations, and conclusions are the author's own.