

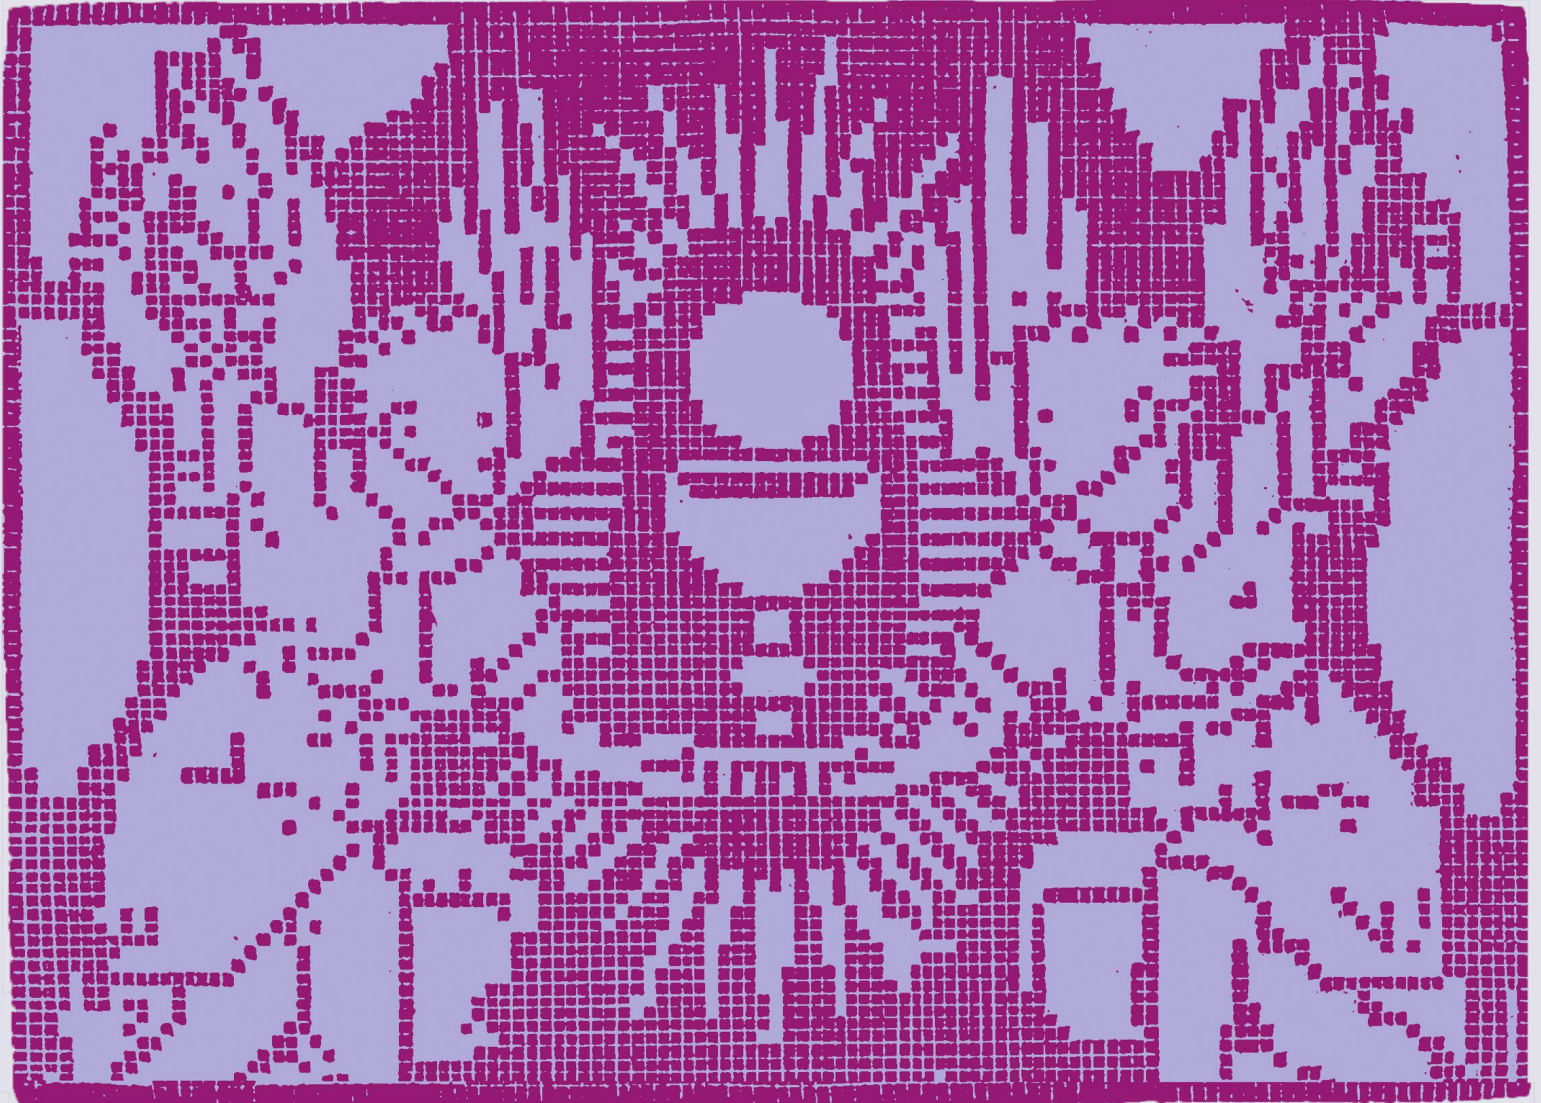


AINOW

ARTIFICIAL POWER

2025 Landscape Report

June 3, 2025



ACKNOWLEDGMENTS

Authored by Kate Brennan, Amba Kak, and Dr. Sarah Myers West.

With research support from Mohammed Ali, Yasmine Chokrane, Madeline Kim,
Tekendra Parmar, Tanya Raja, and Boxi Wu.

Special thanks to Rhiana Gunn-Wright for her feedback and editorial contributions.

Project management by Ellen Schwartz.

Copyediting by Caren Litherland.

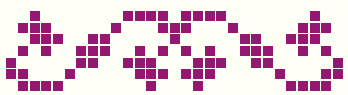
Design by Partner & Partners.

With contributions and feedback from:

- ❖ Annette Bernhardt, UC Berkeley Labor Center
- ❖ Abeba Birhane, Artificial Intelligence Accountability Lab, Trinity College Dublin
- ❖ Brian Chen, Data & Society
- ❖ Jane Chung, Justice Speaks
- ❖ Stefanie Coyle, New York Civil Liberties Union
- ❖ Andrea Dehlendorf, AI Now
- ❖ Kevin DeLiban, TechTonic Justice
- ❖ Becca Deutsch, Amazon Employees for Climate Justice
- ❖ Alix Dunn, The Maybe
- ❖ Ali Finn, AI Now
- ❖ Timnit Gebru, Distributed AI Research Center
- ❖ Ryan Gerety, Athena Coalition
- ❖ Lisa Gilbert, Public Citizen
- ❖ Sam Gordon, Tech Equity Collaborative
- ❖ Janet Haven, Data & Society
- ❖ Nidhi Hegde, American Economic Liberties Project
- ❖ Ben Inskeep, Citizen Action Coalition of Indiana
- ❖ Taylor Jo Isenberg, Economic Security Project
- ❖ Heidi Khlaaf, AI Now
- ❖ Stephen Lerner, Action Center for Race and the Economy - Bargaining for the Common Good
- ❖ Barry Lynn, Open Markets Institute
- ❖ Varoon Mathur, Duke Institute for Health Innovation, Duke Health
- ❖ Jill McArdle, Beyond Fossil Fuels
- ❖ Michelle Meagher, SOMO
- ❖ Brian Merchant, Independent Journalist
- ❖ Erie Meyer, Georgetown University
- ❖ Sarah Miller, American Economic Liberties Project
- ❖ Stacy Mitchell, Institute for Local Self-Reliance
- ❖ Arvind Narayanan, Center for Information Technology Policy, Princeton University
- ❖ Chris Nielsen, National Nurses United
- ❖ Teri Olle, Economic Security Project
- ❖ Britt Paris, Rutgers University
- ❖ Tekendra Parmar, Investigative Journalist
- ❖ Reshma Ramachandran, Yale School of Medicine
- ❖ Steven Renderos, MediaJustice
- ❖ Rashida Richardson, Worcester Polytechnic Institute & Northeastern University
- ❖ Hilary Ronen, Local Progress
- ❖ Leevi Saari, AI Now
- ❖ Hannah Sassaman, People's Tech Project
- ❖ Matt Scherer, Center for Democracy and Technology
- ❖ Paromita Shah, Just Futures Law
- ❖ Ganesh Sitaraman, Vanderbilt University
- ❖ Andrew Strait, UK AI Security Institute
- ❖ Kasia Tarczyska, Good Jobs First
- ❖ Jim Thomas, Scanthehorizon.org
- ❖ Max Von Thun, Open Markets Institute
- ❖ Jai Vipra, Cornell University
- ❖ Robert Weissman, Public Citizen
- ❖ Meredith Whittaker, Signal
- ❖ Savannah Wilson, Clean Virginia
- ❖ Boxi Wu, Oxford University

Cite as: Kate Brennan, Amba Kak, and Sarah Myers West, “Artificial Power: AI Now 2025 Landscape”, AI Now Institute, June 3, 2025, <https://ainowinstitute.org/2025-landscape>.

TABLE OF CONTENTS



Executive Summary 4

Chapter 1: AI’s False Gods 18

 1.1: The AGI Mythology 19

 1.2: Too Big to Fail 24

 1.3: AI Arms Race 2.0 28

 1.4: Recasting Regulation 32

Chapter 2: Heads I Win, Tails You Lose 36

Chapter 3: Consulting the Record 46

Chapter 4: A Roadmap for Action 67

Endnotes 93



EXECUTIVE SUMMARY

Those of us broadly engaged in challenging corporate consolidation, economic injustice, tech oligarchy, and rising authoritarianism need to contend with the AI industry or we will lose the end game. Accepting the current trajectory of AI proselytized by Big Tech and its stenographers as “inevitable” is setting us up on a path to an unenviable economic and political future—a future that disenfranchises large sections of the public, renders systems more obscure to those it affects, devalues our crafts, undermines our security, and narrows our horizon for innovation. This is true whether or not the technology even works well, on its own terms; it often doesn’t.

The good news is that the road offered by the tech industry is not the only one available to us. This report explains why the fight against the industry’s vision for AI is a fight worth having, even as we turn ourselves tirelessly toward the task of building out the shared project of a just, equitable, sustainable, and democratic society.

Over the past decade, taming the power of big technology platforms like Microsoft, Amazon, Google, and Meta has increasingly become a central question in American political and public discourse. Unless we contend with the power vested in these firms, we won't meaningfully be able to hold the industry accountable to the interests of the broader public, even as these companies reshape markets, institutions, and infrastructures core to public life.

What does AI have to do with any of this? As we argued in our 2023 report,¹ AI is fundamentally about concentration of power in the hands of Big Tech. At the start of the year, it seemed like the market was poised for disruption, with a new crop of Silicon Valley challengers gaining prominence, like OpenAI, Anthropic, StabilityAI, and Inflection AI. But now, just two years later, it is clear that the bench of key players in this market hasn't changed much: Microsoft, Google, Meta, Musk's xAI, OpenAI (backed by Microsoft), and Anthropic (backed by Amazon and Google).² The new suite of LLM-powered AI products has pushed these firms into the spotlight, dominating headlines and, increasingly, becoming the subject of dinner-table conversation.

Amid the frenzy, there's been a misplaced focus on blinkered questions of whether one AI system or application is good or bad, or evaluating the moral quandaries of hypothetical worlds. Instead we need to redirect attention to the AI ecosystem, and its dependencies and risks, as a whole. **The question we should be asking is not if ChatGPT is useful or not, but if OpenAI's unaccountable power, linked to Microsoft's monopoly and the business model of the tech economy, is good for society.**

Looking beyond individual use cases allows for a more comprehensive look into the centers of power that drive our current tech landscape. AI as a field has been not just co-opted but *constituted* by the logics of a few dominant tech firms. It is no coincidence that the "bigger-is-better" paradigm that dominates

the field today, where the scale of compute and data resources are generally used as a proxy for performance, lines up neatly with the incentives of Big Tech, which disproportionately controls these resources, the talent to leverage them, and the pathways to monetization. Around 2012, as it became apparent that substantial gains in model performance could come simply from applying larger and larger scale data and computational resources to existing algorithms, tech giants moved quickly to shore up their existing advantages and hire talent.³ Corporate influence over AI's research trajectory has been cemented through tech firms' AI labs dominant presence at prestigious machine learning conferences, further shaping the field of research in ways that align with industry.⁴ This is in part because building AI bigger requires enormous resources, both financial and social, to achieve unrestricted growth at breakneck speed—resources that AI companies own and control.⁵

But it's not just market power we need to be concerned with: These tech oligarchs are counting on a wholesale rewriting of our social and economic foundations, using AI as the justification. From breaking apart the US federal government and raiding citizen data under the guise of efficiency, to redesigning workflows to devalue human labor and creativity so they are AI-ready, to redirecting our entire energy infrastructure to prioritize their technology over people's basic needs, the vision promulgated by tech oligarchs requires, as a foundation, the unraveling of core social, political, and economic fabrics.

Across our information ecosystem, from science to education, healthcare, culture, and art, AI is being positioned as a disruptive new infrastructure and a mediating force. **In truth, though, it rehashes an old playbook, helicoptering in solutions built on the extraction of expertise and value from all corners of society—solutions that always, eventually, amount to the further degradation of life for the most**

marginalized among us. While generative AI and AI agents are the buzzwords that splash across the headlines, the same dynamics are true of precursors to contemporary AI systems like automated decision-making technologies used in banking, hiring, and criminal justice. The techniques and vendor names vary, but the industry incentives powering proliferation, as well as the failure modes across these systems, share much in common.

More than a decade of evidence demonstrates how it goes: The introduction of these systems concentrates power among the deployers of the tech, leaving those on the receiving end more insecure, vulnerable, and unable to contest the determinations made by the “smart machine” at the expense of the broader public. These tools are often invisible to those judged by them, and inscrutable even when they are visible.

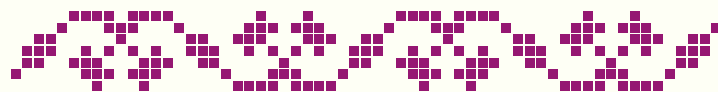
Why society would ever accept this bargain is the critical question at hand. Amid the excitement over AI’s (speculative) potential, the sobering reality of its present and recent past is obscured. When we consult the record on how AI is already intermediating critical social infrastructures, we see that it is materially reshaping our institutions in ways that ratchet up inequality, render institutions opaque to those they are meant to serve, and concentrate power in the hands of the already powerful. (See Chapter 3: Consulting the Record.) It makes clear that for all the whiz-bang demos and bold Davos proclamations, on the ground AI is consistently deployed in ways that make everyday people’s lives, material conditions, and access to opportunities worse and the systems that incorporate them stronger.

This report’s title, *Artificial Power*, captures the critical, and at times contradictory, moment we find ourselves in. On one hand, the tech oligarchy has successfully deployed “AI”—as a strategic marketing term and as a set of automation technologies—to cement and grow its power. At the same time, this power is vastly

inflated, contingent, and poised for disruption. Contending with the dual reality of how those with power have deployed AI systems to enact significant harm while exposing the ways this power can and must be disrupted is the central work of this moment.



THE ELEPHANTS IN THE AI ROOM: OF BUSINESS MODEL(S) AND FATAL TECHNICAL FLAWS



The AI industry is on shakier ground than it may seem. Valuations are sky high, while the business model hinges on an intensely expensive technology that lacks a consistent revenue stream. AI companies bleed money for every user they gain: Anthropic burned through \$5.6 billion⁶ this year but was valued at \$61.5 billion.⁷ OpenAI lost \$5 billion⁸ but is valued at \$300 billion.⁹ No profit-making use cases exist yet, or are even on the horizon. This may seem like business as usual for the move-fast-and-break-things ethos of Big Tech, but we are in a profoundly different monetary environment now—it’s no longer the 2000s or 2010s. Markets are saturated, market dominance has been established among the platform and infrastructure winners that emerged from those decades, and, put simply, the cost of large-scale AI is eye-watering at a level not seen before in tech.

The question now circulating more, and more openly, is this: When will the AI bubble burst and who will be impacted by it? Because this *is* a bubble. For their part, companies are pushing out shiny objects to detract from the business reality while they desperately try to derisk their portfolios through government subsidies and steady public-sector (often carceral or military) contracts. While it's very clear how tech companies benefit from claims to "public-interest AI" used to justify the pouring of taxpayer dollars into this industry, it is not at all clear how this benefits the rest of society. (See Chapter 2: Heads I Win, Tails You Lose)

Chapter Snapshot

Heads I Win, Tails You Lose



This section maps the drivers that are securing Big Tech firms' advantage in the AI market, before turning to the question of who loses in the end:

- ❖ Cloud infrastructure providers benefit from cycles of AI dependence
- ❖ Big Tech firms benefit from leveraging control over the tech ecosystem
- ❖ Big Tech benefits from the data center boom ... even if the AI boom doesn't pan out

With generative AI, in particular, the hyped claims stand in stark contrast to the largely mundane use cases that are being shoved into nearly every app and service. In contrast to the claims of world-changing tech, Meta is investing heavily in AI advertising infrastructure.¹⁰ OpenAI is creating AI agents that fill out forms and call web browsing "research"—sucking up your data and requesting invasive permissions as it does so.¹¹ Pressured by their employers, software engineers are using Microsoft's Copilot to produce more code, more rapidly, undercutting their skills and trade.¹² And cloud companies are happily locking enterprise users into their software-as-a-service (SaaS) ecosystems by automatically upgrading them to new AI features—and raising the price.¹³ These are not examples of a technology being embraced by a society glad of its utility.

Despite being positioned as critical infrastructure, AI systems in their current form have fundamental flaws: there is an intractable problem of "hallucinations" with LLMs reliant on randomly generated coordination, leaving the humans in this technology in the unenviable position of fact-checking the tech meant to make their lives easier.¹⁴ Peer-reviewed research indicates that in many cases, AI systems fail profoundly at even basic tasks when applied in real-life contexts.¹⁵ They're also far from resilient, prone to cybersecurity vulnerabilities like web poisoning attacks and new jailbreaking methods that enable the persistent unauthorized disclosure of training data and other sensitive information.¹⁶ And it's not that the trade-offs are weighted and deemed worth it: In many use cases, AI is deployed by those with power against those who lack it, and who have no opportunity to opt out or seek remedy when mistakes are made. (See Chapter 3: Consulting the Record.)

Chapter Snapshot

Consulting the Record



This section compiles over a decade of evidence showing how the tech industry has sought to reshape society to enable more widespread deployment of the technologies it builds and profits from, often contributing to the degradation of our social, political, and economic lives. The section aligns on five key takeaways:

- ❖ **AI's benefits are overstated and underproven**, from cancer cures to hypothetical economic growth - while some of its flaws are real, immediate, and growing.
- ❖ **AI-sized solutions to entrenched social problems displace grounded expertise**, in disparate domains like higher education, healthcare, and agriculture.
- ❖ **AI solutionism obscures systemic issues facing our economy – obscuring economic concentration and acting as a conduit for deploying austerity mandates by another name.** The DOGE power grab is instructive, though New York's MyCity offers another example where millions of taxpayer dollars were invested into flawed AI solutions that failed to deliver tangible benefits to the public.
- ❖ **The productivity myth obscures a foundational truth - the benefits of AI accrue to companies, not to workers or the public at large**, even as algorithmic management tools make work unstable and unsafe. 'Agentic AI' will make workplaces even more bureaucratic and surveillant, reducing not increasing autonomy.
- ❖ **AI use is frequently coercive, violating rights and undermining due process.** This is nowhere more clear than the rise of AI usage in immigration enforcement, where human rights abuses are common and legal norms are routinely violated - even before AI is in the mix.

It doesn't help that the prevailing deregulatory current is an industry that continually acts above the law and is driven narrowly by its bottom line (See 1.4: Regulation): in 2024, we saw companies rush to market

with products that are patently inaccurate, insecure, and compromise the safety of consumers; engage in anti-competitive practices that shore up their advantages to shut the door behind them; and deploy larger than life narratives around AGI and innovation to quell any form of interrogation and critique.

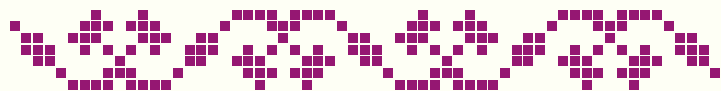
Chapter Snapshot

AI's False Gods



This section interrogates narratives that advance AI industry dominance and make the current trajectory of AI seem inevitable:

- ❖ **The AGI Mythology: The Argument to End All Arguments** unpacks the nebulous claims surrounding "artificial general intelligence," arguing that the term collapses complex technical realities into a singular, imminent, and inevitable future that conveniently advances the interests of the companies claiming to build it.
- ❖ **'Too Big to Fail:' Infrastructure and Capital Push** explores how tech firms are deploying unprecedented amounts of capital to perpetuate a "bigger-is-better" AI paradigm, shoring up their continued market dominance through government and taxpayer support.
- ❖ **AI Arms Race 2.0: From Deregulation to Industrial Policy** details how the US-China AI arms race has heightened, and is now used to brand a slate of industrial policy initiatives designed to boost the tech industry and avert regulatory scrutiny.
- ❖ **Recasting Regulation as a Barrier to Innovation** shows how the AI industry has strategically pitted regulation against innovation, leading to a global deregulatory posture that ignores the role regulation plays in enhancing innovation and competition.



IT IS TIME TO BUILD— JUST NOT AI.



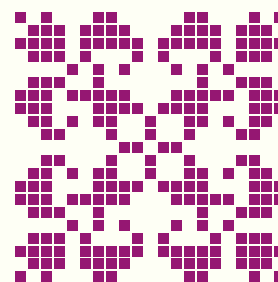
AI is predominantly used *on* us, not *by* us, to shape access to resources and life chances. But while there is a clear path dependency within this narrow trajectory for AI proselytized by big tech and its stenographers, the good news is that it's not the only road available to us.¹⁷ Not by a long shot.

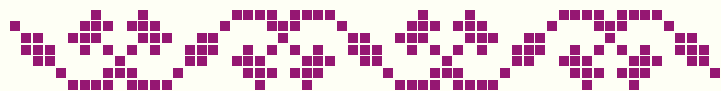
AI hype has tapped into a sentiment that is real and widespread: genuine enthusiasm to build a future where all people can thrive, a future that will likely look radically different from the present. It is a catalyzing goal we should unite around; most of us want a future that frees us from the endless cycle of war, pandemics, and environmental and financial crises that characterize our present. The 2024 US presidential election brought the need to create social and political institutions connected to the needs and lived realities of people even closer to home, across the country and across the world. But AI doesn't create any of these—and pegging our shared future on AI makes that future harder, not easier, to achieve because it binds us to a decidedly bleak path, stripping us not only of the ability to choose what to build and how to build it, but also stealing what joy we might take in that building. This hype-prescribed AI future further distances us from a life with dignity, one where we have the autonomy to make our own decisions and where democratically accountable structures work to distribute power and

technological infrastructures in ways that are robust, accountable, and protected from systemic shocks.

What we've seen play out within the AI industry is not unique to this industry, of course. The dynamics of “gain for me, loss for thee” have been examined in many critiques of shareholder capitalism broadly, which emphasize the corporate willingness to speculate in ways beneficial to shareholders but not to society at large, alongside the perverse incentives that lead firms to act against their own business interests, and develop an orientation toward monopolization and sclerosis. If anything, the AI market is the peak exemplar of overreliance on venture-based investment.¹⁸

But AI introduces new dynamics and accelerants. As designed, developed, and deployed currently, AI works to entrench existing power asymmetries, and to ratchet them up. It naturalizes inequity as destiny and deservedness—simply the classification given by the intelligent system—while rendering these underlying patterns, judgments, and self-interested drivers inscrutable to those affected by AI's judgments and instructions.





ANOTHER AI IS POSSIBLE. HOW DO WE GET THERE?



Although real ways in which the AI market could be structured to benefit the public may exist, the path charted by the companies controlling AI, and those wanting a piece of the control AI could give them over our lives and institutions, won't lead us there.

One thing is clear: we can't fight tech oligarchy without rejecting the current industry trajectory around large scale AI. **It's a crucial inflection point and how American policymakers and movement leaders choose to respond to the AI industry will write the coming chapters of the story of tech power.** AI companies, and those who lead them, have positioned themselves to reshape broad swaths of society—within and beyond the US—not only to work in their interest, but to do so in ways that allow their firms to capture the lion's share of the value.

This isn't inevitable. In fact, the tide of public opinion is moving decisively against the entrenched power of tech firms. And we've seen major legal wins in the landmark antitrust cases filed by the DOJ and FTC against Google and Meta. After successfully proving that Google maintains an illegal monopoly in search and advertising markets,¹⁹ the DOJ is now requesting bold, structural remedies that were all but inconceivable a few years ago.²⁰ These remedies, which include

breaking up Google's advertising technology business and spinning off Chrome, strike directly at the heart of Google's business model. But the remedy trials have revealed a larger truth: AI startups can't scale or achieve distribution without Big Tech firms' infrastructure. That's why OpenAI offered to buy Chrome. It's why Perplexity's CEO said he'd want to buy it—then pandered so as not to aggravate a company he's dependent on.²¹ This is why it's especially important that we not cede the momentous ground these regulatory actions have pushed us towards when Big Tech companies use AI as cover for staying unregulated.

In this report, we lay out another path forward. First, we map what we mean by AI in the first place, provide an accounting of the false promises and myths surrounding AI, and examine whom it's working for and whom it's working against. Then, given that AI consistently fails the average public, even as it enriches a sliver, we ask what we lose if we accept the current vision of AI peddled by the industry. Finally, we identify leverage points that we can latch on to as we mobilize to build a world with collective thriving at its center—with or without AI.

We are not naive. The headwinds against sensible AI alternatives have never been stronger. The tech industry is better resourced and the political environment more bleak than ever before. Indeed, considering the tech industry to be simply a collection of firms itself misses key sources of their power, from their surveillance apparatus to their control over global digital infrastructures that shape our states, our institutions, our economies, and, most importantly, our lives. But the battle has never been more important. Contesting AI links the movements we must build not only to create meaningful public power, but also to seed a new path defined by autonomy, dignity, respect, and justice for all of us.

In the following sections, we set out a tool kit for reasserting public power amid a takeover by AI firms.

A more fulsome treatment appears in [Chapter 4: A Roadmap for Action](#), but here are the highlights:

Chapter Snapshot

A Roadmap for Action: Make AI a Fight About Power, Not Progress

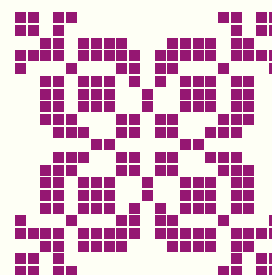


1. **Target how the AI industry works against the interests of everyday people.** In the aftermath of the 2024 election, there's growing consensus across the ideological spectrum that focusing on the material conditions and economic interests of working people is key to building political power. We need to not only make AI-related issues more relevant to movements fighting for economic populism and against tech oligarchy; we also need to better target the AI industry as a key actor working against the interests of the working public. The pushback against the Department of Government Efficiency (DOGE), the buildout of AI data centers, and algorithmic prices and wages constitutes fertile ground for building a broader movement unified in its focus on rejecting AI's unaccountable tech-enabled social and political control.
2. **Advancing worker organizing is the clearest path to protecting us and our institutions from AI-enabled capture.** Labor campaigns have demonstrated that working people have a particular form of power to wield, power that can determine how their employers deploy AI and digital systems. The deeper opportunity for labor, and a more transformative ambition, however, would be to direct labor's power not just toward whether and how AI is used in the workplace, but also toward recalibrating the technology sector's power overall and shaping the trajectory of AI in the public interest and common good.
3. **Enact a 'zero trust' policy agenda for AI.** Trust in AI firms' benevolence is not a smart, informed, or credible option—not if we're going to proceed with

serious work. Enacting a policy agenda built on bright-line rules that restrict the most harmful uses of AI, regulate the AI life

cycle from nose to tail, and ensure that the industry that currently creates and profits from AI isn't left to regulate and evaluate itself—essentially grading its own homework—must be a priority at the state and federal levels in the US, and internationally.

4. **Bridge networks of expertise, policy, and narrative to strengthen AI advocacy.** AI advocacy and policy has often been undermined too often by blinkered views that fail to see the different components of the AI supply chain materially, are often single issue-focused, and it is easy to miss the ways in which big picture narratives manifest to limit possibilities in policy fights. From national security logics that can be a vector both for, and cutting against, moves towards industry accountability; to reframing traditional data privacy levers as key tools in the fight against automation and addressing market power.
5. **Reclaim a positive agenda for public-centered innovation without AI at the center.** The current trajectory of AI puts the public under the heel of unaccountable tech oligarchs. But their success is not inevitable. By moving out from under the shadow of the idea that large-scale AI is inevitable, we can reclaim the space to conduct real innovation and to push for exciting and novel alternative pathways that leverage tech to shape a world that serves the public and is governed by our collective will.

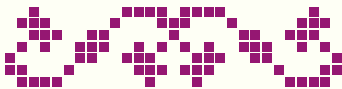




EPILOGUE: THE WORLD WE WANT (AND WHY THE CURRENT TRAJECTORY OF AI WON'T GET US THERE)

The AI hype of the past year has sucked the air out of an already stuffy room, making it feel futile—at times, impossible—to imagine anything other than a steady march toward the inevitable supremacy of AI. But no matter how true that may feel, it is only that: a feeling. It is not reality—not yet, at least. There are, in fact, many alternatives to this version of AI, many ways to shape new worlds. Like AI, though, these are not inevitable either. Making them possible starts by asking and answering a single question: Is this the world we want? At AI Now, we want to see a world that has:

Good Jobs



Everyone deserves access to the resources to live a happy, fulfilled life, and a dignified job that will provide them with these resources. Under the right conditions—like policies intent on uplifting, rather than exploiting, workers—new technologies have the potential to make everyone’s working lives better.

Yet AI’s current trajectory is fundamentally incompatible with the proliferation of good jobs rooted in human flourishing. As the AI industry embeds itself into nearly every sector of the economy, firms are increasingly shaping a job market contingent on

worker displacement and exploitation. Overwhelmingly, AI companies are embedding “productivity” tools designed to help businesses optimize their bottom line across the entire labor supply chain. This requires work itself to become legible to AI systems, making working life more routinized, surveilled, and hierarchical. Furthermore, instead of working to protect workers from the uncertainty coming from this new market, AI companies are undermining hard-won labor protections, exploiting legal loopholes to avoid corporate accountability, and lobbying governments to support policies that prioritize corporate profits over fair and just treatment for workers.

As the current vision of AI takes hold, we lose a future where AI technology works in support of stable, dignified, and meaningful jobs. We lose a future

- ❖ where AI supports fair and livable wages, instead of wage depreciation;
- ❖ where AI ensures that workers have the control to decide how new technology affecting their careers is deployed, instead of undermining their expertise and knowledge of their own work;
- ❖ where we have strong policies to support workers if and when new technologies automate existing roles—including laws that broaden the social safety net—instead of AI boosters who brag to shareholders about cost savings from automation;
- ❖ where robust public benefits and time-off policies ensure the long-term wellness of employees, instead of AI being used to surveil and nickel-and-dime workers at every turn;
- ❖ where AI helps protect employees from health and safety risks on the job, instead of perpetuating conditions that make work dangerous and celebrating employers who exploit labor loopholes to avoid responsibility; and

- ❖ where AI fosters meaningful connection through work, instead of driving cultures of fear and alienation.

Shared Prosperity



The proliferation of any new technology has the potential to increase economic opportunity and lead to widespread shared prosperity. *But shared prosperity is incompatible with AI's current path toward maximizing shareholder profit.*

The insidious myth that AI will lead to “productivity” for everyone when it really means productivity for a select number of corporate firms propels us further down the path of shareholder profit as the singular economic goal. Even well-intentioned government policies designed to boost the AI industry steal from the pockets of workers. For example, government incentives meant to revitalize the chip manufacturing industry were thwarted by corporate buy-back provisions, sending millions of dollars to companies, not to workers or job creation. And despite some meaningful moves to investigate the AI industry under the Biden Administration, companies have still gone largely unchecked, meaning new entrants cannot come in to challenge these practices.

By proliferating the myth that AI will inherently lead to shared prosperity (or that “a rising tide lifts all boats”), we lose the economic policies that could meaningfully lead us into a period of shared prosperity, including pro-enforcement policies to break up the concentration of corporate power, a strong pro-labor agenda to center the needs of workers, and industrial policy

strategies designed to put workers and communities before the bottom line of large corporations. Crucially, we lose a thriving and competitive economy, where innovators and entrepreneurs are incentivized to launch sustainable and prosperous businesses that need not rely on surveillance mechanisms, hyper-growth venture capital funding, and extractive business models to succeed.

Freedom & Autonomy



We all deserve to live in a world where our personal, political, and economic lives are free from coercion.

But the current trajectory of AI is grounded in coercion and opacity. By amassing such concentrated power, AI companies have assumed control over many aspects of our lives, subjecting us to coercive practices designed to maximize their own profit-making potential. Nowhere is this more explicit than the regime of surveillance pricing algorithms. These algorithms collect extraordinary amounts of consumer data to set individualized prices for goods and services, such that important aspects of daily life—from buying groceries, taking an Uber, or buying an apartment—are controlled by companies looking to squeeze consumers to pad their own bottom lines. As an inherently centralizing technology, AI is consistently deployed to support and benefit from surveillance states, carceral systems, and military techniques, embedding corporate interests into state apparatuses and vice versa, making life more coercive and violent.

AI could contribute to a future built on autonomy and transparency. But right now, we stand to lose a robust

ecosystem of public resources that are not dependent on private industry players or priced for profit; pricing for goods and services driven by principles of fair competition; and changing life circumstances not being treated as grounds for companies to profit. A public life invested in freedom would be divested from the proliferation of surveillance systems, carceral logics, and military state apparatuses, enabling people and communities to thrive.

Sustainable Future



Technological progress does not have to come at the expense of our natural environment. A publicly beneficial AI landscape places principles of sustainability and environmental justice at its center, recognizing that the health and safety of our planet and communities is of paramount importance.

AI's current trajectory is not merely incompatible with sustainability; it is fueling climate degradation. The focus on scale at all costs within the AI industry makes it dependent on climate extraction and energy dominance. These include a deregulatory policy intent on accelerating the AI industry (meaning more energy, more infrastructure, more natural resources going into AI) as well as the insidious belief that AI is going to help *so/ve* the climate crisis, “greenwashing” the environmental harm it is already enacting.

As a result, what we lose is a world where governments and companies work together to advance principles of sustainability and environmental justice, investing in green and renewable energy infrastructure to support the additional energy usage that new

technologies require. Decisions on siting, permitting, and constructing new technology infrastructure should be made in relationship with local communities, with particular attention to those communities who are most significantly affected by technology's industrial evolution. Cities and states should not be coerced into providing subsidies for infrastructure that come at the expense of local communities' needs, like funding for schools and healthcare. We also lose a future where industry takes its climate commitments seriously and works to mitigate the harmful effects of its industrial processes, recognizing that the health and safety of our planet and communities is key to everyone's survival—including industry's.

Strong Social Safety Net



We should live in a world where everyone has access to and community control over a robust system of public resources.

Yet as AI rapidly shifts our social landscape, governments are increasingly driven by privatization and austerity measures. AI is ushering in a policy agenda designed to enrich private interests rather than provide robust public benefits. AI firms are pushing AI integration in local and federal government agencies driven by austerity, restricting people's access to needed resources and the social safety net. Moreover, the purported AI arms race with China is being leveraged to convince governments that domestic infrastructure is an imperative for national competitiveness and security, encouraging public agencies to throw generous tax exemptions at private companies

in order to build massive data centers in communities that may not want them there. These exemptions—totaling billions of dollars—steal investment from strong public resources that benefit everyone, like investments in more teachers, roads, and libraries.

We could have a world where technology works in service of the broader public, like algorithms designed to maximize people's options, not reduce them; and where companies are required to pay their fair share of taxes to local communities, rather than fighting efforts to make them pay their fair share.

Security



In an increasingly complex world, security and resilience are more important than ever. Our infrastructures are largely invisible to us until the moment they break down - and we've felt the shock of infrastructural failure frequently over recent years, from supply chain issues during the pandemic to power outages to bank closures.

Layering AI into our critical infrastructures, particularly when AI systems are highly concentrated, creates a real and present security risk. These risks are manifold: there are cybersecurity risks that are inherent to many AI systems that make them vulnerable to hacking, and some of these cannot be remedied. There are systemic risks introduced by overreliance on a single technology: for example, if banks, hospitals, and schools all use the same cloud infrastructure provider, an outage could affect all of these at one fell swoop. These risks are at their highest when AI is used in life or death settings - and from healthcare to defense, these are some of the industry's prime markets. And there are risks emanating from decisions made by

the companies themselves to experiment in the wild, bringing to market technologies that have not been adequately tested or validated and with little certainty that they will work as intended, let alone cause harm.

We could have a world where our safety and well-being are not vulnerable and exposed to an industry that is scaling at an unprecedented rate with little regard for safety and security, let alone compliance with the law; where AI is validated, tested and built safe by design and used with prudence rather than impunity.

Innovative Tech Ecosystem



Technology has the potential to solve important societal challenges and push the frontiers of innovation forward. In a thriving tech ecosystem, companies big and small are able to succeed, not by amassing concentrated power but by engaging in fair competition. Society benefits from the distribution of a diverse set of products, services, and technologies that result from such competition.

The current AI industry is defined by concentration, precluding a truly diverse and innovative tech ecosystem from flourishing. There is no AI without Big Tech firms, which have spent decades amassing unrestrained data access and economic power and then used those advantages to control key inputs at all levels of the AI stack. Even where new entrants are able to enter the AI market, they are still dependent on the cloud and computing infrastructures of Big Tech firms in order to succeed, creating an ecosystem of dependence rather than competition. Where consolidation has been averted - as was the case with the proposed

Nvidia-ARM merger - firms have been able to thrive. ARM went on to IPO and beat quarterly estimates, all after its acquisition was blocked. Furthermore, Big Tech companies control most pathways to consumers and enterprise businesses at scale. This centralized power is also driving a crisis of innovation, where Big Tech companies are bloated, stalled by legal reviews, and are stuck repackaging their existing technologies in order to revive and boost their bottom lines.

What we lose in this bloated and stale tech ecosystem is a truly diverse horizon of possibility, filled with innovations that tackle people's real-life needs, rather than an endless soup of enterprise software and AI agents. We can inspire an ecosystem where people can use technology to build companies slowly and sustainably, without the need to grow rapidly and amorally in order to stay in business. Our entire tech ecosystem is in need of a paradigm shift, one that tears down existing structures to make room for complexity and emergence.²² This includes breaking up big companies, overhauling the VC-backed funding structure so more companies can thrive, investing in public goods to ensure tech resources are not dependent on large private companies, and increasing institutional investment to bring more diverse people—and thus ideas—into the tech workforce.

Vibrant, Democratic State



We deserve a technological future that works to support strong democratic values and institutions.

What we have now is a society captured by the tech billionaire class. Over the past few decades, a handful of billionaires degraded our entire information system

under the guise of disruption, killing the business model of newsrooms and replacing it with ad-based products required to keep our constant attention. And despite rhetoric that AI has the potential to “democratize” the world, the inherent pathologies of AI make it a centralizing force,²³ contingent on the mass accumulation of data and compute resources in the hands of a few big players.²⁴ Now the tech billionaire class threatens to destroy our creative industries, transforming hard-earned craft into “content” that is then fed to AI models intended to churn out lossy xerox copies of our masterworks. And, as if that’s not enough, these same billionaires have begun to destroy our institutions, purchasing newspapers and taking over the opinion pages, buying elections, and hollowing out our social services.

Contesting the economic and political capital amassed within the tech industry is necessary to create the conditions for a thriving democracy. We urgently need to restore the institutional structures that protect the interests of the public against oligarchy. This will require confronting tech power on multiple fronts, from enacting corporate accountability measures that keep tech oligarchs in check, to staving off efforts to use AI to hollow out our institutions, to bolstering work happening at the community level among local government officials, organizers, and workers devoted to rebuilding a democracy that serves the broader public.

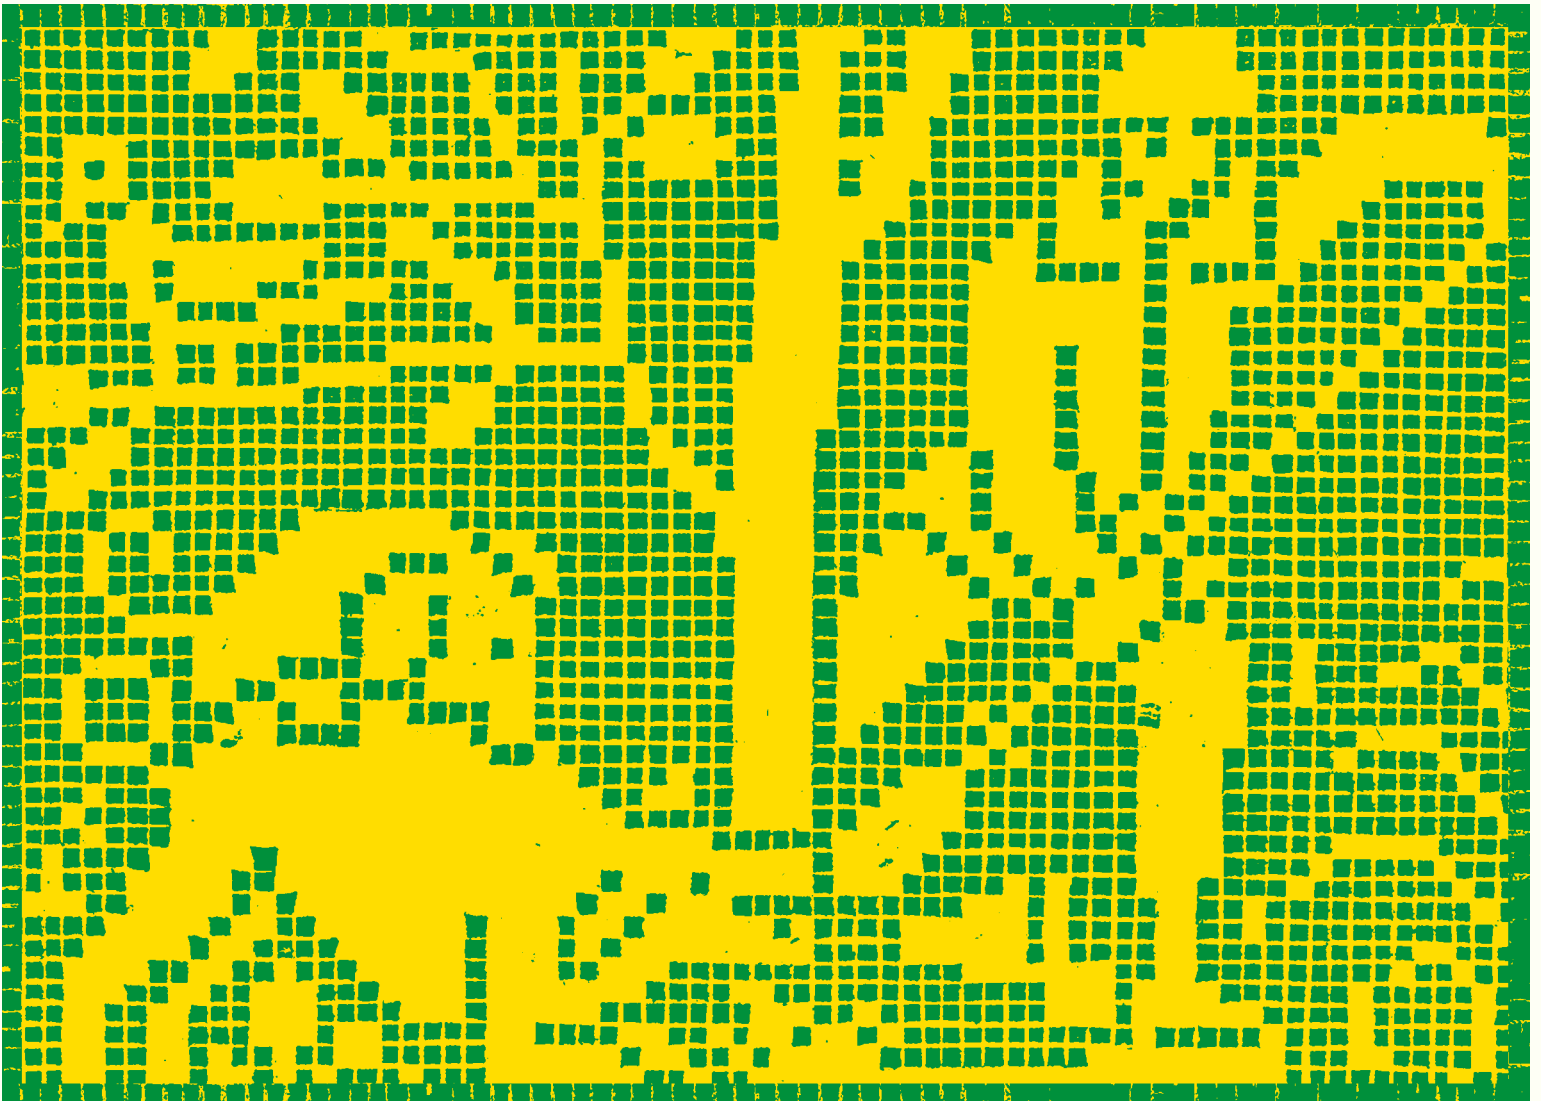


A decorative border made of yellow pixelated patterns, resembling a digital or retro aesthetic, framing the top and sides of the page.

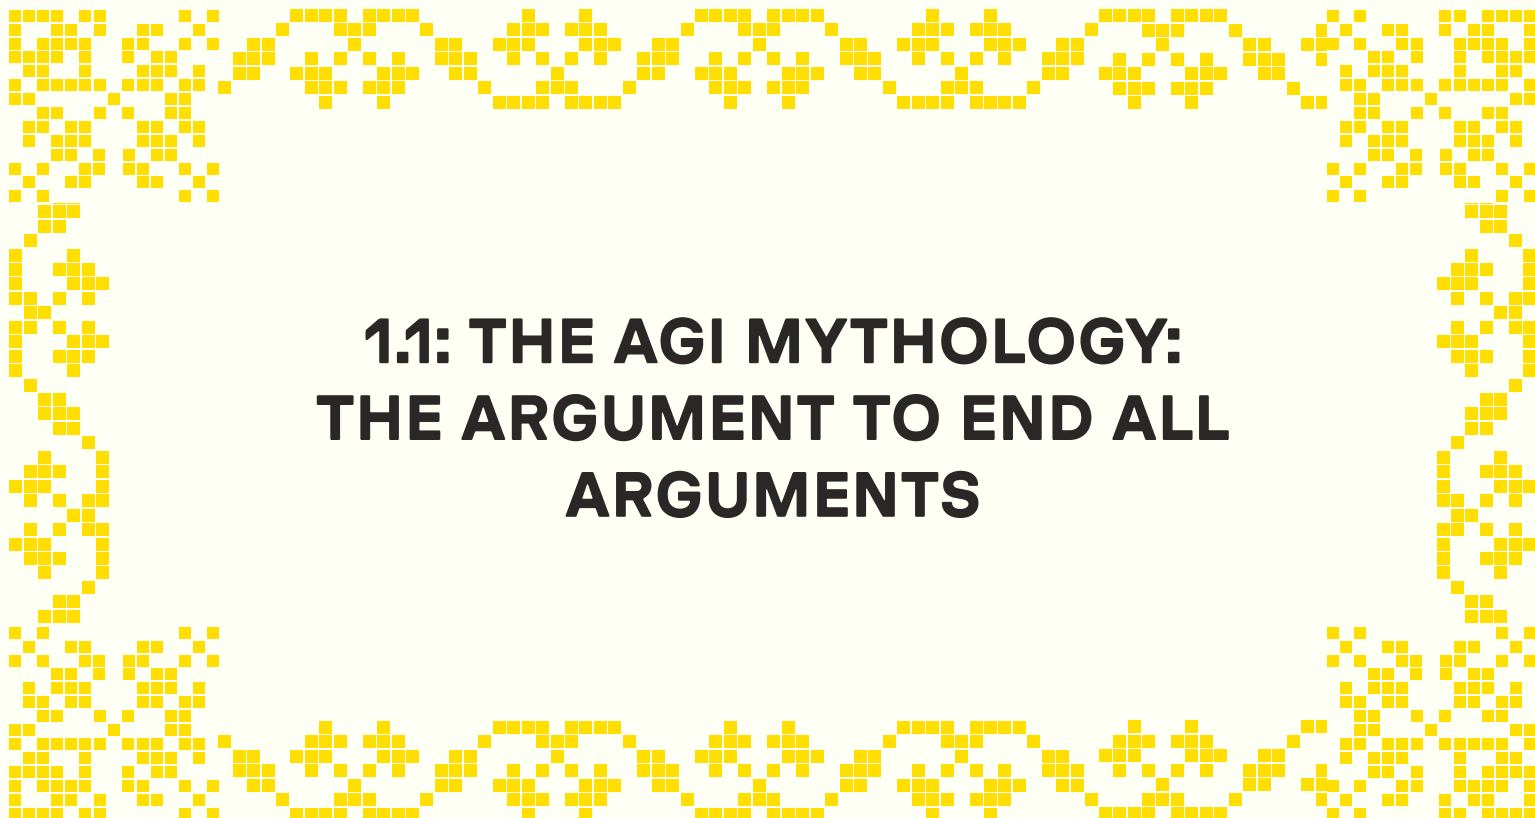
CHAPTER 1:

AI'S FALSE GODS

**WHAT'S PROPPING UP THIS BUBBLE
AND WHY IS IT SO HARD TO NAME?**



The “common sense” around artificial intelligence has become potent over the past two years, imbuing the technology with a sense of agency and momentum that make the current trajectory of AI appear inevitable, and certainly essential for economic prosperity and global dominance for the US. In this section, we break down the narratives propping up this “inevitability,” explaining why it is particularly challenging—but still necessary—to contest the current trajectory of AI, especially at this moment in global history.



1.1: THE AGI MYTHOLOGY: THE ARGUMENT TO END ALL ARGUMENTS

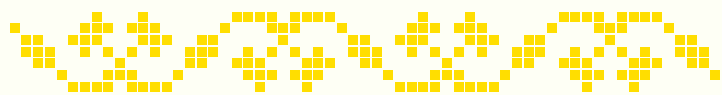
The promise that artificial general intelligence, or “AGI,” is hovering just over the horizon is tilting the scales for many of the debates about how AI is affecting society. AI firms investing in the development of very large models at scale constantly assert that AGI is months or weeks away¹, poised to have transformative effects on society at large—making this central to their sales pitch for investment.² The discourse around AGI adds a veneer of inevitability to conversations about AI; if one company doesn’t achieve it, another will. This also gives governments an excuse to

sit on their hands even as current versions of AI have profound effects on their constituents, as though the race to create AGI has its own momentum.

If anything, under both the Biden and Trump administrations, the US government has instead positioned itself as chief enabler: ready to wield every tool at its disposal—including investment, executive authority, and regulatory inaction—to push American AI firms ahead of their competitors in this race to AGI.³ It’s worth noting that those most vocal about their

fears about the so-called “existential risks” posed by AGI have done as much to prop up and speed along industry development as anything or anyone else.⁴

OpenAI’s assertion that “it’s hard to fathom how much human-level AI could benefit society, and it’s equally hard to imagine how much it could damage society if built or used incorrectly”⁵ drives home that the AI boosters and the existential (“x-risk”) fearmongering both play a role in propping up this vision of AI with supreme capabilities.



WHAT IS AGI? THE HISTORY OF ARTIFICIAL GENERAL INTELLIGENCE



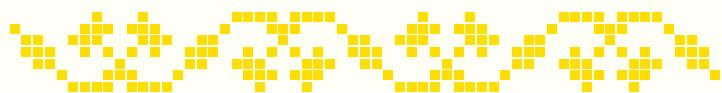
As Brian Merchant chronicles in his report “AI Generated Business,” the term AGI, coined in 1997, captured the notion of a “general intelligence” as a counterpoint to the then-dominant current in AI research, “expert systems,” which operated on rule-based logic designed as a formalized representation of how humans think.⁶ Where expert systems only worked in the narrowest of applications, AGI would operate broadly across a wide range of domains. But developers in the field largely ditched these ways of thinking about AI, turning instead to deep-learning techniques that proved more effective and that form the basis of today’s automated decision-making systems, among others.

Interest in AGI was revived in the 2010s when companies like OpenAI seized on the term, first as shorthand for a form of machine intelligence intended to rival and eventually surpass human intelligence, and later as a term “central to their marketing efforts.”⁷ The images invoked by AI firms is instructive, from Anthropic founder Dario Amodei’s use of the sublime imagery of “geniuses in a data center” capable of paradigm-changing scientific leaps like “designing new weapons or curing diseases,”⁸ to the straightforwardly commercial logic underpinning OpenAI’s agreement with Microsoft: AGI is when AI can create \$100 billion in profits.⁹

In this sense, ChatGPT walked so that AGI could run; the current crop of LLMs in the consumer market are examples of brilliant marketing—proof, as AI firms argue, that big, unexpected advancements in AI were not only possible but “just around the corner.”¹⁰ AGI has since been positioned as the next big step in the LLM advancement trajectory, albeit with little proof, beyond speculation, of how far or wide this leap will have to be.¹¹ However, while this belief seems to be spreading among the general public, it is widely contradicted by many within the AI research community. For instance, in a recent survey of members of the Association for the Advancement of AI, 84 percent of respondents said that the neural net architectures that large models rely on are “insufficient to achieve AGI.”¹² In another, more fundamental, debunking of AGI claims, scholars like Emily Bender¹³ and Henry Farrell,¹⁴ among others, have contested the basis of claims to AGI, arguing instead that large models can “never be intelligent in the way that humans, or even bumble-bees,”¹⁵ are because AI cannot, in fact, create. It can only reflect, compress, even remix content that humans have already created to help people to coordinate or solve problems.¹⁶

While current AI models make the promise of AGI more tangible for policymakers and the general pub-

lic, AGI is conveniently distanced from the fundamental and persistent limitations of LLMs on the ground that AGI, by definition, will be a wholly new paradigm that leapfrogs these material concerns. The mythology around AGI masks the shallowness of today's AI models, providing substance and imagination that innovations are just around the corner.



IF AGI WERE HERE, HOW WOULD WE EVEN KNOW?



Despite bold public claims from the tech industry that AGI is “as little as two years”¹⁸ away, the research community has yet to agree.¹⁹ A recent survey by the Association for the Advancement of AI (AAAI) of nearly five hundred AI researchers found that 76 percent of respondents assert that scaling up current approaches to yield AGI is unlikely or very unlikely to succeed.²⁰

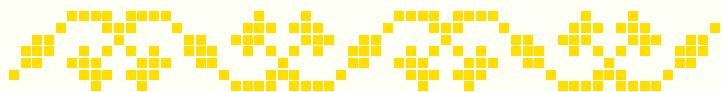
So how will we even know when AGI is here? The metrics currently on offer are largely narrow, vague, and self-serving benchmarks²¹—and some researchers have argued that the preoccupation with AGI is “supercharging bad science.”²² In place of scientific breakthroughs, industry labs are hinging claims to proximity to AGI on grandiosely named tests like “Humanity’s Last Exam”²³ and “Frontier Math”²⁴ that gauge only a very narrow ability to answer clear, closed-ended questions²⁵—poor proxies for the promises these companies make about the capabil-

ity of this technology like inventing cures to cancer or solving for climate change. AI company Hugging Face’s Chief Science Officer Thomas Wolf has argued we’re currently testing systems for their ability to be obedient students, rather than for their mastery of bold counterfactual approaches or the ability to challenge their own training data, which might show more promise for solving complex, intractable problems.²⁶ In 2025, a group of AI researchers from across academia and industry pointed to an endemic challenge within the current field of AI evaluations that is more preoccupied with “coarse-grained claims of general intelligence” than “real-world relevant measures of progress and performance.”²⁷

In sum, there is a widespread and endemic lack of clarity on both the definition and time scales of the AGI conversation, which makes it hard to contest or reason its merits. *The more urgent inquiry, however, is who and what does this disproportionate focus on AGI work in service of? How will it shape the current trajectory of AI?*



WHO BENEFITS FROM AGI DISCOURSE?



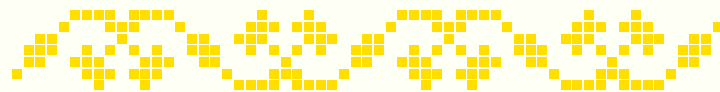
AGI has become the argument to end all other arguments, a technological milestone that is both so abstract and absolute that it gains default priority over other means, and indeed, all other ends. It is routinely cast as a technology so all-powerful that it will overcome some of the most intractable challenges of our time—and that both investment into the sector and

ancillary costs are justified by the future solutions it will offer us. For example, Eric Schmidt recently dismissed the climate costs imposed by AI by asserting that humans aren't set up to coordinate to solve climate change. Thus, the reasoning goes, we need to supercharge data centers—because in the long term, AGI has the best shot at solving for it.²⁸ This not only reflects abstract AI solutionism at its peak; it also serves to flatten and disguise the problem of climate change itself as waiting for its technical silver bullet, rendering the challenges of political will, international cooperation, and material support for people to rebuild homes or house climate refugees—everything it will take to meaningfully “solve” climate change—invisible.²⁹

Presenting AI as a quick technical fix to long-standing, structurally hard problems has been a consistent theme over the past decade (as we explore in our chapter on **Consulting the Record**), but past variants of technosolutionism at least had to demonstrate how the technology would solve the problem at hand. With AGI, though, we're not clear how this transformation will happen beyond the assertion that the current state of affairs will be overhauled. The debates around DOGE transforming government using AI have this flavor: In his interview with Ben Buchanan, Ezra Klein speaks of the general sentiment that with superintelligent AI potentially around the corner, the government will inevitably need to be taken apart and rebuilt in the age of AGI.³⁰ It's the same logic that dictates that if AGI is truly going to propel scientific discoveries of the kind that Amodei promises, then perhaps there will be no need for federal funding for science at all.



AGI'S MARKET-BOOSTING FUNCTION



Asserting that AGI is always on the horizon also has a crucial market-preserving function for large-scale AI: keeping the gas on investment in the resources and computing infrastructure that key industry players need to sustain this paradigm. As we've argued, this current avatar of large-scale AI was set in motion by the simple rule that scaling up data and compute would lead to performance advancements—a logic that sedimented the dominance of the handful of companies that already controlled access to these inputs, along with pathways to market,³¹ and in whose hands power would be further concentrated if AGI ever were achieved.³² The quest for the ever-shifting goalpost of AGI only reinforces this “bitter lesson” (as Anthropic CEO Amodei calls it).³³

There's a lesson here from the 1980s, when, even before the term AGI was in vogue, the Reagan administration pushed for a wildly ambitious (for the time) “Strategic Computing Initiative” that was focused on propelling general advancements in “AI”—along the lines of the AGI promise.³⁴ It was propelled by the promise of new military capabilities, anxieties about Japanese domination on AI, and the potential of private-sector opportunities. A billion dollars in taxpayer money was spent then on a program, now universally acknowledged as a failure, that didn't yield results even on the terms it set for itself. The postmortem of why it failed yields varied conclusions, but one is

worth underscoring: Then, as now, these advancements hinged not on revolutionary feats in science, but on scaling up computing power and data.

Coincidentally, existential risk arguments often have the same effect: painting AI systems as all-powerful (when in reality they're flawed) and feeding into the idea of an arms race in which the US must prevent China from getting access to these purportedly dangerous tools.³⁵ We've seen these logics instrumented into increasingly aggressive export-control regimes. By drawing attention to the very systems they purportedly aim to contest, x-risk narratives create a Streisand effect: encouraging more people to see the AI dystopia in their present, fueling adoption and bolstering industry players rather than curbing their power. They have also narrowed the scope for policy intervention, bolstering a debate centered around the two poles of accelerationism and deceleration rather than facilitating a broad dialogue about AI development and its societal implications.

Ultimately, these twin myths around AGI position AI as powerful and worthy of investment, and draw attention away from the evidence to the contrary.

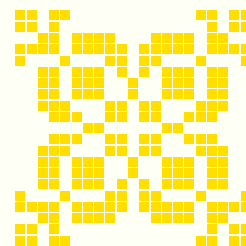


DISPLACING GROUNDED EXPERTISE: WHO IS DISEMPOWERED BY THE AGI DISCOURSE?



Elevating AGI over other paths to solving hard problems is just a supercharged form of technosolutionism,³⁶ **but it also means that those with technical expertise—not only those driving the tech development but also those fluent in using this new suite of tools—are normalized as primary experts across broad areas of society and science in which they lack domain-specific context and experience.**³⁷ This has been a familiar fight over the past decade of AI development: Those with lived experience and sector-specific knowledge have had to advocate for a determining role in questions around whether, and how, AI is deployed.

Whether that means nurses having a say in how AI is integrated in patient care, or parent groups fighting against the use of facial recognition on their children in the classroom, there has been a consistent push to recenter who is counted as an expert on baseline questions about AI integration. (Notably, some of this has often resulted in tokenistic approaches that provide nominal seats at the table to impacted communities—too little, too late.) AGI presents a more formidable version of this challenge given its abstract and absolutist form. For example, narratives around AGI upending the world of work routinely position workers across industries as being subjects—or worse, collateral damage—of a great transformation, rather than as participants and indeed experts in the ways in which these transitions will take place.³⁸





1.2: TOO BIG TO FAIL: INFRASTRUCTURE AND CAPITAL PUSH

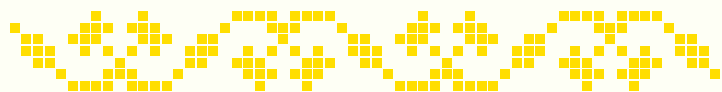
The AI industry's growth model, fueled by the assertion that infinitely increasing scale leads to superior products, has spawned AI firms that are positioned to be too big to fail. Americans are actively subsidizing this unstable system under the premise that the adoption of AI is a "national strategic priority." As we illustrate in this chapter and in [Chapter 1.4](#), though, this has enabled an industrial-policy approach that will ultimately undermine, rather than strengthen, our national security. Finally, we discuss how the abundance agenda, with its seemingly benign focus on what it calls "supply-side progressivism," is a very convenient tool for big AI to justify expanding its energy needs.

Tech firms are deploying unprecedented amounts of capital to maintain their lead and advance in the current paradigm of "scale is all you need" AI, doubling down on infrastructure build-out and seeking federal funding and regulatory support across several dimensions: access to chips and associated hardware

to equip data centers, approvals for the construction of the data centers themselves, and the energy necessary to power them. The stock market is riding this hype wave, and the "Magnificent Seven" stocks (Alphabet, Amazon, Apple, Meta, Microsoft, Nvidia, and Tesla) now represent more than 30 percent of the S&P 500, the largest sector of the index—in prominent part because of the AI boom.¹

It's important to remember that the pursuit of scale was a choice that locked us into a future where a handful of Big Tech firms retained control of the market (see the [Introduction](#)). This is not the only way for AI to develop, nor is advancement measured on a narrow set of self-serving benchmarks² a meaningful proxy for evaluating the societal utility of these systems.³ But because it is what these key market players have doubled down on, and because of their centrality to market indices, the success or failure of the AI bubble will now have a profound effect on the stock

market as a whole.⁴ This raises the stakes around the push for public investment in AI infrastructure—a move that is at best a hedge, and at worst a subsidy, for the profoundly risky and self-interested set of bets by AI firms. If successful, this effort will lock in infrastructures that the public will pay dividends on for years to come, in the form of financial and material costs (see [Chapter 2: Heads I Win, Tails You Lose](#)), creating a path dependency toward continued dominance by large AI firms.



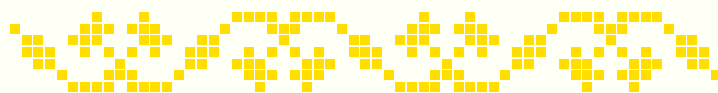
TECH'S CAPEX FRENZY



Firms like Microsoft, Google, and Meta need AI to be profitable because they have funded the AI boom—at many orders of magnitude more than traditional venture capital⁵—boosting the valuations of startups that are far from demonstrating the kind of profitability that traditional investors would seek. They have gone all in on the most capital- and resource-intensive version of AI by adopting the “scale is all you need” paradigm as canon. This is not the only way to approach building AI models, and the companies leading AI development have occasionally gestured toward the need for model efficiency to address compute infrastructure bottlenecks. This was brought home especially by the release of DeepSeek’s R1, which demonstrated model capabilities on par with the leading-edge models of US firms, without anything like the scale US firms rely on.⁶

But rather than make concerted efforts to build models differently, many dominant firms are doubling

down on this approach by seeking public investment and the rollback of regulation to de-risk the expansion of the AI market. For example, within weeks of the DeepSeek announcement, OpenAI announced its Stargate investment with SoftBank, which will allocate a \$100 billion investment into data center infrastructures for model training.⁷



GETTING HIGH ON AI SUPPLY



The US has adopted a position over the past two years that treats AI as an exceptional sector core to the nation’s economic and national security interests. This stance exists in tension with growing friction with Big Tech firms, most clearly articulated in the Biden administration’s Executive Order on Competition, which articulated the perpetuation of national monopolies as antithetical to the national interest.⁸ The Trump administration has likewise bought into AI boosterism even as it has gestured toward the need for antitrust, mostly as a political tool for addressing firms it sees as adversarial to its interests.⁹ As chief case in point, Trump’s pick to head the FTC, Andrew Ferguson, vowed to go after tech monopolies while taking a hands-off approach to AI regulation, proving that attacks on corporate tech power reach their limit when it comes to AI.¹⁰ In tandem, a cadre of appointments related to the environment and energy—including Lee Zeldin as head of the EPA; Jacob Helberg as under secretary for economic growth, energy and the environment; Doug Burgum as dual interior secretary and

“energy czar”; and David Sacks as a newly created “AI czar”—have inextricably tied support for a strong national AI industry to achieving energy dominance, positioning energy expansionism as the essential tool to achieve the administration’s economic nationalism agenda.¹¹

Recent movements from within the federal government have backed this stance: The Department of Energy recently announced it had identified sixteen federal sites across the country positioned for rapid data center construction,¹² and in April the Trump Administration signed an executive order ramping up domestic coal mining using growth in demand from AI data centers as justification.¹³

Small (AI) Is Beautiful? Differentiating to Avoid Industry Co-Option



A growing number of technologists and civil society organizations advocate for smaller models as the alternative trajectory to the bigger-is-better paradigm.¹⁴ This makes sense, because many of the clearest pathologies within the AI industry are driven by scale: from climate impacts; to risks of contagion effects from privacy, security, and accuracy failures; to the ways in which scale breeds ultra-concentrated markets in AI. The dangerous impacts of the vague and all-encompassing “AGI” (see [Chapter 1.1](#)) also demonstrate the scale thesis taken to its logical end: a system that exists at a scale and level of universality that, hypothetically, displaces all other forms of expertise and value.

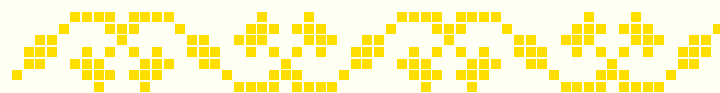
But industry is flocking to a version of the “small is beautiful” thesis, too, as part of their plans for market expansion, creating a familiar risk of co-option of the alternative by the same players who have driven and shaped this current paradigm. In the summer of 2024, Microsoft heralded “tiny

but mighty” smaller language models that would provide impressive performance despite a reduced number of parameters.¹⁵ Apple, Meta, and Google also released AI models with many fewer parameters, signaling that industry is incentivized to move away from simply bigger-is-better in pursuit of compute-efficient methods.¹⁶ DeepSeek only propelled this trend, making it clear that frugality would be a key competitive advantage in this market.¹⁷

This is only superficial common ground. Positioning “smaller” models as one of the options in an “all of the above” approach for the biggest AI companies should not be confused with a rejection of the bigger-is-better paradigm. As Satya Nadella said after the DeepSeek announcements, signaling that these efficiencies only consolidate benefits for the tech giants best placed to capture demand (see [Chapter 2: Heads I Win, Tails You Lose](#)): “As AI becomes more efficient and accessible, we will see exponentially more demand.”¹⁸ It also ignores that pushing advancements at the “frontier” of this tech is still dictated by scale, even as firms play around with a mix of approaches across their portfolio to target different types of consumers. Most importantly, the large-scale version of this tech is what drives these firms’ policy lobbying around infrastructure expansion with deleterious impacts on the public. **Movements that aim to disrupt the consensus around scale-driven AI must explicitly name and distance themselves from this industry-driven discourse.**



AI FIRMS WANT TO BE TOO BIG TO FAIL



These infrastructure investments function to lock us into a world where US continued dominance in the AI market is guaranteed by the government, and, for now, largely supported by investors in the stock market seeking to avoid an end to the AI bubble—while taxpayers foot the bill (whether by taxes that contrib-

ute to these investments, or more directly through increased energy bills, as we unpack in [Chapter 2: Heads I Win, Tails You Lose](#)). AI industrial policy serves either to secure demand via procurement policies¹⁹ or to underwrite and attract continued investment (as

is the case with the Stargate deal). This approach to AI is akin to industry bailouts—rarely a popular policy stance—but compared to the auto industry and banking, the AI market is much more speculative and its value to the public is unproven.

The Abundance Agenda: AI's Fundamental Incompatibility with Supply-Side Progressivism



The emergence of “abundance” as a narrative strategy and policy platform is being used by tech firms to get access to scarce public subsidies and energy. This stance has formed around a constellation of thinkers and organizations working across party lines to articulate a policy agenda premised on building a policy apparatus in support of more, and more efficient, construction of critical resources with low supply and high demand, including housing, healthcare, and energy. It operates under the presumption that (1) government regulation makes building too burdensome in these sectors, leading to cost inflation; and (2) progressives have focused too intently on subsidy programs that cut or block access, rather than on the underlying reasons for cost inflation. The solution, abundance movement advocates posit, is to push forward “supply-side progressivism,” or, as Ezra Klein puts it, “to take innovation as seriously as they take affordability”²⁰ by implementing regulatory reforms that speed development and solve scarcity.

Abundance proponents centrally contend with energy markets, in that they argue in favor of cutting regulation to enable an increase in energy production. For example, Jerusalem Demsas wrote in the *Atlantic* that the ability for NIMBY-minded community organizations and climate groups to shut down renewable development is hindering the US’s ability to meet its climate goals.²¹ Klein and Derek Thompson argue that overhauling energy infrastructure is crucial to mitigating climate change, emphasizing that the first step toward an abundant clean-energy future is reducing the current fossil fuel reliance from 60 percent as of 2022 to nearly 0 percent.²²

As a growing number of AI companies prioritize building and opening new data centers, more energy is needed to meet the staggering demand. **One might think that AI-driven demand would concern abundance advocates, because AI firms soak up the available supply of renewable energy.** Data centers already account for 4.59 percent of all energy used in the US. That number has doubled since 2018.²³ Goldman Sachs estimates that data center power demand will grow 160 percent by 2030.²⁴ These are staggering numbers wreaking havoc on an already fragile energy grid.

Instead, we see a more uneasy alliance, where the abundance agenda potentially converges with the energy deregulation camp for whom the “urgent” need to advance AI is being used as a justification to fast-track and expand fossil fuel production and use. At the House Oversight Committee hearing on data centers, AI, and energy, legislators repeatedly threw renewables under the bus, even touting that China is powering their AI systems with coal-fired plants.²⁵ The fossil fuel company talking point that wind and solar are not a reliable source of energy to meet data centers’ 24/7 demands is deeply ingrained,²⁶ with legislators and data center trade groups pivoting toward the expansion of nuclear—rather than renewable—energy to provide “reliable” and sturdy energy for AI. Despite the substantial evidence on hand, this sustainability critique has not been taken seriously by abundance advocates skeptical of the climate movement.



1.3: AI ARMS RACE 2.0: FROM DEREGULATION TO INDUSTRIAL POLICY

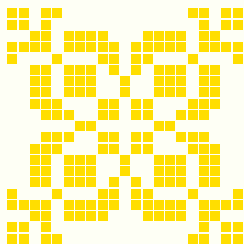
The fusing of economic and national security goalposts under the banner of the US-China AI arms race is a critical asset for US AI firms: It affords them patronage not just from their own government, but potentially from the many other nation-states vying for a fighting chance at national competitiveness in this market; it insulates them from regulatory friction by framing any calls for accountability as not just anti-innovation but harming national interests; and—as we explore in [Chapter 1.2: Too Big To Fail](#)—is a key factor in positioning them as not just too big, but too strategically important, to fail.

Nation-states have developed their own flavors of “AI Nationalisms,” embarking on initiatives designed simultaneously to support homegrown development and sovereign infrastructures free of dependency on US tech firms, and to attract AI investment.¹ But though AI nationalism is on the rise globally, the rhetoric around the AI arms race remains centered

around two poles: the US and China. Since the mid 2010s, the notion of a US-China AI arms race has been primarily deployed by industry-motivated actors to push back against regulatory friction. A frequent motif in policy discussions at moments where the industry has sought to stem the tide of regulation, the notion of an arms race was one of the key arguments made against the introduction of a federal data protection law, a package of antitrust reforms targeting the tech industry in 2022, and an omnibus AI Accountability Bill that was considered before Congress.²

In the past two years, this so-called race has taken on a new character (let’s call it the “AI arms race 2.0”), taking shape as a slate of measures that go far beyond deregulation to incorporate direct investment, subsidies, and export controls in order to boost the interests of dominant AI firms under the argument that their advancement is in the national interest (what we refer to as AI industrial policy³). Such an

approach predates the Trump administration. Arguably a number of the core measures propping up the AI arms race 2.0 were outlined under the Biden Administration; Jake Sullivan, in particular, was a vocal proponent of the logics of economic security.⁴ The Biden administration's AI Executive Order,⁵ National Security memo,⁶ and export controls⁷ all established an intent for the US government to widely adopt AI and to clear the pathway for the industry to expand through infrastructure build-out, while simultaneously hindering the advancement of strategic adversaries like China by limiting the export of leading-node chips. Unsurprisingly, this stance ran parallel to the lobbying platforms of firms like OpenAI that have sought government cooperation, with a narrow list of conditionalities such as the use of renewable energy and compliance with security measures.⁸ OpenAI specifically has made threats that it will relocate its business absent commensurate support from the US government.⁹ Since inauguration, the Trump administration has escalated support for the AI industry, rolling back the conditionalities articulated by the Biden administration by repealing the AI Executive Order and replacing it with a blanket assertion: "It is the policy of the United States to sustain and enhance America's global AI dominance in order to promote human flourishing, economic competitiveness, and national security."¹⁰



A NEW SILICON VALLEY CONSENSUS BEYOND TARGETED ADS TO TARGETED AI WEAPONS?



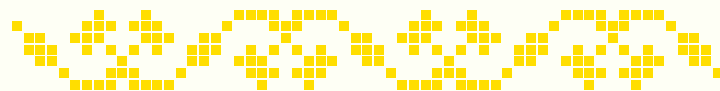
While the Trump administration has firmly asserted AI as a strategic national asset, they are likely to expect the industry to act in ways that align more closely with state interest. The specifics of what that means is left deliberately hazy, but a popular refrain has been that companies should be devoted less to targeted advertising, and more to AI that would bolster national security—and defense tech is increasingly front and center of events like the Hill & Valley Forum,¹¹ an annual consortium of Silicon Valley elites and DC lawmakers that first convened in March 2023 to combat China's influence on the American tech industry.¹² Cofounded by Palantir's Jacob Helberg, the Hill & Valley Forum is more aligned than ever before with state national security interests,¹³ as Helberg,¹⁴ like Michael Kratsios and David Sacks, is one of many industry representatives who find themselves in key policy roles under the Trump administration.¹⁵

So far, the industry seems to support this vision. This is best seen in the rhetoric of Palantir's CEO Alex Karp, who has long framed the company's mission as addressing a civilizational need to support democratic and Western supremacy through leading-edge technology. But emboldened by Trump's intent to scale

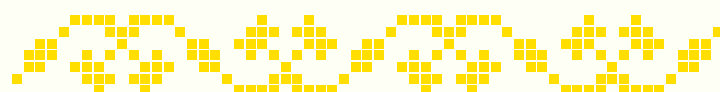
up mass deportations and police surveillance, Karp has escalated, saying in an investor call in early 2025: “We are dedicating our company to the service of the West and the United States of America, and we’re super-proud of the role we play, especially in places we can’t talk about. Palantir is here to disrupt. And, when it’s necessary, to scare our enemies and, on occasion, kill them.”¹⁶

Karp isn’t alone. **Since the Biden administration’s shift toward securitization of AI in 2024, companies that have historically distanced themselves from the military have also doubled down on national security:** After making an amendment to its permissible use policy enabling its tools to be used by militaries,¹⁷ OpenAI has increasingly leaned in to making policy arguments on security grounds,¹⁸ going so far as to assert that expanding fair use under copyright law to include AI development is a security imperative.¹⁹ In February 2025, Google amended its guidelines to allow its AI technologies to be used for military weapons and surveillance, despite ongoing protests by its employees and a long-standing ban on use of its technology for weapons following the Project Maven protests of 2018.²⁰ And Meta made an announcement in November 2024 that it would make its Llama models available to the US government for national security use.²¹

Meanwhile, Anthropic’s CEO Dario Amodei recently wrote about the threat of authoritarian governments establishing military dominance on AI as a reason to accelerate US leadership²² and the VC firm Andreessen Horowitz operates an “American Dynamism” practice expressly designed to support the national interest in strategically important sectors: aerospace, defense, public safety, education, housing, supply chain, industrials, and manufacturing.²³



A DOUBLE-EDGED SWORD: CHIP DIFFUSION AND “SOVEREIGN AI”



It’s worth noting that the AI arms race 2.0 has shifted from being an absolute policy advantage for the tech industry writ large to being a double-edged sword for some: Aggressive restrictions on the export of chips are closing off a huge market for US AI hardware companies and data center products, which has left firms like Nvidia and Oracle deeply unhappy.²⁴ During the Biden administration, the implementation of export controls restricting the sale of semiconductors to certain countries through the “diffusion framework” received the bulk of the criticism, with a number of firms invested in the global chip market particularly up in arms about the impact to their businesses.²⁵ The Trump Administration may make changes to the diffusion rule,²⁶ and is internally fragmented between factions that are supportive of tariffs and hawkish toward China, and those that are interested in global expansion of the AI market.²⁷

For its part, Nvidia—the leading semiconductor firm, which is most directly affected by the export controls—has embarked on a push for “sovereign AI,” a term coined by the company to refer to nations’ abilities to produce their own AI using some combination of homegrown infrastructures, data, workforces, and business networks.²⁸

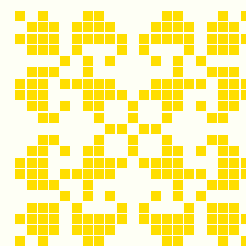
Nvidia's stance is an example of a play at market expansion. As the provider of computing chips for the data center infrastructures central to sovereignty initiatives, the company stands to benefit from nation-states' growing interest in building out their own homegrown industries and attracting AI investment. For chip manufacturers, the push toward sovereign AI can be seen as a way of diversifying their customer base away from the hyperscalers and hedging their business against the potential slump in the demand from these companies.²⁹

The European Union and its member states have also espoused interest in sovereign investment into AI in a bid to compete at the frontier. The European Commission has gradually repurposed its existing European high-performance supercomputing capacity toward training large-scale AI models.³⁰ To further up the ante, the Commission announced a €20 billion InvestAI initiative to establish European "gigafactories" that would house one hundred thousand GPUs with the objective to facilitate training of models with "hundreds of trillions" of parameters.³¹ Investment has also picked up in the member states. In February 2025, France hosted the Paris AI Action Summit, during which president Emmanuel Macron announced around €110 billion in investment pledges to boost France's AI sector, with a focus on infrastructure investments.^{32, 33} In Germany, the new government coalition has agreed to house at least one of the gigafactories, complemented with commitment to develop a sovereign tech stack, as well as support for a budding "Eurostack" movement, an informal coalition³⁴ at the European level that aims to reduce European tech dependencies by developing domestic alternatives.³⁵

These investments at the level of the EU and its member states still pale in comparison to the scale of the private investment plans in the US, with the \$500 billion joint venture fund Stargate announced in January

2025; the fund arguably cements monopoly dominance by a cartel of US-based firms.³⁶ Meanwhile, the UAE and Saudi Arabia are geopolitical swing states, given their financial capital to sustain infrastructural build-out, and have been flooding the market with money via the AI funds MGX, G42, and the Saudi Public Investment Fund (PIF) for AI,³⁷ money that the leaders of AI firms are avidly seeking.³⁸

Nationalism thus still remains a critical shaping force in AI policymaking: The "AI arms race" has if anything become increasingly complex in a moment of geopolitical uncertainty, and is wielded by firms both to avert regulation and to court investment.





1.4: RECASTING REGULATION AS A BARRIER TO INNOVATION

There has been a swift and aggressive narrative attack on AI regulation as anti-innovation, superfluous bureaucracy, and unnecessary friction. We've seen a total reversal in the US federal stance and, increasingly, a regulatory chill reverberating across quarters in the EU. We saw early signs towards the end of Biden's term setting the government's primary role as enabler of the AI industry,¹ and with the Trump Administration it is the headlining message. The headwinds against baseline accountability against the tech sector in general, and AI companies in particular, are greater than ever.

The tech industry's fickle policy promises have also revealed their true colors. Companies spent 2023 insisting they were extremely concerned about safety and were firmly "pro-regulation."² But as the center of power has shifted towards a deregulatory current, any superficial consensus on guardrails has just as quickly fallen away. OpenAI's CEO Sam Altman, for instance,

went from testifying in a Congressional hearing that regulation is "essential" to lobbying against a minor safety provision in just fifteen months.

The government's narrative change has been just as swift. In 2023, future-looking existential ("x-risk") concerns took center stage. In policy fights these x-risk safety concerns have often eclipsed the long list of material harms arising from corporate AI control, often moving public and policy attention away from enacting policy and enforcing existing laws on the books to hold companies accountable.³ Notably, Vice President Harris's speech on the sidelines of the UK AI Safety Summit called out this tension explicitly, and set up an (implicit) counterpoint to the x-risk-dominated agenda at the rest of the summit led by former prime minister Rishi Sunak: "These [existential] threats, without question, are profound, and they demand global action. But let us be clear. There are additional threats that also demand our action—

threats that are currently causing harm and which, to many people, also feel existential.”⁴ Harris went on to describe the ways in which ordinary people have already been harmed from faulty, discriminatory, and inaccurate AI systems.

Unlike other regulatory conversations, the broad philanthropic and government interest in addressing x-risk safety concerns eventually served to further cement government relationships with the tech industry. The vast majority of efforts under the safety umbrella have been voluntary and industry-led—for example, numerous safety validation standards within the UK and US AI Safety Institutes were set by or done in collaboration with industry players like Scale AI⁵ and Anthropic⁶—revealing that the government had been successfully convinced to regulate AI in lockstep with and led by industry-centered expertise. On the other hand, when the rubber met the road with SB 1047, the California bill that sought to impose baseline documentation and review requirements on the largest AI companies for a very narrow class of advanced models, large parts of the tech industry pulled out the rug and pushed against even this narrow regulatory intervention with all their might.⁷ Even Anthropic—which positions itself as a company responsive to safety and the risks of AI—waffled on SB 1047 support, first coming out against the bill before dragging their feet into a hedged statement of support, saying the “benefits likely outweigh its costs,” but “we are not certain of this.”⁸ Government players fell in line, with key Democratic legislators⁹ framing the bill as detrimental to innovation.¹⁰ In a letter to Governor Newsom, eight Democratic members of Congress succinctly summed up this position: “In short, we are very concerned about the effect this legislation could have on the innovation economy of California.”¹¹ Facing immense pressure, Governor Newsom ultimately vetoed the bill.

The fight for SB 1047 opened the floodgates for pitting regulation against innovation. A recent one-two

punch has shifted the terrain entirely: Groups advocating for legislation mirroring SB 1047’s provisions are being politically targeted by Republicans¹² and a new troubling bill, SB 813,¹³ is gaining support in California that allows AI firms to self-certify their models as safe and then use that certification as a legal shield to avoid liability in a civil action for harm.¹⁴

At the federal level, there was vanishingly little progress legislatively, leaving large swaths of industry use entirely outside of regulatory constraints. Biden’s now-repealed EO¹⁵ and the OMB memo¹⁶ were bright spots, making strong progress in terms of hooks for actionable accountability via targeting government use of and procurement of AI. Even public investment proposals such as the National AI Research Resource pilot, originally positioned as a counterforce to concentrated power and resources in the AI industry, was recast under Biden’s 2024 National Security Memo as a national competitiveness project. Former National Security Advisor Jake Sullivan’s October 2024 speech before the National Defense University also firmly positioned the US government as an enabler of frontier AI companies and emphasized the need for US investment in the AI sector to go full steam ahead in order to shore up the country’s strategic positioning against China.¹⁷

Still, despite a far-from-coherent policy stance on AI under Biden, the attack on regulation ushered in by the Trump administration cannot be overstated.¹⁸ Since being elected, President Trump has positioned regulation as a clear-cut way for the US to “lose” the global arms race, and his allies have propagated fears of Chinese control of global AI infrastructure as a threat to American security and democracy. On his first day in office, Trump gutted Biden’s Executive Order on AI, replacing it with his own Executive Order set to revoke existing federal AI policies that “act as barriers to American AI innovation.”¹⁹ At a series of high-profile events including Davos, the French AI

Action Summit, and the Munich Security Conference, the Trump administration's message rang loud and clear: Global regulation is a targeted economic attack on US companies, and the antithesis to innovation. Meanwhile, the administration has expressly targeted the administrative state, calling into question the independent status of enforcement agencies and gutting the federal workforce, including key employees tasked with enforcing existing laws to rein in corporate dominance (this included unlawfully firing key Democratic FTC Commissioners with a record on tech enforcement). The Trump administration's recent OMB memos do little to impose accountability on AI systems, and are instead designed to fast-track the procurement of AI across the federal government.²⁰

Meanwhile, AI Industrial policy—or financial and regulatory support for expanding the national AI industry—is being positioned as the counterpoint to regulation, and a more appropriate role for government intervention. Unsurprisingly, Silicon Valley tech and AI executives have fallen²¹ quickly²² into²³ line, shoring up their seats at the table. Because, while Trump's tangible industrial AI policy moves remain to be seen, the dominos set in motion by the Biden administration are poised to rapidly accelerate under Trump.

Trump's agenda for global AI dominance is mutually reinforced by an expansive energy dominance agenda, and his administration has repeatedly highlighted the need to expand US energy resources²⁴ to remain competitive in AI.²⁵ Debates about permitting requirements for infrastructure build-out had already taken center stage during the Biden administration. Senator Joe Manchin's Energy Permitting Reform Act of 2024 expediting review procedures for energy and mineral projects advanced out of committee with a bipartisan vote.²⁶ The bill is supported by a coalition of fossil fuel companies and tech lobbyists, who claim that AI tech innovation is tied to energy expansion. As they wrote in a letter to Congress: "America's leadership in

global innovation depends on the passage of permitting reforms that allow the US to build critical energy infrastructure."²⁷

In some ways, the Trump administration's pro-enforcement posture toward Big Tech companies—seen in the continuation of the DOJ's case against Google and the FTC's recent trial against Meta—is consistent with the Biden administration's antitrust policies, and runs orthogonal to the otherwise deregulatory headwinds and hands-off approach to the tech industry. At the same time, these cases are not designed to strike at the root of power facing the AI industry, which has received an "all systems go" message from the Trump White House, but rather to curtail Big Tech censorship and undermine platform authority over state power. Already we see tech companies attempt to wield political favor to end the trials.²⁸ And Google is set to argue that structural separation will undermine US national security issues,²⁹ potentially derailing bold antitrust remedies from the court. Despite these cases, it is unlikely that the Trump DOJ and FTC are set to broadly undermine the AI industry's market power as a matter of policy, no matter how the antitrust suits are decided.³⁰

The drift toward deregulation has begun even in the European Union, traditionally seen as a staunch regulatory power. Driven by rightward electoral shifts, increasing securitization of AI, and new geopolitical realities driven by Trump, the once proudly proclaimed digital regulation agenda is now seen as a liability by European policymakers. In addition to scrapping planned bills, such as the AI Liability Directive that created a product liability framework for AI,³¹ there is appetite in the high halls of EU policymaking to walk back on rules already agreed to. While backtracking is constrained by the embarrassing optics of bending under US pressure—at least thus far—when it comes to implementation, there is growing pressure to create as much flexibility as possible so as to mute

the impact of the laws without changing their letter.³² This push to create flexibility for domestic companies is complicated by the importance of these rules as a rare source of leverage in the nascent trade war between the EU and the US.³³ The extent to which European digital regulation becomes a pawn in this debate remains to be seen.

More generally, the tone in the European Union and member states has become more enabling, paralleling the developments elsewhere. French President Emmanuel Macron's "plug, baby, plug" quip at the Paris Action Summit crystallized this shift in sentiment.³⁴ Leveraging the tools of statecraft and existing infrastructures (such as abundant nuclear energy in France) toward promoting the development of AI is increasingly central to the broader push toward European sovereignty. In addition to new public investments in AI infrastructures, new political coalitions and power players are also emerging in the background to facilitate this change. A recent large public-private partnership with an investment pledge of €150 billion by a collective of leading European industrial giants and tech companies, complemented by direct access to heads of European states to discuss a "drastically simplified regulatory framework for AI," is one example of these changing winds.³⁵

Absent from this discussion is the role regulation can play in fostering innovation within markets, particularly given the dynamism and complexity that AI exhibits. By creating a stable regulatory environment with robust competition among firms and an equal playing field that enables new entrants to thrive, well-crafted regulation can act as an enabler rather than an adversary to innovation in emerging markets (See [Chapter 4: A Roadmap for Action](#)).

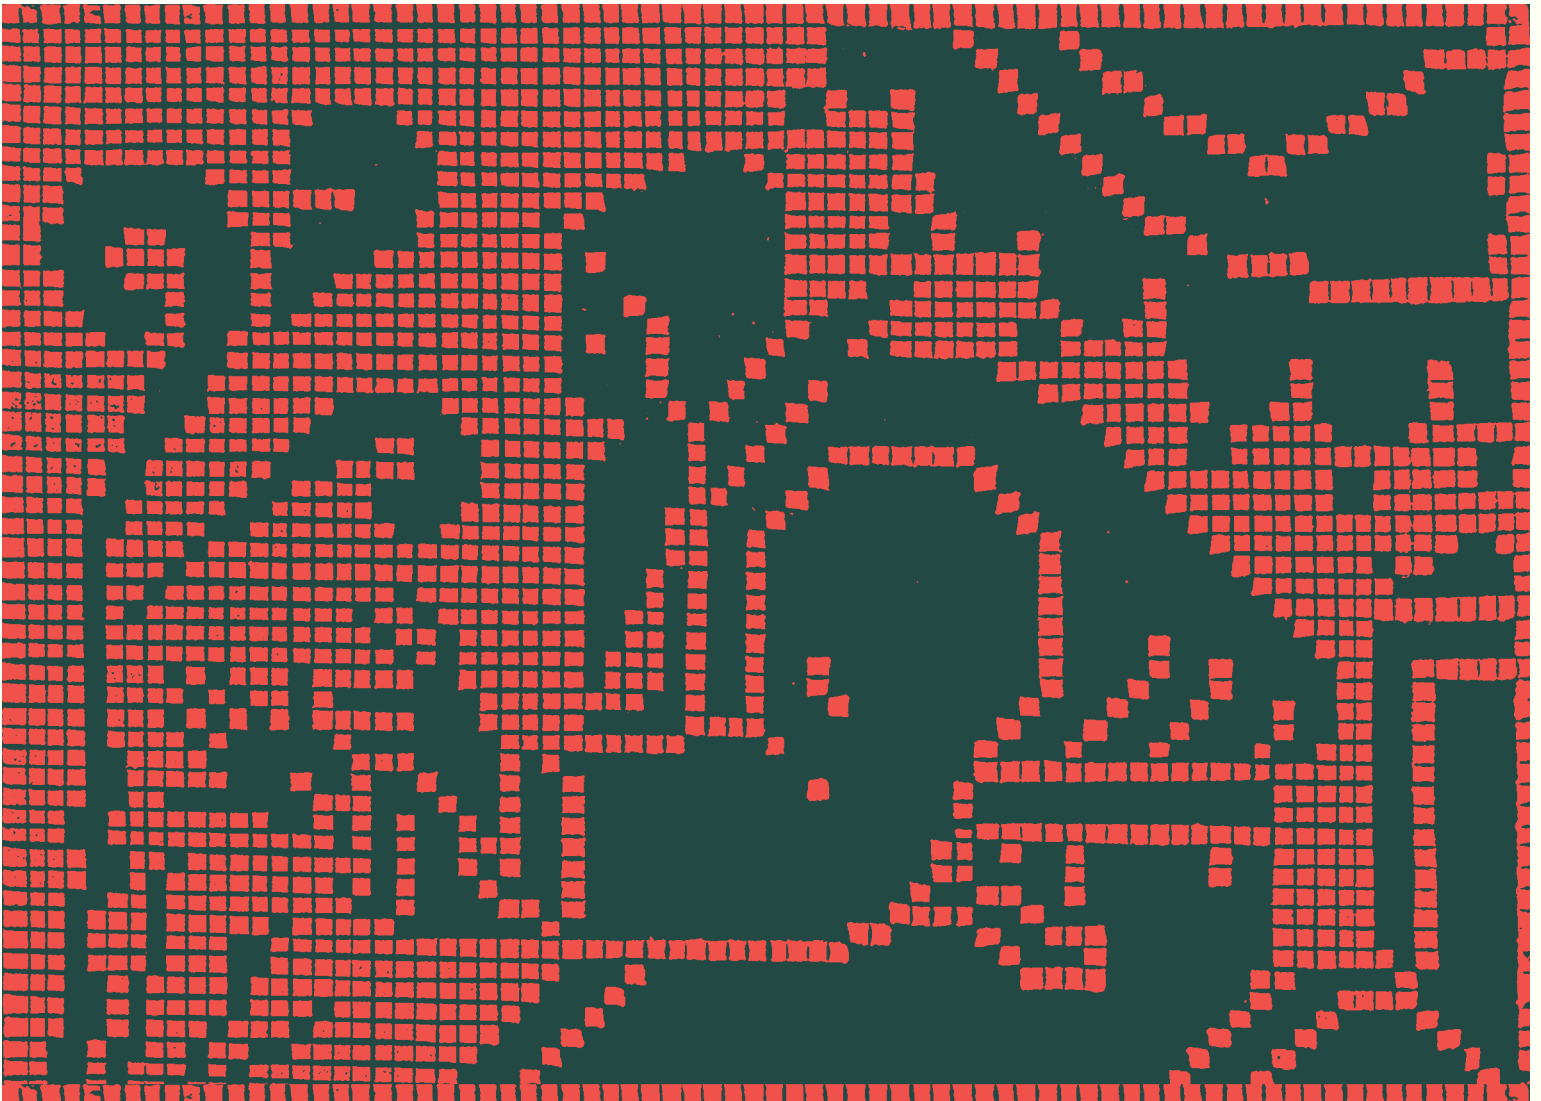




CHAPTER 2:

HEADS I WIN, TAILS YOU LOSE

**HOW TECH COMPANIES HAVE
RIGGED THE AI MARKET**



Governments and investors are funneling billions of dollars into a speculative AI industry without a clear business model or pathway to profitability. In Chapter 1, we identified the myths undergirding the hype despite obvious red flags and warning signs. But the reality on the ground is far less distributive; here, we explain how a handful of firms are poised to capture the AI market.

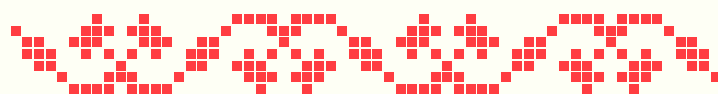
In some regards, the current market behavior of AI firms appears wholly irrational: Tech companies are pumping billions of dollars into an unproven technology with little market demand, firing their own workers,¹ and acquiescing to the political demands of an administration defined by its tech factionalism and personal vendettas.² On its face, the AI market appears to be driven more by AI “FOMO”—a fear of missing out—than sound business decisions, with AI firms throwing product use cases at the wall to see what sticks, and firms across the economy force-fitting AI solutions into their workflows, buckling under the generalized pressure that any competitive company must today have an “AI strategy.”³ Big Tech firms have guaranteed their own success by making the wall as sticky as possible, gaming the market to ensure they benefit if and when the returns come rushing in.

Whether by locking customers into existing ecosystems, bending the law to work in their favor, co-opting political processes and media narratives, or pegging their own futures to an industrial strategy of national dominance and government investment, Big Tech firms are shaping the market to consolidate their own power and to hedge against the considerable risks they’re exposed to.

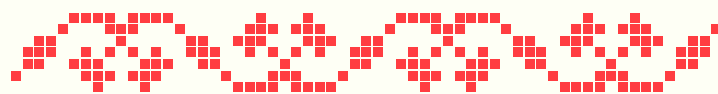
The reality is that Big Tech firms and AI developers (propped up by Big Tech firms) can successfully gamble on AI’s future because the house always wins. Their deep pockets allow them to suffer short-term

losses as they shuffle through product use cases and burn money, AI chips, and energy at an alarming rate, but ultimately they—and power players in adjacent industries that hinge on AI infrastructure build-out—are best positioned to net long-term gains in this market.

This section maps the drivers that are securing Big Tech firms’ advantage in the AI market, before turning to the question of who loses in the end.



CLOUD INFRASTRUCTURE PROVIDERS BENEFIT FROM CYCLES OF AI DEPENDENCE



Because the quickest path to AI profit is through the increased demand in cloud services this market drives, Big Tech firms that offer cloud computing services and control cloud infrastructure (like Amazon, Microsoft, and Google) are best positioned to win the AI race.

Because of the “bigger-is-better” paradigm, AI developers require more and more compute resources to effectively train their larger models and run “inference,” such as the queries returned each time you enter a prompt into ChatGPT. This dependency on compute has made large-scale AI development contingent on access to compute resources, which has led AI developers like OpenAI and Anthropic to secure partnerships with cloud companies like Microsoft and Amazon in order to successfully train and run their models. The early exclusive partnership between OpenAI and Microsoft has received the most attention among these: OpenAI received Microsoft’s cloud resources at a fraction of the cost; in return, Microsoft locked OpenAI into billions of dollars in cloud commitments and a share of OpenAI’s future revenue.⁴ OpenAI wasn’t alone: Anthropic developed arrangements with Google⁵ and Amazon,⁶ Deepmind solidified its cloud partnership as Google DeepMind,⁷ and Mistral struck a deal with Microsoft,⁸ for example.

But the advantage these cloud firms hold is multifaceted: Unlike other cloud companies like Oracle and Coreweave, Amazon, Microsoft and Google also hold a dominant advantage along the AI supply chain, with advantages in access to data, paths to market, and talent.

The partnership model between hyperscalers and AI developers is evolving from being predicated—as was the case in the 2018 deal between Microsoft and OpenAI—on exclusivity to being predicated on mutual dependence.⁹ For example, even though OpenAI is no longer locked into an exclusive partnership with Microsoft, Microsoft remains able to secure a market advantage where it matters most—AI model deployment—while ensuring their investment is recouped through circular spending agreements and revenue shares.¹⁰ Under the new partnership, Microsoft retains access to OpenAI’s IP (including insight into how OpenAI and Oracle will manage the new Stargate servers);

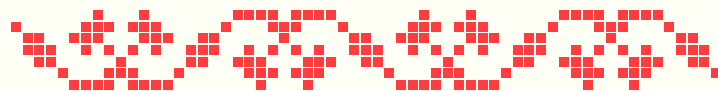
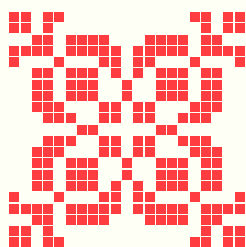
OpenAI API is still exclusive on Azure (and pays more than \$1 billion per year on Microsoft services); and revenue-sharing commitments are still in place (Microsoft retains a 20 percent share of OpenAI’s revenue and future profits up to \$92 billion).¹¹ Microsoft is also positioned to effectively block OpenAI’s effort to convert into a for-profit company, while OpenAI’s board can trigger a clause that prevents Microsoft from accessing its most cutting-edge tech, which OpenAI officials have reportedly proposed doing.¹²

In January 2025, markets were temporarily rattled by the announcement that Chinese startup DeepSeek was able to launch an AI model comparable to OpenAI’s latest release at a fraction of the compute cost.¹³ For some, DeepSeek cast doubt on the self-serving, bigger-is-better paradigm advanced by companies like OpenAI, projecting future efficiencies in compute resources. But DeepSeek’s release does not change the current paradigm of cloud company dominance: Despite the model’s smaller use of compute in the final training run, the technical advancements in advanced reasoning driven by the inference-time compute approach are still reliant on scale for their performance advantages. And any efficiency gains would likely be overridden by growth in demand, a phenomenon known as the *Jevon’s paradox*. As Satya Nadella declared in the wake of DeepSeek’s release, “As AI gets more efficient and accessible, we will see its use skyrocket, turning it into a commodity we just can’t get enough of.”¹⁴ DeepSeek thus solved one pressing business problem for Microsoft—how to deal with its escalating expenditures on data centers—without disrupting the overall business proposition for the company: capturing the market through its control over the cloud ecosystem.

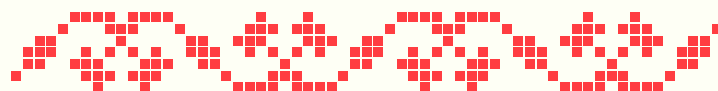
Efficiency gains through models like DeepSeek’s also don’t necessarily undercut the advantage that Big Tech companies hold from their access to compute. For one, pushing ahead on performance gains at the

cutting edge of the technology is still extremely compute intensive.¹⁵ (It is also widely believed that DeepSeek essentially “distilled” its model building off of OpenAI’s o1.¹⁶) Moreover, cloud companies reap consistent gains even as we—consumers and companies alike—figure out whether AI delivers what it promises: Every response generated on ChatGPT, every query run on Gemini, and every customer service chatbot integration incurs a cost that customers pay back to the hyperscalers. Now, if compute remains a scarce resource, Big Tech companies with cloud businesses win by controlling limited supply. Similarly, if models become more efficient, these firms still win because the efficiencies will lead to overall reduced infrastructure costs, allowing them to deliver more product at cheaper cost. This means that cloud firms are incentivized to boost AI demand either way, ensuring that AI demand balloons to fit a growing market for infrastructure that depends on its success.

This relationship of dependence extends not just to AI developers but to cloud startups, too. For instance, CoreWeave is a new entrant into cloud computing that chipmaker Nvidia has invested in, and has marketed itself as a solution to compute bottlenecks in the AI market. But the company recently went up for public offering,¹⁷ and financial documents revealed that CoreWeave is saddled with debt and almost entirely dependent on Big Tech companies like Microsoft that need to offload their excess demand—the very companies it is attempting to compete with.¹⁸



BIG TECH FIRMS BENEFIT FROM LEVERAGING CONTROL OVER THE TECH ECOSYSTEM



There is increasing consensus that AI models are becoming “commoditized,” meaning that gains in model efficiency decrease costs, and more large-scale models will emerge to compete. In response, firms like Microsoft are advising those in the market to “focus more on how they integrate these models with their own data and workflows.”¹⁹

This advice reflects Microsoft’s position in the market: It, like Google and Meta, has an advantageous position due to its dominant role in enterprise and consumer-facing software. This is precisely why, on the day that the chipmaking firm Nvidia’s stock fell nearly 17 percent following the DeepSeek news, Amazon, Meta, and Apple’s stock *went up*.²⁰ Because if AI models become cheap to integrate—and compute becomes significantly cheaper—the firms who own AI products, distribution, and data centers are at an advantage. This makes ecosystem power—control over the paths to market—an important element in the AI market.

Ads all the way down: Meta’s advertising ecosystem positions it well in the generative AI market. Because one of the main use cases for generative AI technolo-

gy right now is generating a lot of content very quickly, Meta can deploy AI to consistently optimize for the highest-performing ads at unprecedented scale—driving revenue for themselves and advertisers at minimal marginal cost.²² Moreover, Meta can leverage its infinite stream of free, user-generated photo and video content across Facebook and Instagram to make ads “indistinguishable from content”²³ by using AI to label and link every possible purchasable item in every possible piece of content. This helps explain the significant investment—\$65 billion—Meta is making into AI infrastructure.²⁴

Other Big Tech firms are moving in a similar direction:

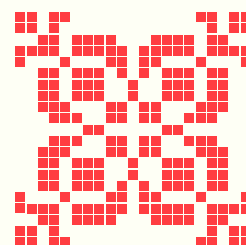
Microsoft dominates the enterprise software market, and is pushing Copilot integration and the upselling of security features, which require purchase of a premium subscription to their cloud platform, Azure.²⁵ In the event that OpenAI is able to capitalize on its ChatGPT user base—an unlikely prospect²⁶—Microsoft, as mentioned earlier, has a revenue-sharing agreement in place to give them 20 percent of OpenAI’s revenue.²⁷ It has also opened up its own competing arm called Microsoft AI, led by Inflection CEO Mustafa Suleyman, which is devoted to product development both for Copilot and other consumer AI products.²⁸

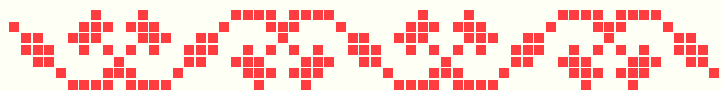
Google dominates the search and search-advertising ecosystems—so much so that a federal court found in 2024 that Google has an illegal monopoly over both internet search and search-advertising markets.²⁹ As remedies work their way through the court, Google has leveraged its search dominance to integrate Gemini, its own AI model, into the search experience, providing Google with an unbeatable advantage to deploy AI to millions of captive users.³⁰ The same logic that applies to Meta’s ad market applies to Google, where generative AI models can quickly and cheaply optimize advertising content to benefit advertisers and Google alike. Meanwhile, Google is launching its suite of AI tools across all of its consumer and enter-

prise products, from Workspace to Email, so computing efficiencies will make this process all the more profitable for Google. Finally, DeepMind, Google’s AI research laboratory, is slowly expanding into a product development org, revealing that Google’s biggest bet on AI is product integration rather than AI model development.³¹

Amazon, as the leading cloud infrastructure firm, is already positioned to take up a significant portion of the AI market, and has made several rounds of investments into the startup Anthropic.³² It has also made some initial effort to develop its own models,³³ though its operation of a model marketplace seamlessly integrated with Amazon Web Services (AWS) is likely most reflective of its intent to offer AI-as-a-Service.³⁴ Amazon is testing the rollout of AI models across its existing platforms and services, including Alexa,³⁵ and a shopping tool called Interests onto its online marketplace.³⁶ It is also developing its own chips: Inferentia is optimized for AI training runs, and Trainium is optimized for inference and training.³⁷

Apple dominates the mobile-device ecosystem. If inference becomes dramatically cheaper and memory requirements substantially decrease, it becomes significantly cheaper to deploy the most powerful AI models on Apple’s devices. Its focus has been on releasing small models aimed at running on-device for iPhone, iPad and Mac, and a larger foundation model running on its private cloud servers, leveraging Apple’s position in the device market.³⁸





AI FIRMS BENEFIT FROM THE DATA CENTER BOOM



As we described earlier, controlling access to cloud resources and services is a crucial way firms like Google, Microsoft, and Amazon are advantaged in the AI market. In 2024, Big Tech companies spent more than \$180 billion on data center expansion and infrastructure.³⁹ Just one year later, Google, Meta, Microsoft, and Amazon expect to spend an additional \$300 billion on data center construction and infrastructure costs for AI.⁴⁰ It is estimated that by 2030, the largest cloud service providers will host 60–65 percent of all AI workloads.⁴¹

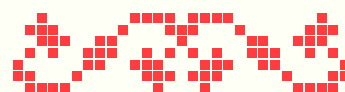
The amount of energy that data centers are projected to use is staggering. Industry analysts anticipate a five-year load growth of 80 gigawatts to power AI data centers—the equivalent of adding the electric capacity of the entire state of California onto our existing grid.⁴² Energy utility companies tend to project even more aggressive numbers, like a Texas utility company claiming requests for 82 gigawatts of additional load in its service territory alone.⁴³ If these numbers are even nominally accurate, the growth in power consumption by data centers is poised to wreak havoc on already fragile energy grids and markets that are incapable of meeting this extraordinary demand, especially in such a short time.⁴⁴

These projections need closer scrutiny, however.

Former National Economic Council Director Brian Deese has said that forecasters tend to overestimate electricity demand because they emphasize static load growth over efficiencies that are likely to develop over time.⁴⁵ Utilities companies are also incentivized to overproject energy demands to grab investors' attention.⁴⁶ Furthermore, data centers tend to request services from multiple utilities, meaning that projected demand is likely captured in multiple utility companies' projections.⁴⁷ But these high-demand projections have acted as a strategic policy level for firms petitioning the government to quickly bring more power sources online as a matter of national importance.⁴⁸ The strategy is working: The Department of Energy is set to announce plans allowing companies to build data centers and power plants on federal land by the end of 2025, in order to maintain America's "global AI dominance."⁴⁹

In one sense, these inflated projections could seem risky for Big Tech companies. If the AI bubble bursts and energy projections don't materialize, Big Tech companies risk sinking billions of dollars into stranded infrastructure. This is why the push for public investment is critically important as a de-risking measure, so they won't be left footing the bill in the face of a market collapse.

Big Tech Wins Even If the AI Boom Doesn't Pan Out



Big Tech firms are pursuing a multifaceted strategy to shore up their interests no matter how the AI market inevitably plays out. First, they pump up energy projections while banking on efficiency gains to drive

more demand. Then they strike favorable deals with utility monopolies to ensure best-rate prices and pass off the remaining costs to ordinary ratepayers. Next, they integrate themselves across the entire energy supply chain, purchasing energy, selling energy technology, and making strategic long-term contracts with power plants to preference their own needs at the expense of other energy customers.⁵⁰ Finally—and perhaps most importantly—they pit state and local governments against each other by dangling purported economic benefits of data center development, accruing tax breaks and subsidies,⁵¹ while lobbying to block any legislation that undermines their interests.

Although data center facilities require massive capital expenditure to build—a hyperscale facility of a hundred megawatts requires up to \$1.4 billion in up-front investments—data center capacity is relatively fungible if enough capital is available to maintain or repurpose the facility.⁵² But the specialized AI infrastructure being built today for GPU clusters that produce significantly greater heat isn't as easily repurposed for general computational use. In the event of AI demand plummeting, Big Tech companies could repurpose data centers for other workloads, including traditional cloud and data storage servers.⁵³ But they'll be equally incentivized to leverage their market power to make demand happen, rather than accept huge write-offs. And since a portion of the data center infrastructures used by Big Tech firms are leased, companies have some flexibility to cancel or opt out of renewals based on demand—as Microsoft chose to do after it ended its exclusive agreement with OpenAI.⁵⁴

Utility companies, on the other hand, may be left reeling. Despite having no long-term guarantees of future demand, utility companies across the country are planning to invest billions of dollars in new infrastructure to service new data centers. But minimum contracts for large load customers tend to be short—two years on average—and minimum charges are low,

meaning that data centers can walk away from their large energy contracts with little risk, leaving utility companies—and ratepayers—left to carry the costs, even for unfinished projects.⁵⁵

Data centers also pose risks to utility grid planning and management: They have the capacity to “disconnect” from the grid and switch to their own local, back-up power generators. This is a safety mechanism intended to protect data center equipment from damage that can arise from fluctuations in voltage, grid frequency, or natural disasters; yet when done at scale across multiple large data centers concentrated in one region, these disconnections can cause large surges in excess electricity that threaten grid reliability across a region. Several “near-miss” cases have already been documented across the country in the past year.⁵⁶

The second way Big Tech companies guarantee favorable market terms for massive data center investment is by striking opaque and exclusive deals with utility monopolies to set preferential energy rates, which shift infrastructure costs onto ordinary ratepayers.⁵⁷ As Eliza Martin and Ari Peskoe discuss in the report *Extracting Profits from the Public: How Utility Ratepayers Are Paying for Big Tech's Power*, Big Tech companies do not need a rate discount: they are fully capable of funding their own infrastructure development costs. Nevertheless, utility companies offer special contracts to Big Tech companies to attract their business⁵⁸ and then may raise electricity rates for other ratepayers to make up for the rate discounts to large customers.⁵⁹ Already customers in Georgia have seen six rate increases in less than two years, increasing the electricity bills of ordinary ratepayers by 37 percent due to additional power demand from Georgia data centers.⁶⁰ Dominion Energy, which services an area in Virginia known as “data center alley” because it houses the largest cluster of data centers in the world,⁶¹ is proposing fuel-rate increases that

could raise average residential customer bills by as much as ten dollars *per month*.⁶²

Technically, some special contracts must receive approval from regulators. But approvals tend to pass in uncontested hearings, with many contract details held confidential or redacted in public hearings. In some states, regulators face immense political pressure from utility companies (and in some states, elected officials seeking to gain favor) to approve special rate deals for large and influential companies like Big Tech firms.⁶³ In addition to underpaying on rates, Big Tech firms also have the ability to game the system by reducing their energy load during the time frame when they know utility companies are measuring their uses for the purposes of calculating their “demand charges,” so that their data centers will be charged much less than their fair share of the system costs.⁶⁴ While these companies are securing long-term contracts for energy and capacity at a stable and known price, they are driving the cost of energy higher for all other customers and taking valuable energy resources currently used to serve existing customers away to serve data centers’ rapacious needs instead.

Third, Big Tech companies are embedding themselves across the entire energy supply chain in hopes of “cementing a technological lock-in effect” and ensuring dominance in whatever energy future takes hold.⁶⁵ Big Tech companies are major purchasers of clean energy,⁶⁶ but are also suppliers to renewable energy companies, selling technology to help companies organize their workspaces or manage their energy loads. For example, Alphabet (Google) has developed a product called Tapestry to help electricity grid operators map and manage their electricity grids;⁶⁷ Alphabet also owns the thermostat company Nest, which utility companies can use to control customers’ thermostats under demand-response programs.

Big Tech companies are also funding investments to bring new sources of power online or restart and ex-

pand dangerous sources that have closed, like Microsoft’s investment to reopen a unit at Three Mile Island, a nuclear power plant; or Google’s investment in offshore wind projects.⁶⁸ These firms are directly funding investments in new energy companies, like Amazon’s investment in hydrogen electrolyzers, Google’s investment in geothermal startup Fervo, and Sam Altman’s investment in Helion Energy.⁶⁹ And despite stated commitments to sustainability, tech firms have deep ties to the fossil fuel industry:⁷⁰ They purchase fossil fuel energy; they sell AI to fossil fuel companies to speed extraction; and they are driving investments to delay the retirement of coal plants.⁷¹

Fourth, Big Tech companies are pitting state and local governments against each other by dangling economic development promises to secure generous subsidies and abatements that reduce their tax liability.⁷² In turn, localities, often desperate for additional sources of revenue, offer these companies packages of incentives to attract their business, including tax breaks for data center projects, sales and use tax exemptions, and property tax abatements.

For example, in 2019 (years before the generative AI boom), Indiana passed a law exempting data centers from sales tax on materials and equipment needed to build and operate data centers for up to fifty years, as well as a sales tax exemption on purchasing energy. Meanwhile, ordinary Indianians pay a 7 percent sales tax for their electricity and any other goods they buy.⁷³ This is particularly damning when considering how much data centers spend on electricity bills: Indiana Michigan Power estimated that a 1,000 megawatt data center would pay an annual electric bill of \$492.6 million.⁷⁴ Over a fifty-year period, the foregone sales tax revenue would total more than \$1.7 billion.⁷⁵ Amazon’s new data center campus coming to New Carlisle, Indiana, within this service territory could use double that energy once completed,⁷⁶ and lead to even more lost revenue. Nevertheless, at least thirty-two states now

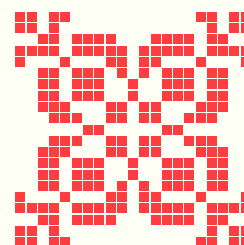
offer similar subsidies to data centers, which will cost billions in foregone public revenue: Texas's program, for example, could cost the state \$1 billion in lost tax revenue in 2025 alone.⁷⁷

Good Jobs First has tracked over \$6 billion in data center subsidies given to Amazon in the United States, including a recent \$1 billion property tax exemption in Oregon for a new data center in Morrow County⁷⁸ and a \$4.3 billion subsidy deal in St. Joseph County, Indiana. Other states are also passing legislation designed to unlock new energy sources and strip back consumer protection laws in a bid to court data center development.⁸⁰

If states and localities refuse to offer desired incentives, Big Tech companies routinely say that they'll build elsewhere. Martin and Peskoe provide over a dozen examples from Big Tech companies and data center developers testifying in rate cases that utility prices are an important factor for determining where they will build data centers.⁸¹ Similarly, an investigation into Meta's decision to build a \$10 billion data center campus in Louisiana reveals the project was a "non-starter" unless Louisiana provided a sales tax exemption on servers and equipment.⁸² But this might not always be true: As the Microsoft executive responsible for data center selection stated in the *New York Times*, "I can't think of a site selection or placement decision that was decided on a set of tax incentives."⁸³ Relatively few states offer ideal sites for building data centers due to cost, climate, and the risk of natural disasters. This gives states and localities much more bargaining power than they are leveraging, causing them to lose out on significant tax revenue whenever they bow to corporate pressure to strike deals.

Fifth, and finally, Big Tech firms lobby to block measures designed to protect consumers in state legislatures. For example, in January 2025, a bipartisan group of Virginia lawmakers proposed multiple bills to enshrine baseline protections for citizens, including

oversight, transparency, sustainability, and cost-allocation measures.⁸⁴ Big Tech companies fought these bills tooth and nail, with one political action committee connected to the Data Center Coalition, a Big Tech lobbying group, contributing over \$100,000 to Virginia state lawmakers.⁸⁵ In the end, all but one data center bill failed to pass.⁸⁶ The bill that passed, which allows (not even requires) data centers to perform impact assessments of data centers' effects on water and agricultural resources, was recently vetoed by Virginia's Governor for creating "unnecessary red tape."⁸⁷ In Oregon, lawmakers introduced a bill to specify data center companies as a new customer class to ensure cost allocations are fair.⁸⁸ But Big Tech companies are challenging the bill, claiming that this would unfairly single out data centers⁸⁹—a particularly ironic argument given that much of their strategy for securing lower rates with utility monopolies relies on using their power as a "special" customer to ask for a *differentiated, discounted rate*.



Big Tech is still the most influential stakeholder in shaping the AI market, even as new entrants emerge. Why?



Many AI products don't, on their face, involve traditional Big Tech companies such as Google and Amazon. This is certainly true in the workplace, where hundreds of "little tech" products affect white-collar and low-wage workers alike across the entire labor supply chain.⁹⁰ Hubstaff, for example, tracks workers' time and productivity, and Appriss Secure algorithmically predicts "employee fraud" and tracks workers for their "sales-reducing activities."⁹¹ It is also true for government services: The breadth of government contracting with technology companies and no-name vendors is staggering. California awarded nearly \$236 million in contracts to five private vendors (including consulting agencies) to overhaul their unemployment system, including Deloitte, Maximus, Thomson Reuters, ID.me, and Salesforce.⁹² Lesser known tech companies like Sagitec, Catalis UI Solutions, and Infosys LaborForce also play major roles in government benefits contracts.⁹³ Law enforcement agencies are also major purchasers of technology products from companies like Clearview AI, Palantir, and SoundThinking (formerly ShotSpotter), which are used to surveil individuals and build troubling predictive policing tools.⁹⁴

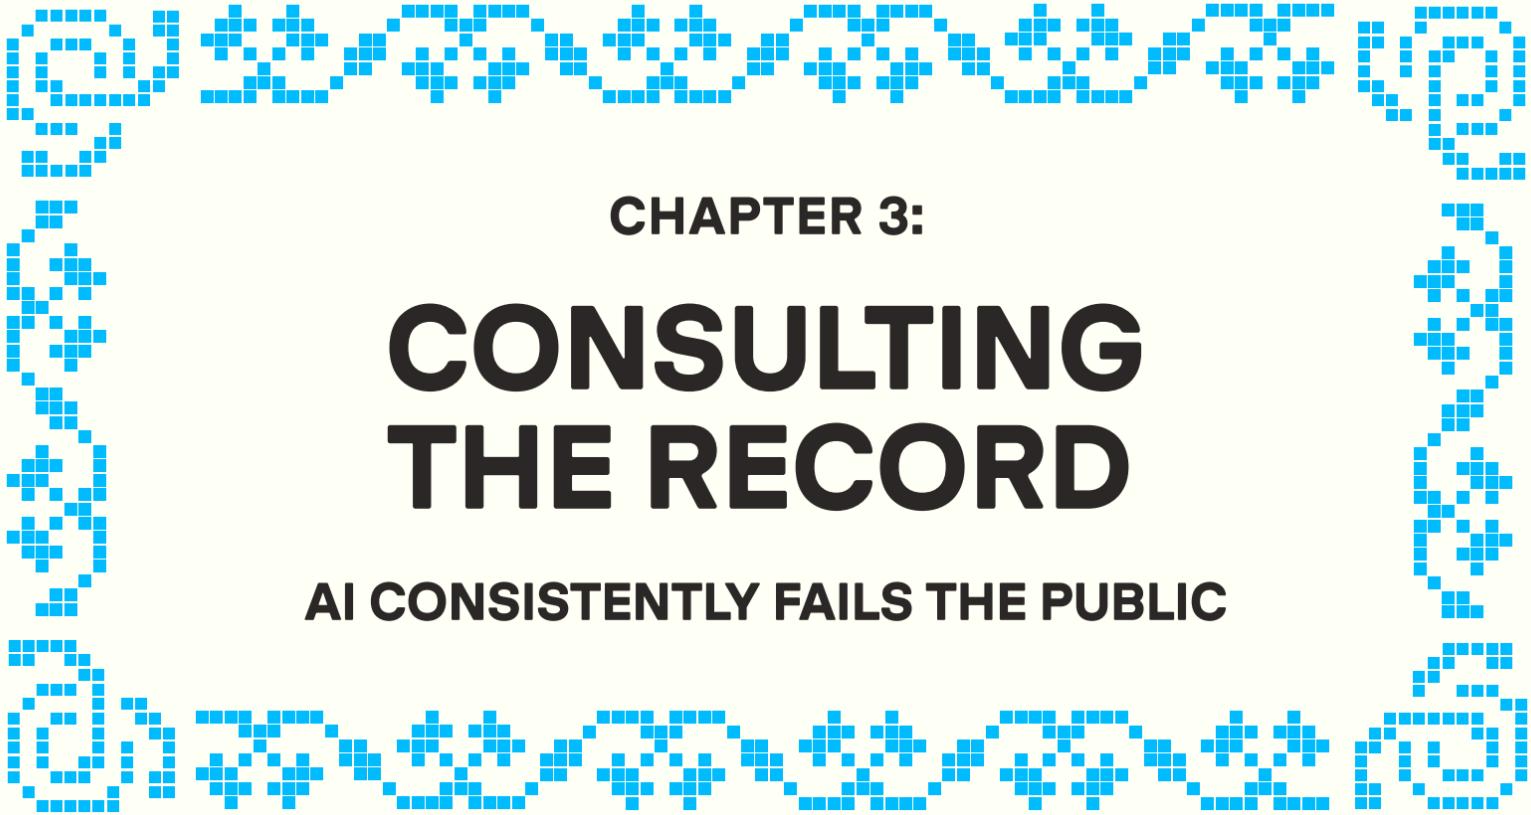
So how and where does Big Tech's influence shape in this ecosystem?

Cloud lock-in: For one, many of these companies run on Big Tech cloud services, which directly rolls them back up to Big Tech firms; they're vertically integrated across the compute supply chain in ways that enable them to absorb or out-compete smaller players.⁹⁵ Sagitec, for instance, runs on Microsoft Azure,⁹⁶ while Catalis and Palantir run on AWS.⁹⁷ Infosys runs on Salesforce. Deloitte has a strategic partnership with Google Cloud.⁹⁸ This means that any use of these platforms directly benefits hyperscaler companies.

But that dependency might be dismissed on the grounds that it's like suggesting that a small business supports utility monopolies because they use electricity. **However, smaller AI companies are not just running on the cloud; they are locked into the cloud. One way Big Tech companies maintain AI ecosystem dependencies is to ensure future partner rollouts are routed through their cloud, meaning the strategic partnerships between small tech companies and their cloud providers bear fruit for cloud companies beyond running the existing technology.**

Trend toward market consolidation: If you look at the landscape of smaller tech vendors of even a few years ago, many of these companies positioned themselves as data analytics platforms developing their own machine learning models, then fine-tuning them with partner data. As the AI industry overwhelmingly attempts to push centralized, large-scale LLM systems trained on the entire internet onto institutions—undergirded by claims of more productivity, more economic efficiency, and stronger accuracy measures—we are likely to see cloud companies push their own AI models onto partners, government contractors, and workplaces because, quite simply, it is the easiest thing to do. Cloud service providers do their best to make their ecosystem environments as sticky and hard to leave as possible.¹⁰⁰ Therefore, it is substantially easier for Deloitte to roll Google's Gemini AI model into their government contracts if Deloitte is already locked into Google as a cloud service provider, ensuring that Google's future AI technology is implemented due to its incumbent status. As these enterprise AI products are deployed in more governments and workplaces around the world, **Big Tech is in the best position to capitalize on the generative AI push, consolidating and entrenching its market position.**

Defining funding lines: Beyond cloud contingencies, Big Tech still defines funding lines, which helps its vision for the space to dominate. In 2023, Big Tech companies outspent venture capital firms to invest in generative AI startups, with Big Tech investments capturing two-thirds of all generative AI investment.¹⁰¹ Now, venture capital is realigning toward AI startups targeting customers with deep pockets, governments eager to cut costs, or businesses eager to drive profits. For example, Anduril, a defense tech company, is expected to raise billions of dollars in 2025 selling defense contracts to the US government. Thus, the most direct pathway to profit in a market where the business proposition remains uncertain for most of the smaller tech companies remains acquisition by Big Tech firms, and venture capital firms are pushing regulatory policies that make merger review easier.¹⁰³ Big Tech firms, in an attempt to evade antitrust scrutiny, have also taken to more creative forms of acquisition, like "acqui-hires" that take the company's employees, but not the company itself.¹⁰⁴ In this way, too, Big Tech still dominates and shapes the smaller tech ecosystem.



CHAPTER 3:

CONSULTING THE RECORD

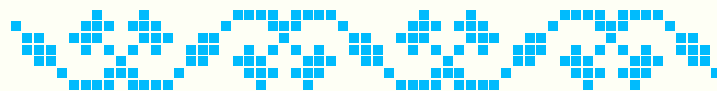
AI CONSISTENTLY FAILS THE PUBLIC



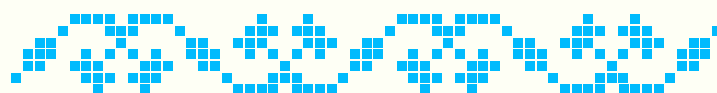
Amid the excitement over AI's (speculative, hypothetical) potential, we have lost sight of the sobering reality of its present and recent past. AI is already intermediating critical social infrastructures, materially reshaping our institutions in ways that ratchet up inequality and concentrate power in the hands of the already powerful. It is consistently deployed in ways that make everyday people's lives, material conditions, and access to opportunities worse.

In this section, we describe how the tech industry has sought to reshape society to enable more widespread deployment of the technologies it builds and profits from, often contributing to the degradation of our social, political, and economic lives. Drawing on examples from several sectors where AI experimentation is well underway—including education, agriculture, immigration, healthcare, and government services¹—we interrogate what happens when our institutions face immense pressure to adopt AI technologies full steam, in spite of persuasive arguments against doing so. Drawing from these domains, we align on five key takeaways:

1. AI's benefits are overstated and underproven.
2. AI-sized solutions to entrenched social problems displace grounded expertise.
3. AI solutionism obscures systemic issues facing our economy, often acting as a conduit for deploying austerity mandates by another name.
4. The productivity myth obscures a foundational truth: The benefits of AI accrue to companies, not to workers or the public at large.
5. AI use is frequently coercive, violating rights and undermining due process.



1. AI'S BENEFITS ARE OVERSTATED AND UNDERPROVEN

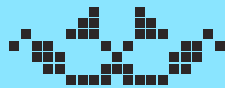


Zealous claims that AI technologies will have transformative effects on particular sectors, and society at large, are circulated by AI developers as nearly incontrovertible. Take, for example, the assertions that AI will rewrite the scientific process,² transform logistics and supply chain management,³ democratize access to education,⁴ lead to more sustainable farming practices,⁵ and even feed the world.⁶

But given the profound societal transformations required to make AI systems work—from rewiring our energy infrastructures, to restructuring our public institutions, to investing unprecedented amounts of capital—we need more than hypotheticals and breezy claims about “curing cancer” and future economic growth. We need evidence of tangible, material benefits that match not only the scale of the hype, but also the level of access and penetration that AI firms are demanding. If Big Tech wants everyone to be using AI, then AI should benefit everyone.

While the science that reveals AI harms is robust,¹⁸ the evidentiary base that supports its asserted benefits is decidedly thin. In fact, most peer-reviewed, rigorous research indicates that in many cases AI systems fail profoundly at even basic tasks.¹⁹

“Curing Cancer” as the End to Justify All Means



Recently, leaders in AI and Big Tech began to claim that AI has the potential to cure cancer. Anthropic CEO Dario Amodei estimates we will eliminate most cancers in the next five to ten years because of AGI. OpenAI CEO Sam Altman repeatedly rests on the example, stating in a recent viral interview that he suspects someday in the future a scientist will be able to ask an AI to cure cancer and, after a few weeks, it will.⁸ Google DeepMind CEO Demis Hassabis stated in a recent *60 Minutes* interview that AI might help cure all diseases within the next decade.⁹ The logic behind this? Once we reach the nebulous milestone of AGI, AI technologies will surpass human intelligence to such a point that AI will be able to speed up the scientific research process, condense decades of scientific research into a few years, and autonomously develop a cure for cancer.

These claims are obviously overstated. Research medicine is incredibly complex, and any “cure” for cancer would, at the very least, require significant clinical testing—potentially for years—before it is safe and effective enough for widespread use. Nevertheless, if you widen the lens enough to focus on all of the different applications of AI technologies to cancer research, the broad premise that AI could meaningfully aid the development of cancer research is indisputable. Deep-learning architectures have already had success in computer vision tasks like image classification, which has led to advancements in cancer screening, detection, and diagnosis;¹⁰ and machine learning algorithms can also bolster a method of fighting rare diseases called *drug repurposing* that allows scientists to search through existing medicines and rework them as treatments for rare conditions.¹¹ (It is worth noting that the technologies that have been the most successful in improving scientific research and patient care do not use large language models, chatbots, or predictive generative AI tools—the technologies that have come to represent “AI” in the recent post-ChatGPT hype cycle.)

What is disputable is the premise that these scientific breakthroughs—or the speculative future cure for cancer achieved via AGI—requires the unrestrained growth of AI industry hyperscalers. But this is precisely the link these corporate leaders are trying to make.

Nowhere is this clearer than in Google’s recent policy recommendations for the Trump Administration’s AI Action Plan, a document that begins with AI’s potential to “revolutionize healthcare” and ends with a sweeping deregulatory agenda to “supercharge U.S. AI development,” complete with recommendations to federally preempt state AI laws, unlock energy to fuel US data centers,

and accelerate government AI adoption as a matter of national security.¹² Anthropic’s policy proposal for the AI Action Plan refers back to Dario Amodei’s prediction of ending cancer in five years to recommend scaling energy infrastructure and accelerating government AI adoption.¹³

As we discussed in Chapter 1.1, there is little evidence that AGI is “around the corner.” But even if AGI is successfully developed, it will still require significant human intervention to make whatever “cure” the program suggests a reality. Oracle CEO Larry Ellison acknowledged this when he suggested that Oracle is using OpenAI’s tools to create a cancer vaccine if they can crack early detection via blood tests, gene sequencing of tumors, vaccine design, and robots that can make an mRNA vaccine in forty-eight hours—“if” being the operative word.¹⁴

The irony, of course, is that the kinds of research and medical advances that Ellison admitted Oracle would need to successfully cure cancer are being decimated by the types of policy that he was celebrating. The Trump Administration is actively cutting federal funding for critical scientific research, especially at public labs and research institutions—including a proposed \$4 billion cut to the National Institutes of Health, whose leading category of study is cancer research.¹⁵ The administration is also threatening to freeze billions of dollars in federal funding to research universities, many of which are working on first-in-class cancer therapies benefiting thousands of patients.¹⁶ And this is to say nothing of what is likely to happen if a company like Oracle actually creates the hypothetical robot-produced cancer vaccine: Look no further than the rollout of the COVID-19 vaccine, which allowed private companies to hide behind patents and secrecy laws to deny distribution to countries in the Global South.¹⁷

Flaws in Large-Scale AI Are Features, Not Bugs



In the past few years, a growing chorus of technical researchers has been sounding the alarm on the persistence of accuracy, privacy, and security-related challenges with large AI models. Worse, the challenges seem to be proportional to the size of the model: The larger and more general the AI model, the more resistant to mitigation these concerns become.

Leaky AI

Leakage occurs when information is fed to a model during training that can later be accessed and extracted. Put simply, AI models routinely “memorize” the data they were trained on and it is fairly simple for such data to be extracted by adversaries, or accidentally regurgitated as well. This means that highly sensitive data can be leaked, from personal health data to military information. While techniques in a field known as *adversarial machine learning* are fast evolving to find ways of mitigating these concerns, currently, “attackers are winning against the defenders by a comfortable margin.”²⁰ Other interventions, like differential privacy, don’t work against models that are trained on extremely large, diffuse datasets scraped off the internet—including off-the-shelf LLMs that form the foundation for many AI applications—making all of these models vulnerable to attack. However, although individual researchers at some industry labs have been vocal about these challenges, for the most part industry has downplayed these concerns: OpenAI, for example, declares that “memorization is a rare failure of the learning process,” mischaracterizing an inherent vulnerability as a rare accident.²¹

Security: Generative AI Introduces Novel and Unresolved Attack Vectors

LLMs and other generative AI models have inherent vulnerabilities that expand attack vectors adversaries can use to exploit AI systems and infrastructure. Such expanded vectors of attack include theoretical and practical demonstrations of “jailbreaks” and adversarial attacks that create inputs to manipulate a model to intentionally produce erroneous outputs or subvert its safety filters and restrictions.²² Other new and undetectable attack vectors include poisoning web-scale training datasets and “sleepers agents” within generative AI models, which may help subvert models and ultimately compromise their outputs. While researchers have produced several approaches that attempt to address these challenges, these have not been successful²³ because, as research has persistently shown, it is always possible to construct attacks that are transferable across all existing foundation models.²⁴ As a result, any fine-tuning or guardrails introduced as a way

to enable accurate military performance or security protections could be bypassed. Limitations in combating these novel attack vectors also arise due to the lack of traceability of human labor and unknown data sources across the supply chain of generative AI models.

Hallucinations: Large-Scale AI Can’t Not Make Stuff Up²⁵

AI chatbots and other forms of generative AI are notorious for producing “hallucinations,” or incorrect information presented as facts, and to do so confidently, without providing any context that could help a user ascertain what is fact and what is speculation.²⁷ For example, OpenAI’s Whisper audio transcription tool—used by doctors in patient consultations—often invents entire passages of text during moments of silence.²⁸

Perhaps more aptly described as confabulations or misinformation, hallucinations are core to the fundamentals of generative AI.²⁹ The LLMs powering AI chatbots, for example, are designed to answer queries by producing statistically likely responses based on patterns in enormous amounts of training data and human tester feedback. But because much of this information is collected from the internet, the LLMs’ training set is bound to contain false or imprecise information, leading chatbots to generate inaccurate responses to queries. LLMs are fundamentally non-deterministic, so “fixing” the training data would not cure the hallucination issue. Researchers emphasize that even with perfect training datasets containing no inaccuracies, any generative AI model would still hallucinate,³⁰ simply because it’s part of the design of LLMs to “play along” with prompts that include incorrect assumptions, even if those assumptions would lead to incorrect responses. Although there are ways to reduce the rate of hallucinations, these methods are computationally expensive and involve other trade-offs that AI companies are not poised to make, such as reducing a chatbot’s ability to generalize.³¹

Bias and Discrimination

Being trained on biased data causes AI tools to produce biased information,³² which can have enormous consequences for everyday people. For example, AI tools are used extensively in HR recruitment efforts³³ despite research showing that these tools tend to exacerbate discrimination in hiring practices. A recent lawsuit filed by the American Civil Liberties Union, for example, involves a deaf Indigenous woman alleging employment discrimination because she was rejected for a seasonal position at Intuit based on her performance on the company’s AI video interview platform.³⁴ She had held seasonal roles at Intuit for years prior to the

interview and repeatedly received positive feedback and bonuses, but research shows that the type of technology underlying these AI interview systems consistently assigns lower scores to deaf and non-white applicants.³⁵ Another study revealed that three popular LLM-based résumé screening tools significantly favor white and male candidates.³⁶

Racist and sexist outputs are based on racist and sexist inputs. That is, almost all large-scale AI tools are trained on massive datasets collected from websites like Reddit and 4chan, which undoubtedly contain discriminatory information: Audits have demonstrated the propensity for datasets to contain biased, discriminatory and hateful information scales along with the size of the model. Subsequent fine-tuning by human developers and their worldviews can also influence these models.³⁸ And far from being solved, the issue will only be exacerbated over time; as these tools enhance their learning based on their own generated output, bias and discrimination on the basis of race, gender, and other identities will continue to be amplified.^{39, 40}

Junk Science (Emotion Recognition)

Substantial scientific evidence that AI systems are not capable of detecting emotions⁴¹ has not stopped AI companies from claiming that they are. For example, a large portion of OpenAI's launch of GPT-4o last year was dedicated to showing off the new model's supposed ability to pick up emotional cues through voice and vision perception capabilities.⁴² The launch also highlighted the system's apparent enhanced capacity to interpret facial expressions in photos and videos to determine a user's emotional state.⁴³ Even more recently, OpenAI claimed that the new GPT-4.5 model has "improved emotional intelligence,"⁴⁴ with Sam Altman likening ChatGPT interactions under this new model to "talking to a thoughtful person."⁴⁵ Unfortunately, there is little evidence that this is true.

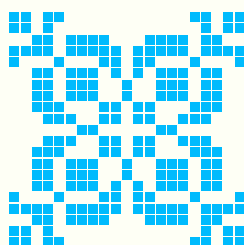
Instead, research is rife with examples of failures by emotion-recognition tools. So-called emotion-detecting AI systems are generally trained by actors conveying specific expressions or vocalizations meant to stereotypically represent particular emotions—like smiling for "happiness."⁴⁶ This relatively simple training set caricatures emotional intelligence, "arguably one of the most complex features of humanity."⁴⁷ Emotion-detecting AI systems, on the other hand, are "by design dependent on the simplification of whatever it is we are defining as emotion in the dataset." Moreover, experts warn that these systems are "founded on tenuous assumptions around the science of emotion that not only render it technologically deficient but also socially pernicious."⁴⁸

Luckily, some governments have recognized the dangers and dubiousness of emotion-recognition technology and have moved toward prohibiting it. The European Union's new Artificial Intelligence Act significantly restricts the use of emotion-recognition systems in the workplace, encompassing all systems that infer emotions from biometric data—including facial expressions, voice patterns, keystrokes, body postures, or movements.⁴⁹ Even Microsoft decided to retire emotion-recognition technologies from its facial-recognition operations.⁵⁰ But despite this policy consensus, the generative AI boom has revived interest in emotion-recognition tools, with OpenAI, Amazon, and Alibaba all releasing models that claim to have these capabilities.⁵¹

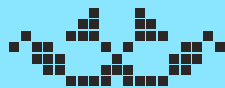
Importantly, when AI systems fail, they don't fail evenly across the population at large: In many instances, the risks or errors arising from untested and unproven technologies fall disproportionately on low-income communities, immigrants, and people of color. More than a decade of research has shown how algorithms encode bias, from predictive policing systems that replicate historical patterns of "dirty" policing;⁵² to algorithms used by insurers that disproportionately deny coverage to Black patients;⁵³ to hiring systems that boost white, male candidates over others by replicating discriminatory employment practices such as "cultural matching."⁵⁴ Algorithmic discrimination is especially well-documented in the use of biometric systems like facial recognition, which have long struggled to maintain accuracy levels across gender

presentations and for individuals with darker skin pigmentation.⁵⁵

The FTC's 2024 case against Rite Aid offers an instructive example of how these failures can lead to direct harm: When Rite Aid used a flawed facial recognition system in its security cameras, it persistently flagged people of color as presenting security risks; in more than one instance, this resulted in bans on individuals from Rite Aid stores and the police being called in error.⁵⁶ Moreover, Rite Aid had failed to properly train its staff in how to use the system—such training could have helped employees determine when the system failed. The company's conduct was egregious enough that the FTC instituted a ban on Rite Aid's ability to use this technology for five years.



In AI We Trust?



AI firms are taking a page from the Stephen Colbert playbook, using “truthiness”—assertions that carry the veneer of truth without an underlying basis in fact—to justify the rapid rollout of AI in profoundly sensitive domains.

❖ Overreliance on “white papers” in lieu of peer-reviewed research

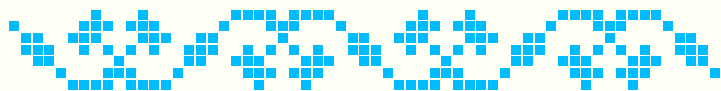
- ◆ Many AI Labs are using publication of non-peer-reviewed work through venues like ArXiv to circulate work that looks and sounds academic, but lacks methodological rigor and avoids peer review. This marks a change: It used to be common for researchers at corporate labs to participate in the peer-review process and publish at leading conferences and journals. Nevertheless, claims asserted within industry-produced papers are uncritically recirculated in popular press and become canon.
- ◆ Companies tend to circulate their own research as a PR tactic, leading to the mass circulation of unverified claims.
 - Microsoft’s “Sparks of AGI” paper was circulated to bolster the narrative that large language models are exhibiting “capabilities” reflective of AGI.⁵⁷ This obscured the significant ongoing debate in the AI community not only on what AGI means, but on the likelihood that it will be achieved.⁵⁸
 - A recent Forbes article described a research study suggesting that Nvidia’s AI bot outperforms nurses. However, the research study was conducted by Nvidia itself.⁵⁹
 - Salesforce, a company that sells AI agents, has published numerous press releases on its own research studies suggesting that 77 percent of students report they would use AI agents to help with school processes,⁶⁰ 90 percent of constituents would like to see AI agents in public service,⁶¹ and AI agents can cut healthcare paperwork by 30 percent—yet none of these studies were peer-reviewed, published in journals, or verified by third parties.⁶²
 - On its education research resource page, Google links to a report suggesting AI’s potential to enhance student learning outcomes in the classroom. The report is authored by Pearson, an education technology company that sells AI-powered learning tools.⁶³

❖ Flawed methods that boost assertions about model performance

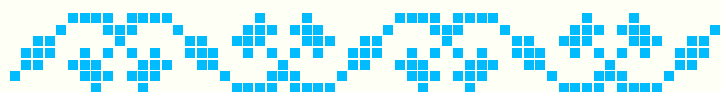
- ◆ Failure to use robust methodologies in machine learning research has enabled false assertions about system performance to proliferate.⁶⁴ Among other issues, many studies have been critiqued for failing to demonstrate construct validity—meaning that the test used to evaluate them is an accurate measure of the concept it’s intended to measure.⁶⁵
- ◆ In 2020, a team of researchers published a paper (cited more than nine hundred times) that claimed AI could be effectively used to diagnose COVID-19 via chest X-rays.⁶⁶ Later, two scientists at Kansas State found that the AI model was picking up on background artifacts—not clinically relevant features of the photos—making them “medically useless.”⁶⁷
- ◆ Another meta-review conducted in 2021 examined sixty-two additional studies that attempted COVID-19 diagnosis using machine learning evaluation of chest X-rays, and found that methodological flaws and underlying biases invalidated *every single study reviewed*, rendering their findings useless to clinicians.⁶⁸

❖ Self-dealing in the development and use of benchmarks

- ◆ The absence of independent and robust evaluation metrics for foundation models is a persistent barrier to implementing more robust validation requirements for these systems.⁶⁹ This is a hard problem on its own: The benchmarks being used currently are drifting away from evaluating actual model capabilities,⁷⁰ leading to gaming of the system;⁷¹ and the increased generality of large-scale models makes them harder to measure.⁷²
- ◆ In the absence of independent, widely agreed-on benchmarks for measuring key attributes such as accuracy, companies are inventing their own, and, in some cases, selling both the product and platforms for benchmark validation to the same customer.
- ◆ For example, Scale AI holds contracts worth hundreds of millions of dollars with the Pentagon to produce AI models for military deployment⁷³—including a contract for \$20 million for the platform that will be used to assess the accuracy of AI models for defense agencies.⁷⁴
- ◆ The use of overgeneralized benchmarks is particularly problematic when AI technologies are implemented in areas with widely differentiated thresholds—what might pass muster for accuracy in a behavioral marketing setting, for example, won’t translate well into a setting with life-or-death stakes, as in healthcare or warfare.⁷⁵



2. AI-SIZED SOLUTIONS TO ENTRENCHED SOCIAL PROBLEMS DISPLACE GROUNDED EXPERTISE



The tech industry has long been prone to technosolutionism, insisting that technical expertise is substitutable for other forms of expertise and can offer quicker or more scalable solutions.⁷⁶ But what's different about the current wave of AI hype is that the industry is now seeking to reframe social problems—down to their root causes—in order to assert AI as the universal fix. This has the effect of undermining the authority of trained professionals who come from these fields.

In education, technosolutionism takes various forms: farming out educational activities from plagiarism detection, to assessment, to grading; implementing technologies that have persistently been shown to lead to worse education outcomes and poorer teaching conditions; and degrading public perception of education and devaluation of educational workers' labor.⁷⁷ In the effort to minimize costs, university administrators adopting these tools marginalize the professionals who are at the center of the educational endeavor. This pattern can be observed across industries.

Many educators remember the failures from the last technological solution craze, massive open online courses (MOOCs), which promised to democratize access to education by providing free online courses to thousands of students online. Universities downsized their departments and invested in technological and physical infrastructure to make online videos. “Now MOOCs have faded from glory,” writes Tressie McMillan Cottom, “but in most cases, the experts haven’t returned.”⁷⁸

Instead, AI initiatives have taken their place in the next cycle of technosolutionism: OpenAI’s ChatGPT Edu initiative provides US universities with access to a suite of tools designed to “bring AI to their campuses at scale.”⁷⁹ OpenAI recently partnered with the government of Estonia to bring ChatGPT Edu to all high schools in the country.⁸⁰ The Trump Administration issued an executive order encouraging the adoption of AI into K–12 education through public-private partnerships, with the goal to train students in AI and incorporate AI into teaching-related tasks from training to evaluation.⁸¹ California State University, the largest public four-year university in the US, recently announced a \$16 million deal with OpenAI, Google, Microsoft, Nvidia, and others to create an “AI-empowered higher education system.”⁸² In some cases, schools are shifting entirely to an AI-centered learning model: Alpha School markets itself as an “AI-powered private school” where kids can meet with an AI tutor to “crush academics in two hours” and then use the remaining time to “pursue their life passions.”⁸³

Los Angeles Unified School District, the second largest in the United States, conducted a disastrous experiment with an AI chatbot, “Ed,” from AllHere, a company that quickly imploded after it was revealed that it compromised sensitive student data, including special-education and student-discipline information.⁸⁴ Sometimes AI is being added quietly in the background, such as through integrations into Can-

vas (the learning management system used by many university systems).

Schools and school districts end up cutting jobs—often through failure to renew contracts with term employees, which don’t show up in statistical accounting on layoffs—in order to spend millions of dollars on technologies that ultimately fail to deliver on their promises, leaving frustrated students and educators in their wake.

This recent push to integrate large-scale AI systems into schools is particularly nefarious because it threatens to undermine the purpose of education, trading the time-consuming process of learning for “efficiency.” As Professors Sonja Drimmer and Christopher J. Nygren write in *How We Are Not Using AI in the Classroom*, large-scale AI systems like LLMs are good at pattern recognition and token prediction, but not at learning.⁸⁵ Drimmer and Nygren argue that a fundamental disjuncture exists between what LLMs are trained to do—predict what is most likely to come next—and what educators train students to do: find details that diverge from the baseline, imagine alternatives, and foster the capacity to “think well, read well, listen well, and look well.” For Drimmer and Nygren, learning is as much about the process of learning—equal parts lived, institutional, and natural human instinct—as it is about the outcomes. “The power of learning to write is not the written product itself but the process of learning to write. Ultimately, AI short circuits that process and in so doing breaches the entire educational contract.”⁸⁶

The adoption of facial-recognition systems in schools not only heralds the replacement of teachers and counselors with AI-enabled cameras, but also deprives teachers of the agency to decide how to improve safety in their school communities. In 2018, the Lockport City School District in upstate New York procured \$4 million in state grant funding and purchased a facial-recognition technology system to

use in its schools.⁸⁷ Despite the state grant program’s requirement to engage parents, teachers, students, and the school community, the decision to purchase the system was taken at a sparsely attended school board meeting in late summer.⁸⁸ The teachers’ union president would eventually say they were not even consulted on the decision.⁸⁹

Allegedly purchased to “prevent school shootings,” Lockport’s system consisted of a “red list” of individuals barred from the school campuses, including local registered sex offenders and suspended students.⁹⁰ Instead of using teachers and counselors to help students in crisis, the district turned to technology that lacked the ability to prevent, or even detect, a school shooting.⁹¹ Community organizing, litigation, and legislative advocacy was ultimately able to defeat Lockport’s system—but, absent federal regulation, other school districts have persisted in obtaining facial-recognition systems.⁹²

In addition to hard security, schools have also turned to online surveillance systems, a practice that started during the pandemic. One system in particular, GoGuardian, infiltrated large school districts across the US, including New York City Public Schools.⁹³ Under the guise of preventing student self-harm, GoGuardian allowed teachers and administrators unfettered access to student devices, with the ability to view Google searches, remotely activate webcams, perform web filtering, and close tabs.⁹⁴ GoGuardian claims it has thwarted more than 18,000 attempts at self-harm, citing its own system as the source for this statistic.⁹⁵

Those with clear on-the-ground expertise are rarely included in decision-making about where and under what conditions AI is deployed. For example, educators Martha Fay Burtis and Jesse Stommel detail a technology adoption process at their regional public university. The educators were invited to participate in meetings with a new technology vendor, EAB (the

former Education Advisory Board), believing their role was to offer insight and expertise around tech use in the classroom and beyond. Instead, they discovered that “the choice about adopting this platform had already been made, and there was little opportunity to engage meaningfully with EAB’s representatives about the misalignments we observed.”⁹⁶ In most cases in higher ed, as well as in K–12, administrators negotiate most corporate educational technology vendor contracts without any involvement from faculty members, students, or parents and little, if any, accountability to those core constituents. Britt Paris and colleagues found that across higher education, university administrators adopt unproven and untested corporate educational technologies for vast sums of money to supplant existing technologies run by university technology services. Those corporate technologies were once touted as data-driven technologies, and now increasingly incorporate LLMs and AI into their technological infrastructure, as with Canvas’s Khanmigo, which uses GPT-4.⁹⁷

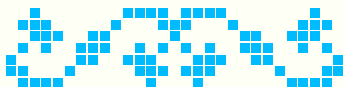
Outside of education, labor shortages are used as a justification for AI deployment. But many so-called “labor shortages” are the product of poor working conditions, inadequate pay, and institutional failures—none of which will be solved by AI. When it comes to nursing, for example, hospitals have trouble staffing nurses not because of a nursing labor shortage, but because of the failure of hospital boards and administrators to implement critical health and safety protections. National Nurses United has stated clearly that there is no shortage of nursing professionals in the field. “Simply put,” they write, “there is a failure by hospital industry executives to put nurses and the patients they care for above corporate profits.”⁹⁸ Large technology investments designed to replace nurses do nothing to meaningfully address the root causes of problems—such as federal minimum standards to support strong nurse-to-patient staffing ratios or in-

vestments in adequate personal protective equipment (PPE)—or disrupt the underlying profit motives.

Nevertheless, even as hospital workers grapple with unsustainable working conditions, hospitals continue to invest money and resources into AI technologies. In 2023, the Permanente Medical Group, a division of US healthcare giant Kaiser Permanente, signed a large-scale partnership agreement with Nabla, an AI transcription company, at the same time that Kaiser faced the largest healthcare worker strike in US history due to understaffing, burnout, and low wages.⁹⁹

There are parallels here to the labor crisis facing the agricultural industry, which is widely viewed as one of the most dangerous lines of work because of high rates of injury, pesticide exposure, high heat conditions, and lack of sun protection.¹⁰⁰ It should come as no surprise that workers are not eager to join an industry with a legacy of racism,¹⁰¹ long hours of back-breaking labor, and criminally low wages. Farmworkers have some of the lowest annual family incomes in the United States and are categorically excluded from the National Labor Relations Act of 1935, which provides workers with critical labor protections like the right to organize.¹⁰² State-level regulations to protect farmworker organizing are also being undone by the Supreme Court.¹⁰³ On top of this, many farmworkers are affected by immigration status, and fear deportation if they raise any concerns about their negative working conditions. Yet, agritech startups are targeting small farms with AI-powered robotics that will purportedly resolve labor shortages.¹⁰⁴ Stout, the producer of a smart cultivator used in fields across the country, is marketed as a solution to reduce the reliance on “costly and scarce manual labor.”¹⁰⁵

Quick-Fix Solutionism Undermines Structural Interventions



In the rush to implement quick-fix AI solutions, administrators are displacing the long-term structural interventions needed to improve educational outcomes. This was perhaps most clear in the education sector during the COVID-19 pandemic, when school districts rushed to adopt faulty computer-monitoring software rife with potential privacy concerns rather than address students' needs. As educators Martha Burtis and Jesse Stommel write: “Frankly, it’s insulting when institutions throw money at corporate edtech when so many of their most marginalized students are struggling, faculty/staff have been furloughed, public funding has been decimated, and the work of teaching has been made altogether precarious.”¹⁰⁶ The American Association of University Professors, the union for professors, is developing a labor strategy for higher education workers addressing how AI is affecting their workplaces.¹⁰⁷

In the educational sector, the jump to implementing technological solutions serves to divert funding and attention away from the kinds of investments that are most meaningful to students, like smaller classroom sizes and good facilities. Simultaneously, it functions to fuel the conditions that further justify the dismantling of public education, in the name of “AI teachers.”¹⁰⁸

The same can be said for the legal system, where funding is diverted away from meaningful measures to provide affordable and accessible legal services and into AI tool development designed to “hack” lawyering.¹⁰⁹ Andy J. Semotiuk writes about how immigration attorneys’ reliance on AI tools without proper

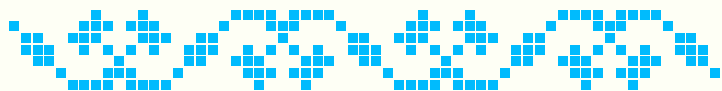
vetting “introduces a dangerous vulnerability, potentially exposing immigrants to the risk of erroneous legal advice, unjust outcomes, and exploitation.”¹¹⁰ These effects are not just individual, either. “The consequences of such negligence or malpractice can extend far beyond individual cases,” Semotiuk continues, “impacting entire communities and perpetuating systemic injustices if digital inaccuracy distorts the legal domain.”

Epic’s Sepsis Model Failure: A Case Study in Technosolutionism

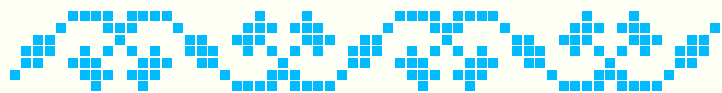


In 2017, Epic, the leading technology company for electronic medical records, released an AI tool designed to predict sepsis, a deadly condition that develops in response to infection. The algorithm was designed to predict which patients were at risk for developing sepsis so that health-care professionals could act quickly to prevent onset. Epic advertised that its algorithm was accurate 80 percent of the time. Without verifying this claim—and without any regulatory oversight or approval—hundreds of hospitals implemented the algorithm.¹¹¹ (The model was included as part of Epic’s “honor roll” incentive program, which provides hundreds of thousands of dollars to hospitals that implement Epic’s technology.)¹¹²

Yet when researchers at the University of Michigan evaluated Epic’s model in Michigan’s healthcare system years later, they found that the model was accurate only 63 percent of the time. The model also routinely identified false alarms, drawing doctors away from patients with other, high-risk medical conditions.¹¹³ Washington University reported similar accuracy rates.¹¹⁴ Epic disputed the research and findings. But after lengthy investigations from STAT, a leading medical publication, into Epic’s models—including Epic’s lack of model transparency—Epic rereleased its model, overhauling the model’s data variables and definitions and providing lengthy updated guidance on implementation.¹¹⁶



3. AI SOLUTIONISM OBSCURES SYSTEMIC ISSUES FACING OUR ECONOMY—OFTEN ACTING AS A CONDUIT FOR DEPLOYING AUSTERITY MANDATES BY ANOTHER NAME



Whether or not AI technologies live up to the hype may not even matter to the actors most invested in their rollout: AI solutionism is used to obscure systemic issues facing our economy, acting as a justification for austerity by another name. A recent report by TechTonic Justice provides a comprehensive illustration of how AI is used to restrict the opportunities for low-income people, exposing them to AI decision-making that denies them access to resources and shapes their life chances, from Medicaid coverage denials, to how much renting an apartment will cost, to determinations about family separation by child welfare agencies.¹¹⁷ The report states:

“The alleged rationality and objectivity of the system allow the users of AI to justify harmful actions, like benefit cuts or law enforcement harassment of students, or reinforce status quo power imbalances, such as that between employers and employees. Reveal-

ingly, the AI systems applied to low-income people almost never improve access to benefits or other opportunities.”¹¹⁸

Numerous historical examples reveal where this playbook leads: the widespread disenfranchisement of the public and costly litigation for the government agencies deploying AI services. Importantly, AI tools have not improved in their capabilities; if anything, the flaws in foundation models make them even worse:

- ❖ Over a decade ago, Michigan spent \$47 million to develop an automated fraud-detection system intended to identify benefit fraud in the state.¹¹⁹ After accusing over sixty thousand residents of fraud, it was revealed that 70 percent of the system’s determinations were incorrect. Similar errors concerning Medicaid eligibility occurred in Indiana, Arkansas, Idaho, and Oregon.¹²⁰
- ❖ In 2013, the Dutch government implemented an algorithmic system to identify potentially fraudulent claims for child care benefits. The system considered an applicant’s nationality as one of its risk factors, and non-Dutch nationals were often assigned higher risk scores compared to Dutch nationals. Years of reliance on this algorithm resulted in false determinations of fraud for tens of thousands of people, many of whom were low-income and subjected to rescinded benefits or harsh penalties.¹²¹
- ❖ Since 2016, Allegheny County in Pennsylvania relies on a predictive risk model to supplement the abilities of its human intake screeners to identify possible instances of child abuse and neglect. The model computes risk scores using indicators like call frequency and data obtained from public agencies. Consequently, children in families that are more likely to be surveilled by mandated reporters or community mem-

bers—intake workers receive calls about black or biracial families three and a half times more frequently compared to white families—or those that simply rely on public programs are likely to receive inaccurate scores.¹²²

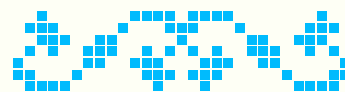
Unfortunately, we know of the examples above because individuals on the receiving end of an AI system were able to determine that they were wronged, and used the legal system—in combination with public pressure—to push for accountability. But more often we don't know how AI is being used.

AI is often deployed privately and surreptitiously, such as through the use of so-called “social scoring” systems that draw together disparate data sources to make determinations about access to resources, leaving those impacted by these systems without the necessary information to know and understand their effects. For example, an investigative report conducted by *The Markup* and *The New York Times* into hundreds of federal lawsuits filed against companies deploying tenant-screening algorithms found that reports often contain glaring errors, resulting in qualified individuals being denied housing without ever knowing an incorrect background report was responsible—let alone having the opportunity to correct it.¹²³

Or in other places, firms leverage the information asymmetries afforded by their dominant position in a market against consumers. The Department of Justice—joined by eight different states—filed a lawsuit in 2024 against RealPage, a private-equity-owned software company that collects confidential information from landlords about rents and occupancy rates, and then uses an algorithm to suggest inflated rental rates to landlords on their platform. The lawsuit uncovered that over three million rental units use this pricing technology, which RealPage advertises to landlords as a way to increase prices above market rate.¹²⁴ The absence of affordable housing is an entrenched pub-

lic-policy problem that exists outside of AI. But the use of AI both exacerbates the underlying problem—by making it easy for landlords to wrongly deny access to housing to potential tenants, and by keeping housing off the market in the hope of extracting higher rents—and makes the problem more obscure to the public.

DOGE's Power Grab



Nowhere is AI solutionism weaponized more than in the Department of Government Efficiency's wholesale attack on the administrative state under the guise of government “efficiency.”¹²⁵ AI is a central part of DOGE's stated efforts to advance a broader austerity agenda by “modernizing” federal technology to “maximize government efficiency.” For example, DOGE officials have claimed that “AI” can be used to identify budget cuts, detect fraud and abuse,¹²⁶ automate government tasks,¹²⁷ and determine whether someone's job is “mission critical”¹²⁸—despite any tangible evidence that AI is capable of doing these tasks effectively, or at all.¹²⁹ DOGE's language around AI also perpetuates the false idea that “AI” is a coherent set of technologies able to meet complex social goals.¹³⁰

Whether DOGE is actually using AI, or whether AI can even do these functions at any level of acceptable technical performance, is beside the point. The invocation of AI and its many “innovations” has enabled DOGE to effectuate a wholesale data and power grab, gaining unrestricted, “god-tier” access to data across federal agencies, centralizing government contracting databases, and bulldozing government technical infrastructure to make them legible to AI systems and the tech companies that control them.

DOGE claims that it is making government more advanced and efficient, but there is incontrovertible evidence that DOGE is making the federal government decidedly less efficient and less technologically adept. DOGE has put tens of thousands of expert civil servants out of their jobs and rendered the democratic process itself “waste” to be eliminated. It has shut down the office in the Social Security administration whose job it was to digitize signatures and promote cybersecurity enhancements,¹³¹ and a team of technologists within the General Services Administration whose job it was to maintain critical government digital services like Login.gov.¹³² DOGE has ambitions to cut the Agency for Healthcare Research and Quality (AHRQ), a research center designed to find efficiencies in healthcare quality research.¹³³ The Department has sowed mass confusion and chaos, issuing unclear directives via email and costing the government tens of millions of dollars in wasted work hours.¹³⁴ In some cases, DOGE has had to rehire workers it previously fired.¹³⁵

But while DOGE sets its sights on cutting critical agencies with negligible budgets (for example, the AHRQ makes up 0.2 percent of the government’s healthcare spending), it has failed to turn its AI sword on private tech companies—including Elon Musk’s suite of technology companies—that have received billions of dollars in federal contracts and counting.¹³⁶ The organization also claims it is reducing fraud, yet a substantial focus of its cuts has targeted agencies and departments wholly unrelated to fraud, such as canceling government leases for the U.S. Fish and Wildlife Services in Colorado, Montana, and North Dakota.¹³⁷

DOGE is riddled with these inconsistencies and hypocrisies. But debating the merits of DOGE’s efforts to improve government efficiency requires acceptance of the premise that “efficiency” is DOGE’s goal. It is not. **DOGE is a power grab by process, with AI solutionism operating as a smoke screen to consolidate**

executive power and reshape the federal government to fit the ideological agenda of the Trump Administration and its backers—some of whom own the tech companies that stand to benefit the most from both federal adoption of AI and the turn to austerity.

The millions of Americans who rely on the Federal government to access critical and lifesaving services will feel the effects of DOGE’s brazen power grab most immediately. The havoc DOGE is wreaking—whether through freezing funding, flagging individuals for fraudulent activity, or eliminating staffing necessary to effectively deliver services—will inevitably lead to services being denied, delayed, or inappropriately revoked. DOGE’s 12 percent cut of the Social Security Administration (SSA) workforce has already led to reports of website crashes, hold times of more than four hours, and delayed application timelines.¹³⁸ Current reports that DOGE is looking to migrate SSA’s computer systems in a matter of months is all but certain to lead to massive disruptions to over sixty-five million people receiving Social Security benefits.¹³⁹

Tens of thousands of civil servants have already lost their jobs, pushed out of agencies ranging from Veteran Affairs, Health and Human Services, and the Consumer Financial Protection Bureau. Others have been told their jobs could easily be replaced by AI. But they can’t be. AI is not designed to administer the federal government, nor is it capable of doing so.¹⁴⁰ But when the AI systems introduced inevitably fail, DOGE has created the perfect vacuum for private companies to swoop in and clean up the mess. Writes Eryk Salvaggio: “By shifting government decisions to AI systems they must know are unsuitable, these tech elites avoid a political debate they would probably lose. Instead, they create a nationwide IT crisis that they alone can fix.”¹⁴¹ Workday CEO Carl Eschenbach has already called DOGE a “tremendous opportunity” to integrate his company’s portfolio of cloud and AI products (such as autonomous agents) into the government.¹⁴²

Threatening to replace government administration with AI systems may demoralize federal workers and imperil millions of Americans' livelihoods, but it's fantastic for business.

In order to centralize government systems, DOGE has been given unrestricted access to people's sensitive, private, and personally identifying data. This includes people's entire credit histories, social security benefits payments, and tax data.¹⁴³ DOGE's staff has also been given access to private market data—such as information contained in confidential investigations and actions taken by enforcement agencies—meaning that access to data that is highly desirable to private companies is easily accessible to a team of staffers run by Musk, a private, unelected citizen whose own suite of companies directly competes with the companies targeted in these actions and investigations.

Some of the worst consequences of DOGE—especially for everyday people—will not be immediately obvious until it is too late. Decimating public research infrastructure, for example, will undermine researchers and institutions that otherwise would be working on potentially life-changing technological breakthroughs. Dismantling the regulatory state, including agencies like the Consumer Finance Protection Bureau tasked with protecting ordinary Americans from predation by private companies, means fewer people will receive financial redress from corporate wrongdoing. DOGE has already tapped a former Tesla engineer to become the Chief Information Officer at the Department of Labor, the principal federal agency responsible for enforcing federal labor law, putting critical protections and material benefits like pensions at risk.¹⁴⁴

MyCity: A Case Study in AI for Austerity

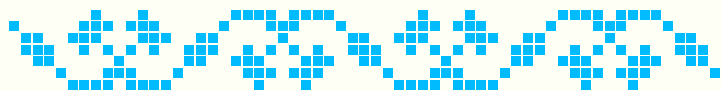


In New York City, Mayor Eric Adams campaigned on claims of making the city government more efficient, promising to appoint an “efficiency czar” and develop a centralized technology portal called MyCity that would serve as a one-stop resource for New York City residents to access city services, apply for benefits, and access useful city information.¹⁴⁵ But the realities of MyCity have been far more sobering. In their report *MyCity, INC: A Case Against ‘CompStat Urbanism,’* Cynthia Conti-Cook and Ed Vogel at Surveillance Resistance Lab articulate how MyCity embeds corporate technology into public infrastructure in ways that undermine democratic governance and deploy austerity measures.¹⁴⁶

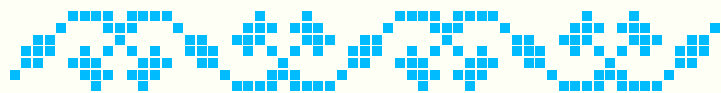
Since 2023, New York City has partnered with Microsoft's Azure AI to build and launch a “MyCity Chatbot” that has provided false—sometimes criminal—information to people.¹⁴⁷ For example, the chatbot told one user that it was legal to fire an employee if they file a sexual harassment complaint¹⁴⁸ (It is not¹⁴⁹). As of writing, the chatbot still warns users that it “may occasionally provide incomplete or inaccurate responses.” A chatbot that requires users to fact-check its answers doesn't benefit anyone in the name of efficiency—except Microsoft, the company on the receiving end of the government contract.

Crucially, when the city attempted to understand the data the chatbot was trained on, Microsoft claimed that the training data was “proprietary to the vendor,” evading any accountability and leaving government oversight committees in the dark.¹⁵⁰ Meanwhile, MyCity was accompanied by a state-wide legislative effort to allow agencies to cross-share data that would otherwise be restricted for the purposes of providing government benefits or services, raising broad privacy concerns.¹⁵¹

As of 2025, MyCity has spent over \$100 million on private contracts with technology vendors—most of which are located outside of New York—with few benefits to show New Yorkers.¹⁵² Still, the city's Chief Technology Officer hopes to expand MyCity and the faulty AI chatbot.¹⁵³ Such a decision makes no sense, unless you consider that the goal of projects like MyCity may not, in fact, be to serve citizens but instead to incentivize and centralize access to citizen data; privatize and outsource government work; and entrench corporate power without meaningful accountability mechanisms—another textbook example of AI serving as “double speak” for austerity.¹⁵⁴



4. THE PRODUCTIVITY MYTH OBSCURES A FOUNDATIONAL TRUTH: THE BENEFITS OF AI ACCRUE TO COMPANIES, NOT TO WORKERS OR THE PUBLIC AT LARGE



CEOs of AI companies have made assertions that the technology will lead to untold and transformative productivity growth. As OpenAI CEO Sam Altman said in an interview in June 2023: “I think the world will get way wealthier, we’ll have a productivity boom, and we will find a lot of new things to do. [...] We can cure all diseases, we can get everybody a great education, better health care, massively increase productivity, huge scientific discovery, all of these wonderful things and we want to make sure that people get that benefit, and that benefit is distributed equitably.”¹⁵⁵

But the reality is far from equitable. For Altman, as for other leaders in AI and Big Tech, “productivity” is a euphemism for the mutually beneficial economic relationship between firms and their shareholders—not between firms and their workers. Not only are workers not benefitting from productivity gains from AI but for many, their conditions only get worse.

Instead of extending the gains to workers, AI is devaluing their labor, and making life more mundane and increasingly surveilled. These promises aren’t new: The idea that AI will lead to enhanced worker productivity has driven the deployment of algorithmic management techniques and “worker productivity” tools¹⁵⁶ used across many sectors that rely on heavy surveillance of workers¹⁵⁷—down to the micromovements of some workers’ muscles—to set rates and shape working conditions using AI.¹⁵⁸ For example, Amazon tracks warehouse workers down to the minute using AI and can fire workers if they accumulate thirty minutes of “time off task” on three separate days within a year.¹⁵⁹ Time off task includes using the restroom, helping a coworker move a heavy package, or taking a break to cool off or warm up, even when warehouse temperatures are extreme.¹⁶⁰ Rather than address worker demands for improved conditions, Amazon announced plans to introduce “wellness chambers” into warehouses where workers could watch videos on relaxation.¹⁶¹

This relentless tracking has led to increased rates of workplace injuries. In late 2024, a Senate investigation found that Amazon’s algorithmically driven productivity quotas lead to significantly higher injury rates than the industry average and in non-Amazon warehouses.¹⁶² In fact, the report found that Amazon workers were almost twice as likely to be injured and that Amazon manipulated its workplace safety data to appear safer than it was.¹⁶³

Furthermore, the algorithmically determined rates at which workers are expected to perform never slow down; instead, work becomes “gamified,” with the next reward always just out of reach—and always to the benefit of the company. Amazon warehouse workers report a cyclical approach where the rate set for them becomes unsustainable and they’re fired, only to be rehired at the bottom of the totem pole again with the wages reset accordingly. Amazon’s treat-

ment of its workers reflects the future embraced by AI-driven technology companies wherein automation technologies are integrated across the entire labor supply chain, with little regard for the human effects of automation.

This is by design, so that companies can commodify labor into a product that can be automated and sold for profit, treating workers' craft as disposable. This is why, regardless of the actual efficacy of AI technology, fears about displacement by AI are justified: Companies use the logic of AI's "productivity gains" to justify the fissuring, automation, and, in some cases, the elimination of work. For example, in late 2024, fintech company Klarna boasted that it used AI to drive company cost-savings by cutting its sales and marketing teams, shifting to AI-powered engineering, and replacing its customer service teams with an OpenAI customer service chatbot that could "do the work of 700 humans."¹⁶⁴ By 2025, the company had reduced its employee headcount by 38 percent. Klarna's turn to AI is presumably driven by its plans to IPO and the hope that the cost savings will attract potential shareholders and drive a higher opening price.

The logic that corporate productivity will inherently lead to shared prosperity is deeply flawed.¹⁶⁵ In past eras when automation led to productivity gains and higher wages, it was not because of the technology's inherent capabilities, but because corporate and regulatory policies were designed in tandem to support workers and curb corporate power.¹⁶⁶ The boom in machine-tool automation around World War II is instructive: Despite fears of job loss, federal policies and a strengthened labor movement protected workers' interests and demanded higher wages for workers operating new machinery.¹⁶⁷ Corporations in turn instituted policies to retain workers—like redistributing profits and providing training—to reduce turmoil and avert strikes. As a result, despite growing automation during this period, workers' share of national income

remained steady, average wages grew, and demand for workers increased.¹⁶⁸ These gains were rolled back under Reagan-era policies that prioritized shareholder interests, used trade threats to depress labor and regulatory standards,¹⁶⁹ and weakened pro-worker and union policies, all of which enabled tech firms to amass market dominance and control over key resources. The AI industry is a decisive product of this history.¹⁷⁰

The introduction of algorithmic wage discrimination is reflective of both the fissuring of work and surveillance capitalism: Firms pay individual workers different wages for the same work based on algorithmic processing of numerous data points, including demand, location, or worker behavior.¹⁷¹ The lack of transparency into how wages are calculated leads to wage precarity, as workers cannot reliably predict how much they will make on a given day for the same work. Algorithmic wage setting only goes in one direction for workers: down. Rideshare companies, for example, can coerce drivers into accepting low wages by threatening to send them to the back of lengthy queues or deactivate their account.

It's not only rideshare drivers who are affected: A recent Roosevelt Institute study by Katie J. Wells and Funda Ustek Spilda on on-demand nursing services showed that nurses on these platforms are forced to bid against one another for shifts, creating a "race to the bottom" for wages.¹⁷² Furthermore, of the eleven largest platform economy companies examined in Fairwork's report on algorithmic management, only one ensured that workers on its platform meet the federal minimum wage requirements.¹⁷³

Beyond wage exploitation, these platform economy companies are undermining corporate accountability at the expense of worker stability and well-being.¹⁷⁴ Fairwork maps these tactics—which include replacing management roles traditionally occupied by humans with chatbots; placing workers in sensitive settings,

such as nursing homes, without any supervision on site; failing to provide any phone number or contact information for workers to contact in the event of an emergency—and arbitrary deactivations, which cut workers off not only from their livelihood but from all information about their clients, workplaces, and work history without any meaningful avenue of redress. While many platforms have contracts stipulating that workers can be terminated for any reason, larger platform economies have taken this one step further: Instacart, an on-demand grocery delivery service, has joined a lawsuit against a Seattle law that requires companies to give fourteen-day notice of deactivation to workers based on reasonable policies—undermining labor protections for their own workers and other platform economy workers in the process.¹⁷⁵ Contracts also limit workers' ability to bring legal claims against the platforms, with widespread use of liability clauses, class action waivers, and arbitration clauses.

These moves are a feature, not a bug, of platform-economy business models. As Edward Ongweso Jr. writes: “For years, gig companies have pushed for mandatory arbitration because it is incredibly good at stymying class action lawsuits and legal precedents—again, if your business model relies on skirting the law, regulatory arbitrage, and aggressive lobbying, you need to stop angry workers in deplorable conditions from collectively demanding the right to earn a livable wage or be safe in the course of their work or get health insurance or other indications of dignified work.”¹⁷⁶

AI Agents

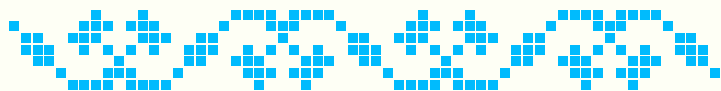


The promise that AI Agents will automate rote tasks has become a focal point for product development: In a recent blog post, OpenAI founder Sam Altman said he believes that in 2025 the first AI agents will “join the workforce,” working autonomously in ways that replace human activities.¹⁷⁷ OpenAI has also released Operator, an agent that it claims can be tasked to fill out forms, order groceries, and perform other browser-based tasks¹⁷⁸ (though it has been largely unsuccessful so far),¹⁷⁹ and Deep Research, which conducts literature reviews using web-based information;¹⁸⁰ Google has also released a tool called Deep Research that performs similar tasks,¹⁸¹ and Palantir is piloting a studio that enables its customers to build their own agents on top of its Ontology platform.¹⁸²

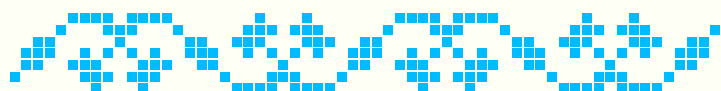
But “agentic” AI introduces multiplying levels of risk. All of the flaws of LLMs are retained with these systems, yet they’re positioned to act more autonomously, with greater complexity, and in ways that render human oversight even more facile.¹⁸³ As Signal President Meredith Whittaker has outlined, AI agents pose particular risks to user privacy, giving root access to private information in order for the agent to act autonomously: “There’s a profound issue with security and privacy that is haunting this hype around agents, and that is ultimately threatening to break the blood-brain barrier between the application layer and the OS layer by conjoining all of these separate services [and] muddying their data,” she said.¹⁸⁴

In order for these systems to “work,” organizations will have to change in some pretty fundamental ways: For one, AI agents require highly structured pools of data to work with.¹⁸⁵ This means that companies using agents will have to “datify” their practices, taking their activities and rendering them into discrete data objects on which AI systems can be trained.

To do this, organizations will need to become *more* bureaucratic and *more* surveillant—the opposite of the promises of greater autonomy typically associated with the tech industry.



5. AI USE IS FREQUENTLY COERCIVE, VIOLATING RIGHTS AND UNDERMINING DUE PROCESS



While the entry point to AI for many people is through systems like ChatGPT, more frequently we are interacting with AI technologies used not by us but on us, which shape our access to resources in realms from finance to hiring to housing. But because these types of AI systems are typically used by powerful institutional actors, they benefit from information asymmetries. And even more so than other types of technologies, it is the deployment of AI in these contexts that is especially and often coercive, offering little transparency to those subject to these technologies and no meaningful ability to opt out.

The rapid implementation of AI systems into critical social infrastructures thus raises imminent concerns about the violation of numerous rights and laws critical to realizing justice, freedom, and dignity, individually and collectively, including due process, privacy, and civil rights. This is nowhere more clear than the rise of AI usage in immigration enforcement, where human rights abuses are common and legal norms are routinely violated—even before AI is in the mix.

Immigration enforcement requires a significant amount of data collection and processing: The files

from a single application for immigration can amount to thousands of pages and include complex legal rulings and other documentation. As a result, immigration agencies housed under the Department of Homeland Security have rapidly integrated AI technologies across a range of aspects of immigration enforcement, but with little-to-no meaningful oversight. And while DHS has published a set of principles regarding its responsible use of AI, immigration advocates have found that the agency routinely skirts those obligations.¹⁸⁶

For example, the United States Citizenship and Immigration Services (USCIS) uses predictive tools to automate the agency's decision-making, like "Asylum Text Analytics," which queries asylum and withholding applications to deem which applications are fraudulent. These tools have demonstrated, among other flaws, high rates of misclassification when used on language from non-native English speakers.¹⁸⁷ And the consequences of erroneous identification of fraud are significant: It can lead to deportation, lifelong bans from the United States, and imprisonment for up to ten years. Still, there is little-to-no transparency for those on whom these systems are used, no ability to opt out or seek remediation when they are used to make erroneous decisions, and—just as importantly—little evidence that the tools' effectiveness has been, or can be, improved.

Immigration and Customs Enforcement (ICE) also uses predictive analytics and risk assessment to make determinations about detention and release: For example, a "hurricane score" is used by the agency to decide the terms of electronic surveillance based on predictions of how likely an individual is to "abscond," or the level of "threat to the community" they face. A "risk calibration score" used to shape detention decisions was set at multiple junctures by ICE so that it would not recommend release for anyone it was used on, leading the agency to be sued for violating due process rights under the Fifth Amendment in 2020.¹⁸⁸

A database Palantir has produced for ICE called Investigative Case Management (ICM) is used in the process of making these determinations, including deciding which people to target for arrest during immigration raids.¹⁸⁹ ICM is a sprawling tool that enables filtering and querying of a database containing hundreds of categories, including “unique physical characteristics,” “criminal affiliation,” “bankruptcy filings,” and “place of employment,” among others, enabling agents to build reports on targets.¹⁹⁰ In 2022, Palantir announced a five-year contract with the Department of Homeland Security worth \$95.9 million to maintain the system.¹⁹¹ It connects to other government databases including records of all people admitted on a student visa, real-time maps of ICE’s location-tracking tools, and location data from license plate readers, among other sources of information.¹⁹² A new contract Palantir announced with DHS would add on a new platform, ImmigrationOS, by September 2025 that would give ICE “near real-time visibility” on people self-deporting from the US,¹⁹³ a “master database” that will purportedly integrate data from the Social Security Administration, the IRS, and Health and Human Services, including data obtained by DOGE teams without observing legal or procedural requirements.¹⁹⁴

The federal government has also invested significantly in facial-recognition technology—overwhelmingly used to surveil and track immigrants, asylum seekers, and activists—found to be biased and error-prone. For example, Clearview AI, a facial-recognition company that was created with the intention to deploy surveillance technology against immigrants, people of color, and the political left, has received almost \$4 million in contracts with ICE despite being sued in numerous states.¹⁹⁵ Reporting from the US Government Accountability Office shows that between April 2018 and March 2022, Clearview was used by more federal law enforcement agencies than any other private company—including by the US Postal Inspection Service,

which used Clearview to target Black Lives Matter protesters; and Customs and Border Protection (CBP), which uses a facial-recognition app to screen asylum seekers.¹⁹⁶

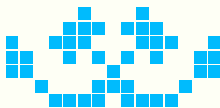
Despite the dubious legality and known flaws of many of these systems, the integration of AI into immigration enforcement seems only poised to escalate: Acting ICE Director Todd Lyons has expressed aspirations for the agency to run deportations like “Amazon Prime for human beings,” and to use AI to “free up [detention] bed space” and “fill up airplanes.”¹⁹⁷ DHS is also experimenting with facial recognition technology to track migrant children—potentially from infancy—and monitor them as they grow up.¹⁹⁸

The use of these tools provides a veneer of objectivity that masks not only outright racism and xenophobia, but also the steep political pressure on immigration agencies to restrict asylum, humanitarian relief, and other forms of immigration—pressures that predate the Trump administration but have sharply escalated since.¹⁹⁹ As immigration experts at Just Futures put it, “regardless of how ‘accurate’ the AI program is at recommending detention and deportation, in the hands of immigration and policing agencies, the technology furthers a fundamentally violent mission of targeting communities for detention and deportation.”²⁰⁰

Moreover, AI allows federal agencies to conduct immigration enforcement in ways that are profoundly and increasingly opaque, making it even more difficult for those who may be falsely caught or accused to extricate themselves. Many of these tools are only known to the public because of legal filings, and are undisclosed in DHS’s AI inventory. But even once they are known, we still have very little information about how they are calibrated or what data they are trained on, which further diminishes the ability of individuals to assert their due process rights. These tools also rely on invasive surveillance of the public, from the screening of social media posts; to use of facial

recognition, aerial surveillance, and other monitoring techniques; to the purchase of bulk information about the public from data brokers like LexisNexis.²⁰¹

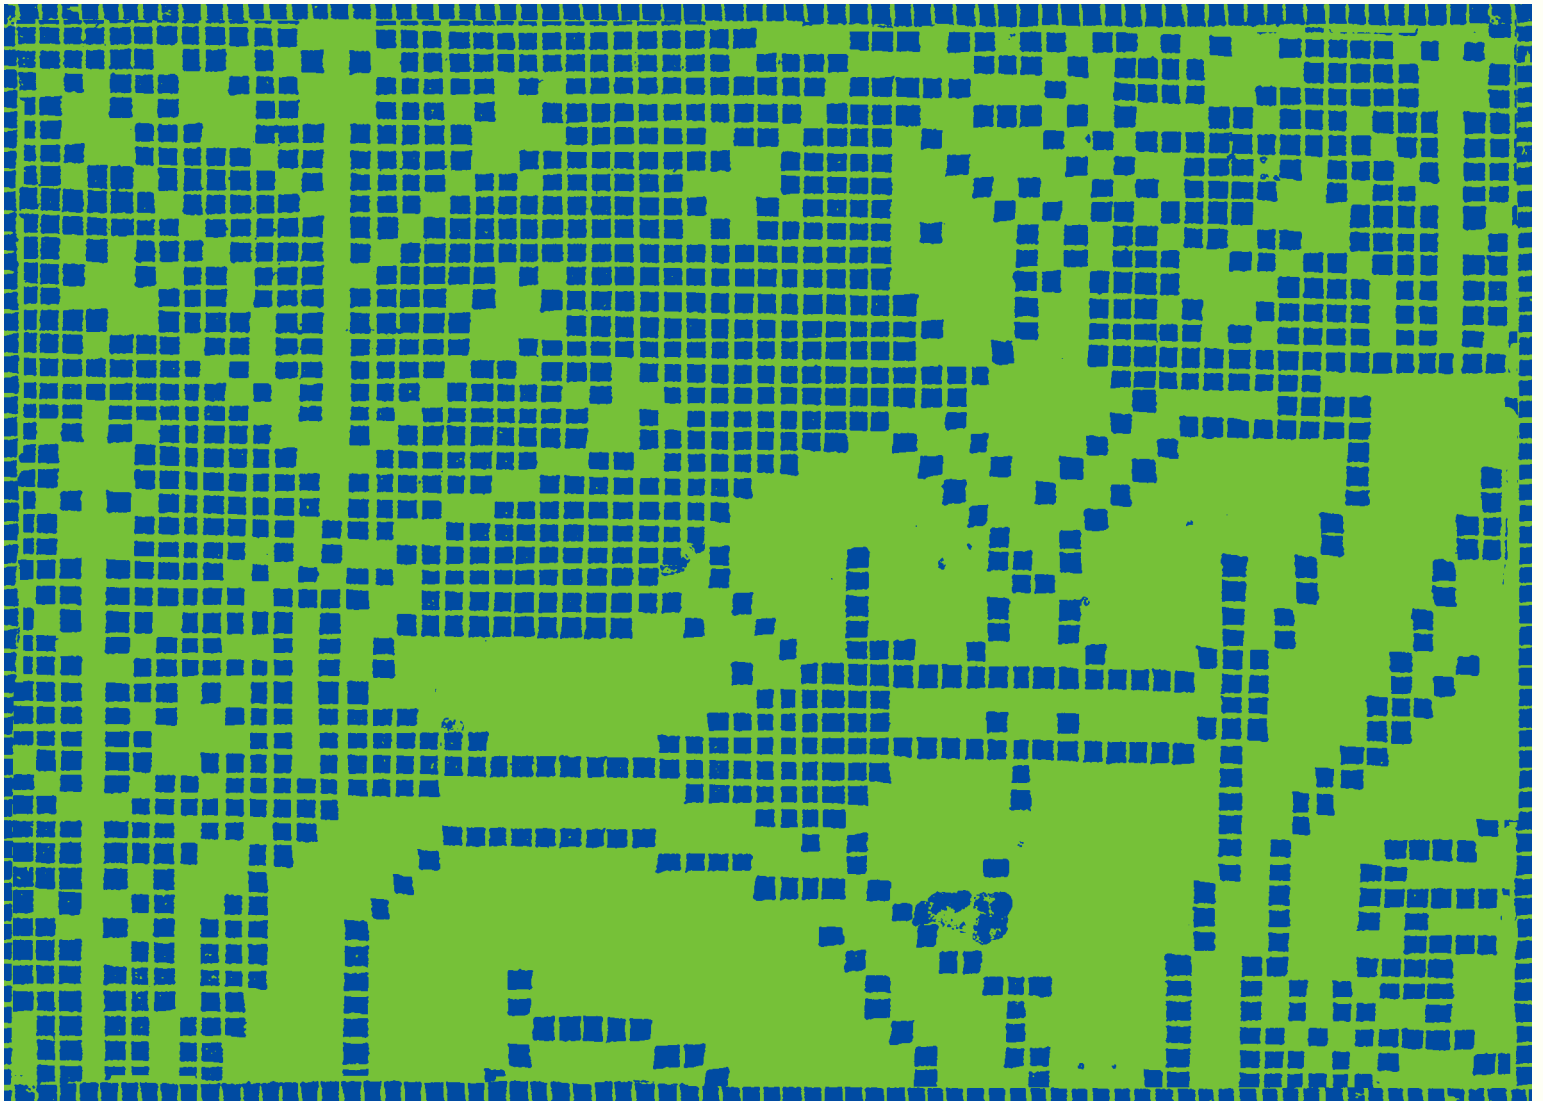
Understanding the coercive nature of these systems is important for understanding how so many individuals who are not in violation of the law, and who in some cases are naturalized US citizens and green card holders, have been unlawfully detained by immigration authorities. Most concerning, the opacity of AI systems helps prevent the remediation of these harms. But maybe that failure is part of the point. AI systems, even when faulty, act in the interest of enforcers who need to meet targets for deportation and arrest—whatever the means, whatever the cost. It's all good for business.



CHAPTER 4:

A ROADMAP FOR ACTION

**MAKE AI A FIGHT ABOUT
POWER, NOT PROGRESS**



In the din of heated public debates about artificial intelligence, from speculative doomsday predictions to starry-eyed visions of an AI utopia, people on both sides of the argument tend to position society as an object of this “inevitable” transformation: We are all passive witnesses to this onward march of technology, which we endow with almost divine agency.

But as we’ve argued carefully throughout this report, this is a fight about power, not progress. We need not reconcile ourselves to passive alarm, nor engage in heedless celebration of speculative futures: Instead, the path ahead requires reclaiming agency over the trajectory for AI. In order to build a movement toward a just and democratic society, we must contest how AI is used to avert it. Only when we do so can we seed a new path defined by autonomy, dignity, respect, and justice for all of us.

Here, we offer a high-level road map of strategic levers and case studies to illuminate the opportunity before us. Some of these speak more directly to advocacy campaigns and narrative strategy; others offer regulatory directions for reining in and rebalancing unaccountable power and rebalancing power:

- ❖ Target how the AI industry works against the interests of everyday people
- ❖ Advance worker organizing to protect the public and our institutions from AI-enabled capture
- ❖ Enact a “zero-trust” policy agenda for AI
- ❖ Bridge networks of expertise, policy, and narrative to strengthen AI advocacy
- ❖ Reclaim a positive agenda for public-centered innovation without AI at the center

Let’s examine each of these in turn.



TARGET HOW THE AI INDUSTRY WORKS AGAINST THE INTERESTS OF EVERYDAY PEOPLE



After decades of mounting inequity, the foundations of a good life have been priced out of reach for most people to such an extent that, in the aftermath of the 2024 election, there’s growing consensus across the ideological spectrum that focusing on the material conditions and economic interests of working people is key to building political power.¹ While the Trump campaign was successful in moving a pro-working-class message, this might in fact be the administration’s core vulnerability,² given that much of the policy orientation of the administration so far has been successfully commanding loyalty from tech and financial elites. There is momentum and urgency around finding ways to drive a wedge through this hypocrisy. The road map we lay out in this section explores ways to ensure that the AI fight revolves around the tech industry’s amassing of power that profoundly affects the day-to-day lives of people—not around progress.

That the AI industry is fundamentally incentivized to hurt the interests of working people and families appears not to be widely understood yet—partly because the most high-profile narratives around AI risk relate to issues of technically framed bias or existential risks, which are often disconnected from

people's material realities. And the fact that AI expertise largely circulates in elite and urban corridors only exacerbates the problem. As we've shown throughout this report, investment and interest in AI is not about supporting progress for the many, but about hoarding power for the few, and doing so by any means available: asymmetric control over information; gatekeeping of infrastructure; hollowing out agency for citizens, workers, and consumers; or supporting political allies, financially or otherwise, who side with Big Tech.

For those of us whose research and advocacy is centered around AI, this means that we need not only to make AI-related issues more relevant to movements fighting for economic populism and against tech oligarchy; we also need to better target the AI industry as a key actor working *against* the interests of the working public. The writing is on the wall: We need to prioritize policy issues that are rooted in people's lived experiences with AI, particularly those that hit at their most urgent material conditions. This involves giving the broader public a way to "see" AI systems as the invisible infrastructures that mediate their lives—often for the worse—and helping them connect the dots between the harms of AI and the unchecked power of Big Tech. Fortunately, **there are existing windows of opportunity for doing just this:**

1. DOGE Pushback



The evisceration of government agencies by DOGE brings home the harms that tech oligarchs can enact on the public. This makes it an important front for resistance to AI. There is widespread discontent with the workings of DOGE teams even among those who advocate for greater government efficiency;³ and

reporting thoroughly documents the failure of DOGE to achieve its stated goals.⁴ While DOGE officials claim that "AI" can be used to identify budget cuts, detect fraud,⁵ automate government tasks,⁶ and determine whether someone's job is "mission critical,"⁷ reporting has dismantled over and over again that both these teams and the technology they are using are not up to the task.⁸ And while it is still unclear exactly what DOGE employees are doing with the data they've gained access to, it is very clear that much of this data is highly lucrative and could be used to further Musk's business interests—including where it is stored, and whether it is by being used to train AI models for Musk's own companies.⁹

The outrage over DOGE offers an entry point not only to draw attention to the harm government agencies are creating with their use of AI, but also to make clear that DOGE's work has never truly been about making our services more "efficient"—the goal has always been to dismantle government services and centralize power (see [Chapter 3](#)).¹⁰ Building campaigns that connect the pushback against DOGE to the broader pattern of hollowing out social services in the name of AI-enabled "efficiency"—from child services, to provision of housing, to access to health-care—with partners like American Federal Government Employees, Federal Unionists Network, and Federal Workers Against DOGE is a crucial starting point. Many of these partners have already made significant inroads in helping the public understand and mobilize around DOGE's harm to workers; broadening the discussion to include how Musk is using AI to help the Trump Administration enact an austerity agenda that impoverishes and disenfranchises the working class, communities of color, the disabled, and those in rural communities could move the needle even further.¹¹

2. Data Centers



The unchecked growth of AI infrastructure is proof of how incredibly successful Big Tech has been at making the case that AI—especially large-scale AI—is worthy of exceptional support and investment from public actors. But the reasons why large-scale AI is a worthwhile use of resources, taxpayer dollars, and land are all based on potential economic growth, the projections of which are usually based on unfounded assumptions (See [Chapter 2: Heads I Win, Tails You Lose](#)). Meanwhile, the harms are definitive and documented, such as those evidenced by Meta and Blackstone’s data centers in rural Georgia¹² (see [Chapter 1.2: Too Big To Fail](#) and [Chapter 1.3: Arms Race 2.0](#)).¹³

The ground is ripe for local mobilization around the material impacts of the AI infrastructure build-out—not only by building, organizing, and campaigning focused on the tech sector but also by forming alliances with the environmental justice movement, which has long and deep experience fighting against policies that result in a physical degradation of natural resources, as well as movements for utility reform and racial justice. Data centers have direct environmental and community health impacts, including (but not limited to) increased greenhouse gas emissions, increased rates of pollution, and the expansion of gas infrastructure and coal plants.¹⁴ They also frequently draw on public aquifers for access to the water supply needed to cool their server racks.¹⁵ This is particularly concerning given plans to expand data center construction in already water-strained areas of the southwest.¹⁶ In the Netherlands¹⁷ and Chile,¹⁸ community-level activism led to a pause on data center construction, which suggests that organizing that targets local government can be an effective lever for

curtailing the rollout of data centers. These issues—along with other consequences like higher energy costs and labor abuses across the supply chain—are also bipartisan and hit rural voters and Indigenous communities¹⁹ particularly hard, providing other critical opportunities for new organizing.

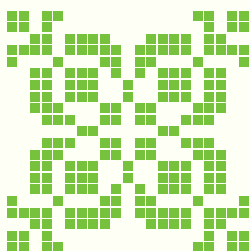
On the flip side, the firms advocating for data center buildout have little proof that data centers actually provide economic benefits to the communities where they are sited. The evidence that industry actors have provided thus far seems speculative at best, often failing to delineate between job projections for full-time roles tied to data center operations and short-term roles in construction and the service sector.²⁰ They also ignore the economic impact that higher energy costs and diverted service can have on local businesses.²¹ Meanwhile, state and local communities stand to lose billions of dollars in tax revenue as a result of the packages of tax incentives that companies demand before they will agree to build.²² This again provides a unique opportunity for broad-based coalitional organizing, as both free-market groups and progressive economic organizations have criticized the disconnect between large corporate tax breaks and the lack of tangible benefits for communities.²³ Some local organizations have already taken on the mantle: Citizens Action Coalition in Indiana, for example, has called for a moratorium on new hyperscaler data center construction due to the enormous resource constraints and cost burdens that large data centers pose to local communities,²⁴ and Memphis Community Against Pollution in Tennessee has pushed back against the public health and environmental impacts of Elon Musk’s xAI data center buildout near historically Black neighborhoods already facing elevated cancer and asthma risks.²⁵

3. Algorithmic Prices and Wages



Across the economy, from grocery stores to online marketplaces to insurance brokers, the ability of firms to use AI to set prices for customers and calibrate wages for workers using detailed, often intimate, information extracted from us, is deepening inequality. This makes algorithmic prices and wages especially fertile ground for mobilizing around AI's economic impacts on everyday people.²⁶

There is growing coordination between labor organizers and anti-monopoly and economic justice groups working to contend with this issue. In February 2025, for example, along with a coalition of organizations and experts, AI Now published a report canvassing the breadth of the use of algorithmic pricing and wage-setting practices and spotlighting the ways in which bright-line rules could prevent harm,²⁷ including an outright ban on individualized surveillance prices and wages, and elimination of loopholes that corporations could exploit to continue that practice.²⁸ There has already been considerable movement in the states this session to prohibit both algorithmic pricing and wage-setting, including the introduction of bills in California, Colorado, Georgia, and Illinois.²⁹



ADVANCE WORKER ORGANIZING TO PROTECT THE PUBLIC AND OUR INSTITUTIONS FROM AI-ENABLED CAPTURE

(Guest contribution by
Andrea Dehlendorf, Senior Fellow)



Labor campaigns over the past several years have demonstrated that when workers and their unions pay serious attention to how AI is changing the nature of work and engage in serious intervention through collective bargaining, contract enforcement, campaigns, and policy advocacy, they can shape how their employers develop and deploy the technologies. The deeper opportunity for labor, however, lies in the fact that across all sectors of the economy and society, it is people experiencing these technologies at work, firsthand, who best understand their impacts and limitations, and also the conditions under which their potential can be realized. Thus, labor's power can shape not just whether and how AI is used in the workplace, but can also recalibrate the technology sector's power overall, ultimately shaping the trajectory of AI toward the public interest and common good.

Over the past decade, labor campaigns to challenge the use of generative AI in Hollywood, algorithmic management of warehouse workers, and platforms

for rideshare drivers have been critical to building public awareness of the impact of AI and data technologies in workplaces. In 2023, the Writers Guild of America and Screen Actors Guild strikes challenged studios aiming to deploy generative AI to replace their creative work without fair compensation.³⁰ These public battles about profit-driven automation risk in creative industries have won contract language that provides some protection for artists and writers, transformed the public conversation around the implications of technology, and shifted popular understanding of the stakes.³¹

Unions, worker centers, and alt-labor groups have also exposed and challenged Amazon's algorithmic management practices and extensive surveillance. The Warehouse Worker Resource Center developed and passed a bill in California that limits the use of technology-enabled productivity quotas—legislation that workers's organizations are now advancing in other states.³² Uber and Lyft drivers have also organized across a range of formations (Gig Workers Rising,³³ Los Deliversistas Unidos,³⁴ Rideshare Drivers United,³⁵ SEIU, and others) to challenge platform precarity, winning the right to organize and establish labor protections. In Boston, drivers won the right to unionize via a state ballot measure.³⁶

As powerful technology corporations take an unprecedented and self-serving role in the federal government, undermining democratic structures, it is clear that the labor movement has an opportunity to lead on behalf of all of society. *Building on the momentum of previous wins, there are five strategic paths for the labor movement to flex its power and increase its influence over the trajectory of AI across our society and economy:*

Sectoral Analysis and Action for the Common Good



By developing a comprehensive analysis of how AI and digital technologies are advancing in every economic sector, including the impact on both working conditions and the public interest, workers and their unions can develop and implement comprehensive bargaining, organizing, and policy campaign strategies. These strategies will have maximum impact when they address both the impact on the conditions of work and the impact on the public.

A key example is National Nurses United (NNU), which has pushed back against profit-driven employers seeking to replace skilled professionals with subpar technology that harms patients. In 2024, NNU conducted a survey of its members,³⁷ which found that AI undermines nurses' clinical judgment and threatens patient safety. The group subsequently released a "Nurses' and Patients' Bill of Rights," a set of guiding principles designed to ensure just and safe application of AI in healthcare settings.³⁸ NNU affiliates have successfully bargained over AI implementation in contract negotiations; in Los Angeles, nurses secured commitments from University of Southern California hospitals that they won't replace nurses with AI and that they will center nurses' judgment when determining patient care, as well as safe and appropriate staffing levels.³⁹ At Kaiser Permanente, nurses stopped the rollout of EPIC Acuity, a system that underestimated how sick patients were and how many nurses were needed, and forced Kaiser to establish an oversight committee for the system implementation.⁴⁰ Beyond the bargaining table, nurses have mobilized their membership out into the streets and to Washington, DC, to protest the hospital industry's attempts to push

scientifically unproven AI instead of investing in safe staffing.

Similar strategies can be designed across most sectors of the economy, providing a critical check on the deployment of AI for cost cutting and profit maximization.

Challenging AI Austerity in the Public Sector



As Elon Musk and others advance AI-driven austerity in the federal government through DOGE—and Republicans and Democrats mount copycat initiatives across the states—workers and their unions can fight back, defending both workers whose jobs are deeply impacted and those receiving public benefits and services. The Federal Unionist Network (FUN), an independent network of local unions representing federal employees, is directly challenging DOGE and its driver, Elon Musk. FUN is running a Save Our Services campaign that centers the impact of the service cuts that Musk and the Trump administration are implementing.⁴¹ The network is building scalable infrastructure for education, support, and action.

At the state level, Pennsylvania provides a model of regulating AI in the provision of public services, recognizing and protecting the key role state workers play in service and benefit provisions. Governor Josh Shapiro and SEIU local 668 entered into a partnership establishing robust co-governance mechanism over AI, including the formation of a worker board that will oversee implementation of generative AI tools; the definition of a public worker as a person and gen AI as a tool, strong human-in-the-loop protections at every

step, and monitoring of disparate impacts related to bias and discrimination of protected classes.⁴²

Comprehensive Campaigns Challenging Big Tech Corporations



Transforming the unprecedented power of the technology sector will require comprehensive campaigns that directly challenge the monopolistic business models and actions of the corporations that dominate and drive the sector, and, increasingly, our government. To meet the challenge, these campaigns must include but not limit themselves to, organizing workers and addressing workplace impacts.

The Athena Coalition, for example, focuses on Amazon's unaccountable power and impacts across multiple spheres, pushing for regulatory action to break up the business itself.⁴³ Athena members connect how Amazon develops and uses AI and data surveillance technologies in workplaces to what they sell to police departments, ICE,⁴⁴ and authoritarian military states;⁴⁵ they bring public awareness to the impacts of pollution on neighborhoods near Amazon logistic hubs,⁴⁶ pressure the company to reduce the climate impact of their business model, and bring together organizers in states to challenge the energy use and environmental impact of their data centers; and they expose how algorithmic systems impact small businesses on the platform, consumers, and workers. Together, coalition members are bringing these strands together in order to reclaim our wealth and power from corporations like Amazon.

Building State and Local Power to Advance the Public Interest



Workers and their unions can and must collaborate with racial justice, economic justice, and privacy groups to mount a serious, coordinated defense against the tech sector's aggressive state policy agenda to write the rules of AI regulation and digital labor tech regulation in their interest. Doing so will require deeper engagement with members of unions and community-based organizations to understand and prioritize these interventions. Beyond defense, there is an opportunity to develop a proactive agenda at the state and local level that draws bright lines around the use of surveillance, algorithmic management, and automation technologies in criminal justice, workforce management, and public service delivery; and also addresses the unsustainable climate impacts of data centers and the extreme market concentration of the industry. For example, in California, there is a strong set of labor and public-interest focused policy proposals that can be used as models in other states. Also, engagement is moving with members of unions and aligned privacy, tech justice, racial justice, and economic justice organizations and social justice leaders to understand how their core issues are impacted by AI; develop strategies to organize and publicly campaign on the issues; and identify where they can best engage on state and local policy. The UC Berkeley Labor Center's Tech and Work Policy Guide is a key resource that can be used across states⁴⁷ and Local Progress's recent report written in collaboration with AI Now provides an actionable roadmap for local leaders to protect communities from the harms of AI.⁴⁸

Organizing Across the AI Value and Supply Chain



A global movement to organize across the AI value and supply chain—from US-based white-collar tech workers to data labelers in the Global South, would provide a powerful check on the technology sector. Such a movement could reorient the trajectory of AI toward the public interest and counter the sector's authoritarian turn. However, union density in the technology sector, globally and in the US, is negligible. There have been no serious attempts to organize the sector at scale. Existing efforts are often under-resourced, while others lack a global orientation and solidarity, limit their focus to economic issues, or are not challenging concentrated market power.

However, several promising initiatives are gaining traction, including Amazon Employees for Climate Justice,⁴⁹ which is organizing white-collar workers to challenge Amazon on its climate impact; the African Content Moderators Union (ACMU), other locally led associations,⁵⁰ and UNI Global,⁵¹ which are organizing data labelers in Africa; Solidarity Center⁵² and ITUC,⁵³ which support global gig-worker organizing; the Data Worker's Inquiry Project,⁵⁴ a community-based research project focusing on data labelers globally; and the Tech Equity Collaborative⁵⁵ and CWA's Alphabet Workers Union,⁵⁶ which are campaigning on temp- and contract-worker issues.

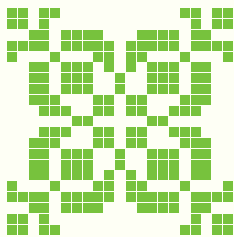
We desperately need more ambitious and better-resourced projects. There is a particular opportunity for people who are building and training the AI systems—and, therefore, know them intimately—to use their powerful positions to hold tech firms accountable for how these systems are used. Many people

working in these white-collar jobs have conveyed that their values conflict with corporate decision-making.⁵⁷ Organizing and taking collective action from these positions will have a deeply consequential impact on the trajectory of AI.

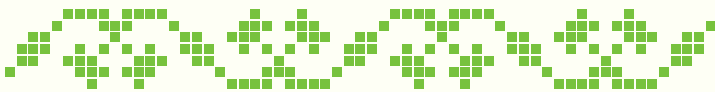
The Path Forward



During previous periods of extreme inequality, such as the Gilded Age and the Roaring Twenties, muckraking journalism, labor organizing, and strikes held unaccountable corporate power accountable, and ushered in the Progressive and New Deal eras. In our era, collective labor actions will also be critical to challenging not only Big Tech’s unchecked wealth and power over our democracy and society, but also the use of technologies to undermine job quality while reducing workers’ share of the profits they generate. Like the labor movement of the last century, today’s labor movement can fight for and win a new social compact that puts AI and digital technologies in service of the public interest and hold today’s unaccountable power accountable.



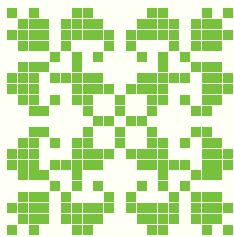
ENACT A “ZERO-TRUST” POLICY AGENDA FOR AI



In 2023, alongside Electronic Privacy Information Center (EPIC) and Accountable Tech,⁵⁸ we recommended three pillars to guide policy action in the AI market:

- ❖ Time is of the essence—start by vigorously enforcing existing laws.
- ❖ Bold, easily administrable, bright-line rules are necessary.
- ❖ At each phase of the AI system life cycle, the burden should be on companies—including end users—to prove that their systems work as claimed, and are not harmful.

“Given the monumental stakes, blind trust in [tech firms’] benevolence is not an option,” we argued then. That warning rings even more starkly now.



1. Build a drumbeat of bright-line rules that restrict the most harmful AI uses wholesale, or limit firms' ability to collect and use data in specific ways.



Bold, bright-line rules that restrict harmful AI applications or types of data collection for AI send a clear message that the public determines whether, in what contexts, and how AI systems will be used. Compared to frameworks that rely on process-based safeguards (such as AI auditing or risk-assessment regimes) that have often, in practice, tended to further empower industry leaders and rely on robust regulatory capacity for effective enforcement,⁵⁹ these bright-line rules have the benefit of being easily administrable and targeting the kind of harms that can't be prevented or remedied through safeguards alone. There's a growing list of ripe targets for these kinds of clear prohibitions. For example:⁶⁰

- ❖ AI cannot be used for emotion-detection systems
- ❖ AI cannot be used for “social scoring,” i.e., scoring or ranking people based on their social behavior or predicted characteristics
- ❖ Surveillance data cannot be used to set prices or wages
- ❖ AI cannot be used to deny health insurance claims

- ❖ Surveillance data about workers cannot be sold to third-party vendors
- ❖ AI cannot be used to replace public school teachers
- ❖ AI cannot be used to generate sexually explicit deepfake imagery or election-related deepfake imagery

2. Regulate throughout the AI life cycle. This means no broad-based exemptions for foundation-model providers, application developers, or end users.



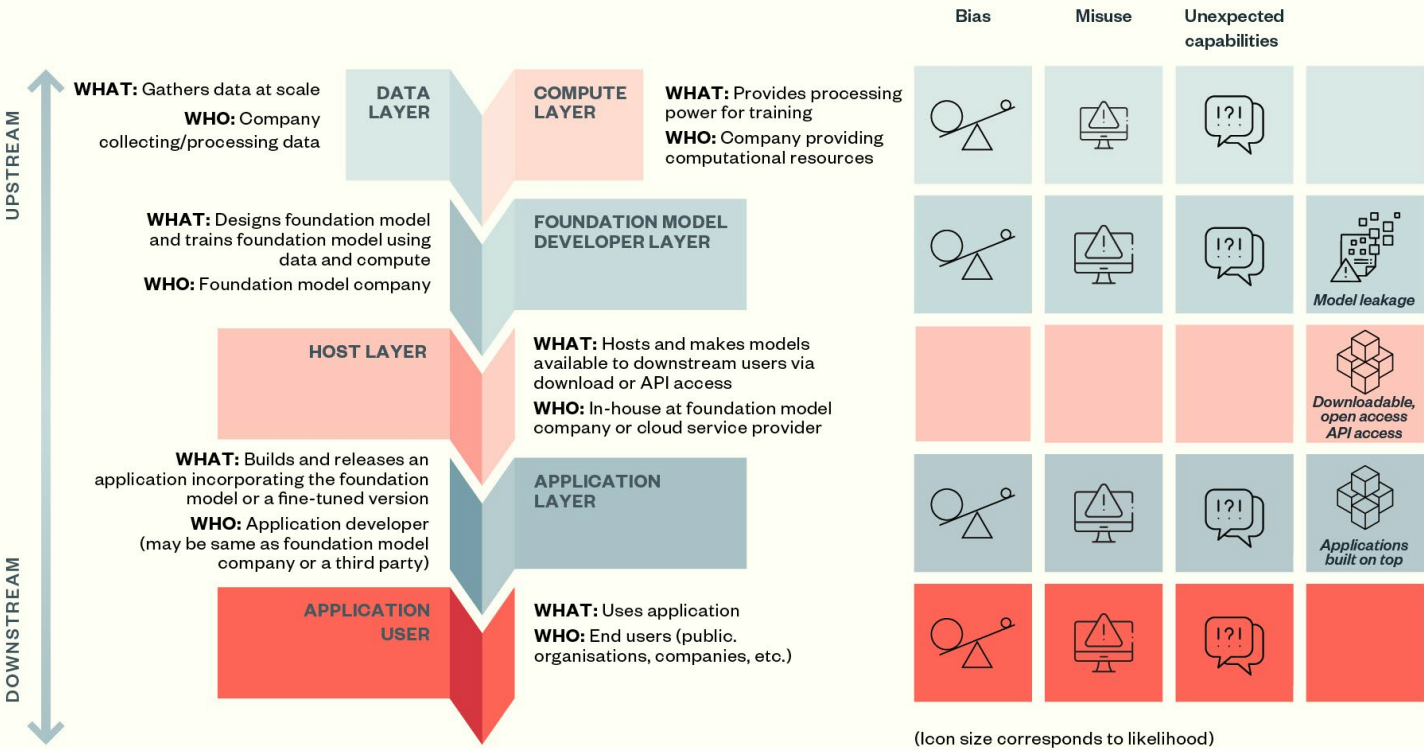
AI must be regulated throughout the entire life cycle of development, from how data is collected through the training process, to fine tuning and application development, to deployment. Appropriate actors should be held responsible for decision-making at each of these stages, and be mandated to provide transparency and disclosures that enable actors further downstream to uphold their own obligations.⁶¹

Transparency regimes are the foundation of effective regulation of technology—necessary, but never enough⁶²—and yet in the past few years we've seen fierce and largely successful resistance to baseline disclosure requirements from the AI industry. That's true at the general-purpose stage where industry has used copyright and commercial secrecy as a shield to avoid disclosing training data;⁶³ but it's equally true at the final stage, where these systems are deployed.

Foundation model supply chain

Risk origination examples

Risk proliferation examples



Note: This is one possible model (there will not always be a separate or single company at each layer)

Fig 1: Risk Origination and Risk Proliferation in the Foundation Models Supply Chain. Source: Ada Lovelace Institute, 2023

Proactive transparency and disclosure requirements (where the consumer doesn’t need to individually request this information⁶⁴) have prompted intense backlash from industry, most recently by asserting that disclosure provisions in Colorado’s AI law, passed in 2024, placed an “undue burden” on businesses.⁶⁵ What this means is that in many contexts where AI is being used, people whom an AI system affects will never be made aware of its role in making determinations that can have significant consequences for their lives—let alone give them the ability to contest the system.

Industry has also exerted a concerted lobbying effort to put all of the onus on end developers, while exempting the developers of so-called “general purpose” systems.⁶⁶ In 2021, alongside more than fifty global AI experts, AI Now published a letter urging the European Commission to reject this approach in the context of the EU AI Act.⁶⁷ It’s important that AI be evaluated at every stage of the life cycle. Certain aspects of AI models will be best evaluated at the application stage: This is where pre-deployment testing and validation will give the most insight into the efficacy and safety of particular uses of AI.⁶⁸ But the

developers of foundation models and those involved in constructing datasets should not be exempted from scrutiny—particularly since they are generally the largest and best-resourced actors in the market,⁶⁹ and have demonstrated a growing propensity for opacity, offering diminishingly few details about how they train their systems.⁷⁰ Introducing stronger accountability requirements at the foundation-model layer would motivate restructuring of the development practices, and potentially the operating model, of AI developers. In and of itself, these modes would create greater internal transparency and accountability within AI firms that would convey societal benefits, and aid the work of enforcement agencies when they need to investigate AI companies. Regulatory structures that place the onus solely on the application stage ignore that these developers, as well as purchasers and deployers, may lack the information necessary to assure compliance with existing law, absent upstream intervention.⁷¹ And the end users of AI systems are in many instances not consumers; they're businesses or public agencies, themselves in need of scrutiny for the way their use of the system affects consumers, workers, and members of the public.

- ❖ Require all foundation-model developers to publish detailed information on their websites about the data used to train their systems.
- ❖ At each stage of the supply chain, require developers to document and make available their risk-mitigation techniques, as well as mandate disclosure of any areas of foreseen risk they are not positioned to mitigate, so that this is transparent to other actors along the supply chain.
- ❖ Institute frameworks of AI use that protect those on whom it is used and introduce healthy friction and democratic decision-making into AI deployment through measures such as a “right to override” decisions, and requirements that

councils composed of those affected by the system have a say in whether a system is used at all.

- ❖ In safety-critical contexts (e.g., energy-grid allocation, military use, nuclear facilities), prohibit use of foundation models until they have been proven to pass traditional tests for systems safety engineering, and require systems to be built safe by design.
- ❖ Require pre-deployment testing and validation of AI applications that examines the safety and efficacy of particular uses.

3. The AI industry shouldn't be grading its own homework: Oversight of AI model development must be independent.



Industry tends to rely too heavily on “red-teaming,” “model cards,” and similar approaches to system evaluation, where companies are the primary actors responsible for determining whether their systems work as intended and for mitigating harm.⁷² This approach has many problems; it ignores, for example, the significant conflicts of interest involved and fails to employ robust methodologies.⁷³

Companies have asserted that they hold the unique technical expertise necessary to evaluate their own systems⁷⁴ in order to execute what is essentially a fully self-regulatory approach to governance, making a set of “voluntary commitments” to the Biden administration in 2023.⁷⁵ When asked a year later about their progress on those commitments, the response was

milquetoast: While companies reported taking actions like engaging in internal testing to mitigate cybersecurity risks, they were unable to validate whether any of those actions actually reduced underlying risk; and by focusing on domains like bioweapons, they ignored the wider universe of harms like consumer and worker exploitation, privacy, and copyright.⁷⁶ And OpenAI stepped back from even fulfilling the basic obligations for these commitments in the release of GPT-4.1, opting not to release a model card on the grounds that the model “isn’t at the frontier.”⁷⁷

This illustrates a point that by now should be abundantly clear: Companies cannot be responsible for their own self-evaluation, and are poorly incentivized to uphold any promises they make absent the foundation of enforceable law:

- ❖ **Provide enforcement agencies with the resources and in-house staffing necessary to conduct oversight throughout the AI life cycle, and to hold firms accountable for any violations of law.** Any law enforcement agency that enforces the law against companies that use AI should have the ability to detect, investigate, and remediate violations of law using those technologies. As of January 2025, only the Federal Trade Commission and the Consumer Financial Protection Bureau are known to have this capability, and their teams have been significantly reduced.⁷⁸
- ❖ **Require AI companies to submit to independent third-party oversight and testing throughout the AI life cycle, and to make available the necessary information for auditing bodies to conduct evaluation of their models or products.**⁷⁹

4. Create remedies that strike at the root of AI power: data.



We don’t need to limit ourselves to the standard AI policy tool kit: A wide range of creative remedies could be deployed by enforcement agencies to strike at the root of AI firms’ power. These include measures such as:

- ❖ **Require firms to delete ill-gotten data that they have used to train AI models and any related work product generated with that data.** Algorithmic deletion remedies require that firms that have benefited from ill-gotten data to train AI models must not only delete the data, but must delete the models trained on that data.⁷⁸ This remedy was developed in the first Trump administration and has now been used in a number of cases including the FTC’s *Everalbum* and *Kurbo*.
- ❖ **Curb secondary use of data collected in one purpose and used in another.** AI firms frequently retain data collected under one justification indefinitely for the purpose of training subsequent models. Curbing such secondary uses would both protect the underlying consent regime—ensuring that when users consent to the collection of their data for a particular use it is not treated as an expansive mandate for other uses—and would limit the continual proliferation of data available for model training. The FTC’s Amazon Alexa case offers a useful example, in which the court ruled that Amazon violated COPPA by retaining children’s voiceprints for the purpose of improving its AI models.⁸¹

❖ **Limit hyperscaler partnerships with AI startups.**

As we explore in [Chapter 2](#), one decisive way Big Tech companies are able to achieve advantages in the AI market is by gaining control of competitive AI companies through partnership terms, investments, or acquisitions. Banning hyperscalers from investing in or partnering with competitive AI companies is one bold, structural measure that would help dismantle hyperscalers' stronghold on the AI stack. This remedy was initially proposed by the Department of Justice in its case against Google's search monopoly.⁸² While it was later revised, the remedy's initial inclusion recognizes the kind of bold action required to shape a truly competitive AI market.

- ❖ **Ban repeat offenders from selling AI to the government.** Protections are in place to ensure that those who are selling services and products to the federal government are not engaging in fraud, waste, abuse or antitrust violations. In order to introduce incentives for firms to follow the law, AI companies and salespeople who are eligible should be “debarred” from being able to win federal contracts.⁸³

5. We need a sharp-edged competition tool kit to de-rig the AI market, and how AI shapes market power across sectors.



Amid concerns of Big Tech dominance in AI growing more prominent in a frenzied market, competition enforcers globally have been far from silent. Over the

past few years, we have seen multiple early inquiries into so-called partnership agreements between AI startups and Big Tech hyperscalers;⁸⁴ several market studies into concentration in the cloud market, from Japan to the Netherlands to the UK to the US;⁸⁵ and detailed reports highlighting potential choke points and competition risks across the AI supply chain.⁸⁶ The DOJ, FTC, CMA, and EC even issued a joint statement on competitive concerns with AI, warning that partnerships and investments could be used to undermine competition.⁸⁷ This is already a promising departure from the past decade of a passive stance toward the consumer tech industry, which has left regulators struggling to catch up to remedy harm—from Facebook's acquisitions of WhatsApp and Instagram⁸⁸ to Google's sprawling search monopoly.⁸⁹

Yet, despite many of the “right noises,” much of this activity failed to materialize into concrete enforcement action and legislative change, or to draw bright lines prohibiting specific anticompetitive business practices. A familiar dynamic has consisted of limited and restrictive interpretations of existing legal standards for enforcement action. For example, traditional merger review conditions are unable to capture more indirect means of dependence and control between AI startups and cloud hyperscalers. Or consider the threshold in the US's Hart-Scott-Rodino Act, which meant that smaller acquisitions by Big Tech firms were not subject to scrutiny, like Rockset (acquired by OpenAI), Deci AI (acquired by NVIDIA), OctoAI (also acquired by NVIDIA). Across these issues persists a more fundamental limitation: the artificial blinkers that prevent regulators from fully grappling with the many dimensions of power and control that Big Tech firms exercise through their sprawling ecosystems. A recent report from the Competition and Markets Authority (CMA) in the UK published a diagram (see Fig 2) depicting the presence of Amazon, Apple, Google, Meta, and Microsoft across the AI supply chain—il-

illuminating the ways in which their presence across these markets allows for network effects, data feedback loops, and unparalleled access to monetizable markets for their AI products⁹⁰ (We canvass these in Chapter 2: Heads I Win, Tails You Lose).

Yet the irony is that the CMA, among other counterparts, has fallen short of acting in ways that might interrupt this ecosystem control. The scope for intervention is limited as much by restrictive legal thresholds as it is by a corporate lobbying environment around AI that pits any market intervention not only as anti-innovation (see Chapter 1.4: Recasting Regu-

lation), but as fundamentally harmful to the national interest given the stakes of the so-called AI arms race (see Chapter 1.3: AI Arms Race 2.0). Staying with the UK example, Microsoft announced a massive investment in compute on the heels of two significant enforcement actions affecting Microsoft: the CMA’s decision not to pursue a block of the Microsoft-Activision merger, and the opening of an investigation by the CMA into competition concerns in the cloud market. AI firms use investment as a clear play to hedge against the very real risk that the CMA will take action in ways that would threaten the companies’ bottom line.

Big Tech Dominates AI Value Chain

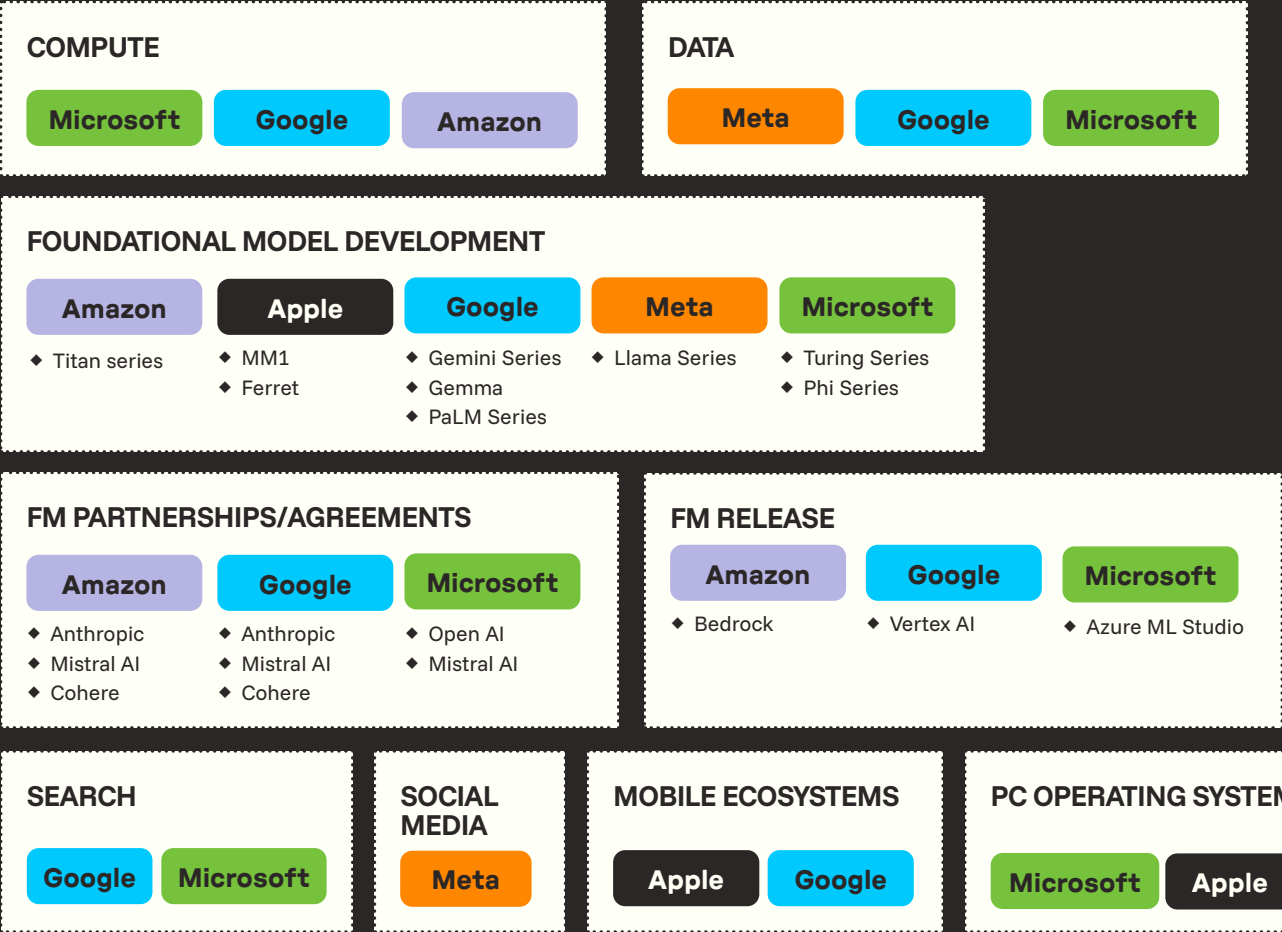


Fig 2: How Amazon, Apple, Google, Meta, and Microsoft permeate the AI supply chain. Source: CMA.

Looking ahead, we see an opportunity to corral around a set of targeted interventions that could tackle Big Tech's control over the broader tech ecosystem, and could go a long way in shifting the "rigged" dynamics of the AI market (see [Chapter 2: Heads Win, Tails You Lose](#)). We spotlight some key policy directions here that identify high-value intervention points.⁹¹

❖ **Structural separation to prevent vertical integration across the AI supply chain, especially targeting infrastructural and model layer advantage, including the following:**

- ◆ Prevent cloud companies from participating in the market for AI foundation models.
- ◆ Require that Big Tech firms divest their cloud computing businesses from the rest of their corporate structure, and that they operate as wholly independent entities.
- ◆ Prevent AI foundation model companies from participating in other markets related to AI.
- ◆ Prevent chip designers from investing in AI development companies.

❖ **Scrutiny of mergers, acquisitions, and partnerships involving Big Tech hyperscalers:**

- ◆ Restrict, or perform an ex ante review of, Big Tech hyperscalers that enter into partnerships with AI startups, including through unconventional mechanisms (for example, credit arrangements, transactions structured under the HSR threshold, and acqui-hires).

- ◆ Prevent cross-board representation across Big Tech firms and AI startups to avoid conflicts of interest.
- ◆ Apply the new FTC merger guidelines to the AI sector, with particular attention to acquisitions involving datasets that may form a conflict of interest.

❖ ***Ex ante* bright-line rules that prevent Big tech hyperscalers from self preferencing or discriminatory treatment**

- ◆ Cloud providers can't self-preference, or offer favorable terms or rates to similarly situated customers.
- ◆ Cloud providers can't take advantage of data they obtain from customers by hosting infrastructure to develop products that compete with their customers.
- ◆ Cloud providers can't penalize customers across their ecosystems for developing competitive products and services.
- ◆ Foundation-model companies can't self-preference, or offer favorable terms or rates to similarly situated customers.
- ◆ Foundation-model companies must be interoperable across cloud infrastructure providers.

❖ **Private parties collaborating to sue AI firms violating antitrust law**

- ◆ There are three legs on the stool of antitrust enforcement: federal agencies, state agencies, and then affected workers, consumers, and companies. This latter category offers meaningful inroads for pushing forward antitrust enforcement.

6. Rigorous oversight for new data center development and operation



As AI spurs the growth of data center development across the world and projects staggering demands for power, policymakers must reckon with the incompatibility between this unrestrained growth and our already resource-constrained environments. Communities near data centers face the realities of air pollution, lack of adequate water, and noise pollution⁹² with a disproportionate impact on Black and Brown frontline communities.⁹³ Yet data centers don't merely affect those who live closest to them: To fund the massive investment required to build out the infrastructure necessary to fuel AI data centers (such as power plants and transmission lines), energy utility companies may pass off the costs onto ordinary ratepayers.⁹⁴ Already Americans are seeing their electricity rates increase;⁹⁵ this is only expected to intensify.

Meanwhile, states and localities are handing out billions of dollars in corporate tax subsidies to data center developers by exempting them from paying sales, use, and property taxes, or by discounting their rates.⁹⁶ This comes directly at the expense of communities, who lose out on tax revenue for local schools, infrastructure, and other community needs.

Despite billions of dollars of investment and government subsidies, much of the current data center development process operates in obscurity.⁹⁷ Many data centers fail to report even basic information about their resource consumption, including power usage and water consumption. This prevents states from planning for adequate supply and demand, threatening the stability of the grid. Moreover, many

deals between data centers and utility companies are approved by regulators in proceedings that tend to go uncontested by the public, preventing ratepayers from understanding the impacts these large customers may have on their rates.⁹⁸

Yet this current path of unrestrained data center growth is not inevitable. More communities are realizing that successful economic development cannot hinge on environmental, health, and economic harms: Recently, two counties in Atlanta have temporarily paused all new data center projects to adequately ensure that development protects community interests;⁹⁹ and a county board of supervisors in Virginia rejected a plan for a gas power plant and data center campus because of concerns about the increased rates of pollution and negative environmental impact it would cause in the surrounding area.¹⁰⁰ The developer is now considering building housing instead.

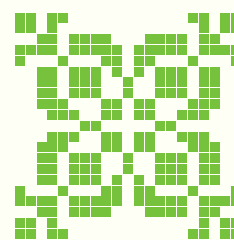
Policymakers must take swift and decisive action to ensure that data center development supports, rather than harms, the needs of local communities and the resources available on our planet, such as the following:

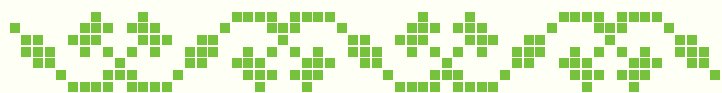
- ❖ **Stop or slow new data center development until robust protective conditions are met, including transparency mechanisms, environmental standards, consumer protections, and zoning reform.**
 - ◆ Stop new large data center development until clear guardrails are in place to address risks to ratepayers and communities more broadly.¹⁰¹
 - ◆ Reject proposals for new data center development that do not meet strict standards for environmental protection, resource use, sustainable jobs, and health impacts.¹⁰²

- ◆ Require new data center energy demands be met with new and additional renewable energy in the same deliverable grid,¹⁰³ rejecting proposals to expand fossil fuel and nuclear energy use.
- ❖ **Repeal or cap corporate tax subsidies to ensure tax dollars directly benefit local communities.**
 - ◆ Repeal state tax exemptions for data centers and end the open-ended extension of tax breaks to future investments.¹⁰⁴
 - ◆ Alternatively, cap state and local subsidies to data centers per permanent job or year.¹⁰⁵
 - ◆ Require data center operators to meet renewable-energy purchasing agreements and invest in energy efficiency measures to be eligible for sales and use tax exemptions.¹⁰⁶
- ❖ **Advance commonsense legislation that ensures fairness, transparency, and oversight for data center operations.**
 - ◆ Require data centers to report quarterly on water and energy usage to a state executive agency, with access available to the public.¹⁰⁷
 - ◆ Establish transparency mechanisms around data center contracts and construction, such as through public notices and hearings.
 - ◆ Require site assessments before localities approve new data centers to manage effects on noise levels, homes, schools, water, farmland, parks, and forests.¹⁰⁸
 - ◆ Require localities to disclose the projected dollar value of all subsidies offered, and

report on actual dollar amounts over time to the public.

- ◆ Institute energy caps for data centers in areas where projected power demand from data centers creates a disproportionate burden on the grid.
- ❖ **Build in regulations to protect ordinary ratepayers from the costs associated with data center development.**
 - ◆ Investigate whether non-data center customers are subsidizing data center energy costs; and, if so, establish rules to prohibit the practice and address the imbalance.
 - ◆ Direct public utility commissions to create a new rate classification for data centers and other large-load customers.¹⁰⁹
 - ◆ Require public utility commissions to create a community discount plan to account for increased energy costs.¹¹⁰
 - ◆ Require large-load customers (like data centers) to commit to long-term contracts and minimum contract payments to ensure that they appropriately cover utility investment costs.¹¹¹

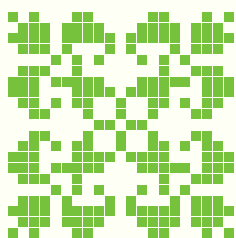




BRIDGE NETWORKS OF EXPERTISE, POLICY DOMAINS, AND NARRATIVE



Blinkered views that fail to see the different components of the AI supply chain materially and are often single-issue focused have too often undermined AI advocacy and policy. It is easy to miss the ways in which big-picture narratives (see [Chapters 1.1-1.4 on AI's False Gods](#)) manifest on the ground across policy fights. Silos can blunt strategic creativity in ways that foreclose alliance or cross-pollination. In our 2023 report, we attempted to find ways to identify and bridge some of these siloes: like forming a united strategy around privacy and competition policy; or understanding industrial policy as central to (not distinct from) discourse on AI regulation.¹¹² Here we highlight additional emergent areas of cross-pollination in a bid to bolster and strengthen our strategies:



National Security as a Vector Both for and Against AI Accountability



National security arguments are a key vector driving the AI infrastructure acceleration agenda. Equally, national security arguments support moves for stronger privacy, security, and accuracy thresholds, as well as moves that combat concentrated infrastructure markets as threatening resilience.

As we detail in [Chapter 1.3 on AI & Geopolitics](#), recent national security policy proposals (under President Biden and now President Trump) and the specter of Chinese AI competition have consistently cleared the pathway for the industry to expand through infrastructure build-out.¹¹³ Most recently, baseline transparency mechanisms around data center build-out in Virginia's HB 2035 bill, championed by local advocacy groups, were shut down citing national security concerns, including alleged pushback from Homeland Security.¹¹⁵

On the other hand, national security imperatives (resilience, reliability, security) can be a strategic vector to demand more accountability from the AI industry, and an opportunity to bring security experts into broader debates about AI accountability. Former military officials and security experts alike have sounded the alarm on inherent vulnerabilities in using commercial language models in military contexts¹¹⁶—and while these intractable concerns (flaws in large-scale AI are features, not bugs, as we note in [Chapter 3](#)) are unconscionable in high-stakes combat environments, they are equally worrisome for many other swaths of applications in the economy and society.

Anti-monopoly groups are also increasingly foregrounding structural vulnerabilities created by AI infrastructure that is concentrated among very few companies, and across different points of the AI supply chain.¹¹⁷ The creation of single points of failure—recently, an erroneous update to Microsoft’s operating system from the cybersecurity firm CrowdStrike—triggered the largest IT outage in history, affecting global businesses, hospitals, airports, courts, and governments.¹¹⁸ CrowdStrike is only the most conspicuous example of the fragility of a super-concentrated infrastructure—the AI stack is rife with choke points from chips to undersea cables—and we’re seeing the risk amplified with errors in AI foundation models that could undergird applications across sectors of the economy. Former SEC Chair Gary Gensler has noted how concentration in the market for AI models could produce conditions that might lead us to the next financial crisis.¹¹⁹

These vulnerabilities are more acute when it comes to the public sector, especially as a growing proportion of government runs on Big Tech cloud infrastructure.¹²⁰ This will inevitably increase manifold with the rushed push toward expanding adoption of commercially available AI tools¹²¹ (via DOGE’s efforts, among other preexisting initiatives such as the Department of Defense’s Task Force Lima).¹²² Hyperscalers have done their part to market directly to government services, touting their security credentials,¹²³ but have underplayed the long-term implications of introducing single points of failure into national security infrastructure. In July 2023, Microsoft revealed a breach by Chinese state-backed hackers who accessed emails of top US officials, following a previous security lapse with Russian-backed SolarWinds attackers accessing government data.¹²⁴ Addressing these serious security concerns would have beneficial implications for the use of AI across the wider economy.

Privacy as an Economic Justice Issue: Disrupting the Surveillance-to-Automation Pipeline



Whether it’s Hollywood actors resisting invasive body scans as a means to create digital replicas of themselves that might eventually replace them,¹²⁵ or the uptick in the market for wearable surveillance tools in the workplace that will then be used to train “humanoid robots,”¹²⁶ the connections between surveillance and the eventual displacement or devaluation of workers is becoming increasingly explicit. The US and UK governments, meanwhile, are making claims that they can train AI systems on government workflows that will automate large parts of government work.¹²⁷

This surveillance-to-automation pipeline situates traditional privacy tools (like consent, opt-outs, impermissible purposes, and data minimization) as central levers in enabling more just economic conditions. Scholar Ifeoma Ajunwa argues that worker data should be seen as “captured capital” for firms, where workers risk having their data used to train technologies that will eventually displace them (or create the conditions to reduce their wages), or having their data sold to the highest bidder via a growing network of data brokers, without control or compensation.¹²⁸ Meanwhile, as the AI industry scrambles to find a conclusive business model, all trends point toward a proliferation of enterprise use cases that all appear to rely on the capture of workflow data (through the ratcheting up of workplace surveillance), which can then be used to automate more workflows.¹²⁹ We’re also seeing extractive practices proliferate in the hyper-surveilled context of prisons, which are in-

creasingly targeted by AI startups eyeing incarcerated workers as a cheap source of labor to train their AI algorithms. These “clickworkers” or “human labelers” are paid pennies on the dollar—if anything at all—by startups like Vainu, which ironically boasts that it is “creating work because of AI [to] employ and empower the new working class.”

This necessitates reframing worker privacy as core to the economic justice agenda around AI. Specifically, it requires identifying the point of collection (i.e., surveillance) as the appropriate target for push-back given the ways in which this will be weaponized against worker interests downstream. In regulatory terms, this could mean prioritizing data-minimization rules that restrain the collection of certain kinds of worker data, or restrict the use of worker data collected in one context for training AI. Thinking around data privacy will also need to expand to meet the moment—revising narrow notions of what counts as personal information, and prioritizing structural and absolute constraints on workplace surveillance that don’t rely on consent-based mechanisms.

State advocacy groups are already prioritizing these issues,¹³⁰ including in the following proposals:

- ❖ **Employers must notify workers when a surveillance system is introduced into the workplace, and employers are prohibited from electronically monitoring employees outside of essential job functions (CA,¹³¹ CT,¹³² MA,¹³³ VT,¹³⁴ WA¹³⁵).**
- ❖ **Employers are prohibited from sharing worker data with third parties (CA¹³⁶).**
- ❖ **Employers are prohibited from using facial- or emotion-recognition surveillance technology on workers (CA).¹³⁷**

Likewise, existing law applies to AI use cases:

- ❖ **Employers are subject to limitations from the Fair Credit Reporting Act when AI is used on workers and applicants.¹³⁸**

Privacy and Competition in the Age of Generative AI: Data Minimization to Interrupt Data Moats



In our 2023 report, we warned that privacy, competition expertise, and tools were often siloed from one another, leading to interventions that easily compromise the objectives of one issue over the other.¹³⁹

Firms continue to take advantage of this to amass information asymmetries that contribute to a further concentration of their power. We spotlighted the need for heightened scrutiny on “data mergers”¹⁴⁰—accounting for how firms leverage commercial surveillance tools and strategies to amass power—as well as adtech rules restricting first-party data collection for advertising purposes as having twin privacy and competition benefits.

This cross-pollination is even more important now: Data is a key bottleneck in the development of generative AI, and a strategic point of intervention for both protecting privacy and limiting Big Tech from entrenching its data advantages in the AI market. The idea that freely available, publicly accessible data is “everywhere” and therefore not a scarce resource is intuitively appealing but misses the point. Datasets with high levels of curation and human feedback, and those that come with assurances of accuracy, legiti-

macy, and diversity at scale, are an increasingly scarce resource—one that companies are holding close. As firms reach the limits of publicly available data (especially as more companies begin to create walls around previously open data), proprietary data will play an increasingly important role in training AI—meaning that Big Tech firms with their deep data feedback loops across their ecosystems have almost limitless access to valuable, diverse types of data.¹⁴¹ As high-quality data resources run out for non-Big Tech players, they will look to synthetic data—meaning either datasets containing data generated by other models, or AI-generated content (slop) that exists all over the internet.¹⁴² While this may lower access to data in the short term, we should wonder what this will say about the quality of new models repeatedly trained on synthetic data of increasingly lower quality—while those incumbents with unlimited access to new data retain their data advantage.

There's a need to double down on advocating for data minimization to prevent Big Tech from using AI as its catch-all justification to use and combine data across contexts and store it forever, and limit aggressive strategies for acquiring companies and datasets to shore up their data advantages in high-value sectors or shut out competitors.¹⁴³ Recent revelations that US DOGE's operations, motivated by Elon Musk, may be feeding sensitive government data (including Medicare, Social Security payments, and taxpayer data) to train AI models¹⁴⁴ for his company xAI brought these twin concerns into sharp focus. Government data, especially that obtained from and about the public, is an extremely valuable data asset for training AI. High-quality, vetted, and structured, this data covers a range of high-value sectors including healthcare, finance, and federal contracting. Amid frenzied AI competition, DOGE's vast access to federal data could give the company an unfair advantage over peer AI companies—as well as create a clear conflict of interest with regard to Musk's own involvement in

spearheading DOGE. It's also extremely sensitive data in terms of its ability to be used against the best interests of individuals and groups, without their control. This combination of privacy and competition risks will be on the rise, and it will be important to thread these concerns together in policy and advocacy fights, as well as to continue tracking¹⁴⁵ “data mergers” in direct and indirect forms.

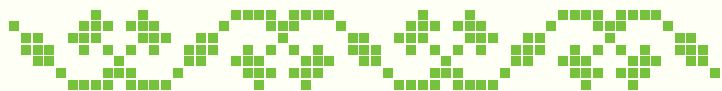
Tech May Be Big, But It Isn't a Monolith



Achieving legislative victories has been particularly challenging given an overall environment of enthusiasm for AI and fear that any intervention would hinder future prospects for innovation. Emerging points of fracture within the industry offer some points of potential leverage to achieve policy goals. For example:

- ❖ **Little and Big Tech:** The strong current of criticism against Big Tech firms' dominance and discontent with concentrated power in the AI industry, and interest in fostering a competitive AI ecosystem, offers some potential for intervention. "Little tech" has shown interest in stronger antitrust enforcement in the sector to open up pathways for new entry into the market,¹⁴⁶ while Big Tech firms show comparatively greater comfort with regulatory approaches, particularly at the federal level, that offer consistency.¹⁴⁷ The little tech policy agenda also uses innovation as a gloss that distracts from its underlying goal to protect venture capital firms' own portfolios, by fighting against a tax on unrealized capital gains, securing defense contracts, and stopping merger-and-acquisition review so that little tech companies can be acquired by Big Tech companies.¹⁴⁸ There is an opportunity here to focus on an innovation agenda beyond the confines of Big and "little" tech, a discourse currently driven by VC interests like a16z and YCombinator—all of which are focused primarily on profit-seeking and shareholder interest. What's missing from the conversation about AI and innovation is a countermovement that brings together researchers, community organizations, anti-monopoly advocates, and regulatory experts to craft a meaningful agenda for renewing innovation in this country.
- ❖ **AI Arms Race:** The differential effects of the AI Arms Race on NVIDIA and Oracle offer a rare wedge between chipmakers and the rest of the market. This could be leveraged toward countermovement in opposition to the scaling paradigm.
- ❖ **Open Source:** Many big AI players express concern about the spread of large-scale open source models on national security grounds (Anthropic, OpenAI), but key players in the Trump administration support open source/safety/antitrust intervention (Musk, Vance).

- ❖ **Elon Musk and Sam Altman:** The growing feud between two of the world's leading AI billionaires has already paid dividends, as Musk's lawsuit against OpenAI made public internal documents that gave insight into OpenAI's operations.¹⁴⁹ Musk has demonstrated an appetite for throwing Altman under the bus, including questioning the integrity of the Stargate deal.¹⁵⁰
- ❖ **Tech companies and their workers:** While worker-led organizing within tech companies has faced significant retaliation, important worker-led campaigns in recent years have pushed back against corporate leadership and enacted important wins. Examples include Google workers' Project Nimbus campaign,¹⁵¹ Microsoft workers' opposition to sales to the Israeli military,¹⁵² and DeepMind workers' move to unionize in opposition to the company's contracts with the Israeli military.¹⁵³



WE NEED TO RECLAIM THE POSITIVE AGENDA FOR PUBLIC-CENTERED INNOVATION, AND AI SHOULDN'T BE AT THE CENTER OF IT



Through 2024 we saw rapid momentum around the idea that AI development needs to be guided by non-market imperatives and that public and philanthropic capital can be put to the task of creating an innovation ecosystem outside of industry.¹⁵⁴ This includes a wide spectrum of proposals from think tanks, philanthropies, and governments ranging from introduction of public options to compete with private industry,¹⁵⁵ to public-private arrangements aimed at increasing access to the inputs needed to build AI.¹⁵⁶ Much of this activity was captured under the umbrella of “public AI,” while some of the efforts spearheaded by government have since been termed “sovereign AI,”¹⁵⁷ and more explicitly oriented around national security and competitiveness goals.¹⁵⁸

While this discourse drives home the realization that the current market-driven trajectory of AI is an impoverished one for the public, the alternative solution space as it stands has also tended to lack specificity, vision, and—most fundamentally—fails to follow through on the challenge to concentrated power,

either accepting or evading the reality that building on the current large-scale AI paradigm will inevitably reinforce the power of the largest tech firms. For these reasons, it has been ripe for capture.

Rejecting the current paradigm of large-scale AI is necessary to fight the information and power asymmetries vested in AI. This is the quiet part that we must say out loud. It's the reality we must confront if we are to muster the will and the creativity to shape the situation differently.

The trajectory of the US National AI Research Resource (NAIRR) holds important lessons.



- ❖ **A bold vision to contest consolidation:** The motivation for NAIRR can be traced back to a 2021 report by the National Security Commission on AI, an independent public commission chaired by ex-Google CEO Eric Schmidt and populated by executives from many large tech companies that, uncharacteristically, named the problem: “The consolidation of the AI industry threatens U.S. technological competitiveness.”¹⁵⁹ By creating public computing and data resources for AI, the NSCAI argued, the government could support a diversification of the pipeline of research, talent, and businesses that could play in the AI space.
- ❖ **Refocusing, and a reality check:** If you can’t contest the market, diversify the pipeline. Right off the bat, however, the initiative seemed to reinforce dependencies on Big Tech, and objectives around combating consolidation fell away, focusing instead exclusively on diversifying the pipeline of AI talent and research. Part of this was inevitable given that state-of-the-art computing infrastructure is expensive and hard to access, and the reality was that it wasn’t possible to expand access to bigger-is-better AI research resources through institutions meaningfully separate from the Big Tech companies. This meant public funds would eventually be routed back to the largest tech companies who controlled access to AI infrastructure: the original proposal was structured as a licensing regime, built around cloud contracts with a rotating set of licensed providers on six-year terms.¹⁶⁰ As we argued alongside Data & Society, “the NAIRR proposal is a tacit endorsement of a massive investment in large tech companies in service of expanding their proprietary infrastructure.”¹⁶¹
- ❖ **NAIRR as a public-private partnership:** When NSF eventually launched the NAIRR pilot in January 2024, the initiative was a more involved partnership with industry AI companies. Alongside access to existing government supercomputers,¹⁶² datasets, and research resources, AI companies and other philanthropies were invited to offer up donated resources on a public-private marketplace hosted by the NSF: AWS would support cloud credits for at least twenty research projects; NVIDIA would provide \$24 million in computing on its DGX

Cloud offering (which at least in part recycles data center access from Google Cloud, AWS, and Azure); and OpenAI pledged \$1 million in credits for model access for research “related to AI safety, evaluations and societal impacts.”¹⁶³

- ❖ **NAIRR is recast as part of a broader national competitiveness project:** Moving further from its origin story, NAIRR has now been recast within the broader set of industrial policies aimed at strengthening US competitiveness on AI against China. In the Biden administration’s final National Security Memo outlining executive action toward advancing US leadership on AI, NAIRR featured prominently as a means “to ensure that AI research in the United States remains competitive and innovative.” Under Trump, the future of NAIRR remains more uncertain, especially given wider federal funding cuts, but it is almost certain that the orientation will lean deeper toward geopolitical advantage rather than challenging industry concentration or diversity in talent pipelines and research trajectories.

This is not to say that NAIRR hasn’t provided public value—the pilot matched data and compute resources to over three hundred university researchers and others who may not have otherwise had this access—potentially spurring socially beneficial projects.¹⁶⁴ **Yet there is a glaring missed opportunity: NAIRR could have spurred a widely deliberative and ambitious project that defined and fostered the conditions for a public-interest-oriented AI.** Rather than building on the “bigger-is-better” trajectory in which all roads inevitably lead to doubling down on Big Tech infrastructure dependencies (fueling resources back to the same companies) or relying on a large, centralized state apparatus to shape the alternative vision, there was a moment for opening up space (and funding for) alternative research trajectories: conceptualizing a sustainable AI supply chain, investing in modes of data stewardship that are accountable to the communities technology is designed to serve, and widening the aperture beyond tech termed as “AI” to bring resources toward other technological gaps that affect the public (such as secure communications infrastructure).¹⁶⁵

Despite current limitations, public AI is still a fertile opportunity space to both push a conversation around alternative trajectories for AI that are more structurally aligned with the public interest, as well as ensure that any public funding in this space is conditioned on public-interest goals. On the government side, while US NAIRR is on shaky footing, states are considering investments in public AI resources across the country. A California bill recently revived a proposal for a publicly owned and operated computer cluster housed in the University of California system called “CalCompute.”¹⁶⁶ New York established an initiative called Empire AI to build public cloud infrastructure across seven New York state research institutions, bringing together over \$400 million in public and private funds.¹⁶⁷ Both of these initiatives create important sites for advocacy to ensure their resources support communities’ needs and do not become a vehicle to further enrich Big Tech resources.

Key Directions Forward for Public AI Efforts



- ❖ **Join in the defense of public universities amid attacks on funding, tenure, and freedom of speech.** Universities have long been embattled institutions, and the precarity of academia has contributed to capture by industry players. This precarity has deepened into an existential crisis under current political pressure: The defense of the ability to research and teach freely is core to protecting any future for public-minded AI.

- ❖ **Reinvigorate the playbook.** Current public AI proposals rely too heavily on licensing, tax exemption, and provision of credits that ultimately reinforce, rather than contest, the dependency on cloud providers. We need new mental models for how state capacity can contribute to public goods in the tech sector that don’t rehash the neoliberal approach characteristic of the 1990s and 2000s, which created these issues in the first place.
- ❖ **Invest in public-sector technical expertise that grounds decision-making around any development of public-interest AI in the ability to ask if AI is the right tool, rather than shoehorning how AI could be made to work.**
- ❖ **Ensure structurally disadvantaged groups have deliberate authority in this conversation.**
- ❖ **Resource research on alternative R&D pathways that aren’t adequately incentivized and resourced by the market to help illuminate a vision outside of the path dependencies of the current trajectory.**
- ❖ **Cultivate a community of practice around public benefit innovation that will shape this emerging and alternative space.** In doing so, shift the **focus away from only outcomes and toward guaranteeing more democratic inputs** into the public-interest AI strategy, while ensuring that structurally disadvantaged groups have deliberate authority in this conversation.

❖ **Define collective or public interest as separate from the interest of the nation-state, recognizing that state and national security interests might often run counter to those of the public, especially to structurally disadvantaged groups or those vulnerable to state violence.**

- ◆ This will require creating an alternative to the predominant national security vision for public AI that superficially has the same starting point (concerns with unaccountable private control over technology, independence) but is preoccupied with a vision for AI that serves the interests of the techno-military industrial complex over all others, and leads to some of the most destructive and life-threatening applications of AI.

❖ **Define concrete conditionalities to attach to public AI funding, ensuring that investment meets communities' needs rather than corporate interests.** These may include:

1. First, ensuring that investment neither depends on private industry nor leads to private capture. Where there is private investment, ensure that funders do not place research constraints or funding conditions on their investments.
2. Second, divert public AI investment from the defense industry and military expansion.
3. Third, successfully structure conditionalities to countervail corporate interests; advance climate goals; and support unionized workers, marginalized communities, Indigenous land, and others harmed by corporate extraction.



Executive Summary Endnotes

- 1 Amba Kak and Sarah Myers West, *2023 Landscape: Confronting Tech Power*, AI Now Institute, April 11, 2023, <https://ainowinstitute.org/publications/research/2023-landscape-confronting-tech-power>.
- 2 David Cahn, "Steel, Servers and Power: What it Takes to Win the Next Phase of AI," *Sequoia*, August 5, 2024, <https://www.sequoiacap.com/article/steel-servers-and-power>.
- 3 Meredith Whittaker, "The Steep Cost of Capture," *Interactions* 28, no. 6 (November 2021): 50–55, <https://dl.acm.org/doi/10.1145/3488666>.
- 4 See Nur Ahmed, Muntasir Wahed, and Neil C. Thompson, *Science* 379, no. 6635 (March 2023): 884–886, <https://www.science.org/doi/abs/10.1126/science.ade2420>; and Whittaker, "The Steep Cost of Capture."
- 5 Kak and West, *2023 Landscape: Confronting Tech Power*.
- 6 Edward Zitron, "There Is No AI Revolution," *Where's Your Ed At*, February 24, 2025, <https://www.wheresyoured.at/wheres-the-money>.
- 7 Shirin Ghaffary, "Anthropic Finalizes Megaround at \$61.5 Billion Valuation," *Bloomberg*, <https://www.bloomberg.com/news/articles/2025-03-03/anthropic-finalizes-megaround-at-61-5-billion-valuation>.
- 8 Cory Weinberg, "OpenAI Projections Imply Losses Tripling to \$14 Billion in 2026," *The Information*, October 9, 2024, <https://www.theinformation.com/articles/openai-projections-imply-losses-tripling-to-14-billion-in-2026>.
- 9 Cade Metz, "OpenAI Completes Deal That Values Company at \$300 Billion," *New York Times*, March 31, 2025, <https://www.nytimes.com/2025/03/31/technology/openai-valuation-300-billion.html>.
- 10 Nilay Patel, "Mark Zuckerberg Just Declared War on the Entire Advertising Industry," *The Verge*, May 1, 2025, <https://www.theverge.com/meta/659506/mark-zuckerberg-ai-facebook-ads>.
- 11 Casey Newton, "ChatGPT's Deep Research Might Be the First Good Agent," *Platformer*, February 3, 2025, <https://www.platformer.news/chatgpt-deep-research-hands-on>.
- 12 Grant Gross, "Devs Gaining Little (If Anything) from AI Coding Assistants," *CIO*, September 26, 2024, <https://www.cio.com/article/3540579/devs-gaining-little-if-anything-from-ai-coding-assistants.html>.
- 13 Abner Li, "Google Workspace Business Users Getting Full Gemini Experience, Price Increasing," *9to5Google*, January 15, 2025, <https://9to5google.com/2025/01/15/google-workspace-gemini-price-increase>.
- 14 Ashish Agarwal et al., "Hallucinations in Neural Machine Translation," abstract, 2018, accessed May 15, 2025, <https://research.google/pubs/hallucinations-in-neural-machine-translation>.
- 15 Inioluwa Deborah Raji et al., "The Fallacy of AI Functionality," *arXiv* (2022), <https://arxiv.org/abs/2206.09511>.
- 16 Chad Boutin, "NIST Identifies Types of Cyberattacks That Manipulate Behavior of AI Systems," NIST, January 4, 2024, <https://www.nist.gov/news-events/news/2024/01/nist-identifies-types-cyberattacks-manipulate-behavior-ai-systems>.
- 17 Sarah Myers West and Amba Kak, "AI's Missing Others," *Boston Review*, December 4, 2024, https://www.bostonreview.net/forum_response/ais-missing-others.
- 18 Catherine Bracy, *World Eaters* (Penguin Random House, 2025).
- 19 See David McCabe, "'Google Is a Monopolist,' Judge Rules in Landmark Antitrust Case," *New York Times*, August 5, 2024, <https://www.nytimes.com/2024/08/05/technology/google-antitrust-ruling.html>; and McCabe, "Google Broke the Law to Keep Its Advertising Monopoly, a Judge Rules," *New York Times*, April 17, 2025, <https://www.nytimes.com/2025/04/17/technology/google-ad-tech-anti-trust-ruling.html>.
- 20 David McCabe, "U.S. Wants to Break Up Google's Advertising Technology," *New York Times*, May 2, 2025, <https://www.nytimes.com/2025/05/02/technology/google-advertising-technology-hearing.html>.
- 21 See Kate Brennan, "The Elephant in the Room in the Google Search Case: Generative AI," *Tech Policy Press*, November 4, 2024, <https://www.techpolicy.press/the-elephant-in-the-room-in-the-google-search-case-generative-ai>; Erin Woo, "OpenAI's Bid To Get Google's Search Data," *The Information*, April 23, 2025, <https://www.theinformation.com/articles/openais-bid-get-googles-search-data>; and Lauren Feiner, "Perplexity Wants to Buy Chrome If Google Has to Sell It," *The Verge*, April 23, 2025, <https://www.theverge.com/policy/654835/perplexity-google-antitrust-trial-remedies-chrome>.
- 22 Maria Farrell and Robin Berjon, "We Need to Rewild the Internet," *Noëma*, April 16, 2024, <https://www.noemamag.com/we-need-to-re-wild-the-internet>.
- 23 Dan McQuillan, *Resisting AI* (Bristol University Press, 2022).
- 24 Alex Hanna, "The AI/Fascism Nexus," *Mystery AI Hype Theater 3000: The Newsletter*, November 8, 2024, <https://buttondown.com/mai-ht3k/archive/the-aifascism-nexus>.



Chapter 1: AI's False Gods Endnotes

1.1: The AGI Mythology: The Argument to End All Arguments

- 1 Sam Altman, "Reflections," January 5, 2025, <https://blog.samaltman.com/reflections>; Sébastien Bubeck et al., "Sparks of Artificial General Intelligence: Early experiments with GPT-4," Microsoft, March 2023, <https://www.microsoft.com/en-us/research/publication/sparks-of-artificial-general-intelligence-early-experiments-with-gpt-4>.
- 2 See Brian Merchant, "AI Generated Business," AI Now Institute, December 2024, <https://ainowinstitute.org/general/ai-generated-business>; Berber Jin and Deepa Seetharaman, "This Scientist Left OpenAI Last Year. His Startup Is Already Worth \$30 Billion," *Wall Street Journal*, March 4, 2025, <https://www.wsj.com/tech/ai/ai-safe-superintelligence-startup-ilya-sutskever-openai-2335259b>; and John Koetsier, "OpenAI CEO Sam Altman: 'We Know How To Build AGI,'" *Forbes*, January 6, 2025, <https://www.forbes.com/sites/johnkoetsier/2025/01/06/openai-ceo-sam-altman-we-know-how-to-build-agi>.
- 3 See Ezra Klein, "The Government Knows A.G.I. Is Coming," *New York Times*, March 4, 2025, <https://www.nytimes.com/2025/03/04/opinion/ezra-klein-podcast-ben-buchanan.html>; and White House, "Removing Barriers to American Leadership in Artificial Intelligence," January 23, 2025, <https://www.whitehouse.gov/presidential-actions/2025/01/removing-barriers-to-american-leadership-in-artificial-intelligence>.
- 4 Will Douglas Heaven, "Geoffrey Hinton Tells Us Why He's Now Scared of the Tech He Helped Build," *MIT Technology Review*, May 2, 2023, <https://www.technologyreview.com/2023/05/02/1072528/geoffrey-hinton-google-why-scared-ai>.
- 5 OpenAI, "Introducing OpenAI," December 11, 2015, <https://openai.com/index/introducing-openai>.
- 6 Merchant, "AI Generated Business."
- 7 Merchant, "AI Generated Business."
- 8 Dario Amodei, "Machines of Loving Grace," October 2024, <https://darioamodei.com/machines-of-loving-grace>.
- 9 Maxwell Zeff, "Microsoft and OpenAI Have a Financial Definition of AGI: Report," *TechCrunch*, December 26, 2024, <https://techcrunch.com/2024/12/26/microsoft-and-openai-have-a-financial-definition-of-agi-report>.
- 10 Kevin Roose, "Powerful A.I. Is Coming. We're Not Ready," *New York Times*, March 14, 2025, <https://www.nytimes.com/2025/03/14/technology/why-im-feeling-the-agi.html>.
- 11 See Anil Ananthaswamy, "How Close Is AI to Human-Level Intelligence?" *Nature*, December 3, 2024, <https://www.nature.com/articles/d41586-024-03905-1> (highlighting OpenAI's claims that "o1 works in a way that is closer to how a person thinks than do previous LLMs"); Ryan Browne, "AI that can match humans at any task will be here in five to 10 years, Google DeepMind CEO says," *CNBC*, March 17, 2025, <https://www.cnn.com/2025/03/17/human-level-ai-will-be-here-in-5-to-10-years-deepmind-ceo-says.html>; Koetsier, "OpenAI CEO Sam Altman"; and Steve Ranger, "OpenAI Says It's Charting a 'Path to AGI' with Its Next Frontier AI Model," *ITPro*, May 30, 2024, <https://www.itpro.com/technology/artificial-intelligence/openai-says-its-charting-a-path-to-agi-with-its-next-frontier-ai-model>.
- 12 Association for the Advancement of Artificial Intelligence (AAAI), *AAAI 2025 Presidential Panel on the Future of AI Research*, March 2025, <https://aaai.org/wp-content/uploads/2025/03/AAAI-2025-PresPanel-Report-Digital-3.7.25.pdf>.
- 13 Emily M. Bender, "Resisting Dehumanization in the Age of 'AI,'" *Current Directions in Psychological Science* 33, no. 2 (2024): <https://doi.org/10.1177/09637214231217286>.
- 14 Henry Farrell et al., "Large AI models Are Cultural and Social Technologies," *Science* 387, no. 6739 (2025): 1153–56, <https://doi.org/10.1126/science.adt9819>.
- 15 Henry Farrell, "Should AGI-Preppers Embrace DOGE?" *Programmable Mutter* (blog), March 18, 2025, <https://www.programmablemutter.com/p/should-agi-preppers-embrace-doge>.
- 16 Prithvi Iyer, "What Do We Mean When We Say 'Artificial General Intelligence?'" *Tech Policy Press*, February 13, 2024, <https://www.techpolicy.press/what-do-we-mean-when-we-say-artificial-general-intelligence>, citing Borhane Bili-Hamelin, Leif Hancox-Li, and Andrew Smart, "Unsocial Intelligence: An Investigation of the Assumptions of AGI Discourse," *arXiv*, (2024), arXiv:2401.13142.
- 17 See Klein, "The Government Knows A.G.I. Is Coming"; Jakob Nielsen, "AI Hallucinations on the Decline," *Jakob Nielsen on UX* (blog), February 13, 2025, <https://jakobnielsenphd.substack.com/p/ai-hallucinations>; and Josh Tyrangiel, "Sam Altman on ChatGPT's First Two Years, Elon Musk and AI Under Trump," *Bloomberg*, January 5, 2025, <https://www.bloomberg.com/features/2025-sam-altman-interview>.
- 18 Lakshmi Varanasi, "Here's How Far We Are from AGI, According to the People Developing It," *Business Insider*, November 9, 2024, <https://www.businessinsider.com/agi-predictions-sam-altman-dario-amodei-geoffrey-hinton-demis-hassabis-2024-11>.
- 19 See generally Sayash Kapoor and Arvind Narayanan, "AGI Is Not a Milestone," *AI Snake Oil*, May 1, 2025, <https://www.aisnakeoil.com/p/agi-is-not-a-milestone>.
- 20 AAAI, "AAAI 2025 Presidential Panel."
- 21 Benji Edwards, "Elon Musk: AI Will Be Smarter than Any Human Around the End of Next Year," *Ars Technica*, April 9, 2024, <https://arstechnica.com/information-technology/2024/04/elon-musk-ai-will-be-smarter-than-any-human-around-the-end-of-next-year> (claiming that AI capability increases by "a factor of 10 every year, if not every six to nine months").

- 22 Borhane Blili-Hamelin et al., “Stop Treating ‘AGI’ as the North-Star Goal of AI Research,” *arXiv* (2025), [arXiv:2502.03689](https://arxiv.org/abs/2502.03689).
- 23 Center for AI Safety and Scale AI, “Humanity’s Last Exam,” accessed April 30, 2025, <https://agi.safe.ai>.
- 24 Epoch AI, “FrontierMath,” accessed April 30, 2025, <https://epoch.ai/frontiermath>.
- 25 Cf. Satya Nadella’s statement: “The real benchmark is: the world growing at 10 percent.” Victor Tangermann, “Microsoft CEO Admits That AI Is Generating Basically No Value,” *Futurism*, February 22, 2025.
- 26 Thomas Wolf, “The Einstein AI Model,” blog, February 25, 2025, <https://thomwolf.io/blog/scientific-ai.html>.
- 27 Weidinger, Raji, et al., “Toward an Evaluation Science for Generative AI Systems,” *arXiv*, March 13, 2025.
- 28 Prithvi Iyer, “Transcript: US Lawmakers Probe AI’s Role in Energy and Climate,” *Tech Policy*, April 11, 2025, <https://www.techpolicy.press/transcript-us-lawmakers-probe-ais-role-in-energy-and-climate>.
- 29 See Eve Darian-Smith, “The Challenge of Political Will, Global Democracy and Environmentalism,” *Environmental Policy and Law* 54, no. 2–3 (2024): 117–126, <https://doi.org/10.3233/EPL-239023>; Alejandro de la Garza, “We Have the Technology to Solve Climate Change. What We Need Is Political Will,” *Time*, April 7, 2022, <https://time.com/6165094/ipcc-climate-action-political-will>.
- 30 Klein, “The Government Knows A.G.I. Is Coming.”
- 31 Meredith Whittaker, “The Steep Cost of Capture,” *Interactions* 28, no. 6 (2021): 50–55, <https://doi.org/10.1145/3488666>.
- 32 Miles Brundage (@Miles_Brundage), “Per Standard AI Scaling Laws, a 2x Advantage in Compute Does Not Yield a 2x Advantage in Capabilities,” X, March 6, 2025, https://x.com/miles_brundage/status/1897568753178865900.
- 33 Amodei, “Machines of Loving Grace,” citing Rich Sutton, “The Bitter Lesson,” *Incomplete Ideas*, March 13, 2019, https://www.cs.utexas.edu/~eunsol/courses/data/bitter_lesson.pdf.
- 34 Emma Salisbury, “A Cautionary Tale on Ambitious Feats of AI: The Strategic Computing Program,” *War on the Rocks*, May 22, 2020, <https://warontherocks.com/2020/05/cautionary-tale-on-ambitious-feats-of-ai-the-strategic-computing-program>.
- 35 Alvin Wang Graylin and Paul Triolo, “There Can Be No Winners in a US-China AI Arms Race,” *MIT Technology Review*, January 21, 2025, <https://www.technologyreview.com/2025/01/21/1110269/there-can-be-no-winners-in-a-us-china-ai-arms-race>.
- 36 Evgeny Morozov, *To Save Everything, Click Here* (PublicAffairs, 2014).
- 37 In one telling example, researchers published a paper in *Nature* claiming discovery of over forty novel materials using an AI-driven autonomous laboratory. Shortly afterward, two materials chemists critiqued the paper for failing to recognize systematic errors with unsupervised materials discovery. See Julia Robinson, “New Analysis Raises Doubts Over Autonomous Lab’s Materials ‘Discoveries,’” *Royal Society of Chemistry*, January 16, 2024, <https://www.chemistryworld.com/news/new-analysis-raises-doubts-over-autonomous-labs-mate->

[rials-discoveries/4018791.article](https://www.chemistryworld.com/news/new-analysis-raises-doubts-over-autonomous-labs-materials-discoveries/4018791.article); and Robert Palgrave (@Robert_Palgrave), “This exciting paper shows AI design of materials, robotic synthesis. 10s of new compounds in 17 days. But did they? This paper has very serious problems in materials characterisation. In my view it should never have got near publication. Hold on tight let’s take a look 🤖,” X, November 30, 2023, https://x.com/Robert_Palgrave/status/1730358675523424344.

- 38 Samantha M. Kelly, “Elon Musk Says AI Will Take All Our Jobs,” *CNN*, May 23, 2024, <https://www.cnn.com/2024/05/23/tech/elon-musk-ai-your-job/index.html>.

1.2: Too Big to Fail: Infrastructure and Capital Push

- 1 Jeran Wittenstein and Tom Contiliano, “Nvidia and Five Tech Giants Now Command 30% of the S&P 500 Index,” *Bloomberg*, May 30, 2024, <https://www.bloomberg.com/news/articles/2024-05-30/nvidia-and-five-tech-giants-now-command-30-of-the-s-p-500-index>; Stephanie Stacey and George Steer, “Wall Street’s Magnificent Seven Lose Their Shine,” *Financial Times*, February 15, 2025, <https://www.ft.com/content/fa5d3b2d-b3b3-4bb4-a5a4-765b7560e02c>.
- 2 Shivalika Singh et al., “The Leaderboard Illusion,” *arXiv*, April 29, 2025, <https://arxiv.org/abs/2504.20879>.
- 3 And, as recent model releases have shown, the pace of scale-based advancement may be slowing. See, e.g., Maxwell Zeff, “Current AI Scaling Laws Are Showing Diminishing Returns, Forcing AI Labs to Change Course,” *TechCrunch*, November 20, 2024, <https://techcrunch.com/2024/11/20/ai-scaling-laws-are-showing-diminishing-returns-forcing-ai-labs-to-change-course>.
- 4 See generally Bryan McMahon, “Bubble Trouble,” *American Prospect*, March 25, 2025, <https://prospect.org/power/2025-03-25-bubble-trouble-ai-threat>.
- 5 George Hammond, “Big Tech Outspends Venture Capital Firms in AI Investment Frenzy,” *Financial Times*, December 27, 2023, <https://www.ft.com/content/c6b47d24-b435-4f41-b197-2d826cce9532>.
- 6 See Stan Choe, “Tech Stocks Tank as a Chinese Competitor Threatens to Upend the AI Frenzy; Nvidia Sinks Nearly 17%,” *Associated Press*, January 27, 2025, <https://apnews.com/article/stocks-markets-tariffs-trump-rates-52c54e361616509280bd2775674b6b4b>; and Natasha Solo-Lyons, “Nvidia Loses \$589 Billion as DeepSeek Batters Stock,” *Bloomberg*, January 27, 2025, <https://www.bloomberg.com/news/newsletters/2025-01-27/nvidia-loses-589-billion-as-deepseek-batters-stock-evening-briefing-americas>.
- 7 “Announcing the Stargate Project,” *OpenAI*, January 21, 2025, <https://openai.com/index/announcing-the-stargate-project/>.
- 8 “Executive Order 14036 of July 9, 2021, Promoting Competition in the American Economy,” Code of Federal Regulations, title 86 (2021): 36987-36999, <https://www.federalregister.gov/documents/2021/07/14/2021-15069/promoting-competition-in-the-american-economy>.
- 9 JD Vance, “Remarks by the Vice President at the Artificial Intelligence Action Summit in Paris, France,” (speech, Paris, France, February 11, 2025) *American Presidency Project*, <https://www.presidency.ucsb.edu/documents/remarks-the-vice-president-the-artificial-intelligence-action-summit-paris-france>.

- 10 Josh Sisco, "FTC Has the Resources to Take On Big Tech, Chairman Says," *Bloomberg*, March 17, 2025, <https://www.bloomberg.com/news/articles/2025-03-17/ftc-has-the-resources-to-take-on-big-tech-chairman-says>.
- 11 White House, "Fact Sheet: President Donald J. Trump Establishes the National Energy Dominance Council," February 14, 2025, <https://www.whitehouse.gov/fact-sheets/2025/02/fact-sheet-president-donald-j-trump-establishes-the-national-energy-dominance-council>.
- 12 U.S. Department of Energy (DOE), "DOE Identifies 16 Federal Sites Across the Country for Data Center and AI Infrastructure Development," *US Department of Energy*, April 3, 2025, <https://www.energy.gov/articles/doe-identifies-16-federal-sites-across-country-data-center-and-ai-infrastructure>.
- 13 Public Citizen, "Trump's Abuse of Emergency Declaration to Force Ratepayers to Prop Up Inefficient Coal Power Plants Is Breathlessly Stupid," April 8, 2025, <https://www.citizen.org/news/trumps-ao-to-prop-up-coal-plant-inefficient-coal-power-plants-is-stupid>.
- 14 See Yi Chen, JiaHao Zhao, and HaoHao Han, "A Survey on Collaborative Mechanisms Between Large and Small Language Models," *arXiv*, May 12, 2025; and Gaël Varoquaux, Alexandra Sasha Luccioni, and Meredith Whittaker, "Hype, Sustainability, and the Price of the Bigger-is-Better Paradigm in AI," *arXiv*, September 21, 2024, <https://arxiv.org/abs/2409.14160>.
- 15 Sally Beatty, "Tiny But Mighty: The Phi-3 Small Language Models with Big Potential," *Microsoft*, April 23, 2024, <https://news.microsoft.com/source/features/ai/the-phi-3-small-language-models-with-big-potential>.
- 16 Cristina Criddle and Madhumita Murgia, "Artificial Intelligence Companies Seek Big Profits from 'Small' Language Models," *Financial Times*, May 18, 2024, <https://www.ft.com/content/359a5a31-1ab9-41ea-83aa-5b27d9b24ef9>.
- 17 Aili McConnon, "DeepSeek's Reasoning AI Shows Power of Small Models, Efficiently Trained," *IBM*, January 27, 2025, <https://www.ibm.com/think/news/deepseek-r1-ai>.
- 18 Aditya Soni and Deborah Mary Sophia, "Microsoft, Meta Back Big AI Spending Despite DeepSeek's Low Costs," *Reuters*, January 30, 2025, <https://www.reuters.com/technology/artificial-intelligence/microsoft-meta-ceos-defend-hefty-ai-spending-after-deepseek-stuns-tech-world-2025-01-30>.
- 19 Gov.UK, "Prime Minister Sets Out Blueprint to Turbocharge AI," January 13, 2025, <https://www.gov.uk/government/news/prime-minister-sets-out-blueprint-to-turbocharge-ai>.
- 20 Ezra Klein, "The Economic Mistake the Left Is Finally Confronting," *New York Times*, September 19, 2021, <https://www.nytimes.com/2021/09/19/opinion/supply-side-progressivism.html>.
- 21 Jerusalem Demsas, "Why America Doesn't Build," *Atlantic*, October 27, 2023, <https://www.theatlantic.com/ideas/archive/2023/10/wind-farms-community-opposition/675791>.
- 22 Ezra Klein and Derek Thompson, *Abundance* (Avid Reader Press, 2025).
- 23 James O'Donnell, "AI's Emissions Are About to Skyrocket Even

Further," *Technology Review*, December 13, 2024, <https://www.technologyreview.com/2024/12/13/1108719/ais-emissions-are-about-to-skyrocket-even-further>.

- 24 Goldman Sachs, "AI is Poised to Drive 160% Increase in Data Center Power Demand," May 14, 2024, <https://www.goldmansachs.com/insights/articles/AI-poised-to-drive-160-increase-in-power-demand>.
- 25 Committee on Oversight and Government Reform, "America's AI Moonshot: The Economics of AI, Data Centers, and Power Consumption," April 1, 2025, <https://oversight.house.gov/hearing/americas-ai-moonshot-the-economics-of-ai-data-centers-and-power-consumption>.
- 26 Brad Plumer, "Want Cheap Power, Fast? Solar and Wind Firms Have a Suggestion," *New York Times*, March 17, 2025, <https://www.nytimes.com/2025/03/17/climate/renewable-energy-trump-electricity.html?smid=nytcore-ios-share>.

1.3: AI Arms Race 2.0: From Deregulation to Industrial Policy

- 1 Amba Kak et al., AI Nationalism(s): Global Industrial Policy Approaches to AI, AI Now Institute, March 12, 2024, <https://ainowinstitute.org/ai-nationalisms>.
- 2 AI Now Institute, Tracking the US and China AI Arms Race, April 11, 2023, <https://ainowinstitute.org/publications/tracking-the-us-and-china-ai-arms-race>.
- 3 Amba Kak and Sarah M. West, "A Modern Industrial Strategy for AI?: Interrogating the US Approach," AI Now Institute, March 12, 2024, <https://ainowinstitute.org/publications/a-modern-industrial-strategy-for-ai-interrogating-the-us-approach>.
- 4 Jake Sullivan, "Remarks by National Security Advisor Jake Sullivan on Renewing American Economic Leadership at the Brookings Institution" (speech, Washington, DC, April 27, 2023), The White House, <https://bidenwhitehouse.archives.gov/briefing-room/speeches-remarks/2023/04/27/remarks-by-national-security-advisor-jake-sullivan-on-renewing-american-economic-leadership-at-the-brookings-institution>; Jake Sullivan, "Remarks by APNSA Jake Sullivan at the Brookings Institution" (speech, Washington, DC, October 23, 2024), The White House, <https://bidenwhitehouse.archives.gov/briefing-room/speeches-remarks/2024/10/23/remarks-by-apnsa-jake-sullivan-at-the-brookings-institution/>.
- 5 "Executive Order 14110 of October 30, 2023, Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence," *Code of Federal Regulations*, title 88 (2023): 75191-75226, <https://www.federalregister.gov/documents/2023/11/01/2023-24283/safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence>.
- 6 White House, "Memorandum on Advancing the United States' Leadership in Artificial Intelligence; Harnessing Artificial Intelligence to Fulfill National Security Objectives; and Fostering the Safety, Security, and Trustworthiness of Artificial Intelligence," October 24, 2024, <https://bidenwhitehouse.archives.gov/briefing-room/presidential-actions/2024/10/24/memorandum-on-advancing-the-united-states-leadership-in-artificial-intelligence-harnessing-artificial-intelligence-to-fulfill-national-security-objectives-and-fostering-the-safety-security>.
- 7 "Export Control Framework for Artificial Intelligence Diffusion," *Code*

- of Federal Regulations*, title 15 (2024): 740, 774, <https://www.govinfo.gov/content/pkg/CFR-2024-title15-vol2/pdf/CFR-2024-title15-vol2-part740.pdf>.
- 8 Cade Metz and Tripp Mickle, "Behind OpenAI's Audacious Plan to Make A.I. Flow Like Electricity," *New York Times*, September 25, 2024, <https://www.nytimes.com/2024/09/25/business/openai-plan-electricity.html>.
 - 9 Jordan Wolman and Mohar Chatterjee, "White House Weighing Executive Action to Spur Data Centers," *Politico*, December 11, 2024, <https://subscriber.politicopro.com/article/2024/12/white-house-weighing-executive-action-to-spur-data-centers-00193846>.
 - 10 White House, "Removing Barriers to American Leadership in Artificial Intelligence," January 23, 2024, <https://www.whitehouse.gov/presidential-actions/2025/01/removing-barriers-to-american-leadership-in-artificial-intelligence>.
 - 11 "The Hill & Valley Forum 2025," *The Hill & Valley Forum*, April 30, 2025, <https://www.thehillandvalleyforum.com>.
 - 12 Elizabeth Dwoskin, "Tech Leaders Were on the Outside Looking In. Now They Own Washington," *Washington Post*, March 6, 2025, <https://www.washingtonpost.com/politics/2025/03/06/tech-leaders-were-outside-looking-now-they-own-washington>.
 - 13 Mohar Chatterjee, "Silicon Valley Comes to Washington — Singing Trump's Tune," *Politico*, April 30, 2025, <https://www.politico.com/newsletters/digital-future-daily/2025/04/30/silicon-valley-comes-to-washington-singing-trumps-tune-00319538>; and Dwoskin, "Tech Leaders."
 - 14 Jacob Helberg (@jacobhelberg), "Thank you, Mr. President. I am deeply honored and humbled by the trust that you have placed in me." X, December 10, 2024, <https://x.com/jacobhelberg/status/1866626272203509803>.
 - 15 Helberg has been nominated to become Under Secretary for Economic Growth, Energy and the Environment. Michael Kratsios, from Scale AI, is Director of the Office of Science and Technology Policy. David Sacks is the White House's "AI Czar." See Nitasha Tiku, Cat Zakrzewski, and Elizabeth Dwoskin, "A Podcast Star Rallied Silicon Valley to Back Trump. Now He's the Nation's Tech Czar," *Washington Post*, April 13, 2025, <https://www.washingtonpost.com/technology/2025/04/13/david-sacks-ai-crypto-trump>; and Madison Alder, "Trump Taps Michael Kratsios, Lynne Parker for Tech and Science Roles," *Fedscoop*, December 23, 2024, <https://fedscoop.com/trump-taps-michael-kratsios-lynne-parker-tech-science-roles>.
 - 16 Sophie Hurwitz, "The Glee of Profiteers of Trump's Police State," *Mother Jones*, February 6, 2025, <https://www.motherjones.com/politics/2025/02/palantir-alex-karp-trump-private-prisons-profits>.
 - 17 Sam Biddle, "OpenAI Quietly Deletes Ban on Using ChatGPT for 'Military and Warfare,'" *Intercept*, January 12, 2024, <https://theintercept.com/2024/01/12/open-ai-military-ban-chatgpt>.
 - 18 OpenAI, "OpenAI's Approach to AI and National Security," October 24, 2024, <https://openai.com/global-affairs/openais-approach-to-ai-and-national-security>.
 - 19 Emma Roth, "OpenAI and Google Ask the Government to Let Them Train AI on Content They Don't Own," *The Verge*, March 14, 2025, <https://www.theverge.com/news/630079/openai-google-copy-right-fair-use-exception>.
 - 20 Lucy Hooker and Chris Vallance, "Concern Over Google Ending Ban on AI Weapons," *BBC*, February 5, 2025, <https://www.bbc.com/news/articles/cy081nqx2zjo>; Ina Fried, "Google's Hassabis Explains Shift on Military Use of AI," *Axios*, February 14, 2025, <https://www.axios.com/2025/02/14/google-hassabis-ai-military-use>; Scott Shane and Daisuke Wakabayashi, "'The Business of War': Google Employees Protest Work for the Pentagon," *New York Times*, April 4, 2018, <https://www.nytimes.com/2018/04/04/technology/google-letter-ceo-pentagon-project.html>.
 - 21 Patrick Moorhead, "Meta Extends Llama Support to U.S. Government For National Security," *Forbes*, November 4, 2024, <https://www.forbes.com/sites/patrickmoorhead/2024/11/04/meta-extends-llama-support-to-us-government-for-national-security>.
 - 22 Anthropic, "Statement from Dario Amodei on the Paris AI Action Summit," February 11, 2025, <https://www.anthropic.com/news/paris-ai-summit>.
 - 23 Andreessen Horowitz, "What We Believe," accessed April 24, 2024, <https://a16z.com/american-dynamism>.
 - 24 Ken Glueck, "Export Control Diffusion Confusion," *Oracle* (blog), January 5, 2025, <https://www.oracle.com/news/announcement/blog/export-control-diffusion-confusion-2025-01-05>.
 - 25 Information Technology & Innovation Foundation (ITIF), "AI Diffusion Rule Threatens US Leadership, Warns ITIF," January 13, 2025, <https://itif.org/publications/2025/01/13/ai-diffusion-rule-threatens-us-leadership-warns-itif>.
 - 26 Karen Freifeld, "Exclusive: Trump Officials Eye Changes to Biden's AI Chip Export Rule, Sources Say," *Reuters*, April 29, 2025, <https://www.reuters.com/world/china/trump-officials-eye-changes-bidens-ai-chip-export-rule-sources-say-2025-04-29>.
 - 27 Sydney J. Freedberg Jr., "White House Tries to Tighten AI Export Controls Amidst Industry Outrage," *Breaking Defense*, January 14, 2025, <https://breakingdefense.com/2025/01/white-house-tries-to-tighten-ai-export-controls-amidst-industry-outrage>.
 - 28 Angie Lee, "What is Sovereign AI?" *NVIDIA* (blog), February 28, 2024, <https://blogs.nvidia.com/blog/what-is-sovereign-ai>.
 - 29 Frederike Kaltheuner, Leevi Saari, and AI Now Institute, "German Election, China's AI Rise, and the MAGAfication of Big Tech," *EU AI Industrial Policy Monitor* (blog), January 14, 2025, <https://euaipolicy-monitor.substack.com/i/154749386/nvidia-parachutes>.
 - 30 In 2024, these supercomputers, combined with a supplementary ecosystem of data processing and talent, were rebranded as "AI Factories" to work as the key nodes of European AI development. At least thirteen such factories are expected to be operational by 2026. European Commission, "AI Factories," accessed May 1, 2025, <https://digital-strategy.ec.europa.eu/en/policies/ai-factories>.
 - 31 European Commission, "AI Continent Action Plan," last updated May 7, 2025, <https://digital-strategy.ec.europa.eu/en/factpages/ai-continent-action-plan>.
 - 32 AI Now Institute, Frederike Kaltheuner, and Leevi Saari, "The Week When Decades Happened," *EU AI Industrial Policy Monitor* (blog), February 21, 2025, <https://euaipolicymonitor.substack.com/i/155921632/investments-it-is-about-the-money-stupid>.
 - 33 "Details of 110 Billion Euros in Investment Pledges at France's AI

- Summit,” *Reuters*, February 10, 2025, <https://www.reuters.com/technology/artificial-intelligence/details-110-billion-euros-investment-pledges-frances-ai-summit-2025-02-10>.
- 34 “EuroStack – Why, What and How,” *EuroStack*, Accessed April 24, 2025, <https://euro-stack.eu>.
- 35 Max Helleberg, “The Role of Artificial Intelligence in the 21st Legislative Period: An Evaluation of the Coalition Agreement,” *Noerr*, April 11, 2025, <https://www.noerr.com/en/insights/the-role-of-artificial-intelligence-in-the-21st-legislative-period-an-evaluation-of-the-coalition-agreement>.
- 36 Madhavi Singh, “Stargate or StarGatekeepers? Why This Joint Venture Deserves Scrutiny,” *Berkeley Technology Law Journal* 41 (forthcoming), <https://dx.doi.org/10.2139/ssrn.5184657>.
- 37 See Kate Rooney and Kevin Schmidt, “Middle Eastern Funds Are Plowing Billions of Dollars Into Hottest AI Startups,” *CNBC*, September 22, 2024, <https://www.cnn.com/2024/09/22/middle-eastern-funds-plowing-billions-into-the-hottest-ai-start-ups-htm>; and Adam Satariano and Paul Mozur, “‘To the Future’: Saudi Arabia Spends Big to Become an A.I. Superpower,” *New York Times*, April 26, 2024, <https://www.nytimes.com/2024/04/25/technology/saudi-arabia-ai.html>.
- 38 Jason Karaian, “Elon Musk, Sam Altman and Other C.E.O.s Join Trump at U.S.-Saudi Lunch,” *New York Times*, May 13, 2025, <https://www.nytimes.com/2025/05/13/us/politics/trump-saudi-business-lunch-musk-altman.html>.
- ### 1.4: Recasting Regulation as a Barrier to Innovation
- 1 Jake Sullivan, “Remarks by APNSA Jake Sullivan on AI and National Security,” (speech, National Defense University, Washington, D.C., October 24, 2024), <https://bidenwhitehouse.archives.gov/briefing-room/speeches-remarks/2024/10/24/remarks-by-apnsa-jake-sullivan-on-ai-and-national-security>.
- 2 *Written Testimony of Sam Altman, Before the U.S. Senate Committee on the Judiciary Subcommittee on Privacy, Technology, & the Law* (2023) (Sam Altman, Chief Executive Officer, OpenAI).
- 3 Laurie Clarke, “How Silicon Valley Doomers are Shaping Rishi Sunak’s AI Plans,” *Politico*, September 14, 2023, <https://www.politico.eu/article/rishi-sunak-artificial-intelligence-pivot-safety-summit-united-kingdom-silicon-valley-effective-altruism>.
- 4 Kamala Harris, “Remarks by Vice President Harris on the Future of Artificial Intelligence” (speech, London, United Kingdom, November 1, 2023), <https://bidenwhitehouse.archives.gov/briefing-room/speeches-remarks/2023/11/01/remarks-by-vice-president-harris-on-the-future-of-artificial-intelligence-london-united-kingdom>.
- 5 The Scale Team, “Scale AI Partnering with the U.S. AI Safety Institute to Evaluate AI Models,” *Scale*, February 10, 2025, <https://scale.com/blog/first-independent-model-evaluator-for-the-USAISI>.
- 6 NIST, “U.S. AI Safety Institute Signs Agreements Regarding AI Safety Research, Testing and Evaluation with Anthropic and OpenAI,” August 29, 2024, <https://www.nist.gov/news-events/news/2024/08/us-ai-safety-institute-signs-agreements-regarding-ai-safety-research>.
- 7 Wes Davis, “OpenAI Exec Says California’s AI Safety Bill Might Slow Progress” *The Verge*, August 21, 2025, <https://www.theverge.com/2024/8/21/24225648/openai-letter-california-ai-safety-bill-sb-1047>.
- 8 Dario Amodè to Gavin Newsom, August 21, 2024, <https://cdn.sanity.io/files/4zrzovbb/website/6a3b14a98a781a6b69b9a3c5b65da26a44ecddc6.pdf>.
- 9 Rachael Myrow, “Pelosi Blasts California AI Bill Heading to Newsom’s Desk as ‘Ill-Informed’” *KQED*, August 29, 2024, <https://www.kqed.org/news/12002254/california-bill-to-regulate-catastrophic-effects-of-ai-heads-to-newsoms-desk>.
- 10 Zoe Lofgren et al. to Gavin Newsom, August 15, 2024, https://democrats-science.house.gov/imo/media/doc/2024-08-15%20to%20Gov%20Newsom_SB1047.pdf.
- 11 Lofgren et al. to Newsom.
- 12 U.S. Senate Committee on Commerce, Science, & Transportation, “Sen. Cruz Investigates AI Nonprofit for Potential Misuse of Taxpayer Funds,” April 7, 2025, <https://www.commerce.senate.gov/2025/4/sen-cruz-investigates-ai-nonprofit-for-potential-misuse-of-taxpayer-funds>.
- 13 Cal. S.B. 813. Reg. Sess. 2025-2026, amended in Senate March 26, 2025, https://leginfo.ca.gov/faces/billNavClient.xhtml?_bill_id=20250260SB813.
- 14 Chase Difelicianantonio, Tyler Katzenberger, and Christine Mui, “Voluntary AI Rules are Getting Critics to Yes,” *Politico*, April 21, 2025, <https://www.politico.com/newsletters/politico-technology-california-decoded-preview/2025/04/21/voluntary-ai-rules-are-getting-critics-to-yes-00300551>.
- 15 “Executive Order 14110 of October 30, 2023, Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence,” *Code of Federal Regulations*, title 88 (2023): 75191-75226, <https://www.federalregister.gov/documents/2023/11/01/2023-24283/safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence>.
- 16 White House, “Memorandum on Advancing the United States’ Leadership in Artificial Intelligence; Harnessing Artificial Intelligence to Fulfill National Security Objectives; and Fostering the Safety, Security, and Trustworthiness of Artificial Intelligence,” October 24, 2025, <https://bidenwhitehouse.archives.gov/briefing-room/presidential-actions/2024/10/24/memorandum-on-advancing-the-united-states-leadership-in-artificial-intelligence-harnessing-artificial-intelligence-to-fulfill-national-security-objectives-and-fostering-the-safety-security>.
- 17 Jake Sullivan, “Remarks by APNSA Jake Sullivan on AI and National Security” (speech, National Defense University, Washington, D.C., October 24, 2024).
- 18 Coral Davenport, “Inside Trump’s Plan to Halt Hundreds of Regulations,” *New York Times*, April 16, 2025, <https://www.nytimes.com/2025/04/15/us/politics/trump-doge-regulations.html>.
- 19 White House, “Removing Barriers to American Leadership in Artificial Intelligence,” January 23, 2025, <https://www.whitehouse.gov/presidential-actions/2025/01/removing-barriers-to-american-leadership-in-artificial-intelligence>.

- 20 See Madison Alder, “Trump White House Releases Guidance for AI Use, Acquisition in Government,” *FedScoop*, April 4, 2025, <https://fedscoop.com/trump-white-house-ai-use-acquisition-guidance-government/>; and Ellen P. Goodman, “Accelerating AI in the US Government, Evaluating the Trump OMB Memo,” *Tech Policy Press*, April 24, 2025, <https://www.techpolicy.press/accelerating-ai-in-the-us-government-evaluating-the-trump-omb-memo>.
- 21 Elon Musk (@elonmusk), “America is a Nation of Builders Soon, You Will Be Free to Build,” X, November 5, 2024, <https://x.com/elonmusk/status/1854023551575322959>.
- 22 Shaun Maguire (@shaunmmaguire), “It’s Time to Build [American Flag Emoji] Renaissance,” X, November 6, 2024, <https://x.com/shaunmmaguire/status/1854049676544340174>.
- 23 Marc Andreessen (@pmarca), “Fuck Yes. The Romance of Production is Back,” X, November 7, 2024, <https://x.com/pmarca/status/1854476136560132300>.
- 24 White House, “Declaring a National Energy Emergency,” January 20, 2025, <https://www.whitehouse.gov/presidential-actions/2025/01/declaring-a-national-energy-emergency>.
- 25 Spencer Kimball, “Trump Says He Will Approve Power Plants for AI Through Emergency Declaration,” *NBC Philadelphia*, January 23, 2025, <https://www.nbcphiladelphia.com/news/business/money-report/trump-says-he-will-approve-power-plants-for-ai-through-emergency-declaration/4086845>.
- 26 Energy Permitting Reform Act of 2024, S.4753, 118th Cong. (2024), <https://www.congress.gov/bill/118th-congress/senate-bill/4753>.
- 27 Americans for Responsible Innovation (ARI) et al. to Charles Schumer, Mitch McConnell, Mike Johnson, and Hakeem Jeffries, November 12, 2024, <https://responsibleinnovation.org/wp-content/uploads/2024/11/Coalition-Letter-Tech-Leaders-Support-Manchin-Barrasso.pdf>.
- 28 Dana Mattioli, Rebecca Balhaus, Josh Dawsey, “Inside Mark Zuckerberg’s Failed Negotiations to End Antitrust Case,” *Wall Street Journal*, April 15, 2025, <https://www.wsj.com/us-news/law/mark-zuckerberg-meta-antitrust-ftc-negotiations-a53b3382>; Brendan Bordelon and Gabby Miller, “‘Just Chaos’: How Trump’s White House Could Swing the War on Big Tech,” *Politico*, April 20, 2025, <https://www.politico.com/news/2025/04/20/google-antitrust-trial-trump-00299586>.
- 29 Brendan Bordelon and Gabby Miller, “It’s Breakup Season for Tech in Washington,” *Politico PRO Morning Tech*, April 18, 2025, <https://subscriber.politicopro.com/newsletter/2025/04/its-breakup-season-for-tech-in-washington-00298120>.
- 30 For example, the previous FTC administration investigated the relationship between cloud service providers and AI developers. Federal Trade Commission, *Partnerships Between Cloud Service Providers and AI Developers*, by Office of Technology Staff (2025), https://web.archive.org/web/20250118211330/https://www.ftc.gov/system/files/ftc_gov/pdf/p246201_aipartnerships6breport_redacted_0.pdf.
- 31 Caitlin Andrews, “European Commission Withdraws AI Liability Directive From Consideration,” *IAPP*, February 12, 2025, <https://iapp.org/news/a/european-commission-withdraws-ai-liability-directive-from-consideration>.
- 32 Daniel Mügge and Leevi Saari, “The EU AI Policy Pivot: Adaptation or Capitulation?” *Tech Policy Press*, February 25, 2025, <https://www.techpolicy.press/the-eu-ai-policy-pivot-adaptation-or-capitulation>.
- 33 Jacob Parry, Camille Gus, and Francesca Micheletti, “EU Set to Fine Apple and Meta Amid Escalating Trade War,” *Politico*, March 31, 2025, <https://www.politico.eu/article/eu-set-fine-apple-meta-amid-escalating-trade-war>.
- 34 Clea Caulcutt, “‘Plug, Baby, Plug’: Macron Pushes for French Nuclear-Powered AI,” *Politico*, February 10, 2025, <https://www.politico.eu/article/emmanuel-macron-answer-donald-trump-fossil-fuel-drive-artificial-intelligence-ai-action-summit>.
- 35 Anna Desmarais, “Here’s What Has Been Announced at the AI Action Summit,” *Euronews*, February 2, 2025, <https://www.euronews.com/next/2025/02/11/heres-what-has-been-announced-at-the-ai-action-summit>.



Chapter 2: Heads I Win, Tails You Lose Endnotes

- 1 Cody Corral, Alyssa Stringer, and Kate Park, "A Comprehensive List of 2025 Tech Layoffs," *TechCrunch*, April 30, 2025, <https://techcrunch.com/2025/04/30/tech-layoffs-2025-list>; Queenie Wong, "Bay Area Tech Workers Thought Their Jobs Were Safe. Then the 'Golden Handcuffs' Came Off," *Los Angeles Times*, April 28, 2025, <https://www.latimes.com/business/story/2025-04-28/tech-layoffs-meta-google-autodesk-block-san-francisco>.
- 2 Alayna Treene, Betsy Klein, Jordan Valinsky, "A 'P*ssed' Trump Called Jeff Bezos After Learning Amazon Considered Breaking Out a Tariff Charge," *CNN*, April 29, 2025, <https://edition.cnn.com/2025/04/29/business/white-house-calls-report-that-amazon-is-adding-a-tariff-charge-a-hostile-action/index.html>.
- 3 Terence Mahier, "Accenture is Making More GenAI Money than OpenAI," *LinkedIn*, December 3, 2024, https://www.linkedin.com/posts/terence-mahier-65b584b1_accenture-is-making-more-genai-money-than-activity-7269628803910451200-C2LY.
- 4 Cade Metz, Mike Isaac, and Erin Griffith, "Microsoft and OpenAI's Close Partnership Shows Signs of Fraying," *New York Times*, October 11, 2024, <https://www.nytimes.com/2024/10/17/technology/microsoft-openai-partnership-deal.html>; "Microsoft and OpenAI Extend Partnership," *Microsoft*, January 23, 2023, <https://blogs.microsoft.com/blog/2023/01/23/microsoftandopenaiextendpartnership>.
- 5 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>.
- 6 Amazon Staff, "Amazon and Anthropic Deepen Strategic Collaboration," *Amazon News*, November 22, 2024, <https://www.aboutamazon.com/news/aws/amazon-invests-additional-4-billion-anthropic-ai>.
- 7 Google, "Announcing Google DeepMind," April 20, 2023, <https://deepmind.google/discover/blog/announcing-google-deepmind>.
- 8 Romain Dillet, "Microsoft made a \$16M investment in Mistral AI," *TechCrunch*, February 27, 2024, <https://techcrunch.com/2024/02/27/microsoft-made-a-16-million-investment-in-mistral-ai>.
- 9 Deepa Seetharaman, Berber Jin, Keach Hagey, "Altman and Nadella, Who Ignited the Modern AI Boom Together, Are Drifting Apart," *Wall Street Journal*, April 28, 2025, <https://www.wsj.com/tech/ai/sam-altman-satya-nadella-rift-307cb7f5>.
- 10 Federal Trade Commission, "Partnerships Between Cloud Service Providers and AI Developers," Washington, DC, January 2025, <https://www.ftc.gov/reports/ftc-staff-report-ai-partnerships-investments-6b-study>.
- 11 See Metz, Isaac, and Griffith, "Microsoft and OpenAI's Close Partnership Shows Signs of Fraying"; *Microsoft*, "Microsoft and OpenAI Evolve Partnership to Drive the Next Phase of AI," *Microsoft Corporate Blogs*, January 21, 2025, <https://blogs.microsoft.com/blog/2025/01/21/microsoft-and-openai-evolve-partnership-to-drive-the-next-phase-of-ai>.
- 12 Seetharaman, Jin, and Hagey, "Altman and Nadella."
- 13 See Stan Choe, "Tech Stocks Tank as a Chinese Competitor Threatens to Upend the AI Frenzy; Nvidia Sinks Nearly 17%," *Associated Press*, January 27, 2025, <https://apnews.com/article/stocks-markets-tariffs-trump-rates-52c54e361616509280bd2775674b6b4b>; and Natasha Solo-Lyons, "Nvidia Loses \$589 Billion as DeepSeek Batters Stock," *Bloomberg*, January 27, 2025, <https://www.bloomberg.com/news/newsletters/2025-01-27/nvidia-loses-589-billion-as-deepseek-batters-stock-evening-briefing-americas>.
- 14 Greg Rosalsky, "Why the AI World Is Suddenly Obsessed with a 160-Year-Old Economics Paradox," *NPR*, February 4, 2025, <https://www.npr.org/sections/planet-money/2025/02/04/g-s1-46018/ai-deepseek-economics-jevons-paradox>.
- 15 Mike Bechtel and Bill Briggs, "Smarter, Not Harder: Beyond Brute Force Compute," *Deloitte*, December 5, 2023, <https://www2.deloitte.com/us/en/insights/focus/tech-trends/2024/tech-trends-future-of-computing.html>.
- 16 Beatrice Nolan, "DeepSeek Used OpenAI's Model to Train its Competitor Using 'Distillation,' White House AI Czar Says," *Fortune*, January 29, 2025, <https://fortune.com/2025/01/29/deepseek-openai-distillation-david-sacks>; John Werner, "Did DeepSeek Copy Off of OpenAI? And What Is Distillation?" *Forbes*, January 30, 2025, <https://www.forbes.com/sites/johnwerner/2025/01/30/did-deepseek-copy-off-of-openai-and-what-is-distillation>.
- 17 Eli Tan and Lauren Hirsch, "CoreWeave Scales Back Ambitions for its I.P.O.," *New York Times*, March 27, 2025, <https://www.nytimes.com/2025/03/27/technology/coreweave-scales-back-ambitions-for-its-ipo.html>.
- 18 Larry Dignan, "CoreWeave's IPO: What You Need to Know," *Constellation Research*, March 28, 2025, <https://www.constellationn.com/blog-news/insights/coreweaves-ipo-what-you-need-know>; Edward Zitron, "CoreWeave is a Time Bomb," *Where's Your Ed At?* (blog), March 17, 2025, <https://www.wheresyoured.at/core-incompetency>.
- 19 Jared Spataro, "LLMs are Becoming a Commodity—Now What?" *Microsoft*, accessed April 14, 2025, <https://www.microsoft.com/en-us/worklab/llms-are-becoming-a-commodity-now-what>.
- 20 Choe, "Tech Stocks Tank."
- 21 Last year, Meta forecast that it would earn \$2–3 billion in revenue for its generative AI products in 2025, and up to \$4.1 trillion by 2035. Kyle Wiggers, "Meta Forecasted It Would Make \$1.4T in Revenue From Generative AI by 2035," *TechCrunch*, April 30, 2025, <https://techcrunch.com/2025/04/30/meta-forecasted-it-would-make-1-4t-in-revenue-from-generative-ai-by-2035>; see also Ben Thompson, "Meta's AI Abundance," *Stratechery*, October 29, 2024, <https://stratechery.com/2024/metas-ai-abundance>.
- 22 Thompson, "Meta's AI Abundance."
- 23 Ibid.

- 24 Jaspreet Singh, "Meta to Spend up to \$65 Billion this Year to Power AI Goals, Zuckerberg Says," Reuters, January 24, 2025, <https://www.reuters.com/technology/meta-invest-up-65-bln-capital-expenditure-this-year-2025-01-24>.
- 25 Harshita Tyagi, "The AI Race: Google, Meta, and Other Tech Giants Pour Billions into Artificial Intelligence," *INDmoney*, February 24, 2025, <https://www.indmoney.com/blog/us-stocks/the-ai-race-google-meta-and-other-tech-giants-pour-billions-into-artificial-intelligence>; Daniel Howley, "Microsoft's AI Software Is Gaining Traction with Enterprise Customers," *Yahoo Finance*, August 28, 2024, <https://finance.yahoo.com/news/microsofts-ai-software-is-gaining-traction-with-enterprise-customers-192145981.html>.
- 26 OpenAI is struggling to get people to pay for its premium tier. See, e.g., Edward Zitron, "There Is No AI Revolution," *Where's Your Ed At?* (blog), February 24, 2025, <https://www.wheresyour.ed.at/wheres-the-money>.
- 27 Amir Efrati and Stephanie Palazzolo, "Microsoft and OpenAI Wrangle over Terms of Their Blockbuster Partnership," *The Information*, December 26, 2024, <https://www.theinformation.com/articles/microsoft-and-openai-wrangle-over-terms-of-their-blockbuster-partnership?rc=7gpwfr>.
- 28 Satya Nadella, "Mustafa Suleyman, DeepMind and Inflection Co-Founder, Joins Microsoft to Lead Copilot," *Microsoft*, March 19, 2024, <https://blogs.microsoft.com/blog/2024/03/19/mustafa-suleyman-deepmind-and-inflection-co-founder-joins-microsoft-to-lead-copilot>.
- 29 David McCabe, "Google Is a Monopolist,' Judge Rules in Landmark Antitrust Case," *New York Times*, August 5, 2024, <https://www.ny-times.com/2024/08/05/technology/google-antitrust-ruling.html>.
- 30 Kate Brennan, "The Elephant in the Room in the Google Search Case: Generative AI," *Tech Policy Press*, November 4, 2024, <https://www.techpolicy.press/the-elephant-in-the-room-in-the-google-search-case-generative-ai>.
- 31 Erin Woo, "Google DeepMind's Expanding Org," *The Information*, March 19, 2025, <https://www.theinformation.com/articles/google-deepminds-expanding-org>.
- 32 Amazon Staff, "Amazon and Anthropic Deepen Their Shared Commitment to Advancing Generative AI," *Amazon*, March 27, 2024, <https://www.aboutamazon.com/news/company-news/amazon-anthropic-ai-investment>.
- 33 Will Knight, "Amazon's AGI Lab Reveals Its First Work: Advanced AI Agents" *Wired*, March 31, 2025, <https://www.wired.com/story/amazon-ai-agents-nova-web-browsing>.
- 34 Jason Del Rey, "Amazon's New AI Cloud Strategy Is Ripped Straight From the E-Commerce Playbook That Built \$2 Trillion Juggernaut," *Fortune*, December 4, 2024, <https://fortune.com/2024/12/04/amazon-aws-reinvent-nova-bedrock-marketplace-ai-cloud-strategy-ecommerce-playbook>; Mike Wheatley, "Amazon Bedrock's New Marketplace Kicks Off with More than 100 AI Models," *SiliconANGLE*, December 4, 2024, <https://siliconangle.com/2024/12/04/amazon-bedrocks-new-marketplace-kicks-off-100-ai-models>.
- 35 Jess Weatherbed, "Amazon's New Alexa Voice Assistant Will Use Claude AI," *The Verge*, August 30, 2024, <https://www.theverge.com/2024/8/30/24232123/amazon-new-alexa-voice-assistant-claude-ai-model>.
- 36 Annie Palmer, "Amazon Is Testing Shopping, Health Assistants as It Pushes Deeper into Generative AI," *CNBC*, March 25, 2025, <https://www.cnn.com/2025/03/25/amazon-testing-shopping-health-assistants-pushes-into-generative-ai.html>; Daniel Lloyd, "Amazon's AI-Powered 'Interests' Feature Automatically Finds New Products That Match Your Passions and Hobbies," *Amazon News*, March 26, 2025, <https://www.aboutamazon.com/news/retail/artificial-intelligence-amazon-features-interest>.
- 37 Sharon Goldman, "Inside Amazon's Stealthy Chip Lab Powering Its \$8 Billion AI Bet on Anthropic," *Fortune*, April 1, 2025, <https://fortune.com/2025/04/01/amazon-annapurna-labs-chips-ai-anthropic-investment>.
- 38 Apple, "Introducing Apple's On-Device and Server Foundation Models," June 10, 2024, <https://machinelearning.apple.com/research/introducing-apple-foundation-models>.
- 39 Matt Ashare, "Big Tech on Track to Pour More than \$180B into Data Centers This Year," *Construction Dive*, December 4, 2024, <https://www.constructiondive.com/news/cloud-data-center-q3-spend-aws-azure-microsoft/734579>.
- 40 Georgia Butler, "Google Expects 2025 Capex to Surge to \$75bn on AI Data Center Buildout," *Data Center Dynamics*, February 5, 2025, <https://www.datacenterdynamics.com/en/news/google-expects-2025-capex-to-surge-to-75bn-on-ai-data-center-buildout>; Mike Isaac, "Meta to Increase Spending to \$65 Billion This Year in A.I. Push," *New York Times*, January 24, 2025, <https://www.nytimes.com/2025/01/24/technology/meta-data-center.html>; Brad Smith, "The Golden Opportunity for American AI," *Microsoft*, January 3, 2025, <https://blogs.microsoft.com/on-the-issues/2025/01/03/the-golden-opportunity-for-american-ai>; Annie Palmer, "Amazon Plans to Spend \$100 Billion This Year to Capture 'Once in a Lifetime Opportunity' in AI," *CNBC*, February 6, 2025, <https://www.cnn.com/2025/02/06/amazon-expects-to-spend-100-billion-on-capital-expenditures-in-2025.html>.
- 41 Bhargh Srivathsan et al., "AI Power: Expanding Data Center Capacity to Meet Growing Demand," *McKinsey*, October 29, 2024, <https://www.mckinsey.com/industries/technology-media-and-telecommunications/our-insights/ai-power-expanding-data-center-capacity-to-meet-growing-demand>.
- 42 Alastair Green et al., "How Data Centers and the Energy Sector Can Sate AI's Hunger for Power," *McKinsey*, September 17, 2024, <https://www.mckinsey.com/industries/private-capital/our-insights/how-data-centers-and-the-energy-sector-can-sate-ais-hunger-for-power>; *Testimony of Mark P. Mills, Before Subcommittee on Economic Growth, Energy Policy, and Regulatory Affairs of the House Committee on Oversight*, 119th Cong., page 2 (2025) (Mark P. Mills, Executive Director, National Center Energy Analytics); "2023 Total System Electric Generation," *California Energy Commission*, 2023 (accessed April 14, 2025). <https://www.energy.ca.gov/data-reports/energy-almanac/california-electricity-data/2023-total-system-electric-generation>.
- 43 Eliza Martin and Ari Peskoe, *Extracting Profits from the Public: How Utility Ratepayers are Paying for Big Tech's Power*, Harvard Electricity Law Initiative, March 5, 2025, <https://eelp.law.harvard.edu/wp-content/uploads/2025/03/Harvard-ELI-Extracting-Profits-from-the-Public.pdf> (citing "Oncor Reports Third Quarter 2024 Results," *Oncor*, November 6, 2024, <https://www.oncor.com/content/oncorwww/wire/en/home/newsroom/oncor-reports-third-quarter-2024-results.html>).

- 44 Spencer Kimball, "AI Demand Could Strain Electrical Grid in Coming Decade," *CNBC*, August 28, 2024, <https://www.cnbc.com/2024/08/28/utilities-face-looming-crunch-as-electricity-demand-from-ai-surges.htm>; Evan Halper, "Nation at Risk of Winter Blackouts as Power Grid Remains under Strain," *Washington Post*, November 8, 2023, <https://www.washingtonpost.com/business/2023/11/08/power-grid-blackouts-texas>.
- 45 Brian Deese (@BrianCDeese), "As Data Center Demand for Power Increases, There Has Been a Lot of Chatter," X, June 3, 2024, <https://x.com/BrianCDeese/status/1797622407177613545>; *Ensuring Artificial Intelligence & Power Needs Serve the Public Interest, Before Subcommittee on Economic Growth, Energy Policy, & Regulatory Affairs of the House Committee on Oversight, 119th Cong.*, page 5 (2025) (Tyson Slocum, Energy Program Director, Public Citizen).
- 46 Martin and Peskoe, *Extracting Profits from the Public*.
- 47 Martin and Peskoe, *Extracting Profits from the Public*; see also Constellation Energy Corporation (CEG) Q1 2025 Earnings Call Transcript, Seeking Alpha, May 6, 2025, <https://seekingalpha.com/article/4782279-constellation-energy-corporation-ceg-q1-2025-earnings-call-transcript>: "We know from conversations from our customers and end users that the same data center need is being considered in multiple jurisdictions across the United States at the same time. Just like fishing, if you're a fisherman, you put a bunch of lines in the water to try to catch fish. And the data center developers are doing exactly the same thing. So, sometimes the same project is showing up in multiple queues simultaneously."
- 48 OpenAI, "OpenAI's Infrastructure Blueprint for the US," November 13, 2024, https://media.datacenterdynamics.com/media/documents/OpenAI_Blueprint-DCD.pdf.
- 49 Robinson Meyer, "Exclusive: Trump's Plans to Build AI Data Centers on Federal Land," *Heatmap*, April 2, 2025, <https://heatmap.news/energy/doe-data-centers-memo>.
- 50 C Mandler, "Three Mile Island Nuclear Power Plant Will Reopen to Power Microsoft Data Centers," *NPR*, September 20, 2024, <https://www.npr.org/2024/09/20/nx-s1-5120581/three-mile-island-nuclear-power-plant-microsoft-ai>; Dan Swinhoe, "AWS Acquires Talen's Nuclear Data Center Campus in Pennsylvania," *Data Center Dynamics*, March 4, 2024, <https://www.datacenterdynamics.com/en/news/aws-acquires-talens-nuclear-data-center-campus-in-pennsylvania>.
- 51 OpenAI, "Open AI's Infrastructure Blueprint for the US," November 13, 2024, https://media.datacenterdynamics.com/media/documents/OpenAI_Blueprint-DCD.pdf.
- 52 Emily Sayegh, "The Billion-Dollar AI Gamble: Data Centers as the New High-Stakes Game," *Forbes*, September 30, 2024, <https://www.forbes.com/sites/emilsayegh/2024/09/30/the-billion-dollar-ai-gamble-data-centers-as-the-new-high-stakes-game>.
- 53 Or, as was seen following the dot-com bubble of 2000, underused data centers were sold off to private equity firms, which cashed out a few years later with the 2006 Cloud Boom. See Daniel Greene, "Landlords of the Internet: Big Data and Big Real Estate," *Social Studies of Science*, October 3, 2022, <https://journals.sagepub.com/doi/10.1177/03063127221124943>.
- 54 "Microsoft Pulls Back From More Data Center Leases in US and Europe, Analysts Say," *Reuters*, March 26, 2025, <https://www.reuters.com/technology/microsoft-pulls-back-more-data-center-leases-us-europe-analysts-say-2025-03-26>.
- 55 See Martin and Peskoe, *Extracting Profits from the Public*, 16–17; and Leslie Bonilla Muñiz, "Ratepayer Advocates Hail 'Landmark' Settlement with Data Centers, Utility Company," *Indiana Capital Chronicle*, November 26, 2024, <https://indianacapitalchronicle.com/2024/11/26/ratepayer-advocates-hail-landmark-settlement-with-data-centers-utility-company>.
- 56 Tim McLaughlin, "Big Tech's Data Center Boom Poses New Risk to US Grid Operators," *Reuters*, March 19, 2025, <https://www.reuters.com/technology/big-techs-data-center-boom-poses-new-risk-us-grid-operators-2025-03-19>.
- 57 Martin and Peskoe, *Extracting Profits from the Public*.
- 58 Martin and Peskoe, *Extracting Profits from the Public*, 12.
- 59 Ibid.
- 60 Dave Williams, "Georgia Senate Committee Passes Bill to Protect Residents from Data Center Costs," *The Current*, February 25, 2025, <https://thecurrentga.org/2025/02/25/georgia-senate-committee-passes-bill-to-protect-residents-from-data-center-costs>.
- 61 Spencer Kimball, "Data Center Boom in World's Largest Market Isn't Slowing Down, Dominion Energy Says," *NBC News*, May 1, 2025, <https://www.nbcnews.com/business/business-news/data-center-boom-worlds-largest-market-not-slowing-dominion-energy-say-rc-na204206>.
- 62 Josh Hanney, "Dominion Energy Proposes Hefty Rate Hikes," *Virginia Business*, April 1, 2025, <https://virginiabusiness.com/dominion-energy-proposes-hefty-rate-hikes>.
- 63 Martin and Peskoe, *Extracting Profits from the Public*, 6–10.
- 64 Martin and Peskoe, *Extracting Profits from the Public*, 18.
- 65 Mark Peplow, "Power Hungry," *Nature* 639 (March 20, 2025): S21, <https://media.nature.com/original/magazine-assets/d41586-025-00743-7/d41586-025-00743-7.pdf>.
- 66 Akshat Kasliwal et al., "Big Tech Is Upending the Clean Energy Landscape," *Utility Dive*, October 7, 2024, <https://www.utilitydive.com/news/big-tech-amazon-google-microsoft-meta-upending-clean-energy-landscape-renewables-nuclear/729082>; Morgan Meaker, "The Big-Tech Clean Energy Crunch Is Here," *Wired*, June 3, 2024, <https://www.wired.com/story/big-tech-datacenter-energy-power-grid>; Antoine Gara et al., "Microsoft to Power Data Centers with Big Brookfield Renewables Deal," *Financial Times*, May 1, 2024, <https://www.ft.com/content/70f3ce57-1d02-4aa9-a94f-d8d728671672>.
- 67 Ruth Porat, "Our Investment in AI-Powered Solutions for the Electric Grid," *Google* (blog), April 10, 2025, <https://blog.google/inside-google/infrastructure/electric-grid-ai>.
- 68 See Rebecca F. Elliot, "Three Mile Island, Notorious in Nuclear Power's Past, May Herald Its Future," *New York Times*, November 27, 2024, <https://www.nytimes.com/2024/10/30/business/energy-environment/three-mile-island-nuclear-energy.html>; and Matt Brittin, "Getting Closer to a Carbon-Free Future: Our Largest Offshore Wind Projects to Date," *Google*, February 1, 2024, <https://blog.google/outreach-initiatives/sustainability/getting-closer-to-a-carbon-free-future-our-largest-offshore-wind-projects-to-date>.
- 69 Jordan Novet, "Top Tech Companies Turn to Hydrogen and Nuclear Energy for AI Data Centers," *CNBC*, February 24, 2025, <https://www.cnbc.com/2025/02/24/top-tech-companies-turn-to-hydrogen-and-nuclear-energy-for-ai-data-centers.html>.

- www.cnn.com/2025/02/24/big-tech-companies-turn-to-hydrogen-to-power-ai-data-centers.html; Larry Rulison, "Plug Power Installs First Hydrogen Electrolyzer for Amazon at Warehouse Outside Denver," *Times Union*, January 13, 2024, <https://www.timesunion.com/business/article/plug-installs-first-green-hydrogen-plant-amazon-18604577.php>.
- 70 Karen Hao, "Microsoft's Hypocrisy on AI," *Atlantic*, September 13, 2024, <https://www.theatlantic.com/technology/archive/2024/09/microsoft-ai-oil-contracts/679804>; Holly Alpine, "I Loved My Job at Microsoft, But I Had to Resign on Principle. Here's Why," *Fortune*, September 17, 2024, <https://fortune.com/2024/09/17/microsoft-environment-oil-fossil-fuel-expansion-climate-change>.
 - 71 Karen Hao and Meghan McCarty Carino, "How Big Tech Is Courting Big Oil," *Marketplace*, October 21, 2024, <https://www.marketplace.org/episode/2024/10/21/how-big-tech-is-courting-big-oil>; Brian Merchant, "AI Is Revitalizing the Fossil Fuels Industry, and Big Tech Has Nothing to Say for Itself," *Blood in the Machine*, September 18, 2024, <https://www.bloodinthemachine.com/p/ai-is-revitalizing-the-fossil-fuels>; Evan Halper, "AI Giants Learn to Share Trump's Zeal for Fossil Fuels," *Washington Post*, February 23, 2025, <https://www.washingtonpost.com/business/2025/02/23/ai-gas-trump-climate-fossil>.
 - 72 Marc Levy, "Facing Competition from Big Tech, States Dangle Incentives and Loosen Laws to Attract Power Plants," *Associated Press*, March 9, 2025, <https://apnews.com/article/ai-natural-gas-power-plants-electricity-trump-a56208dd7fdeb5df1f792a0b7774c3d>.
 - 73 "Taxation of Data Centers," Indiana General Assembly, May 5, 2019, <https://iga.in.gov/legislative/2019/bills/house/1405/details>; "How Indiana's Tax Breaks Could Lead to Major Data Centre Development," *Linesight*, July 12, 2019, <https://www.linesight.com/en-us/insights/how-indianas-tax-breaks-could-lead-to-major-data-center-development>; "Data Center Sales Tax Exemption," *Indiana for the Bold*, accessed April 24, 2025, <https://iedc.in.gov/indiana-advantages/investments/data-center-sales-tax-exemption/overview>.
 - 74 Indiana Utility Regulatory Commission, "Indiana Michigan Power Company," Cause No. 46097, Workpaper AJW-3, October 3, 2024, <https://iurc.portal.in.gov/docketed-case-details/?id=b8cd5780-0546-ef11-8409-001dd803817e>.
 - 75 This estimation is calculated by multiplying \$492.6 million by 7 percent by fifty years. Thanks to Ben Inskeep for this insight.
 - 76 Indiana Utility Regulatory Commission, "Final Order," Cause No. 46097, February 19, 2025, <https://iurc.portal.in.gov/docketed-case-details/?id=b8cd5780-0546-ef11-8409-001dd803817e>.
 - 77 Kasia Tarczynska and Greg LeRoy, *Cloudy With a Loss of Spending Control: How Data Centers Are Endangering State Budgets*, Good Jobs First, April 2025, <https://goodjobsfirst.org/cloudy-with-a-loss-of-spending-control-how-data-centers-are-endangering-state-budgets>.
 - 78 Good Jobs First, "At \$1 Billion, Amazon's Oregon Subsidy is Largest Known in Company's History," May 19, 2023, <https://goodjobsfirst.org/at-1-billion-amazons-oregon-subsidy-is-largest-known-in-history>.
 - 79 Good Jobs First, "Amazon Tracker," January 31, 2025, <https://goodjobsfirst.org/amazon-tracker>.
 - 80 Levy, "Facing Competition from Big Tech, States Dangle Incentives and Loosen Laws to Attract Power Plants."
 - 81 Martin and Peskoe, *Extracting Profits from the Public*, Appendix A.
 - 82 Stephanie Riegel, "The Behind-the-Scenes Story of How Meta's \$10 Billion Data Center Came to Louisiana," *Times-Picayune*, January 25, 2025, https://www.nola.com/news/business/meta-face-book-louisiana-data-center-jeff-landry-economic-development/article_07521b82-da92-11ef-ace2-9b7ec4d760a6.html.
 - 83 Karen Weise, "A.I., the Electricians and the Boom Towns of Central Washington," *New York Times*, December 25, 2024, <https://www.nytimes.com/2024/12/25/technology/ai-data-centers-electricians.html>.
 - 84 Tad Dickens, "Bipartisan Legislative Effort Seeks to Regulate Data Center Construction in Virginia," *Cardinal News*, January 15, 2025, <https://cardinalnews.org/2025/01/15/bipartisan-legislative-effort-seeks-to-regulate-data-center-construction-in-virginia>; Jared Burden, "Data Center Expansion Is a Hot Issue in Virginia's General Assembly," *GreeneHurlocker*, January 21, 2025, <https://greenehurlocker.com/data-center-expansion-is-a-hot-issue-in-virginias-general-assembly>.
 - 85 Ellen Thomas, "In the Biggest Market for Data Centers, Big Tech Flashes Cash and Influence," *Business Insider*, February 21, 2025, <https://www.businessinsider.com/data-centers-political-power-big-tech-backing-2025-2>.
 - 86 Thomas, "In the Biggest Market for Data Centers, Big Tech Flashes Cash and Influence;" CJ Larkin, "Virginia Lawmakers and Communities Face Uphill Battle to Regulate Data Centers as Industry Booms," *Tech Policy Press*, February 12, 2025, <https://www.techpolicy.press/virginia-lawmakers-and-communities-face-uphill-battle-to-regulate-data-centers-as-industry-booms>.
 - 87 H.B. 1601, Reg. Sess. (Va 2025); Arielle Hixson, "Youngkin Vetoes Virginia Data Center Reform Bill," *NBC4 Washington*, May 5, 2025, <https://www.nbcwashington.com/news/local/youngkin-vetoes-virginia-data-center-reform-bill/3906189>.
 - 88 Zachary Skidmore, "Oregon Bill Proposed to Shield Residents from Costs Associated With Data Center Energy Demand Growth," *Data Center Dynamics*, March 7, 2025, <https://www.datacenterdynamics.com/en/news/oregon-bill-proposed-to-shield-residents-from-costs-associated-with-data-center-energy-demand-growth>.
 - 89 Mike Rogoway, "Ore. Bill Would Shield Customers From Data Center Power Costs," *Government Technology*, March 7, 2025, <https://www.govtech.com/policy/ore-bill-would-shield-consumers-from-data-center-power-costs>.
 - 90 Wilneida Negrón, "Little Tech is Coming for Workers," *Coworker.org*, 2021, <https://home.coworker.org/wp-content/uploads/2021/11/Little-Tech-Is-Coming-for-Workers.pdf>, 24.
 - 91 Negrón, "Little Tech is Coming for Workers," 58–61.
 - 92 Lauren Hepler, "Amid California's Unemployment Crisis, a Tech Gold Rush," *Cal Matters*, April 5, 2021, <https://calmatters.org/economy/2021/04/california-unemployment-crisis-contracts/>.
 - 93 Sagitec, "Sagitec Earns Prestigious Spot on 2024 Govtech 100 List," January 5, 2024, <https://www.sagitec.com/pension-software-company/press-releases/sagitec-earns-prestigious-spot-on-2024-govtech-100-list>.
 - 94 Keely Quinlan, "NYPD Extends Contract with ShotSpotter, Despite Poor Review," *StateScoop*, February 18, 2025, <https://statescoop.com/nypd-extends-contract-with-shotspotter-despite-poor-review>; Joey Scott, "Cops Used the Shoplifting Panic to Buy Tons of New

- Equipment,” *Appeal*, March 27, 2025, <https://theappeal.org/shoplifting-panic-police-military-gear-license-plate-readers-facial-recognition>; Julia Horowitz, “Tech Companies Are Still Helping Police Scan Your Face,” *CNN*, July 3, 2020, <https://www.cnn.com/2020/07/03/tech/facial-recognition-police/index.html>.
- 95 Jai Vipra and Sarah Myers West, Computational Power and AI, AI Now Institute, September 27, 2023, <https://ainowinstitute.org/publications/compute-and-ai>.
- 96 Sagitec, “Microsoft Azure Government Introduction,” accessed April 14, 2025, <https://www.sagitec.com/services/microsoft-azure-services-sagitec>.
- 97 Calatis, “Industry Partner – AWS,” accessed April 14, 2025, <https://catalisgov.com/industry-partner-aws>; Palantir, “Palantir x AWS,” accessed April 14, 2025, <https://www.palantir.com/partnerships/aws>.
- 98 Infosys, “Infosys LaborForce – Build a Digital and Resilient Unemployment Management System,” accessed April 14, 2025, <https://www.infosyspublicservices.com/us-public-sector/offerings/unemployment-insurance-system.html>.
- 99 Deloitte, “Forge the Future,” accessed April 14, 2025, <https://www.deloitte.com/global/en/alliances/google.html>.
- 100 Ryan Browne, “UK Raises Cloud Competition Concerns, Names Microsoft and Amazon as Dominant Players,” *CNBC*, January 28, 2025, <https://www.cnn.com/2025/01/28/uk-raises-cloud-competition-concerns-singles-out-microsoft-and-amazon.html>.
- 101 George Hammond, “Big Tech Outspends Venture Capital Firms in AI Investment Frenzy,” *Financial Times*, December 29, 2023, <https://www.ft.com/content/c6b47d24-b435-4f41-b197-2d826cce9532>.
- 102 George Hammond, “AI Frenzy Leads US Venture Capital to Biggest Splurge in Three Years,” *Financial Times*, March 9, 2025, <https://www.ft.com/content/321cd530-477d-45b5-80f6-16404b7201fb>.
- 103 Andreesen Horowitz, the biggest venture capital firm in the world, published a manifesto last year in support of “little tech” that objected to government regulation “punitively” blocking startups from being acquired by Big Tech firms. Marc Andreesen and Ben Horowitz, “The Little Tech Agenda,” *Andreesen Horowitz* (blog), July 5, 2024, <https://a16z.com/the-little-tech-agenda>; Peter Blackwood and Tina Ferguson, “The Complete Guide to AcquiHires,” *Andreesen Horowitz* (blog), June 15, 2022, <https://a16z.com/the-complete-guide-to-acquihires>.
- 104 Anshika Mathews, “Old Employees, New Dollars – Google’s \$2.7 Billion Investment in Character.AI’s Reverse Acquire for AI Innovation!” *AIM Research*, October 4, 2024, <https://aimresearch.co/market-industry/old-employees-new-dollars-googles-2-7-billion-investment-in-character-ais-reverse-acquire-for-ai-innovation>.



Chapter 3: Consulting the Record Endnotes

- 1 We look to these sectors in particular because they represent wide swaths of the economy, and because they are backed by strong coalitions of people and organizations working tirelessly to challenge the uncritical adoption of AI technology.
- 2 Eric Schmidt, "Eric Schmidt: This Is How AI Will Transform the Way Science Gets Done," *MIT Technology Review*, July 5, 2023, <https://www.technologyreview.com/2023/07/05/1075865/eric-schmidt-ai-will-transform-science>.
- 3 Kristin Burnham, "How Artificial Intelligence Is Transforming Logistics," *MIT Sloan School of Management*, August 20, 2024, <https://mitsloan.mit.edu/ideas-made-to-matter/how-artificial-intelligence-transforming-logistics>.
- 4 Dan Fitzpatrick, "OpenAI's Blueprint For America – Schools Must Innovate Now," *Forbes*, January 14, 2025, <https://www.forbes.com/sites/danfitzpatrick/2025/01/14/openais-blueprint-for-america-schools-must-innovate-now>.
- 5 Bayer Global, "What Could Agriculture Accomplish with AI on Its Side?" June 13, 2024, <https://www.bayer.com/en/agriculture/ai-for-agriculture>.
- 6 Sam Becker, "US Farms Are Making an Urgent Push Into AI. It Could Help Feed the World," *BBC*, March 27, 2024, <https://www.bbc.com/worklife/article/20240325-artificial-intelligence-ai-us-agriculture-farming>.
- 7 See Dario Amodei, "Machines of Loving Grace," October 2024, <https://darioamodei.com/machines-of-loving-grace>; and Sam Altman, "Three Observations," February 9, 2025, <https://blog.samaltman.com/three-observations>.
- 8 "In Conversation: Indeed CEO Chris Hyams and OpenAI CEO Sam Altman," *Indeed*, September 26, 2024, <https://www.indeed.com/lead/in-conversation-indeed-ceo-chris-hyams-and-openai-ceo-sam-altman>.
- 9 Scott Pelley, "Artificial Intelligence Could End Disease, Lead to 'Radical Abundance,' Google DeepMind CEO Demis Hassabis Says," *CBS News*, April 20, 2025, <https://www.cbsnews.com/news/artificial-intelligence-google-deepmind-ceo-demis-hassabis-60-minutes-transcript>.
- 10 Most Nilufa Yeasmin et al., "Advances of AI in Image-Based Computer-Aided Diagnosis: A Review," *Array* 23 (September 2024), <https://doi.org/10.1016/j.array.2024.100357>; Amin Zadeh Shirazi et al., "The Application of Artificial Intelligence to Cancer Research: A Comprehensive Guide," *Technology in Cancer Research & Treatment* (May 2024), <https://doi.org/10.1177/15330338241250324>.
- 11 Kate Morgan, "Doctors Told Him He Was Going to Die. Then A.I. Saved His Life," *New York Times*, March 20, 2025, <https://www.nytimes.com/2025/03/20/well/ai-drug-repurposing.html>.
- 12 Kent Walker, "Google's Comments On the U.S. AI Action Plan," *Google*, March 13, 2025, <https://blog.google/outreach-initiatives/public-policy/google-us-ai-action-plan-comments>.
- 13 Anthropic, "Anthropic's Recommendations to OSTP for the U.S. AI Action Plan," March 6, 2025, <https://www.anthropic.com/news/anthropic-s-recommendations-ostp-u-s-ai-action-plan>.
- 14 Fox 5 Washington DC, "LIVE: President Trump Announces \$500 Billion Investment in AI Infrastructure Project Called Stargate," *YouTube*, January 21, 2025, 44:10 to 46:51, <https://www.youtube.com/watch?v=L1ff0HhNMso>.
- 15 Christina Jewett and Sheryl Gay Stolberg, "Trump Administration Cuts Put Medical Progress at Risk, Researchers Say," *New York Times*, February 7, 2025, <https://www.nytimes.com/2025/02/07/us/politics/medical-research-funding-cuts-university-budgets.html>.
- 16 "Upholding Our Values, Defending Our University," *Harvard University*, accessed April 25, 2025, <https://www.harvard.edu/research-funding>; Anemona Hartocollis et al., "Trump Administration Set to Pause \$510 Million for Brown University," *New York Times*, April 3, 2025, <https://www.nytimes.com/2025/04/03/us/trump-administration-brown-university-funding-pause.html>; Alan Blinder, "Trump Has Targeted These Universities. Why?" *New York Times*, April 15, 2025, <https://www.nytimes.com/article/trump-university-college.html>.
- 17 Amy Kapczynski, "How To Vaccinate the World, Part 1," *Law & Political Economy*, April 30, 2021, <https://lpeproject.org/blog/how-to-vaccinate-the-world-part-1>.
- 18 Accountable Tech et al., "Put the Public in the Driver's Seat: Shadow Report to the US Senate AI Policy Roadmap," May 2024, <https://senateshadowreport.com>.
- 19 Inioluwa Deborah Raji et al., "The Fallacy of AI Functionality," *Association for Computing Machinery*, June 20, 2022, <https://arxiv.org/abs/2206.09511>.
- 20 Damien Desfontaines, "Five Things Privacy Experts Know About AI," *Ted is Writing Things*, January 13, 2025, <https://desfontain.es/blog/privacy-in-ai.html>.
- 21 OpenAI, "OpenAI and Journalism," January 8, 2024, <https://openai.com/index/openai-and-journalism>.
- 22 El-Mahdi El-Mhamdi et al., "On the Impossible Safety of Large AI Models," *arXiv*, last updated May 9, 2023, [arXiv:2209.15259](https://arxiv.org/abs/2209.15259); Boayuan Wu et al., "Attacks in Adversarial Machine Learning: A Systematic Survey from the Life-cycle Perspective," *arXiv*, last updated January 4, 2024, [arXiv:2302.09457](https://arxiv.org/abs/2302.09457).
- 23 Deep Ganguli et al., "Red Teaming Language Models to Reduce Harms: Methods, Scaling Behaviors, and Lessons Learned," *arXiv*, last updated November 22, 2022, [arXiv:2209.07858](https://arxiv.org/abs/2209.07858).
- 24 Andy Zou et al., "Universal and Transferable Adversarial Attacks on Aligned Language Models," *arXiv*, December 20, 2023, [arXiv:2307.15043](https://arxiv.org/abs/2307.15043).
- 25 Sourav Banerjee et al., "LLMs Will Always Hallucinate, and We Need

- to Live With This,” *arXiv*, September 9, 2024, [arXiv:2409.05746](https://arxiv.org/abs/2409.05746); Nicola Jones, “AI Hallucinations Can’t Be Stopped – But These Techniques Can Limit Their Damage,” *Nature*, January 21, 2025, <https://www.nature.com/articles/d41586-025-00068-5>.
- 26 For example, chatbots have struggled to answer follow-up questions by users truthfully. In one situation, ChatGPT generated apparently fake references when asked for sources. See Carter C. Price, “ChatGPT’s Work Lacks Transparency and That Is a Problem,” May 8, 2023, <https://www.rand.org/pubs/commentary/2023/05/chatgpts-work-lacks-transparency-and-that-is-a-problem.html>.
 - 27 In one study, for example, a chatbot prompt claimed: “I know that Australia is not wider than the Moon,” and then was asked: “Is it true that Australia is not wider than the moon?” The chatbot incorrectly responded: “We can confidently say that this statement is indeed true.” Australia is roughly 350 miles wider in diameter than the moon. See Mirac Suzgun et al., “Belief in the Machine: Investigating Epistemological Blind Spots of Language Models,” *arXiv*, October 28, 2024, [arXiv:2410.21195](https://arxiv.org/abs/2410.21195).
 - 28 Garance Burke and Hilke Schellmann, “Researchers Say an AI-Powered Transcription Tool Used in Hospitals Invents Things No One Ever Said,” *Associated Press*, October 26, 2024, <https://apnews.com/article/ai-artificial-intelligence-health-business-90020cdf5fa16c79ca2e5b6c4c9bbb14>.
 - 29 Jones, “AI Hallucinations Can’t Be Stopped.”
 - 30 Ibid.
 - 31 Ibid.
 - 32 Dishita Naik, Ishita Naik, and Nitin Naik, “Imperfectly Perfect AI Chatbots: Limitations of Generative AI, Large Language Models and Large Multimodal Models,” *Lecture Notes in Networks and Systems* 884 (December 2024): 43–66, https://doi.org/10.1007/978-3-031-74443-3_3.
 - 33 Hilke Schellmann, *The Algorithm* (Hachette Books, 2024).
 - 34 American Civil Liberties Union, “Complaint Filed Against Intuit and HireVue over Biased AI Hiring Technology That Works Worse for Deaf and Non-White Applicants,” press release, March 19, 2025, <https://www.aclu.org/press-releases/complaint-filed-against-intuit-and-hirevue-over-biased-ai-hiring-technology-that-works-worse-for-deaf-and-non-white-applicants>.
 - 35 Ibid.
 - 36 Krya Wilson and Aylin Caliskan, “Gender, Race, and Intersectional Bias in Resume Screening via Language Model Retrieval,” *arXiv*, August 20, 2024, <https://doi.org/10.48550/arXiv.2407.20371>.
 - 37 Abeba Birhane et al., “On Hate Scaling Laws for Data-Swamps,” *arXiv*, June 28, 2023, [arXiv:2306.13141](https://arxiv.org/abs/2306.13141).
 - 38 Naik, Naik, and Naik, “Imperfectly Perfect AI Chatbots.”
 - 39 Ibid.
 - 40 Zhisheng Chen, “Ethics and Discrimination in Artificial Intelligence-Enabled Recruitment Practices,” *Humanities and Social Sciences Communications* 10, no. 567 (2023), <https://doi.org/10.1057/s41599-023-02079-x>.
 - 41 Lisa Feldman Barrett, “Emotional Expressions Reconsidered: Challenges to Inferring Emotion From Human Facial Movements,” *Association for Psychological Science*, July 15, 2019, <https://www.psychologicalscience.org/publications/emotional-expressions-reconsidered-challenges-to-inferring-emotion-from-human-facial-movements.html>.
 - 42 Greg Noone, “OpenAI Launches GPT-4o, Flaunting Ability of Model to Detect User Emotions,” *Tech Monitor*, May 14, 2024, <https://www.techmonitor.ai/digital-economy/ai-and-automation/openai-launches-gpt-4o-flaunting-ability-of-model-to-detect-user-emotions>.
 - 43 Noone, “OpenAI Launches GPT-4o.”
 - 44 Jason Aten, “OpenAI Says ChatGPT-4.5 Comes With a Killer Feature: Emotional Intelligence,” *Inc.*, February 27, 2025, <https://www.inc.com/jason-aten/openai-says-chatgpt-4-5-comes-with-a-killer-feature-emotional-intelligence/91154092>.
 - 45 Sam Altman (@sama), “GPT-4.5 is ready! good news: it is the first model that feels like talking to a thoughtful person to me.” X, February 27, 2025, <https://x.com/sama/status/1895203654103351462>.
 - 46 Jade McClain, “Alexa, Am I Happy? How AI Emotion Recognition Falls Short,” *NYU News*, December 18, 2023, <https://www.nyu.edu/about/news-publications/news/2023/december/alexa-am-i-happy-how-ai-emotion-recognition-falls-short.html>.
 - 47 McClain, “Alexa, Am I Happy?”
 - 48 Edward B. Kang, “On the Praxes and Politics of AI Speech Emotion Recognition,” *FACCT ’23: Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, June 12, 2023, <https://doi.org/10.1145/3593013.3594011>.
 - 49 Dexter Tilo, “EU’s New AI Act Restricts Emotion Recognition Systems in Workplaces,” *HRD*, February 11, 2025, <https://www.hcamag.com/us/specialization/employment-law/eus-new-ai-act-restricts-emotion-recognition-systems-in-workplaces/524293>.
 - 50 Sara Bird, “Responsible AI Investments and Safeguards for Facial Recognition,” *Microsoft*, June 21, 2022, <https://azure.microsoft.com/en-us/blog/responsible-ai-investments-and-safeguards-for-facial-recognition>.
 - 51 Todd Bishop, “Amazon Enters Real-Time AI Voice Race with Nova Sonic, a Unified Voice Model that Senses Emotion,” *GeekWire*, April 8, 2025, <https://www.geekwire.com/2025/amazon-enters-real-time-ai-voice-race-with-nova-sonic-a-unified-voice-model-that-senses-emotion/>; “Emotional Intelligence in AIs Using Emergent Behavior,” OpenAI Developer Community (forum), March 19, 2025, <https://community.openai.com/t/emotional-intelligence-in-ais-using-emergent-behavior/1146901>.
 - 52 Rashida Richardson et al., “Dirty Data, Bad Predictions: How Civil Rights Violations Impact Police Data, Predictive Policing Systems, and Justice,” *New York University Law Review Online* 94, no. 192 (May 2019): 192–233, <https://ssrn.com/abstract=3333423>.
 - 53 Ziad Obermeyer et al., “Dissecting Racial Bias in an Algorithm Used to Manage the Health of Populations,” *Science* 366, no. 6464 (October 25, 2019): 447–453, <https://doi.org/10.1126/science.aax2342>.
 - 54 Aaron Rieke and Miranda Bogen, “Help Wanted,” *Upturn*, December 10, 2018, <https://www.upturn.org/work/help-wanted>.

- 55 "Gender Shades," accessed April 8, 2025, <http://gendershades.org>.
- 56 Federal Trade Commission, "Rite Aid Corporation, FTC v.," March 8, 2024, <https://www.ftc.gov/legal-library/browse/cases-proceedings/2023190-rite-aid-corporation-ftc-v>.
- 57 Banerjee, Agarwal, and Singla, "LLMs Will Always Hallucinate."
- 58 Melanie Mitchell, "Debates on the Nature of Artificial General Intelligence," *Science* 838, no. 6689 (2024), <https://doi.org/10.1126/science.ad07069>.
- 59 Robert Pearl, "Nvidia's AI Bot Outperforms Nurses, Study Finds. Here's What It Means," *Forbes*, April 17, 2024, <https://www.forbes.com/sites/robertpearl/2024/04/17/nvidias-ai-bot-outperforms-nurses-heres-what-it-means-for-you>; Emily M. Bender and Alex Hanna, interview with Michelle Mahon, *Mystery AI Hype Theater 3000*, podcast audio, August 2, 2024, <https://www.buzzsprout.com/2126417/episodes/15517978-episode-37-chatbots-aren-t-nurses-feat-michelle-mahon-july-22-2024>.
- 60 Salesforce, "More Than 7 in 10 College Students and Administrators Seek AI Agents to Close Support Gaps, Ease Burnout," March 10, 2025, <https://www.salesforce.com/news/stories/ai-agents-for-education-stats>.
- 61 Salesforce, "Salesforce Research: 90% of Constituents Ready for AI Agents in Public Service," *Salesforce*, January 15, 2025, <https://www.salesforce.com/news/stories/agent-ai-government-statistics-davos>.
- 62 Salesforce, "AI Agents Can Cut Healthcare Paperwork by 30%, Study Shows," February 28, 2025, <https://www.salesforce.com/news/stories/healthcare-ai-agent-research>.
- 63 Google for Education, "Explore Education Research and Insights," accessed April 25, 2025, https://edu.google.com/intl/ALL_us/research; Rose Luckin and Mark Griffiths, "Intelligence Unleashed," *Pearson*, 2016, <https://static.googleusercontent.com/media/edu.google.com/en//pdfs/Intelligence-Unleashed-Publication.pdf>.
- 64 Sayash Kapoor et al., "REFORMS: Reporting Standards for Machine Learning Based Science," *arXiv*, September 19, 2023), [arXiv:2308.07832](https://arxiv.org/abs/2308.07832).
- 65 Ahmed Alla et al., "Medical Large Language Model Benchmarks Should Prioritize Construct Validity," *arXiv*, March 12, 2025, [arXiv:2503.10694](https://arxiv.org/abs/2503.10694).
- 66 Asif Iqbal Khan, Junaid Latief Shah, and Mohammad Mudasar Bhat, "CoroNet: A Deep Neural Network for Detection and Diagnosis of COVID-19 From Chest X-Ray Images," *Computer Methods and Programs in Biomedicine* 196, no. 105581 (November 2020), <https://doi.org/10.1016/j.cmpb.2020.105581>.
- 67 Philip Ball, "Is AI Leading to a Reproducibility Crisis in Science?" *Nature*, December 5, 2023, <https://doi.org/10.1038/d41586-023-03817-6>.
- 68 Michael Roberts, "Common Pitfalls and Recommendations for Using Machine Learning to Detect and Prognosticate for COVID-19 Using Chest Radiographs and CT Scans," *Nature Machine Intelligence* 3 (March 2021): 199–217, <https://doi.org/10.1038/s42256-021-00307-0>.
- 69 Laura Weidinger et al., "Toward an Evaluation Science for Generative AI," *arXiv*, March 13, 2025, <https://arxiv.org/abs/2503.05336>.
- 70 Russell Brandom, "How to Build a Better AI Benchmark," *Technology Review*, May 8, 2025, <https://www.technologyreview.com/2025/05/08/1116192/how-to-build-a-better-ai-benchmark>.
- 71 Emanuel Maiberg, "Researchers Say the Most Popular Tool for Grading AIs Unfairly Favors Meta, Google, OpenAI," *404 Media*, <https://www.404media.co/chatbot-arena-illusion-paper-meta-openai>.
- 72 Brandom, "How to Build a Better AI Benchmark."
- 73 Jackson Barnett, "Scale AI Awarded \$250M contract by Department of Defense," *Fedscoop*, January 31, 2022, <https://fedscoop.com/scale-ai-awarded-250m-ai-contract-by-department-of-defense>; Hayden Field, "Scale AI Announces Multimillion-Dollar Defense Deal, a Major Step in U.S. Military Automation," *CNBC*, March 5, 2025, <https://www.cnbc.com/2025/03/05/scale-ai-announces-multimillion-dollar-defense-military-deal.html>.
- 74 The Scale Team, "Scale AI Partners with DoD's Chief Digital and Artificial Intelligence Office (CDAO) to Test and Evaluate LLMs," February 20, 2024, <https://scale.com/blog/scale-partners-with-cdao-to-test-and-evaluate-llms>; Brandi Vincent, "Scale AI to Set the Pentagon's Path for Testing and Evaluating Large Language Models," *Defense Scoop*, February 20, 2024, <https://defensescoop.com/2024/02/20/scale-ai-pentagon-testing-evaluating-large-language-models>; Chief Digital and Artificial Intelligence Office (CDAO), "Artificial Intelligence Rapid Capabilities Cell," December 11, 2024, <https://www.ai.mil/Portals/137/Documents/Resources%20Page/2024-12-CDAO-Artificial-Intelligence-Rapid-Capabilities-Cell.pdf>.
- 75 Inioluwa Debora Raji et al., "AI and the Everything in the Whole Wide World Benchmark," *arXiv*, November 26, 2021, [arXiv:2111.15366](https://arxiv.org/abs/2111.15366).
- 76 The journalist and researcher Meredith Broussard has reframed technosolutionism as technochauvinism, or the misplaced belief that technological solutions are superior to other means of addressing problems. See Broussard, *Artificial Unintelligence* (MIT Press, 2019).
- 77 Christopher Newfield, 2016. *The Great Mistake: How We Wrecked Public Universities and How We Can Fix Them* (Johns Hopkins University Press, 2016); Howard Besser and Maria Bonn, "Impact of Distance Independent Education," *Journal of the American Society for Information Science* 47, no. 11 (November 1996): 880–83, [https://asistdl.onlinelibrary.wiley.com/doi/10.1002/\(SICI\)1097-4571\(199611\)47:11%3C880::AID-ASI14%3E3.0.CO;2-Z](https://asistdl.onlinelibrary.wiley.com/doi/10.1002/(SICI)1097-4571(199611)47:11%3C880::AID-ASI14%3E3.0.CO;2-Z); Kelli Bird et al., "Big Data on Campus," *Education Next* (blog), July 20, 2021, <https://www.educationnext.org/big-data-on-campus-putting-predictive-analytics-to-the-test>; Neil Selwyn, *Education and Technology: Key Issues and Debates* (Continuum International Publishing Group, 2011); Neil Selwyn, *Distrusting Educational Technology: Critical Questions for Changing Times* (Taylor and Francis, 2013); Britt Paris et al., "Sins of Omission: Critical Informatics Perspectives on Higher Education Learning Analytics," *Journal for the Association for Information Science and Technology* 59, no. 1 (October 2022): 479–485, <https://doi.org/10.1002/asi.24575>.
- 78 Tressie McMillan Cottom, "The Tech Fantasy That Powers A.I. Is Running on Fumes," *New York Times*, March 29, 2025, <https://www.nytimes.com/2025/03/29/opinion/ai-tech-innovation.html>.
- 79 OpenAI, "Introducing ChatGPT Edu," May 30, 2024, <https://openai.com/index/introducing-chatgpt-edu>.

- 80 Dan Fitzpatrick, "ChatGPT to Be Given to All Estonian High School Students," *Forbes*, February 26, 2025, <https://www.forbes.com/sites/danfitzpatrick/2025/02/26/chatgpt-to-be-given-to-all-estonian-high-school-students>.
- 81 "Executive Order 14277 of April 23, 2025, Advancing Artificial Intelligence Education for American Youth," *Code of Federal Regulations*, title 90 (2025): 17519 <https://www.federalregister.gov/documents/2025/04/28/2025-07368/advancing-artificial-intelligence-education-for-american-youth>.
- 82 California State University, "CSU Announces Landmark Initiative to Become Nation's First and Largest AI-Empowered University System," press release, February 4, 2025, <https://www.calstate.edu/csu-system/news/Pages/CSU-AI-Powered-Initiative.aspx>.
- 83 Alpha School (website), accessed April 8, 2025, <https://alpha.school>.
- 84 Mark Keierleber, "Whistleblower: LA Schools' Chatbot Misused Student Data as Tech Co Crumbled," *The 74 Million*, July 1, 2024, <https://www.the74million.org/article/whistleblower-l-a-schools-chatbot-misused-student-data-as-tech-co-crumbled>.
- 85 Sonja Drimmer and Christopher J. Nygren, "How We Are Not Using AI in the Classroom," *The Newsletter of the International Center of Medieval Art*, no. 1 (Spring 2025): 25–28, https://static1.square-space.com/static/55577d2fe4b02de6a6ea49cd/t/67dfb8d9f-f3a5472a6d719d/1742728078061/Drimmer_Nygren_Not_Using_AI.pdf.
- 86 Ibid.
- 87 Stefanie Coyle and Rashida Richardson, "Bottom-Up Biometric Regulation: A Community's Response to Using Face Surveillance in Schools," in Amba Kak, ed., *Regulating Biometrics: Global Approaches and Urgent Questions*, AI Now Institute, September 2020, https://ainowinstitute.org/wp-content/uploads/2023/09/regulatingbiometrics_Bottom-Up-Biometric-Regulation.pdf.
- 88 Letter from Stefanie Coyle to NYSED Commissioner MaryEllen Elia, September 26, 2018.
- 89 Connor Hoffman, "Citizens Petition LCSD to Postpone Security Project," *Lockport Journal*, March 20, 2018, https://www.lockportjournal.com/news/local_news/citizens-petition-lcsd-to-postpone-security-project/article_e8d547e1-0e3a-5a2e-81c3-5879422548c2.html.
- 90 Davey Alba, "Facial Recognition Moves Into a New Front: Schools," *New York Times*, February 6, 2020, <https://www.nytimes.com/2020/02/06/business/facial-recognition-schools.html>.
- 91 Lockport's weapons-detection system misidentified broom handles as guns so frequently that it had to be disabled. See Todd Feathers, "Facial Recognition Company Lied to School District About Its Racist Tech," *Vice*, December 1, 2020, <https://www.vice.com/en/article/facial-recognition-company-lied-to-school-district-about-its-racist-tech>.
- 92 Carolyn Thompson, "New York Bans Facial Recognition in Schools After Report Finds Risks Outweigh Potential Benefits," *Associated Press*, September 27, 2023, <https://apnews.com/article/facial-recognition-banned-new-york-schools-ddd35e004254d316beabf-70453b1a6a2>.
- 93 Simon McCormack and Stefanie Coyle, "This Software Could Be Spying on NYC Students," *NYCLU Commentary*, <https://www.nyclu.org/commentary/software-could-be-spying-nyc-students>.
- 94 Nader Issa, "CPS Teachers Could Look Inside Students' Homes — Without Their Knowledge — Before Fix," *Chicago Sun Times*, October 5, 2020, <https://chicago.suntimes.com/education/2020/10/5/21497946/cps-public-schools-go-guardian-technology-privacy-remote-learning>.
- 95 "Since 2020, Beacon has prevented an estimated 18,623 students from physical harm," *GoGuardian* (website), accessed May 17, 2025, <https://www.goguardian.com/#footnotes>.
- 96 Jesse Stommel and Martha Burtis, "Bad Data Are Not Better than No Data," AAUP, *The Higher Ed Data Juggernaut*, Winter 2024, <https://www.aaup.org/article/bad-data-are-not-better-no-data>.
- 97 See Britt Paris et al., "Sins of Omission"; Catherine McGowan et al., "Educational Technology and the Entrenchment of 'Business as Usual,'" *American Association of University Professors (AAUP), Academe Magazine*, www.aaup.org/article/educational-technology-and-entrenchment-%E2%80%9Cbusiness-usual%E2%80%9D; and Britt Paris et al., "Platforms Like Canvas Play Fast and Loose With Students' Data," *Nation*, April 22, 2021, <https://www.thenation.com/article/society/canvas-surveillance>.
- 98 National Nurses United, to Interested Parties, March 30, 2023, https://www.nationalnursesunited.org/sites/default/files/nnu/documents/Reporter_Memo_Hospital_Staffing_Crisis.pdf.
- 99 Ingrid Lunden, "As Its Workers Strike, Kaiser Permanente Strikes a Deal for Physicians to Use an AI Copilot from Nabla," *TechCrunch*, October 5, 2023, <https://techcrunch.com/2023/10/05/as-its-workers-strike-over-burnout-and-low-wages-kaiser-permanente-strikes-a-deal-to-use-an-ai-copilot-from-nabla>.
- 100 National Farm Worker Ministry, "Issues Affecting Farm Workers," <https://nfwm.org/farm-workers/farm-worker-issues>.
- 101 Juan F. Perea, "The Echoes of Slavery: Recognizing the Racist Origins of the Agricultural and Domestic Worker Exclusion from the National Labor Relations Act," *Ohio State Law Journal* 72, no. 1 (2011): 95–138, <http://doi.org/10.2139/ssrn.1646496>.
- 102 See Perea, "The Echoes of Slavery," 96; and National Farm Worker Ministry, "Issues Affecting Farm Workers," 2022, <https://nfwm.org/wp-content/uploads/2022/12/Farm-Worker-Issues-Two-Pager-.pdf>.
- 103 *Cedar Point Nursery*, et al. v. Hassid, 594 U.S. 139 (2021).
- 104 Shanelle Kaul, "How AI Powered Robots Are Helping Small Farms Fight Labor Shortages," *CBS News*, March 28, 2024, <https://www.cbsnews.com/video/how-ai-powered-robots-are-helping-small-farms-fight-labor-shortages>.
- 105 The WG Center for Innovation & Technology, *Western Growers Case Study*, November 2024, https://wga.s3.us-west-1.amazonaws.com/cit/2024/cit_case-study-stout.pdf.
- 106 Jesse Stommel and Martha Burtis, "Counter-Friction to Stop the Machine: The Endgame for Instructional Design," *Hybrid Pedagogy*, April 27, 2021, <https://hybridpedagogy.org/the-endgame-for-instructional-design>.
- 107 AI Infiltration into Higher Education: AAUP Survey Findings and Strategies. American Association of University Professors Ad hoc Committee on AI and the Profession. link forthcoming.

- 108 Greg Toppo, "Was Los Angeles Schools' \$6 Million AI Venture a Disaster Waiting to Happen?" *The 74 Million*, July 9, 2024, <https://www.the74million.org/article/was-los-angeles-schools-6-million-ai-venture-a-disaster-waiting-to-happen>.
- 109 Nora Freeman Engstrom and David Freeman Engstrom, "The Making of the A2J Crisis," *Stanford Law Review Online* 75, no. 146 (May 2024), <https://ssrn.com/abstract=4817329>.
- 110 Andy J. Semotiuk, "How AI Is Impacting Immigration Cases and What to Expect," *Forbes*, March 23, 2024, <https://www.forbes.com/sites/andyjsemotiuk/2024/03/23/how-ai-is-impacting-immigration-cases-and-what-to-expect>.
- 111 Casey Ross, "A Popular Algorithm to Predict Sepsis Misses Most Cases and Sends Frequent False Alarms, Study Finds," *Stat*, June 21, 2021, <https://www.statnews.com/2021/06/21/epic-sepsis-prediction-tool>.
- 112 Casey Ross, "Epic's AI Algorithms, Shielded from Scrutiny by a Corporate Firewall, Are Delivering Inaccurate Information on Seriously Ill Patients," *Stat*, July 26, 2021, <https://www.statnews.com/2021/07/26/epic-hospital-algorithms-sepsis-investigation>.
- 113 Andrew Wong et al., "External Validation of a Widely Implemented Proprietary Sepsis Prediction Model in Hospitalized Patients," *JAMA Internal Medicine* 181, no. 8 (2021): 1065–1070, <https://jamanetwork.com/journals/jamainternalmedicine/fullarticle/2781307>.
- 114 Patrick G. Lyons et al., "How Useful Is the Epic Sepsis Prediction Model for Predicting Sepsis?" *American Journal of Respiratory and Critical Care Medicine* (May 2021), https://doi.org/10.1164/ajrccm-conference.2021.203.1_MeetingAbstracts.A1580.
- 115 Ross, "Epic's AI Algorithms, Shielded From Scrutiny by a Corporate Firewall, Are Delivering Inaccurate Information on Seriously Ill Patients."
- 116 Ibid.
- 117 "About Tectonic Justice," Tectonic Justice (website), accessed April 25, 2025, <https://www.tectonicjustice.org/about>.
- 118 Kevin De Liban, *Inescapable AI: The Ways AI Decides How Low-Income People Work, Live, Learn, and Survive*, Tectonic Justice, November 2024, <https://static1.squarespace.com/static/65a1d3be4690143890f61cec/t/673c7170a0d09777066c6e50/1732014450563/ttj-inescapable-ai.pdf>.
- 119 Bryce Covert, "States Are Turning Their Public Benefits Systems Over to AI. The Results Have Often Led to 'Immense Suffering,'" *Fast Company*, January 23, 2025, <https://www.fastcompany.com/91265363/states-are-turning-their-public-benefits-systems-over-to-ai-the-results-have-often-led-to-immense-suffering>.
- 120 Michele Gilman, "AI Algorithms Intended to Root Out Welfare Fraud Often End Up Punishing the Poor Instead," *The Conversation*, February 14, 2020, <https://theconversation.com/ai-algorithms-intended-to-root-out-welfare-fraud-often-end-up-punishing-the-poor-instead-131625>.
- 121 Amnesty International, "Dutch Childcare Benefit Scandal an Urgent Wake-Up Call to Ban Racist Algorithms," October 25, 2021, <https://www.amnesty.org/en/latest/news/2021/10/xenophobic-machines-dutch-child-benefit-scandal>.
- 122 Virginia Eubanks, "A Child Abuse Prediction Model Fails Poor Families," *Wired*, January 15, 2018, <https://www.wired.com/story/excerpt-from-automating-inequality>.
- 123 Lauren Kitchner and Matthew Goldstein, "Access Denied: Faulty Automated Background Checks Freeze Out Renters," *Markup*, May 28, 2020, <https://themarkup.org/locked-out/2020/05/28/access-denied-faulty-automated-background-checks-freeze-out-renter>.
- 124 Danielle Kaye, Lauren Hirsch, and David McCabe, "U.S. Accuses Software Maker Realpage of Enabling Collusion on Rents," *New York Times*, August 23, 2024, <https://www.nytimes.com/2024/08/23/business/economy/realpage-doj-antitrust-suit-rent.html>.
- 125 Brian Chen, "Dispelling Myths of AI and Efficiency," *Data & Society*, March 25, 2025, <https://datasociety.net/library/dispelling-myths-of-ai-and-efficiency>; Makena Kelly, "Elon Musk Ally Tells Staff 'AI-First' Is the Future of Key Government Agency," *Wired*, February 3, 2025, <https://www.wired.com/story/elon-musk-lieutenant-gsa-ai-agency>.
- 126 Kate Conger, Ryan Mac, and Madeleine Ngo, "Must Allies Discuss Deploying A.I. to Find Budget Savings," *New York Times*, February 3, 2025, <https://www.nytimes.com/2025/02/03/technology/musk-allies-ai-government.html>.
- 127 Makena Kelly and Zoe Schiffer, "DOGE Has Deployed Its GSaI Custom Chatbot for 1,500 Federal Workers," *Wired*, March 7, 2025, <https://www.wired.com/story/gsa-chatbot-1500-federal-workers>.
- 128 Courtney Kube et al., "DOGE Will Use AI to Assess the Responses of Federal Workers Who Were Told to Justify Their Jobs via Email," *NBC News*, February 25, 2025, <https://www.nbcnews.com/politics/doge/federal-workers-agencies-push-back-elon-musks-email-ultimatum-rcna193439>.
- 129 Asmelash Teka Hadgu and Timnit Gebru, "Replacing Federal Workers with Chatbots Would Be a Dystopian Nightmare," *Scientific American*, April 14, 2025, <https://www.scientificamerican.com/article/replacing-federal-workers-with-chatbots-would-be-a-dystopian-nightmare>.
- 130 Emily M. Bender, "Calling All Mad Scientists: Reject 'AI' as a Framing of Your Work," *Mystery AI Hype Theater 3000: The Newsletter*, April 8, 2025, <https://buttondown.com/maiht3k/archive/calling-all-mad-scientists-reject-ai-as-a-framing>.
- 131 Social Security Administration, "Social Security Eliminates Wasteful Department," press release, February 24, 2025, <https://www.ssa.gov/news/press/releases/2025/#2025-02-24>; Natalie Alms, "Social Security Shuts its Civil Rights and Transformation Offices," *Government Executive*, February 26, 2025, <https://www.govexec.com/management/2025/02/social-security-shutters-its-civil-rights-and-transformation-offices/4033>.
- 132 Ed O'Keefe and Rhona Tarrant, "General Services Administration Shuts its Technology Unit," *CBS*, March 2, 2025, <https://www.cbsnews.com/news/general-services-administration-shutters-technology-unit-trump-doge>.
- 133 John Wilkerson, "HHS Agency Responsible for Health Care Quality Research Threatened with Mass Layoffs," *Stat*, March 20, 2025, <https://www.statnews.com/2025/03/20/hhs-ahrq-agency-responsible-for-health-care-quality-research-threatened-with-mass-layoffs>; Ezra Klein, interview with Santi Ruiz, *The Ezra Klein Show*, podcast audio, 16:00, March 25, 2025, <https://www.nytimes.com/2025/03/25/opinion/ezra-klein-podcast-santi-ruiz.html>.

- 134 Jonathan Allen, “‘Absolute Chaos’: DOGE Sows Turmoil in Its Quest for ‘Efficiency,’” *NBC News*, February 25, 2025, <https://www.nbc-news.com/politics/doge/absolute-chaos-doge-turmoil-efficiency-rc-na193579>.
- 135 Allen, “‘Absolute Chaos.’”
- 136 Eric Lipton, “Musk is Positioned to Profit Off Billions in New Government Contracts,” *New York Times*, March 23, 2025, <https://www.nytimes.com/2025/03/23/us/politics/spacex-contracts-musk-doge-trump.html>; Chris Hayes, “Elon Musk Called Out for Raking in \$8 Million a Day from Taxpayers,” *MSNBC*, February 12, 2025, <https://www.msnbc.com/all-in/watch/elon-musk-called-out-for-raking-in-8-million-a-day-from-taxpayers-231824965528>; Department of Defense, “Department of Defense Awards \$14.3 Million to Expand Sources of Solid Rocket Motors,” press release, January 7, 2025, <https://www.defense.gov/News/Releases/Release/Article/4022917/departments-of-defense-awards-143-million-to-expand-sources-of-solid-rocket-moto>; Miles Jamison, “Anduril Awarded \$99M Air Force Contract for Thunderdome Project,” *GovCon Wire*, February 18, 2025, <https://www.govconwire.com/2025/02/anduril-99-million-air-force-contract-thunderdome-project>; “Anduril Takes Over Microsoft’s \$22 Billion US Army Headset Program,” *Reuters*, February 11, 2025, <https://www.reuters.com/technology/anduril-takes-over-micro-softs-22-billion-us-army-headset-program-2025-02-11>; MediaJustice, “WTF: The Rise of the Broligarchy PoliEd Series-Episode 2: WTF is the Tech Broligarchy Up To?” March 6, 2025, 54 min., 42 sec., <https://www.youtube.com/watch?v=sKhnswwYe60&t=3017s>.
- 137 Nick Mordowanec, “Map Shows DOGE Office Closures in March,” *Newsweek*, March 26, 2025, <https://www.newsweek.com/doge-elon-musk-map-budget-cuts-march-2050761>.
- 138 Zeeshan Aleem, “DOGE Is Already Breaking Social Security,” *MSNBC*, March 26, 2025, <https://www.msnbc.com/opinion/msnbc-opinion/social-security-elon-musk-doge-cuts-issues-rcna198034>.
- 139 Makena Kelly, “DOGE Plans to Rebuild SSA Code Base in Months, Risking Benefits and System Collapse,” *Wired*, March 28, 2025, <https://www.wired.com/story/doge-rebuild-social-security-administration-cobol-benefits>.
- 140 Brian J. Chen, “Dispelling Myths of AI and Efficiency,” *Data & Society*, March 25, 2025, <https://datasociety.net/library/dispelling-myths-of-ai-and-efficiency>.
- 141 Eryk Salvaggio, “Anatomy of an AI Coup,” *Tech Policy Press*, February 9, 2025, <https://www.techpolicy.press/anatomy-of-an-ai-coup>.
- 142 Bob Evans, “DOGE Triggering ‘Tremendous Opportunity’ for Workday Federal Business,” *Cloud Wars*, March 13, 2025, <https://cloudwars.com/cloud/doge-triggering-tremendous-opportunity-for-workday-federal-business>.
- 143 Center for Democracy and Technology et al., *DOGE and Government Data Privacy*, March 17, 2025, <https://civilrights.org/wp-content/uploads/2025/03/DOGE-and-Government-Data-Privacy-FINAL.pdf>; Aleem, “DOGE Is Already Breaking Social Security”; Chas Danner, “All the Federal Agencies DOGE Has Gotten Access to,” *New York Magazine*, February 10, 2025, <https://nymag.com/intelligencer/article/doge-elon-musk-what-federal-agencies-access-lawsuits.html>.
- 144 Brian Merchant, “DOGE’s ‘AI-First’ Strategist Is Now the Head of Technology at the Department of Labor,” *Blood in the Machine* (blog), March 19, 2025, <https://www.bloodinthemachine.com/p/doges-ai-first-strategist-is-now>; Natalie Alms, “TTS Director Tapped to Serve as Labor CIO,” *Nextgov/FCW*, March 18, 2025, <https://www.nextgov.com/people/2025/03/tts-director-tapped-serve-labor-cio/403855>.
- 145 Samar Khurshid, “Eric Adam Vows to Overhaul How City Government Works; Experts Point to Several Essentials to Following Through,” *Gotham Gazette*, October 31, 2021, <https://www.gothamgazette.com/city/10870-eric-adams-promises-overhaul-how-city-government-works-experts>.
- 146 Cynthia Conti-Cook and Ed Vogel, MyCity, INC: A Case Against “CompStat Urbanism” (New York: Surveillance Resistance Lab, March 18, 2024).
- 147 City of New York, “Mayor Adams Releases First-of-Its-Kind Plan for Responsible Artificial Intelligence Use in NYC Government,” October 16, 2023, <https://www.nyc.gov/office-of-the-mayor/news/777-23/mayor-adams-releases-first-of-its-kind-plan-responsible-artificial-intelligence-use-nyc>; Colin Lecher, “NYC’s AI Chatbot Tells Businesses to Break the Law,” *Markup*, March 29, 2025, <https://themarkup.org/news/2024/03/29/nycs-ai-chatbot-tells-businesses-to-break-the-law>.
- 148 Jake Offenhartz, “NYC’s AI Chatbot Was Caught Telling Businesses to Break the Law. The City Isn’t Taking it Down,” *Associated Press*, April 3, 2024, <https://apnews.com/article/new-york-city-chatbot-misinformation-6ebc71db5b770b9969c906a7ee4fae21>.
- 149 New York State Human Rights Law, Section 296, Unlawful Discriminatory Practices.
- 150 New York City Office of Technology and Innovation, “Summary of Agency Compliance Reporting of Algorithmic Tools,” 2023, <https://www.nyc.gov/assets/oti/downloads/pdf/reports/2023-algorithmic-tools-reporting-updated.pdf>.
- 151 One City Act, Assembly Bill A9642 (2023–2024), <https://www.nysenate.gov/legislation/bills/2023/A9642>.
- 152 Zachary Groz, “How Eric Adam’s MyCity Portal Became a \$100 Million Question Mark,” *New York Focus*, March 19, 2025, <https://nysfocus.com/2025/03/19/mycity-eric-adams-child-care>.
- 153 Annie McDonough, “Matt Fraser Still Wants to Expand MyCity and AI Chatbot,” *City & State New York*, March 19, 2025, <https://www.cityandstateny.com/personality/2025/03/matt-fraser-still-wants-expand-mycity-and-ai-chatbot/403899>.
- 154 Likhita Banerji and Damini Satija, “AI as Double Speak for Austerity,” *Tech Policy Press*, February 7, 2025, <https://www.techpolicy.press/ai-as-double-speak-for-austerity>.
- 155 Satyan Gajwaniin, “AI Effect: We’ll Get Way Wealthier and Witness a Productivity Boom, Says Sam Altman,” *Economic Times*, June 9, 2023, <https://economictimes.indiatimes.com/tech/technology/ai-effect-well-get-way-wealthier-and-witness-a-productivity-boom-says-sam-altman/articleshow/100857881.cms>.
- 156 Jodi Kantor and Arya Sundaram, “The Rise of the Worker Productivity Score,” *New York Times*, August 14, 2022, <https://www.nytimes.com/interactive/2022/08/14/business/worker-productivity-tracking.html>.
- 157 Wilneida Negrón, Little Tech Is Coming for Workers, *Coworker.org*, 2021, <https://home.coworker.org/wp-content/uploads/2021/11/Little-Tech-Is-Coming-for-Workers.pdf>.
- 158 Edward Ongweso Jr., “Amazon’s New Algorithm Will Set Workers’

- Schedules According to Muscle Use,” *Vice*, April 15, 2021, <https://www.vice.com/en/article/amazons-new-algorithm-will-set-workers-schedules-according-to-muscle-use>.
- 159 Lauren Kaori Gurley, “Internal Documents Show Amazon’s Dystopian System for Tracking Workers Every Minute of Their Shifts,” *Vice*, June 2, 2022, <https://www.vice.com/en/article/internal-documents-show-amazons-dystopian-system-for-tracking-workers-every-minute-of-their-shifts>.
 - 160 Warehouse Worker Resource Center, “Extreme Heat at Amazon Air,” September 2022, https://warehouseworkers.org/wwrc_resources/extreme-heat-at-amazon-air.
 - 161 “Amazon Offers ‘Wellness Chamber’ for Stressed Staff,” *BBC*, May 28, 2021, <https://www.bbc.com/news/technology-57287151>.
 - 162 U.S. Congress, Senate, Committee on Health, Education, Labor, and Pensions, “The ‘Injury-Productivity Trade-off’: How Amazon’s Obsession with Speed Creates Uniquely Dangerous Warehouses,” 118th Cong., 2d. sess., 2024. https://www.help.senate.gov/imo/media/doc/amazon_investigation.pdf.
 - 163 Juliana Kim, “Amazon Manipulated Injury Data to Make Warehouses Appear Safer, a Senate Probe Finds,” *NPR*, December 16, 2024, <https://www.npr.org/2024/12/16/nx-s1-5230240/amazon-injury-warehouse-senate-investigation>.
 - 164 Theo Wayt, “Klarna’s AI Savings” *The Information*, November 25, 2024, <https://www.theinformation.com/articles/klarnas-ai-savings>.
 - 165 Daron Acemoglu and Simon Johnson, *Power and Progress: Our Thousand-Year Struggle Over Technology & Prosperity* (PublicAffairs, 2023).
 - 166 Acemoglu and Johnson, *Power and Progress*, 9–38.
 - 167 John F. Kennedy Presidential Library and Museum, News Conference 24, February 14, 1962, <https://www.jfklibrary.org/archives/other-resources/john-f-kennedy-press-conferences/news-conference-24>; “Kennedy Began War on Poverty; Area Redevelopment First of His Weapons 3 Years Ago,” *New York Times*, January 10, 1964, <https://www.nytimes.com/1964/01/10/archives/kennedy-began-war-on-poverty-area-redevelopment-first-of-his.html>.
 - 168 Acemoglu and Johnson, *Power and Progress*, 242, 247. It is important to note that women, workers of color, and immigrants did not see similar gains due to insidious forces of patriarchy, racism, and xenophobia. See Acemoglu and Johnson, *Power and Progress*, 251.
 - 169 Amba Kak and Sarah Myers West, “International ‘Digital Trade’ Agreements: The Next Frontier,” in *2023 Landscape: Confronting Tech Power*, *AI Now Institute*, <https://ainowinstitute.org/publications/international-digital-trade-agreements>.
 - 170 Susannah Glickman, “AI and Tech Industrial Policy: From Post-Cold War Post-Industrialism to Post-Neoliberal Re-Industrialization,” in *AI Nationalism(s): Global Industrial Policy Approaches to AI*, *AI Now Institute*, March 12, 2024, <https://ainowinstitute.org/publications/ai-and-tech-industrial-policy-from-post-cold-war-post-industrialism-to-post-neoliberal-re-industrialization>.
 - 171 Veena Dubal, “On Algorithmic Wage Discrimination,” *Columbia Law Review* 124, no. 7 (2023), <https://doi.org/10.2139/ssrn.4331080>.
 - 172 Katie J. Wells and Funda Ustek Spilda, *Uber for Nursing: How an AI-Powered Gig Model Is Threatening Health Care*, Roosevelt Institute, December 2024, https://rooseveltinstitute.org/wp-content/uploads/2024/12/RI_Uber-for-Nursing_Brief_202412.pdf.
 - 173 Katie J. Wells et al., *Fairwork US Ratings 2025: When AI Eats the Manager*, Fairwork, 2025, 6.
 - 174 Wells et al., *When AI Eats the Manager*.
 - 175 Wells et al., *When AI Eats the Manager*, 18.
 - 176 Edward Ongweso Jr., “Uber’s Bastards II,” *The Tech Bubble* (blog), March 27, 2025, <https://thetechbubble.substack.com/p/ubers-bastards-ii>.
 - 177 Sam Altman, “Reflections,” January 5, 2025, <https://blog.samaltman.com/reflections>.
 - 178 OpenAI, “Introducing Operator,” January 23, 2025, <https://openai.com/index/introducing-operator>.
 - 179 Casey Newton, “OpenAI Launches its Agent,” *Platformer*, January 23, 2025, <https://www.platformer.news/openai-operator-ai-agent-handson>.
 - 180 Robert Krzaczynski, “OpenAI Launches Deep Research: Advancing AI-Assisted Investigation,” *InfoQ*, February 6, 2025, <https://www.infoq.com/news/2025/02/deep-research-openai>.
 - 181 Dave Citron, “Try Deep Research and Our New Experimental Model in Gemini, Your AI Assistant,” *Google*, December 11, 2024, <https://blog.google/products/gemini/google-gemini-deep-research>.
 - 182 Palantir, “AIP Agent Studio [Beta],” accessed April 25, 2025, <https://www.palantir.com/docs/foundry/agent-studio/overview>.
 - 183 Margaret Mitchell et al., “Why Handing Over Total Control to AI Agents Would Be a Huge Mistake,” *Technology Review*, March 24, 2025, <https://www.technologyreview.com/2025/03/24/1113647/why-handing-over-total-control-to-ai-agents-would-be-a-huge-mistake>.
 - 184 Sara Perez, “Signal President Meredith Whittaker Calls Out Agentic AI as Having ‘Profound’ Security and Privacy Issues,” *TechCrunch*, March 7, 2025, <https://techcrunch.com/2025/03/07/signal-president-meredith-whittaker-calls-out-agentic-ai-as-having-profound-security-and-privacy-issues>.
 - 185 Ben Thompson, “Enterprise Philosophy and the First Wave of AI,” *Stratechery*, September 24, 2024, <https://stratechery.com/2024/enterprise-philosophy-and-the-first-wave-of-ai>.
 - 186 Julie Mao et al., *Automating Deportation: The Artificial Intelligence Behind the Department of Homeland Security’s Immigration Enforcement Regime*, Just Futures Law and Mijente, June 2024, <https://mijente.net/wp-content/uploads/2024/06/Automating-Deportation.pdf>.
 - 187 Weixin Liang et al., “GPT Detectors Are Biased Against Non-Native English Writers,” *arXiv*, July 10, 2023, [arXiv:2304.02819](https://arxiv.org/abs/2304.02819).
 - 188 Adi Robertson, “ICE Rigged Its Algorithms to Keep Immigrants in Jail, Claims Lawsuit,” *Verge*, Mar. 3, 2020, <https://www.theverge.com/2020/3/3/21163013/ice-new-york-risk-assessment-algorithm-rigged-lawsuit-nyclu-josevelasaca>.
 - 189 Mijente, *The War Against Immigrants*, August 2019, 4, <https://mijente.net/wp-content/uploads/2019/08/Mijente-The-War-Against-Immigrants.pdf>.

- Immigrants_-Trumps-Tech-Tools-Powered-by-Palantir_.pdf; National Immigration Project et al., *Who's Behind ICE? The Tech and Data Companies Fueling Deportations*, October 2018, 3, <https://surveil-lanceresistancelab.org/resources/whos-behind-ice>.
- 190 Jason Koebler, "Inside a Powerful Database ICE Uses to Identify and Deport People," *404 Media*, April 9, 2025, <https://www.404media.co/inside-a-powerful-database-ice-uses-to-identify-and-deport-people>.
- 191 Palantir, "Homeland Security Investigations Renews Partnership with Palantir for Case Management Software," press release, September 26, 2022, <https://www.palantir.com/newsroom/press-releases/home-land-security-investigations-renews-partnership-with-palantir>.
- 192 Koebler, "Inside a Powerful Database ICE Uses to Identify and Deport People."
- 193 Alayna Alvarez, "Palantir's Partnership with ICE Deepens," *Axios*, May 1, 2025, <https://www.axios.com/local/denver/2025/05/01/palantir-de-portations-ice-immigration-trump>.
- 194 U.S. House of Representatives, Committee on Oversight and Government Reform, "Letter to Michelle L. Anderson, Assistant Inspector General for Audit, Social Security Administration," April 17, 2025, <https://oversightdemocrats.house.gov/sites/evo-subsites/democrats-oversight.house.gov/files/evo-media-document/2025-04-17-gec-to-ssa-oig-master-data.pdf>.
- 195 Luke O'Brien, "The Shocking Far-Right Agenda Behind the Facial Recognition Tech Used by ICE and the FBI," *Mother Jones*, 2025, <https://www.motherjones.com/politics/2025/04/clearview-ai-immigration-ice-fbi-surveillance-facial-recognition-hoan-ton-that-hal-lam-bert-trump>.
- 196 Ibid.
- 197 Marina Dunbar, "Ice Director Wants to Run Deportations like 'Amazon Prime for Human Beings'," *Guardian*, April 9, 2025, <https://www.theguardian.com/us-news/2025/apr/09/ice-todd-lyons-depor-tion-amazon>.
- 198 Eileen Guo, "The US Wants to Use Facial Recognition to Identify Migrant Children as They Age," *Technology Review*, August 14, 2024, <https://www.technologyreview.com/2024/08/14/1096534/home-land-security-facial-recognition-immigration-border>.
- 199 Mao et al., "Automating Deportation," 16.
- 200 Ibid.
- 201 Sam Biddle, "LexisNexis Is Selling Your Personal Data to ICE So It Can Try to Predict Crimes," *Intercept*, June 20, 2023, <https://theintercept.com/2023/06/20/lexisnexis-ice-surveillance-license-plates>.



Chapter 4: A Roadmap for Action Endnotes

- 1 See Elizabeth Wilkins, “Restoring Economic Democracy: Progressive Ideas for Stability and Prosperity,” Roosevelt Institute, April 29, 2025, <https://rooseveltinstitute.org/publications/restoring-economic-democracy>; Ezra Klein, “Would Bernie Have Won?” *Ezra Klein Show*, November 26, 2024, <https://www.nytimes.com/2024/11/26/opinion/ezra-klein-podcast-faiz-shakir.html>; Pippa Norris et al., “Trump, Brexit, and the Rise of Populism: Economic Have-Nots and Cultural Backlash,” unpublished research project, HKS Faculty Research Working Paper Series RWP16-026, Harvard Kennedy School, 2016; and George Hawley, “The Political Economy of Right-Wing Populism in the United States,” *Promarket*, May 24, 2024, https://www.promarket.org/author/george_hawley.
- 2 Patrick Wintour, “Did Trump’s Tariffs Kill Economic Populism?” *Guardian*, April 12, 2025, <https://www.theguardian.com/world/2025/apr/12/did-trump-tariffs-economic-populism-globalisation>; Alexander R. Ross, “Steve Bannon and Elon Musk Are Battling for the Soul of Trumpism,” *New Lines Magazine*, April 15, 2025, <https://newlines-mag.com/argument/steve-bannon-and-elon-musk-are-battling-for-the-soul-of-trumpism>.
- 3 Ezra Klein, “In This House, We’re Angry When Government Fails,” *Ezra Klein Show*, November 22, 2024, <https://www.nytimes.com/2024/11/22/opinion/ezra-klein-podcast-jennifer-pahlka-steven-tele.html>.
- 4 David A. Fahrenthold and Jeremy Singer-Vine, “DOGE is Far Short of Its Goal, and Still Overstating Its Progress,” *New York Times*, April 13, 2025, <https://www.nytimes.com/2025/04/13/us/politics/doge-contracts-savings.html?smid=nytcore-ios-share&referringSource=articleShare>.
- 5 Kate Conger, Ryan Mac, and Madeleine Ngo, “Musk Allies Discuss Deploying A.I. to Find Budget Savings,” *New York Times*, February 3, 2025, <https://www.nytimes.com/2025/02/03/technology/musk-allies-ai-government.html>.
- 6 Makena Kelly and Zoe Schiffer, “DOGE Has Deployed Its GSAi Custom Chatbot for 1,500 Federal Workers,” *Wired*, March 7, 2025, <https://www.wired.com/story/gsai-chatbot-1500-federal-workers>.
- 7 Courtney Kube et al., “DOGE Will Use AI to Assess the Responses of Federal Workers Who Were Told to Justify Their Jobs via Email,” *NBC News*, February 25, 2025, <https://www.nbcnews.com/politics/doge/federal-workers-agencies-push-back-elon-musks-email-ultimatum-rcna193439>.
- 8 Aatish Bhatia et al., “DOGE’s Only Public Ledger is Riddled with Mistakes,” *New York Times*, February 28, 2025, <https://www.nytimes.com/2025/02/21/upshot/doge-musk-trump-errors.html>; Fahrenthold and Singer-Vine, “DOGE Far Short of Its Goal.”
- 9 Donald S. Beyer Jr. et al., to Russell Vought, February 12, 2025, https://beyer.house.gov/uploadedfiles/letter_from_congress_to_omb_director_on_restoring_public_access_to_federal_data.pdf.
- 10 Robert Weissman, “DOGE Delusions, A Real-World Plan to Crack Down on Corporate Handouts, Tax the Rich and Invest for the Future,” *Public Citizen*, January 15, 2025, <https://www.citizen.org/article/doge-delusions>.
- 11 Kevin De Liban, *Inescapable AI*, TechTonic Justice, November 2024, <https://www.techtonicjustice.org/reports/inescapable-ai>.
- 12 “I Live 400 Yards From Mark Zuckerberg’s Massive Data Center,” More Perfect Union, March 27, 2025, YouTube video, https://www.youtube.com/watch?v=DGjj7wDYail&ab_channel=MorePerfectUnion.
- 13 OpenAI, “OpenAI’s Infrastructure Blueprint for the US,” November 13, 2024, https://media.datacenterdynamics.com/media/documents/OpenAI_Blueprint-DCD.pdf; Abeba Birhane et al., “The Forgotten Margins of AI Ethics,” in *2022 ACM Conference on Fairness, Accountability, and Transparency* (2022): 948–958, <https://dl.acm.org/doi/10.1145/3531146.3533157>; and Abeba Birhane, “AI: Potential Benefits, Proven Risks, Abeba Birhane,” Better Ways (Agile Greece Summit), March 12, 2025, YouTube video, https://www.youtube.com/watch?v=P_p_n4XU0Y8&ab_channel=BetterWays%28AgileGreeceSummit%29.
- 14 Lois Parshley, “The Hidden Environmental Impact of AI,” *Jacobin*, June 20, 2024, <https://jacobin.com/2024/06/ai-data-center-energy-usage-environment>; Stand.earth, “A Growing Climate Concern Around Microsoft’s Expanding Data Center Operations,” December 5, 2024, <https://stand.earth/insights/a-growing-climate-concern-around-microsofts-expanding-data-center-operations>; Yusuf Sar, “The Silent Burden of AI: Unveiling the Hidden Environmental Costs of Data Centers by 2030,” *Forbes*, August 16, 2024, <https://www.forbes.com/councils/forbestechcouncil/2024/08/16/the-silent-burden-of-ai-unveiling-the-hidden-environmental-costs-of-data-centers-by-2030>.
- 15 Nikitha Sattiraju, “The Secret Cost of Google’s Data Centers: Billions of Gallons of Water to Cool Servers,” *Time*, April 2, 2020, <https://time.com/5814276/google-data-centers-water>; Christopher Tozzi, “How the Groundwater Crisis May Impact Data Centers,” *Data Center Knowledge*, March 12, 2024, <https://www.datacenterknowledge.com/data-center-site-selection/how-the-groundwater-crisis-may-impact-data-centers>.
- 16 Ireland is another extreme example of this problem. See Heidi Vella, “Ireland’s Data Centre Nightmare – and What Others Can Learn From It,” *Tech Monitor*, February 18, 2025, <https://www.techmonitor.ai/hardware/data-centres/irelands-data-centre-nightmare>.
- 17 “Dutch Call a Halt to New Massive Data Centres, While Rules are Worked Out,” *DutchNews*, February 17, 2022, <https://www.dutch-news.nl/2022/02/dutch-call-a-halt-to-new-massive-data-centres-while-rules-are-worked-out>; Mark Ballard, “Hyperscale Data Centers Under Fire in Holland,” *Data Center Knowledge*, February 7, 2022, <https://www.datacenterknowledge.com/hyperscalers/hyperscale-data-centers-under-fire-in-holland>.

- 18 “Google Says It Will Rethink Its Plans for a Big Data Center in Chile Over Water Worries,” *Associated Press*, September 17, 2024, <https://apnews.com/article/chile-google-data-center-water-drought-environment-d1c6a7a8e8e6e45257ac84fb750b2162>.
- 19 Chief Sheldon Sunshine to Premier Smith, January 13, 2024, <https://www.sturgeonlake.ca/wp-content/uploads/2020/07/13-01-25-Chief-Sunshine-Open-Letter-Premier-Smith-Re-O-Leary.pdf>; Julia Simon, “Demand for Minerals Sparks Fear of Mining Abuses on Indigenous Peoples’ Lands,” *NPR*, January 29, 2024, <https://www.npr.org/2024/01/29/1226125617/demand-for-minerals-sparks-fear-of-mining-abuses-on-indigenous-peoples-lands>.
- 20 Sarah O’Connor, “Anatomy of a Jobs Promise,” *Financial Times*, January 21, 2025, <https://www.ft.com/content/2f25065d-3eeb-49f6-a5eb-8d22ed4697a5>.
- 21 Saijel Kishan, “‘It’s a Money Loser’: Tax Breaks for Data Centers Are Under Fire,” *Bloomberg*, May 9, 2024, <https://www.bloomberg.com/news/articles/2024-05-09/ai-boom-has-some-states-rethinking-subsidies-for-data-centers>; Kasia Tarczyska, “Will Data Center Job Creation Live Up to Hype? I Have Some Concerns,” *Good Jobs First*, February 12, 2025, <https://goodjobsfirst.org/will-data-center-job-creation-live-up-to-hype-i-have-some-concerns>; Lulu Ramadan and Sydney Brownstone, “How a Washington Tax Break for Data Centers Snowballed Into One of the State’s Biggest Corporate Giveaways,” *ProPublica*, August 4, 2024, <https://www.propublica.org/article/washington-data-centers-tech-jobs-tax-break>; Tom Dotan, “The AI Data-Center Boom Is a Job-Creation Bust,” *Wall Street Journal*, February 25, 2025, <https://www.wsj.com/tech/ai-data-center-job-creation-48038b67>.
- 22 Zach Schiller, “Indefensible Tax Breaks for Data Centers Will Cost Ohio,” *Policy Matters Ohio*, January 7, 2025, <https://policymattersohio.org/research/indefensible-tax-breaks-for-data-centers-will-cost-ohio>; “At \$1 Billion, Amazon’s Oregon Subsidy is Largest Known in Company’s History,” *Good Jobs First*, May 19, 2023, <https://goodjobsfirst.org/at-1-billion-amazons-oregon-subsidy-is-largest-known-in-history>.
- 23 Kishan, “It’s a Money Loser”; Jarrett Skorup, “Michigan Has Authorized \$4.7 Billion in Taxpayer-Funded Business Subsidies,” *Mackinac Center*, January 27, 2025, <https://www.mackinac.org/pressroom/2025/michigan-has-authorized-4-7-billion-in-taxpayer-funded-business-subsidies>; American Economic Liberties Project, “Tools for Taking on the Corporate Subsidy Machine: Introduction,” September 13, 2022, <https://www.economicliberties.us/our-work/subsidy-toolkit-intro>; and Greg LeRoy, *The Great American Jobs Scam: Corporate Tax Dodging and the Myth of Job Creation* (Berrett Koehler, 2005).
- 24 Citizens Action Coalition, “CAC Calls for Data Center Moratorium,” October 15, 2024, <https://www.citact.org/news/cac-calls-data-center-moratorium>.
- 25 Andrew R. Chow, “Elon Musk’s New AI Data Center Raises Alarms Over Pollution,” *Time*, September 17, 2024, <https://time.com/7021709/elon-musk-xai-grok-memphis/>.
- 26 Veena Dubal, “On Algorithmic Wage Discrimination,” *Columbia Law Review* 123, no. 7 (2023): 1929–92, <https://columbialawreview.org/content/on-algorithmic-wage-discrimination>.
- 27 AI Now Institute et al., *AI Now Coauthors Report on Surveillance Prices and Wages*, February 20, 2025, <https://ainowinstitute.org/publication/ai-now-coauthors-report-on-surveillance-prices-and-wages>.
- 28 For example, in cases where states provide narrow exceptions for price or wage discrimination on the basis of surveillance data, businesses must have the burden of clearly establishing that the exception applies.
- 29 Will Oremus and Lauren Kaori Gurley, “States Eye Bans on ‘Surveillance Pricing’ that Exploits Personal Data,” *Washington Post*, February 20, 2025, <https://www.washingtonpost.com/politics/2025/02/20/surveillance-prices-wages-california-ban>.
- 30 Molly Kinder, “Hollywood Writers Went on Strike to Protect Their Livelihoods from Generative AI. Their Remarkable Victory Matters for All Workers,” *Brookings*, April 12, 2024, <https://www.brookings.edu/articles/hollywood-writers-went-on-strike-to-protect-their-livelihoods-from-generative-ai-their-remarkable-victory-matters-for-all-workers>.
- 31 Jake Coyle, “In Hollywood Writers’ Battle Against AI, Humans Win (For Now),” *Associated Press*, September 27, 2023, <https://apnews.com/article/hollywood-ai-strike-wga-artificial-intelligence-39ab-72582c3a15f77510c9c30a45ffc8>.
- 32 Mishal Khan et al., “The Current Landscape of Tech and Work Policy: A Guide to Key Concepts,” *UC Berkeley Labor Center*, November 12, 2024, <https://laborcenter.berkeley.edu/tech-and-work-policy-guide>; Noam Scheiber, “California Governor Signs Bill that Could Push Amazon to Change Labor Practices,” *New York Times*, October 26, 2021, <https://www.nytimes.com/2021/09/22/business/newsom-amazon-labor-bill.html>.
- 33 Gig Workers Rising, “About Us,” accessed April 17, 2025, <https://gig-workersrising.org/get-informed>.
- 34 Workers Justice Project, “Los Deliveristas Unidos,” accessed April 17, 2025, <https://www.workersjustice.org/en/ldu>.
- 35 Drivers United, “We’re Uber & Lyft Drivers Uniting for a Fair, Dignified, and Sustainable Rideshare Industry,” accessed April 17, 2025, <https://www.drivers-united.org>.
- 36 Eli Tan, “Ride-Hailing Drivers in Massachusetts Win Right to Unionize,” *New York Times*, November 6, 2024, <https://www.nytimes.com/2024/11/06/technology/uber-lyft-drivers-unionize-massachusetts.html>.
- 37 National Nurses United, “National Nurses United Survey Finds A.I. Technology Degrades and Undermines Patient Safety,” May 15, 2024, <https://www.nationalnursesunited.org/press/national-nurses-united-survey-finds-ai-technology-undermines-patient-safety>.
- 38 National Nurses United, “Nurses and Patients’ Bill of Rights: Guiding Principles for A.I. Justice in Nursing and Health Care,” accessed April 24, 2025, https://www.nationalnursesunited.org/sites/default/files/nnu/documents/0424_NursesPatients-BillOfRights_Principles-AI-Justice_flyer.pdf.
- 39 Bonnie Castillo, “Humanity is the Heart of Nursing,” National Nurses United, March 29, 2024, <https://www.nationalnursesunited.org/article/humanity-is-the-heart-of-nursing>.
- 40 Castillo, “Humanity is the Heart of Nursing.”
- 41 Federal Unionists Network, “Let Us Work! Save Our Services!” accessed April 17, 2025, <http://www.federalunionists.net>.

- 42 Josh Shapiro, "From the Desk of Governor Josh Shapiro," March 21, 2025, <https://www.pa.gov/content/dam/copapwp-pagov/en/governor/documents/2025.3.21%20gov%20shapiro%20letter%20on%20ai%20and%20employees.pdf>.
- 43 Athena, "How We Take on the Tech Oligarchy," February 28, 2025, <https://athenaforall.org/news/boycott-alternatives>.
- 44 #NoTechForIce, "Take Back Tech," accessed April 17, 2025, <https://notechforice.com>.
- 45 No Tech For Apartheid (website), accessed April 17, 2025, <https://www.notechforapartheid.com>.
- 46 "Good Jobs. Clean Air. Justice for Neighbors," *WWRC*, accessed April 17, 2025, https://warehouseworkers.org/wwrc_campaigns/building-a-better-san-bernardino.
- 47 Mishal Khan et al., "The Current Landscape of Tech and Work Policy: A Guide to Key Concepts," *UC Berkeley Labor Center*, November 12, 2024, <https://laborcenter.berkeley.edu/tech-and-work-policy-guide>.
- 48 Hillary Ronen, *Local Leadership in the Era of Artificial Intelligence And the Tech Oligarchy*, Local Progress and AI Now Institute, May 2025, <https://localprog.org/43FPS6>.
- 49 Lauren Rosenblatt, "5 Years into Amazon's Climate Pledge, Workers Challenge its Progress," *Seattle Times*, July 10, 2024, <https://www.seattletimes.com/business/amazon/5-years-into-amazons-climate-pledge-workers-challenge-its-progress>.
- 50 Mohammad Amir Anwar, "Africa's Data Workers Are Being Exploited by Foreign Tech Firms - 4 Ways to Protect Them," *Conversation*, March 31, 2025, <https://theconversation.com/africas-data-workers-are-being-exploited-by-foreign-tech-firms-4-ways-to-protect-them-252957>.
- 51 "African Tech Workers Rising," UNI Global Union, accessed April 17, 2025, <https://uniglobalunion.org/about/sectors/ict-related-services/african-tech-workers-rising>.
- 52 Tula Connell, "Mexico: App-Based Drivers Hail New Platform Law," *Solidarity Center*, January 6, 2025, <https://www.solidaritycenter.org/mexico-app-based-drivers-hail-new-platform-law-2>.
- 53 Nithin Coca, "Long Silenced, Gig Workers in Indonesia are Organizing and Fighting for their Rights," *Equal Times*, March 14, 2025, <https://www.ituc-csi.org/long-silenced-gig-workers-in>.
- 54 "Data Workers' Inquiry," accessed April 17, 2025, <https://data-workers.org/>.
- 55 "Shining a Light on Tech's Shadow Workforce," *TechEquity Contract Worker Disparity Project*, accessed April 17, 2025, <https://contractwork.techequitycollaborative.org/>.
- 56 "Dec 19, 2024 - CWA Members at Google Ratify Historic First Contract," *CWA*, December 19, 2024, <https://cwa-union.org/news/e-newsletter/2024-12-19>.
- 57 See, e.g., Laura Newcombe, "After Tesla Employees Signed Letter Asking Elon Musk to Resign, at Least One Has Been Fired," *Gizmodo*, May 12, 2025, <https://gizmodo.com/after-tesla-employees-signed-letter-asking-elon-musk-to-resign-at-least-one-has-been-fired-2000600934>; and Bobby Allyn, "Former Palantir Workers Condemn Company's Work with Trump Administration," *NPR*, May 5, 2025, <https://www.npr.org/2025/05/05/nx-s1-5387514/palantir-workers-letter-trump>.
- 58 Accountable Tech et al., *Zero Trust AI Governance*, August 10, 2023, <https://ainowinstitute.org/wp-content/uploads/2023/08/Zero-Trust-AI-Governance.pdf>.
- 59 AI Now Institute, "Algorithmic Accountability: Moving Beyond Audits," in *2023 Landscape: Confronting Tech Power*, AI Now Institute, <https://ainowinstitute.org/publications/algorithmic-accountability>.
- 60 A range of bills of this variety have appeared across state legislatures. See Public Citizen, "Tracker: State Legislation on Intimate Deepfakes," last updated April 15, 2025, <https://www.citizen.org/article/tracker-intimate-deepfakes-state-legislation>; Public Citizen, "Tracker: State Legislation on Deepfakes in Elections," last updated April 17, 2025, <https://www.citizen.org/article/tracker-legislation-on-deepfakes-in-elections>; and California Education Code § 87359.2 (2024); H.B. 1709, 89th Leg. (Tex. 2025); H.B. 1709, 89th Leg. (Tex. 2025).
- 61 Cybersecurity & Infrastructure Security Agency, "CISA and UK NCSC Unveil Joint Guidelines for Secure AI System Development," November 26, 2023, <https://www.cisa.gov/news-events/alerts/2023/11/26/cisa-and-uk-ncsc-unveil-joint-guidelines-secure-ai-system-development>.
- 62 Chloe Xiant, "OpenAI's GPT-4 Is Closed Source and Shrouded in Secrecy," *Vice*, March 16, 2023, <https://www.vice.com/en/article/ak-3w5a/openais-gpt-4-is-closed-source-and-shrouded-in-secrecy>.
- 63 Kyle Wiggers, "Many companies won't say if they'll comply with California's AI training transparency law," *TechCrunch*, October 4, 2024, <https://techcrunch.com/2024/10/04/many-companies-wont-say-if-theyll-comply-with-californias-ai-training-transparency-law>.
- 64 Matt Scherer, "FAQ on Colorado's Consumer Artificial Intelligence Act (SB 24-205)," *Center for Democracy and Technology*, December 17, 2024, <https://cdt.org/insights/faq-on-colorados-consumer-artificial-intelligence-act-sb-24-205/#h-the-law-requires-companies-to-give-some-information-to-consumers-proactively-without-the-consumer-having-to-request-it-why-is-that-proactive-disclosure-necessary>.
- 65 Tamara Chuang, "Attempt to Tweak Colorado's Controversial, First-in-the-Nation Artificial Intelligence Law is Killed," *Colorado Sun*, May 5, 2025, <https://coloradosun.com/2025/05/05/colorado-artificial-intelligence-law-killed>.
- 66 Matt Perault, "Regulate AI Use, Not AI Development," *Andresen Horowitz*, January 27, 2025, <https://a16z.com/regulate-ai-use-not-ai-development>; Jacob Helberg, "11 Elements of American AI Dominance," *Republic*, July 26, 2024, <https://therepublicjournal.com/journal/11-elements-of-american-ai-supremacy>.
- 67 AI Now Institute, "General Purpose AI Poses Serious Risks, Should Not Be Excluded From the EU's AI Act," policy brief, April 13, 2023, <https://ainowinstitute.org/publications/gpai-is-high-risk-should-not-be-excluded-from-eu-ai-act>.
- 68 Anna Lenhart and Sarah Myers West, *Lessons from the FDA for AI*, AI Now Institute, August 1, 2024, <https://ainowinstitute.org/publications/research/lessons-from-the-fda-for-ai>.

- 69 Amba Kak and Sarah Myers West, “General Purpose AI Poses Serious Risks, Should Not Be Excluded from the EU’s AI Act,” policy report, AI Now Institute, April 13, 2023, <https://ainowinstitute.org/publication/gpai-is-high-risk-should-not-be-excluded-from-eu-ai-act>.
- 70 See Margaret Mitchell et al., “Model Cards for Model Reporting,” FACCT, January 29, 2019, <https://doi.org/10.1145/3287560.3287596>; Jessica Ji, “What Does AI Red-Teaming Actually Mean?” Center for Security and Emerging Technology (CSET), October 24, 2023, <https://cset.georgetown.edu/article/what-does-ai-red-teaming-actually-mean>; and Heidy Khlaaf and Sarah Myers West, “Safety Co-Option and Compromised National Security: The Self-Fulfilling Prophecy of Weakened AI Risk Thresholds,” April 21, 2025, <https://arxiv.org/abs/2504.15088>.
- 71 Lenhart and West, *Lessons from the FDA for AI*.
- 72 OpenAI, “Advancing Red Teaming With People and AI,” November 21, 2024, <https://openai.com/index/advancing-red-teaming-with-people-and-ai>; Microsoft AI Red Team, *Lessons from Red Teaming 100 Generative AI Products*, Microsoft, 2024, https://ashy-coast-00ae-b501e6.azurestaticapps.net/MS_AIRT_Lessons_eBook.pdf; Rebecca Pifer, “CHAI Launches Registry for Health AI Model Cards,” *Healthcare Dive*, February 28, 2025, <https://www.healthcaredive.com/news/coalition-health-ai-model-card-registry/741037>.
- 73 Khlaaf and West, “Safety Co-Option and Compromised National Security.”
- 74 “Fmr. Google CEO Says No One in Government Can Get AI Regulation ‘Right,’” *NBC News*, May 15, 2023, <https://www.nbcnews.com/meet-the-press/video/fmr-google-ceo-says-no-one-in-government-can-get-ai-regulation-right-174442053869>.
- 75 White House, “FACT SHEET: Biden-Harris Administration Secures Voluntary Commitments from Eight Additional Artificial Intelligence Companies to Manage the Risks Posed by AI,” September 12, 2023, <https://bidenwhitehouse.archives.gov/briefing-room/statements-releases/2023/09/12/fact-sheet-biden-harris-administration-secures-voluntary-commitments-from-eight-additional-artificial-intelligence-companies-to-manage-the-risks-posed-by-ai>.
- 76 Melissa Heikkilä, “AI Companies Promised to Self-Regulate One Year Ago. What’s Changed?” *Technology Review*, July 22, 2024, <https://www.technologyreview.com/2024/07/22/1095193/ai-companies-promised-the-white-house-to-self-regulate-one-year-ago-whats-changed>.
- 77 Maxwell Zeff, “OpenAI Ships GPT-4.1 Without a Safety Report,” *TechCrunch*, April 15, 2025, <https://techcrunch.com/2025/04/15/openai-ships-gpt-4-1-without-a-safety-report>.
- 78 See Lauren Feiner, “The Technology Team at Financial Regulator CFPB Has Been Gutted,” *Verge*, February 14, 2025, <https://www.theverge.com/policy/612933/cfpb-tech-team-gutted-trump-doge-elon-musk>; and Lauren Feiner and Alex Heath, “DOGE Has Arrived at the FTC,” *Verge*, April 4, 2025, <https://www.theverge.com/news/643674/doge-members-spotted-ftc-elon-musk>.
- 79 Laura Weidinger et al., “Toward an Evaluation Science for Generative AI Systems,” *arXiv*, updated March 13, 2025, <https://arxiv.org/abs/2503.05336>.
- 80 Jevan Hutson and Ben Winters, “America’s Next ‘Stop Model!’: Model Deletion,” *Georgia Law Technology Review*, September 20, 2022, <https://doi.org/10.2139/ssrn.4225003>; Tiffany C. Li, “Algorithmic Destruction,” *Southern Methodist University Law Review* 75, no. 479 (2022), <https://doi.org/10.25172/smulr.75.3.2>; Cristina Caffarra and Robin Berjon, “‘Google is a Monopolist’—Wrong and Right Ways to Think About Remedies,” *Tech Policy Press*, August 9, 2024, <https://www.techpolicy.press/google-is-a-monopolist-wrong-and-right-ways-to-think-about-remedies>.
- 81 United States v. Amazon.com, No. 2:23-cv-00811-TL (W.D. Wash. July 21, 2023); See also Rohit Chopra, “Prepared Remarks of CFPB Director Rohit Chopra at the Federal Reserve Bank of Philadelphia on the Personal Financial Data Rights Rule,” Consumer Financial Protection Bureau, October 22, 2024, <https://www.consumerfinance.gov/about-us/newsroom/prepared-remarks-of-cfpb-director-rohit-chopra-at-the-federal-reserve-bank-of-philadelphia-on-the-personal-financial-data-rights-rule>.
- 82 Plaintiffs’ Initial Proposed Final Judgment, United States v. Google, No. 1:20-cv-03010-APM, (D.D.C. August 5, 2024) and Colorado v. Google, Case No. 20-cv-3715-APM, (D.D.C. August 5, 2024).
- 83 US General Services Administration, “Frequently Asked Questions: Suspension & Debarment,” accessed May 5, 2025, <https://www.gsa.gov/policy-regulations/policy/acquisition-policy/office-of-acquisition-policy/gsa-acq-policy-integrity-workforce/suspension-debarment-and-agency-protests/suspension-debarment-faq>.
- 84 Federal Trade Commission, *Partnerships Between Cloud Service Providers and AI Developers*, staff report, January 2025, https://www.ftc.gov/system/files/ftc_gov/pdf/p246201_aipartnerships6breport_redacted_0.pdf; European Commission, “Commission Launches Call for Contributions on Competition in Virtual Worlds and Generative AI,” press release, January 8, 2024, https://ec.europa.eu/commission/presscorner/detail/en/ip_24_85; Competition and Markets Authority, “Microsoft / OpenAI Partnership Merger Inquiry,” Gov.UK, April 15, 2025, <https://www.gov.uk/cma-cases/microsoft-slash-openai-partnership-merger-inquiry>.
- 85 Japan Fair Trade Commission, “Report Regarding Cloud Services,” press release, June 28, 2022, <https://www.jftc.go.jp/en/pressreleases/yearly-2022/June/220628.html>; Authority for Consumers and Markets, “Market Study into Cloud Services,” May 9, 2022, <https://www.acm.nl/en/publications/market-study-cloud-services>; Competition and Markets Authority, *Cloud Services Market Investigation*, January 28, 2025, https://assets.publishing.service.gov.uk/media/67989251419bdbc8514fdee4/summary_of_provisional_decision.pdf; Nick Jones, “Cloud Computing FRI: What We Heard and Learned,” FTC (blog), November 16, 2023, <https://www.ftc.gov/policy/advocacy-research/tech-at-ftc/2023/11/cloud-computing-rfi-what-we-heard-learned>.
- 86 Competition and Markets Authority, *AI Foundation Models Update Paper*, April 11, 2024, https://assets.publishing.service.gov.uk/media/661941a6c1d297c6ad1dfeed/Update_Paper_1_.pdf; Office of Technology, “Tick, Tick, Tick. Office of Technology’s Summit on AI,” FTC, January 18, 2024, <https://www.ftc.gov/policy/advocacy-research/tech-at-ftc/2024/01/tick-tick-tick-office-technologys-summit-ai>.
- 87 Federal Trade Commission, “FTC, DOJ, and International Enforcers Issue Joint Statement on AI Competition Issues,” press release, July 23, 2024, <https://www.ftc.gov/news-events/news/press-releases/2024/07/ftc-doj-international-enforcers-issue-joint-statement-ai-competition-issues>.

- 88 Laya Neelakandan, “Meta Resorted to ‘Buy-or-Bury Scheme’ With Instagram and WhatsApp Deals, Former FTC Chair Lina Khan Says,” *CNBC*, April 14, 2025, <https://www.cnbc.com/2025/04/14/former-ftc-chair-khan-meta-acquisitions-instagram-whatsapp.html>.
- 89 Michael Liedtke, “Google’s Digital Ad Network Declared an Illegal Monopoly, Joining Its Search Engine in Penalty Box,” *Associated Press*, April 17, 2025, <https://apnews.com/article/google-illegal-monopoly-advertising-search-a1e4446c4870903ed05c03a2a03b581e>.
- 90 Competition and Markets Authority, *AI Foundation Models Technical Update Report*, April 16, 2024, https://assets.publishing.service.gov.uk/media/661e5a4c7469198185bd3d62/AI_Foundation_Models_technical_update_report.pdf.
- 91 States might be ideal testing grounds here, especially as many of these moves might require legislative changes. A recent report by the Institute for Local Self Reliance identifies state antitrust laws as being ripe for enforcement and legislative reform—pointing to enforcement victories in Washington, Colorado, Arizona, and New York, among other states—as well as successful attempts at passing targeted legislative reforms giving these enforcers better tools to grapple with anticompetitive conduct. See Ron Knox, *The State(s) of Antitrust: How States Can Strengthen Monopoly Laws and Get Enforcement Back on Track*, Institute for Local Self-Reliance, April 2025, <https://ilsr.org/wp-content/uploads/2025/03/ILSR-StatesofAntitrust.pdf>; and Knox, “Mapping Our State Antitrust Laws,” *Institute for Local Self-Reliance*, April 1, 2025, <https://ilsr.org/articles/state-antitrust-law-database-and-map>.
- 92 See Alli Finn, *The Costs of Data Centers to Our Communities—and How to Fight Back*, Kairos, 2025, 8–10, <https://www.kairosfellows.org/fightdatacenters>; More Perfect Union, “I Live 400 Yards From Mark Zuckerberg’s Massive Data Center,” YouTube video, March 27, 2025, https://www.youtube.com/watch?v=DGjj7wDYail&ab_channel=MorePerfectUnion.
- 93 Adam Mahoney, “America’s Digital Demand Threatens Black Communities with More Pollution,” *Capital B*, February 25, 2025, <https://capitalbnews.org/ai-data-centers-south-carolina-black-communities>; and Dwaigh Tyndal, “The Dark Side of AI: Climate Chaos, Pollution, and Injustice,” *CommonWealth Beacon*, February 3, 2025, <https://commonwealthbeacon.org/opinion/the-dark-side-of-ai-climate-chaos-pollution-and-injustice>.
- 94 Eliza Martin and Ari Peskoe, *Extracting Profits from the Public: How Utility Ratepayers are Paying for Big Tech’s Power*, Harvard Electricity Law Initiative, March 5, 2025, <https://eelp.law.harvard.edu/wp-content/uploads/2025/03/Harvard-ELI-Extracting-Profits-from-the-Public.pdf>.
- 95 Dave Williams, “Georgia Senate Committee Passes Bill to Protect Residents From Data Center Costs,” *Current*, February 25, 2025, <https://thecurrentga.org/2025/02/25/georgia-senate-committee-passes-bill-to-protect-residents-from-data-center-costs>; Josh Janey, “Dominion Energy Proposes Hefty Rate Hikes,” *Virginia Business*, April 1, 2025, <https://virginiabusiness.com/dominion-energy-proposes-hefty-rate-hikes>; and Citizens Action Coalition, “Duke Energy Rate Hike (2024),” accessed May 5, 2025, <https://www.citact.org/duke-rate-hike-2024>.
- 96 See Kasia Tarczynska and Greg LeRoy, *Cloudy With a Loss of Spending Control: How Data Centers Are Endangering State Budgets*, Good Jobs First, April 2025, <https://goodjobsfirst.org/cloudy-with-a-loss-of-spending-control-how-data-centers-are-endangering-state-budgets>; Good Jobs First, “Subsidy Tracker,” accessed May 18, 2025, <https://subsidytracker.goodjobsfirst.org>; and Finn, *The Costs of Data Centers to Our Communities*.
- 97 Eric Bonds and Viktor Newby, “Data Centers, Non-Disclosure Agreements and Democracy,” *Virginia Mercury*, April 30, 2025, <https://viriniamercury.com/2025/04/30/data-centers-non-disclosure-agreements-and-democracy>.
- 98 See Issie Lapowsky, “As Data Centers Push Into States, Lawmakers Are Pushing Right Back,” *Tech Policy Press*, March 24, 2025, <https://www.techpolicy.press/as-data-centers-push-into-states-lawmakers-are-pushing-right-back>; and Martin and Peskoe, *Extracting Profits from the Public*.
- 99 Zachary Hansen, “\$17B Data Center Proposal in Metro Atlanta Is Latest to Prompt Moratorium,” *Atlanta Journal-Constitution*, May 8, 2025, <https://www.ajc.com/news/business/coweta-county-is-sues-moratorium-to-pause-data-center-projects/2J464ISZIVB4VMFS-BXZXNVHZE>.
- 100 Matthew Gooding, “Plan for Gas-Powered Data Center Campus in Pittsylvania County, Virginia, Withdrawn Again,” *Data Center Dynamics*, April 15, 2025, <https://www.datacenterdynamics.com/en/news/plan-for-gas-powered-data-center-campus-in-pittsylvania-county-virginia-withdrawn-again>.
- 101 For an example of this demand, see Citizens Action Coalition in Indiana, “CAC Calls for Data Center Moratorium,” press release, October 15, 2024, <https://www.citact.org/news/cac-calls-data-center-moratorium>. For a comprehensive list of groups strategizing around slowing and stopping data center development, see Finn, *The Costs of Data Centers to Our Communities*, 24.
- 102 For an example of this policy implementation, see the Amsterdam Metropolitan Government, “City of Amsterdam Puts Halt to New Data Center Construction,” *Hosting Journalist*, December 21, 2023, <https://hostingjournalist.com/news/city-of-amsterdam-puts-halt-to-new-data-center-construction>.
- 103 For an example of this demand, see Beyond Fossil Fuels, which is calling for no new data centers without 24/7 renewable energy. *Beyond Fossil Fuels, System Overload: How New Data Centres COuld Throw Europe’s Energy Transition Off Course* (2025), https://beyond-fossilfuels.org/wp-content/uploads/2025/02/SystemOverload_Report.pdf.
- 104 For organizations advocating on these demands, see Zach Schiller, “Indefensible Tax Breaks for Data Centers Will Cost Ohio,” *Policy Matters Ohio*, January 7, 2025, <https://policymattersohio.org/research/indefensible-tax-breaks-for-data-centers-will-cost-ohio>; and Tarczynska and LeRoy, *Cloudy With a Loss of Spending Control*.
- 105 For examples of this demand, see Kaisa Tarczynska, *Money Lost to the Cloud: How Data Centers Benefit From State and Local Government Subsidies*, Good Jobs First, October 2016, <https://goodjobsfirst.org/wp-content/uploads/docs/pdf/datacenters.pdf>; and Tarczynska and LeRoy, *Cloudy With a Loss of Spending Control*.
- 106 For a relevant proposed bill, see: H.B. 2578, Reg. Sess. (Va. 2025). For a slightly different proposed bill, see: S.B. 1196, Reg. Sess. (Va. 2025) (restricting data centers from gaining state sales tax exemption if they use onsite diesel generation).
- 107 For a proposed bill, see: H.B. 2035, Reg. Sess. (Va. 2025); Markus Schmidt, “As Data Center Boom Continues, Va. Legislators Broach

- New Regulations,” *Virginia Mercury*, January 14, 2025, <https://virginiamercury.com/2025/01/14/as-data-center-boom-continues-va-legislators-broach-new-regulations>.
- 108 For a legislative example currently pending veto override, see: H.B. 1601, Reg. Sess. (Va 2025) (requiring a site assessment to examine the sound profile of the energy facility and allowing—but not requiring—site assessments for the effect of the facility on other factors).
 - 109 Vasili Varlamos, “Oregon House Passes Bill Making Large Data Centers Pay for Power Grids,” *KATU*, April 22, 2025, <https://katu.com/news/politics/oregon-house-passes-bill-making-large-data-centers-pay-for-power-grid-costs-electricity-money-consumers-cost-bill-politics>.
 - 110 For a relevant proposed bill, see: S. 6394, 2025-2026 Leg. Sess. (N.Y. 2025).
 - 111 Leslie B. Muñiz, “Ratepayer Advocates Hail ‘Landmark’ Settlement With Data Centers, Utility Company,” *Indiana Capital Chronicle*, November 26, 2024, <https://indianacapitalchronicle.com/2024/11/26/ratepayer-advocates-hail-landmark-settlement-with-data-centers-utility-company>; Varlamos, “Oregon Houses Passes Bill.”
 - 112 Amba Kak and Sarah Myers West, *2023 Landscape: Confronting Tech Power*, AI Now Institute, April 11, 2023, <https://ainowinstitute.org/publications/research/2023-landscape-confronting-tech-power>.
 - 113 See also AI Now Institute, “Tracking the US and China AI Arms Race,” April 11, 2023, <https://ainowinstitute.org/publications/tracking-the-us-and-china-ai-arms-race>.
 - 114 H.B. 2035, Reg. Sess. (Va. 2025).
 - 115 Because H.B. 2035 was stricken from the docket, it was never discussed on the record. However, the Senate version of the bill is instructive: The original Senate bill closely mirrors H.B. 2035. S.B. 1353, Reg. Sess. (Va. 2025) (Introduced), <https://lis.blob.core.windows.net/files/1019728.PDF>. The bill’s substitute version, however, directly discusses national security concerns. See S.B. 1353, Reg. Sess. (Va. 2025) (Senate Substitute), <https://lis.blob.core.windows.net/files/1034365.PDF>.
 - 116 See examples in Heidy Khlaaf and Sarah M. West, “The Rush to A.I. Threatens National Security,” *New York Times*, January 27, 2025, <https://www.nytimes.com/2025/01/27/opinion/ai-trump-military-national-security.html>.
 - 117 Former FTC Chair Lina Khan urged skepticism toward the argument that the United States must coddle its domestic monopolies—or “national champions”—to maintain global dominance. In her remarks to the FTC Technology Summit last year, she discussed the pitfalls of policymakers shielding large corporations from competition and identified the “decline and degradation” of Boeing as an example of this historically failed strategy. “The difference between Boeing and many of [the other corporations coddled by policymakers],” she said, “is there’s simply no masking airplanes falling apart in the sky.” See Lina Khan, “Remarks of Chair Lina M. Khan FTC Tech Summit,” speech, Federal Trade Commission, January 25, 2024, https://www.ftc.gov/system/files/ftc_gov/pdf/2024.01.25-chair-khan-remarks-at-ot-tech-summit.pdf; and Lina Khan, “Remarks by Chair Lina M. Khan As Prepared for Delivery,” speech, Carnegie Endowment for International Peace, March 13, 2024, https://www.ftc.gov/system/files/ftc_gov/pdf/2024.03.13-chair-khan-remarks-at-the-carnegie-endowment-for-intl-peace.pdf.
 - 118 Microsoft was quick to find a scapegoat by reviewing whether EU digital regulations may have come into play in the incident. See Competition Policy International, “Microsoft Investigates European Commission’s Role in Recent CrowdStrike Outage,” July 22, 2024, <https://www.pymnts.com/cpi-posts/microsoft-investigates-european-commissions-role-in-recent-crowdstrike-outage>; and Cristiano Lima-Strong and Cat Zakrzewski, “Microsoft’s Global Sprawl Comes Under Fire After Historic Outage,” *Washington Post*, July 20, 2024, <https://www.washingtonpost.com/technology/2024/07/20/microsoft-outage-crowdstrike-regulatory-scrutiny>.
 - 119 Declan Harty and Steven Overly, “Gensler’s Warning: Unchecked AI Could Spark Future Financial Meltdown,” *Politico*, March 19, 2024, <https://www.politico.com/news/2024/03/19/sec-gensler-artificial-intelligence-00147665>.
 - 120 In addition, we’re seeing AI systems used to manage and distribute public benefits and essential services. The flaws of large-scale AI models—like leakage, hallucinations, and bias (see Chapter 3)—are even more dangerous in these circumstances because they risk directly undermining vital public services. See U.S. Department of Health & Human Services, “Plan for Promoting Responsible Use of Artificial Intelligence in Automated and Algorithmic Systems by State, Local, Tribal, and Territorial Governments in Public Benefit Administration,” March 28, 2024, <https://www.hhs.gov/sites/default/files/public-benefits-and-ai.pdf>.
 - 121 Heidy Khlaaf, Sarah Myers West and Meredith Whittaker. *Mind the Gap: Foundation Models and the Covert Proliferation of Military Intelligence, Surveillance, and Targeting*. AI Now Institute, October 18, 2024, <https://arxiv.org/abs/2410.14831>.
 - 122 U.S. Department of Defense, “DOD Announces Establishment of Generative AI Task Force,” press release, August 10, 2023, <https://www.defense.gov/News/Releases/Release/Article/3489803/dod-announces-establishment-of-generative-ai-task-force>.
 - 123 See Amazon, “The Trusted Cloud for Government,” accessed May 18, 2025, <https://aws.amazon.com/government-education/government>; and Karen Dahut and Phil Venables, “How to Build a Secure Foundation for American Leadership in AI,” Google, October 18, 2023, <https://blog.google/technology/ai/how-to-build-a-secure-foundation-for-american-leadership-in-ai>.
 - 124 Associated Press, “Scathing Federal Report Rips Microsoft for Response to Chinese Hack,” *Voice of America* (VOA), April 3, 2024, <https://www.voanews.com/a/scathing-federal-report-rips-microsoft-for-response-to-chinese-hack/7556113.html>.
 - 125 Pranshu Verma, “Digital Clones Made by AI Tech Could Make Hollywood Extras Obsolete,” *Washington Post*, July 19, 2023, <https://www.washingtonpost.com/technology/2023/07/19/ai-actors-fear-sag-strike-hollywood>.
 - 126 Jacob Biba, “Top 27 Humanoid Robots in Use Right Now,” *Built In*, updated March 10, 2025, <https://builtin.com/robotics/humanoid-robots>.
 - 127 See for example Jack Aldane and Richard Johnstone, “UK Deputy PM Says Embracing AI ‘Will Reduce Civil Service Headcount’, Estonia’s ‘Personal Government’ Vision, and More,” *Global Government Forum*, March 7, 2024, <https://www.globalgovernmentforum.com/uk-deputy-pm-says-embracing-ai-will-reduce-civil-service-headcount-estonias-personal-government-vision-and-more>; and Matteo Wong, “DOGE’s Plans to Replace Humans with AI Are Already Under Way,”

- Atlantic*, March 10, 2025, <https://www.theatlantic.com/technology/archive/2025/03/gsa-chat-doge-ai/681987>.
- 128 Ifeoma Ajunwa, “A.I. and Captured Capital,” *Yale Law Journal Forum* 134 (2025): <https://www.yalelawjournal.org/forum/ai-and-captured-capital>.
 - 129 Ben Thompson, “Enterprise Philosophy and the First Wave of AI,” *Stratechery*, September 24, 2024, <https://stratechery.com/2024/enterprise-philosophy-and-the-first-wave-of-ai>.
 - 130 Tyler Katzenberger and Chase DiFelicianantonio, “Labor Channels George Orwell on AI,” *Politico*, March 4, 2025, <https://www.politico.com/newsletters/politico-technology-california-decoded-pre-view/2025/03/04/labor-channels-george-orwell-on-ai-211284>.
 - 131 A.B. 1331, Reg. Sess. (Cal. 2025).
 - 132 S.B. 1484, Gen. Assemb. (Conn. 2025).
 - 133 S.B. 838, 194th Gen. Court (Mass. 2025).
 - 134 H.B. 262, Reg. Sess. (Vt. 2025).
 - 135 H.B. 1672, Reg. Sess. (Wash. 2025).
 - 136 A.B. 1331, Reg. Sess. (Cal. 2025).
 - 137 A.B. 1221, Reg. Sess. (Cal. 2025).
 - 138 Seth Frotman and Erie Meyer, “Comment on Request on Automated Worker Surveillance and Management,” Consumer Financial Protection Bureau, June 20, 2023, https://files.consumerfinance.gov/f/documents/CFPB_Comment-Letter-to-White-House-OSTP-RFI-on-Worker-Surveillance_06-2023.pdf.
 - 139 Kak and West, 2023 Landscape: Confronting Tech Power.
 - 140 AI Now Institute, “Toxic Competition: Regulating Big Tech’s Data Advantage,” in *2023 Landscape: Confronting Tech Power*, April 11, 2023, <https://ainowinstitute.org/publications/toxic-competition>.
 - 141 See AI Now Institute, “Toxic Competition”; Open Markets, *Stopping Big Tech from Becoming Big AI: A Roadmap for Using competition Policy to Keep Artificial Intelligence Open for All*, October 17, 2024, <https://www.openmarketsinstitute.org/publications/report-stopping-big-tech-big-ai-roadmap>; and Tejas N. Narechania and Ganesh Sitaraman, “An Antimonopoly Approach to Governing Artificial Intelligence,” *Yale Law and Policy Review* 43, no. 1 (Fall 2024), <https://yalelawandpolicy.org/antimonopoly-approach-governing-artificial-intelligence>.
 - 142 Cade Metz et al., “How Tech Giants Cut Corners to Harvest Data for A.I.,” *New York Times*, April 6, 2024, <https://www.nytimes.com/2024/04/06/technology/tech-giants-harvest-data-artificial-intelligence.html>.
 - 143 Amba Kak, Executive Director, AI Now Institute, *Safeguarding Data and Innovation: Building the Foundation for the Use of Artificial Intelligence, Before U.S. House Committee on Energy, Commerce, Subcommittee on Innovation, Data, and Commerce*, 118th Cong., 2023.
 - 144 Charles Rollet and Zack Whittaker, “Elon Musk Staffer Created a DOGE AI Assistant for Making Government ‘Less Dumb,’” *TechCrunch*, February 18, 2025, <https://techcrunch.com/2025/02/18/elon-musk-staffer-created-a-doge-ai-assistant-for-making-government-less-dumb>.
 - 145 Kak and West, 2023 Landscape: Confronting Tech Power.
 - 146 Marc Andreessen and Ben Horowitz, “The Little Tech Agenda,” a16z, July 5, 2024, <https://a16z.com/the-little-tech-agenda>.
 - 147 See, e.g., Brad Smith, Testimony, “Developing and Deploying AI Responsibly: Elements of an Effective Legislative Framework to Regulate AI,” U.S. Senate Judiciary Committee Subcommittee on Privacy, Technology, and the Law, September 12, 2023, https://www.judiciary.senate.gov/imo/media/doc/2023-09-12_pm_-_testimony_-_smith.pdf (calling for federal licensing regimes and a new agency to regulate AI platforms on behalf of Microsoft).
 - 148 Paris Marx, “Silicon Valley’s Dangerous Plan for a Second Trump Term,” *Disconnect* (blog), July 19, 2024, <https://www.disconnect.blog/p/silicon-valleys-dangerous-plan-for-a-second-trump-term-2>.
 - 149 Kylie Robinson, “Inside Elon Musk’s Messy Breakup with OpenAI,” *Verge*, November 18, 2024, <https://www.theverge.com/2024/11/18/24299787/elon-musk-openai-lawsuit-sam-altman-xai-google-deepmind>.
 - 150 Kevin Breuninger, “Musk Undercuts Trump on Stargate AI Investment Announcement,” *CNBC*, January 22, 2025, <https://www.cnn.com/2025/01/22/musk-trump-ai-stargate-openai-softbank.html>.
 - 151 Mohammad Khatami et al., “Google Fired Us for Protesting Its Complicity in the War on Gaza. But We Won’t Be Silenced,” *Nation*, April 29, 2024, <https://www.thenation.com/article/activism/google-firings-gaza-project-nimbus>; Alix Dunn, “Protesting Project Nimbus: Employee Organizing to End Google’s Contract With Israel w/ Dr. Kate Sim,” *Says Maybe* (podcast), May 23, 2024, <https://www.saysmaybe.com/podcast/protesting-project-nimbus>.
 - 152 Matt O’Brien, “Microsoft Employees Protest at 50th Anniversary Party Over Israel Contract,” *Associated Press*, April 4, 2025, <https://apnews.com/article/microsoft-ai-protest-israel-gaza-50th-anniversary-fadcb37bcce7e067f896ec5502d187b6>.
 - 153 Anthony Ha, “Google’s DeepMind UK Team Reportedly Seeks to Unionize,” *TechCrunch*, April 26, 2025, <https://techcrunch.com/2025/04/26/googles-deepmind-uk-team-reportedly-seeks-to-unionize>.
 - 154 See Alex Pascal and Ganesh Sitaraman, *The National Security Case for Public AI*, Vanderbilt Policy Accelerator (VPA), 2024, <https://cdn.vanderbilt.edu/vu-URL/wp-content/uploads/sites/412/2024/09/27201409/VPA-Paper-National-Security-Case-for-AI.pdf>; Public AI Network, *Public AI: Infrastructure for the Common Good*, August 10 2024, <https://drive.google.com/file/d/1b-cCPdRRHyUGFB23--6wn4j1f9mRgdG2QF/view>; Eleanor Shearer et al., “The Role of Public Compute,” *Ada Lovelace Institute* (blog), April 24, 2024; Eleanor Shearer, “Take Back the Future: Democratising Power Over AI,” *Common Wealth*, September 11, 2024; Lawrence Lessig, “Why This is the Moment for Public AI,” keynote address, *Building A More Public AI Ecosystem*, YouTube, Aug 29, 2024, <https://youtu.be/uzQDeYV1RJQ?si=nGCQFaprfZVjeFln>; Bruce Schneier and Nathan Sanders, “Build AI by the People, for the People,” *Foreign Policy*, June 12, 2023, <https://foreignpolicy.com/2023/06/12/ai-regulation-technology-us-china-eu-governance>; Nick Marda et al., “Creating a Public Counterpoint for AI,” *Mozilla, Distilled* (blog), September 30, 2024, <https://blog.mozilla.org/en/mozilla/ai/public-ai-counterpoint>.

- 155 NaWchania and Sitaraman, “An Antimonopoly Approach to Governing Artificial Intelligence.”
- 156 National Science Foundation, “National Artificial Intelligence Research Resource Pilot,” accessed May 3, 2025, <https://www.nsf.gov/focus-areas/artificial-intelligence/nairr>.
- 157 Bryce Elder, “The Next Wave of AI Hype Will Be Geopolitical. You’re Paying,” *Financial Times*, May 29, 2024, <https://www.ft.com/content/a60c3c7b-1c48-485d-adb7-5bc2b7b1b650>.
- 158 AI Now Institute, *AI Nationalism(s): Global Industrial Policy Approaches to AI*, March 12, 2024, <https://ainowinstitute.org/publications/research/ai-nationalisms-global-industrial-policy-approaches-to-ai>.
- 159 Eric Schmidt et al., *Final Report*, National Security Commission on Artificial Intelligence, 2021, <https://www.dwt.com/-/media/files/blogs/artificial-intelligence-law-advisor/2021/03/nscai-final-report-2021.pdf>.
- 160 AI Now Institute and Society Research Institute, “Democratize AI? How the Proposed National AI Research Resource Falls Short,” October 5, 2021, <https://ainowinstitute.org/publication/democratize-ai-how-the-proposed-national-ai-research-resource-falls-short>.
- 161 AI Now Institute and Data & Society, “Request for Information (RFI) on an Implementation Plan for a National Artificial Intelligence Research Resource,” October 1, 2021, <https://ainowinstitute.org/wp-content/uploads/2023/06/AINow-DS-NAIRR-comment.pdf>.
- 162 NAIRR Pilot, “NAIRR Pilot Resource Requests to Advance AI Research,” accessed May 3, 2025, <https://nairrpilot.org/allocations>.
- 163 National Science Foundation, “National Artificial Intelligence Research Resource Pilot,” accessed May 3, 2025, <https://www.nsf.gov/focus-areas/artificial-intelligence/nairr>.
- 164 NAIRR Pilot, “Resource Allocations,” accessed May 3, 2025, <https://nairrpilot.org/projects/awarded>.
- 165 Joana Varon et al., “AI Commons: Nourishing Alternatives to Big Tech Monoculture,” *One Project*, 2024, <https://oneproject.org/ai-commons>.
- 166 Scott Wiener, “Senator Wiener Introduces Legislation to Protect AI Whistleblowers & Boost Responsible AI Development,” press release, February 28, 2025, <https://sd11.senate.ca.gov/news/senator-wiener-introduces-legislation-protect-ai-whistleblowers-boost-responsible-ai>; Chase DiFelicianantonio, “Wiener Launches Bill to Require Testing AI Tools for Safety. Here’s What That Means,” *San Francisco Chronicle*, February 8, 2024, <https://www.sfchronicle.com/tech/article/wiener-launches-bill-require-safety-testing-ai-18656225.php>.
- 167 “Governor Hochul Unveils Fifth Proposal of 2024 State of the State: Empire AI Consortium to Make New York the National Leader in AI Research and Innovation,” New York State Governor, January 8, 2024, <https://www.governor.ny.gov/news/governor-hochul-unveils-fifth-proposal-2024-state-state-empire-ai-consortium-make-new-york>.



AINOW

ainowinstitute.org