

Executive Summary

Artificial intelligence¹ is captivating our attention, generating both fear and awe about what's coming next. As increasingly dire prognoses about AI's future trajectory take center stage in the headlines about generative AI, it's time for regulators, and the public, to ensure that there is nothing about artificial intelligence (and the industry that powers it) that we need to accept as given. This watershed moment must also swiftly give way to action: to galvanize the considerable energy that has already accumulated over several years towards developing meaningful checks on the trajectory of AI technologies. This must start with confronting the concentration of power in the tech industry.

The AI Now Institute was founded in 2017, and even within that short span we've witnessed similar hype cycles wax and wane: when we wrote the 2018 AI Now report, the proliferation of facial recognition systems already seemed well underway, until pushback from local communities pressured government officials to pass bans in cities across the United States and around the world.² Tech firms were associated with the pursuit of broadly beneficial innovation,³ until worker-led organizing, media investigations, and advocacy groups shed light on the many dimensions of tech-driven harm.⁴

These are only a handful of examples, and what they make clear is that **there is nothing about artificial intelligence that is inevitable.** Only once we stop seeing AI as synonymous with progress can we establish popular control over the trajectory of these technologies and meaningfully confront their serious social, economic, and political impacts—from exacerbating patterns of inequality in housing,⁵ credit,⁶ healthcare,⁷ and education⁸ to inhibiting workers' ability to organize⁹ and incentivizing content production that is deleterious to young people's mental and physical health.¹⁰

In 2021, several members of AI Now were asked to join the Federal Trade Commission (FTC) to advise the Chair's office on artificial intelligence.¹¹ This was, among other things, a recognition of the growing centrality of AI to digital markets and the need for regulators to pay close attention to potential harms

¹ The term 'artificial intelligence' has come to mean many different things over the course of its history, and may be best understood as a marketing term rather than a fixed object. See for example: Michael Atleson, "Keep your AI Claims in Check", Federal Trade Commission, February 27, 2023.; Meredith Whittaker, "Signal, and the Tech Business Model Shaping Our World", Conference on Steward-Ownership 2023., Annie Lowery, "AI Isn't Omnipotent. It's Janky", The Atlantic, April 3, 2023.

² Meredith Whittaker, Kate Crawford, Roel Dobbe, Genevieve Fried, Elizabeth Kazianas, Varoon Mathur, Sarah Myers West, Rashida Richardson, Jason Schultz, Oscar Schwartz, *AI Now 2018 Report*, AI Now Institute, December 2018.

Tom Simonite, "Face Recognition Is Being Banned—But It's Still Everywhere," *Wired*, December 22, 2021.

³ See Jenna Wortham, "Obama Brought Silicon Valley to Washington," *New York Times*, October 25, 2016.; and Cecilia Kang and Juliet Eilperin, "Why Silicon Valley Is the New Revolving Door for Obama Staffers," *Washington Post*, February 28, 2015.

⁴ Varoon Mathur, Genevieve Fried, and Meredith Whittaker, "AI in 2019: A Year in Review," *Medium*, October 9, 2019.

⁵ Robert Bartlett, Adair Morse, Richard Stanton, and Nancy Wallace, "Consumer-Lending Discrimination in the FinTech Era," *Journal of Financial Economics* 143, no. 1 (January 1, 2022): 30–56.

⁶ Christopher Gilliard, "Prepared Testimony and Statement for the Record," Hearing on "Banking on Your Data: The Role of Big Data in Financial Services," House Financial Services Committee Task Force on Financial Technology, 2019.

⁷ Ziad Obermeyer, Brian Powers, Christine Vogeli, and Sendhil Mullainathan, "Dissecting Racial Bias in an Algorithm Used to Manage the Health of Populations," *Science* 366, no. 6464 (October 25, 2019): 447–53.

⁸ Rashida Richardson and Marci Lerner Miller, "The Higher Education Industry Is Embracing Predatory and Discriminatory Student Data Practices," *Slate*, January 13, 2021.

⁹ *Ibid.*

¹⁰ See Zach Praiss, "New Poll Shows Dangers of Social Media Design for Young Americans, Sparks Renewed Call for Tech Regulation," *Accountable Tech*, March 29, 2023; and Tawnell D. Hobbs, Rob Barry, and Yoree Koh, "'The Corpse Bride Diet': How TikTok Inundates Teens with Eating-Disorder Videos," *Wall Street Journal*, December 17, 2021.

¹¹ Federal Trade Commission, "FTC Chair Lina M. Khan Announces New Appointments in Agency Leadership Positions," press release, November 19, 2021.

to consumers and competition. Our experience within the US government helped clarify the path for the work ahead.

ChatGPT was unveiled during the last month of our time at the FTC, unleashing a wave of AI hype that shows no signs of letting up. This underscored the importance of addressing AI's role and impact, not as a philosophical futurist exercise but as something that is being used to shape the world around us here and now. We urgently need to be learning from the "move fast and break things" era of Big Tech; we can't allow companies to use our lives, livelihoods, and institutions as testing grounds for novel technological approaches, experimenting in the wild to our detriment. Happily, we do not need to draft policy from scratch: artificial intelligence, the companies that produce it, and the affordances required to develop these technologies already exist in a regulated space, and companies need to follow the laws already in effect. This provides a foundation, but we'll need to construct new tools and approaches, built on what we already have.

There is something different about this particular moment: it is primed for action. We have abundant research and reporting that clearly documents the problems with AI and the companies behind it. This means that more than ever before, we are prepared to **move from identifying and diagnosing harms to taking action to remediate them**. This will not be easy, but now is the moment for this work. This report is written with this task in mind: we are drawing from our experiences inside and outside government to outline an agenda for how we—as a group of individuals, communities, and institutions deeply concerned about the impact of AI unfolding around us—can meaningfully confront the core problem that AI presents, and one of the most difficult challenges of our time: **the concentration of economic and political power in the hands of the tech industry—Big Tech in particular**.

There is no AI without Big Tech.

Over the past several decades, a handful of private actors have accrued power and resources that rival nation-states while developing and evangelizing artificial intelligence as critical social infrastructure. AI is being used to make decisions that shape the trajectory of our lives, from the deeply impactful, like what kind of job we get and how much we're paid; whether we can access decent healthcare and a good education; to the very mundane, like the cost of goods on the grocery shelf and whether the route we take home will send us into traffic.

Across all of these domains, the same problems show themselves: the technology doesn't work as claimed, and it produces high rates of error or unfair and discriminatory results. But the visible problems are only the tip of the iceberg. The opacity of this technology means we may not be informed when AI is in use, or how it's working. This ensures that we have little to no say about its impact on our lives.

This is underscored by a core attribute of artificial intelligence: it is foundationally reliant on resources that are owned and controlled by only a handful of Big Tech firms.

The dominance of Big Tech in artificial intelligence plays out along three key dimensions:

1. **The Data Advantage:** Firms that have access to the widest and deepest swath of behavioral data insights through surveillance will have an edge in the creation of consumer AI products. This is reflected in the acquisition strategies adopted by tech companies, which have of late focused on expanding this data advantage. Tech companies have amassed a tremendous degree of economic power, which has enabled them to embed themselves as core infrastructure within a number of industries, from health to consumer goods to education to credit.
2. **Computing Power Advantage:** AI is fundamentally a data-driven enterprise that is heavily reliant on substantial computing power to train, tune, and deploy these models. This is expensive and runs up against material dependencies such as chips and the location of data centers that mean efficiencies of scale apply, as well as labor dependencies on a relatively small pool of highly skilled tech workers that can most efficiently use these resources.¹² Only a handful of companies actually run their own infrastructure – the cloud and compute resources foundational to building AI systems. What this means is that even though “AI startups” abound, they must be understood as barnacles on the hull of Big Tech – licensing server infrastructure, and as a rule competing with each other to be acquired by one or another Big Tech firm. We are already seeing these firms wield their control over necessary resources to throttle competition. For example, Microsoft recently began penalizing customers for developing potential competitors to GPT-4, threatening to restrict their access to Bing search data.¹³
3. **Geopolitical Advantage:** AI systems (and the companies that produce them) are being recast not just as commercial products but foremost as strategic economic and security assets for the nation that need to be boosted by policy, and never restrained. The rhetoric around the US-China AI race has evolved from a sporadic talking point to an increasingly institutionalized stance (represented by collaborative initiatives between government, military, and Big Tech companies) that positions AI companies as crucial levers within this geopolitical fight. This narrative conflates the continued dominance of Big Tech as synonymous with US economic prowess, and ensures the continued accrual of resources and political capital to these companies.

To understand how we got here, we need to look at how tech firms presented themselves in their incipency: their rise was characterized by marketing rhetoric promising that commercial tech would serve the public interest, encoding democratic values like freedom, democracy, and progress. But what’s clear now is that the companies developing and deploying AI and related technologies are motivated by the same things that—structurally and necessarily—motivate all corporations: growth, profit, and rosy market valuations. This has been true from the start.

¹² For example, Microsoft is even rationing access to server hardware internally for some of its AI teams to ensure it has the capacity to run GPT-4. See Aaron Holmes and Kevin McLaughlin, “Microsoft Rations Access to AI Hardware for Internal Teams,” *The Information*.

¹³ Leah Nylen and Dina Bass, “Microsoft Threatens to Restrict Data in Rival AI Search,” March 24, 2023.

Why “Big Tech”?

In this report, we pay special attention to policy interventions that target large tech companies. The term “Big Tech” became popular around 2013¹⁴ as a way to describe a handful of US-based megacorporations, and while it doesn’t have a definite composition, today it’s typically used as shorthand for Google, Apple, Facebook, Amazon, and Microsoft (often abbreviated as GAFAM), and sometimes also includes companies like Uber or Twitter.

It’s a term that draws attention to the unique scale at which these companies operate: the network effects, data, and infrastructural advantages they have amassed. Big Tech’s financial leverage has allowed these firms to consolidate this advantage across sectors from social media to healthcare to education and across media (like the recent pivot to virtual and augmented realities), often through strategic acquisitions. They seek to protect this advantage from regulatory threats through lobbying and similar non-capital strategies that leverage their deep pockets.¹⁵ Following on from narratives around “Big Tobacco,” “Big Pharma,” and “Big Oil,” this framing draws upon lessons from other domains where consolidation of power in industries has led to movements to reassert public accountability. (As one commentator puts it, “society does not prepend the label ‘Big’ with a capital B to an industry out of respect or admiration. It does so out of loathing and fear – and in preparation for battle.”¹⁶) Recent name changes, like Google to Alphabet or Facebook to Meta, also make Big Tech helpful terminology to capture the sprawl of these companies and their continually shifting contours.¹⁷

Focusing on Big Tech is a useful prioritization exercise for tech policy interventions for several reasons:

- **Tackling challenges that either originate from or are exemplified by Big Tech companies can address the root cause of several key concerns:** invasive data surveillance, the manipulation of individual and collective autonomy, the consolidation of economic power, and exacerbation of patterns of inequality and discrimination, to name a few.
- **The Big Tech business and regulatory playbook has a range of knock-on effects on the broader ecosystem, incentivizing and even compelling other companies to fall in line.** Google and Facebook’s adoption of the behavioral advertising business model that effectively propelled commercial surveillance into becoming the business model of the internet is just one example of this.

¹⁴ Nick Dyer-Witheford and Alessandra Mularoni, “Framing Big Tech: News Media, Digital Capital and the Antitrust Movement,” *Political Economy of Communication* 9, no. 2 (2021): 2–20.

¹⁵ Zephyr Teachout and Lina Khan, “Market Structure and Political Law: A Taxonomy of Power,” *Duke Journal of Constitutional Law & Public Policy* 9, no. 1 (2014): 37–74.

¹⁶ Will Oremus, “Big Tobacco. Big Pharma. Big Tech?” *Slate*, November 17, 2017.

¹⁷ Kean Birch Kean and Kelly Bronson, “Big Tech,” *Science as Culture* 31, no. 1 (January 2, 2022): 1–14.

- **Growing dependencies on Big Tech across the tech industry and government make them a single point of failure.** A core business strategy for these firms is to make themselves infrastructural, and much of the wider tech ecosystem relies on them in one way or another, from cloud computing to advertising ecosystems and, increasingly, to payments. This makes these companies both a choke point and a single point of failure. We're also seeing spillover into the public sector. While a whole spectrum of vendors for AI and tech products sells to government agencies, the dependence of government on Big Tech affordances came into particular focus during the height of the pandemic, when many national governments needed to rely on Big Tech infrastructure, networks, and platforms for basic governance functions.

Finally, this report takes aim not just at the pathologies associated with these companies, but also at the broader narratives that justify and normalize them. From unrestricted innovation as a social good to digitization, to data as the only way to see and interpret the world, to platformization as necessarily beneficial to society and synonymous with progress—and regulation as chilling this progress—these narratives pervade the tech industry (and, increasingly, government functioning as well).

Strategic priorities

Where do we go from here? Across the chapters of this report, we offer a set of approaches that, in concert, will collectively enable us to confront the concentrated power of Big Tech. Some of these are bold policy reforms that offer bright-line rules and structural changes. Others aren't in the traditional domain of policy at all, but acknowledge the importance of non regulatory interventions such as collective action, worker organizing, while acknowledging the role public policy can play in bolstering or kneecapping these efforts. We also identify trendy policy responses that seem positive on their surface, but because they fail to meaningfully address power discrepancies should be abandoned. The primary jurisdictional focus for these recommendations is the US, although where relevant we point to policy windows or trends in other jurisdictions (such as the EU) with necessarily global impacts.

Four strategic priorities emerge as particularly crucial for this moment:

1. **Employ strategies that place the burden on companies to demonstrate that they are not doing harm, rather than on the public and regulators to continually investigate, identify, and find solutions for harms after they occur.**

Investigative journalism and independent research has been critical to tech accountability: the hard work of those testing opaque systems has surfaced failures that have been crucial for establishing evidence for tech-enabled harms. But, as we outline in the section on [Algorithmic Accountability](#), as a policy response, audits and similar accountability frameworks dependent on third-party evaluation play

directly into the tech company playbook by positioning responsibility for identifying and addressing harms outside of the company.

The finance sector offers a useful corollary for thinking this through. Much like AI, the actions taken by large financial firms have diffuse and unpredictable effects on the broader financial system and the economy at large. It's hard to predict any particular harm these may cause, but we know the consequences can be severe, and the communities hit hardest are those that already experience significant inequality. After multiple crisis cycles, there's now widespread consensus that the onus needs to be on *companies* to demonstrate that they are mitigating harms and to comply with regulations, rather than on the broader public to root these out.

The tech sector, likewise has diffuse and unpredictable effects not only on our economy, but our information environment and labor market, among many other things. We see value in a due-diligence approach that requires firms to demonstrate their compliance with the law rather than turn to regulators or civil society to show where they haven't complied—similar in orientation to how we already regulate many goods that have significant public impact, like food and medicine. And we need structural curbs like bright lines and no-go zones that identify types of use and domains of implementation that should be barred in any instance, as many cities have already established by passing bans on facial recognition. For example, in the chapter on [Algorithmic Management](#) we identify emotion recognition as a type of technology that should never be deployed, but particularly in the workplace: aside from the clear concerns about its use of pseudoscience and accompanying discriminatory effects, it is fundamentally unethical for employers to seek to draw inferences about their employees' inner state to maximize their profit. And in [Biometric Surveillance](#), we identify the absence of such bright-line measures as the animating force behind a slow creep of facial recognition and other surveillance systems into domains like cars and virtual reality.

We also need to lean further toward scrutiny of harms before they happen rather than waiting to rectify harms after they've already occurred. We discuss what this might look like in the context of merger reviews in the [Toxic Competition](#) section, advocating for an approach to merger reviews that looks to predict and prevent abusive practices before they manifest, and in [Antitrust](#), we break down how needed legal reforms would render certain kinds of mergers invalid in the first place, and put the onus on companies to demonstrate they aren't anti-competitive.

2. Break down silos across policy areas, so we're better prepared to address where advancement of one policy agenda impacts others. Firms play this isolation to their advantage.

One of the primary sources of Big Tech power is the expansiveness of their reach across markets, with digital ecosystems that stretch across vast swathes of the economy. This means that effective tech policy must be similarly expansive, attending to how measures adopted in the advancement of one policy agenda ramify across other policy domains. For example, as we underscore in the section on [Toxic Competition](#), legitimate concerns about third-party data collection must be addressed in a way that doesn't inadvertently enable further concentration of power in the hands of Big Tech firms.

Disconnection between the legal and policy approaches to privacy on the one hand and competition on the other have enabled firms to put forward self-regulatory measures like Google's Privacy Sandbox in

the name of privacy that ultimately will lead to the depletion of both privacy and competition by strengthening Google's ability to collect information on consumers directly while hollowing out its competitors. These disconnects can also prevent progress in one policy domain from carrying over to another. Despite years of carefully accumulated evidence on the fallibility of AI-based content filtration tools, we're seeing variants of the magical thinking that AI tools will be able to scan effectively for illegal content, crop up once again in encryption policy with the EU's recent "chat control" client-side scanning proposals.¹⁸

Policy and advocacy silos can also blunt strategic creativity in ways that foreclose alliance or cross-pollination. We've made progress on this front in other domains, ensuring for example that privacy and national security are increasingly seen as consonant, rather than mutually exclusive, objectives. But AI policy has been undermined too often by a failure to understand AI materially, as a composite of data, algorithmic models, and large-scale computational power. Once we view AI this way, we can understand data minimization and other approaches that limit data collection not only as protecting consumer privacy, but as mechanisms that help mitigate some of the most egregious AI applications, by reducing firms' data advantage as a key source of their power and rendering certain types of systems impossible to build. It was through data protection law that Italy's privacy regulator was the first to issue a ban on ChatGPT¹⁹ and, the week before that, Amsterdam's Court of Appeal ruled automated firing and opaque algorithmic wages to be illegal.²⁰ FTC officials also recently called for leveraging antitrust as a tool to enhance worker power, including to push back against worker surveillance.²¹ This opens up space for advocates working on AI-related issues to form strategic coalitions with those that have been leveraging these policy tools in other domains. This multivariate approach has the added advantage of necessitating that those focused on AI-related issues form strategic coalitions with those that have been leveraging these policy tools in other domains.

Throughout this report, we attempt to establish links between related, but often siloed domains: data protection or competition reform as AI policy (see section on [Data Minimization](#)); [Antitrust](#)], or AI policy as industrial policy (see section on [Algorithmic Accountability](#)).

3. Identify when policy approaches get co-opted and hollowed out by industry, and pivot our strategies accordingly.

The tech industry, with its billions of dollars and deep political networks, has been both nimble and creative in its response to anything perceived as a policy threat. There are relevant lessons from the European experience around the perils of shifting from a "rights-based" regulatory framework, as in the GDPR, to a "risk-based" approach, as in the upcoming AI Act and how the framing of "risk" (as opposed to rights) could tip the playing field in favor of industry-led voluntary frameworks and technical standards.²²

Responding to the growing chorus calling for bans on facial recognition technologies in sensitive social domains, several tech companies pivoted from resisting regulation to claiming to support it, something they often highlighted in their marketing. The fine print showed that what these companies actually

¹⁸ Ross Anderson, "Chat Control or Child Protection?", *University of Cambridge Computer Lab*, October 13, 2022.

¹⁹ Clothilde Goujard, "Italian Privacy Regulator Bans ChatGPT" *Politico*, March 31, 2023.

²⁰ Worker Info Exchange, "Historic Digital Rights Win for WIE and the ADCU Over Uber and Ola at the Amsterdam Court of Appeals", April 4, 2023

²¹ Elizabeth Wilkins, "Rethinking Antitrust", March 30, 2023

²² Fanny Hidvegi and Daniel Leufer, "The EU should regulate AI on the basis of rights, not risks", *Access Now*, February 17, 2021.

supported were soft moves positioned to undercut bolder reform. For example, Washington state's widely critiqued facial recognition law passed with Microsoft's support. The bill prescribed audits and stakeholder engagement, a significantly weaker stance than banning police use which is what many advocates were calling for (see section on [Biometrics](#)).

For example, mountains of research and advocacy demonstrate the discriminatory impacts of AI systems and the fact that these issues cannot be addressed solely at the level of code and data. While the AI industry has accepted that bias and discrimination is an issue, companies have also been quick to narrowly cast bias as a technical problem with a technical fix.

Civil society responses must be nimble in responding to Big Tech subterfuge, and we must learn to recognize such subterfuge early. We draw from these lessons when we argue that there is disproportionate policy energy being directed toward AI and algorithmic audits, impact assessments, and "access to data" mandates. Indeed, such approaches have the potential to eclipse and nullify structural approaches to curbing the harms of AI systems (see section on [Algorithmic Accountability](#)). In an ideal world, such transparency-oriented measures would live alongside clear standards of accountability and bright-line prohibitions. But this is not what we see happening. Instead, a steady stream of proposals position algorithmic auditing as the *primary* policy approach toward AI.

Finally, we also need to stay on top of companies' moves to evade regulatory scrutiny entirely: for example, firms have been seeking to introduce measures in global trade agreements (see section on [Global Digital Trade](#)) that would render regulatory efforts seeking accountability by signatory countries presumptively illegal. And companies have sought to use promises of AI magic as a means of evading stronger regulatory measures, such as by clinging to the familiar false argument that AI can provide a fix for unsolvable problems, such as in content moderation.²³

4. Move beyond a narrow focus on legislative and policy levers and embrace a broad-based theory of change.

To make progress and ensure the longevity of our wins, we must be prepared for the long game, and author strategies that keep momentum going in the face of inevitable political stalemates. We can learn from ongoing organizing in other domains, from climate advocacy (see section on [Climate](#)) that identifies the long-term nature of these stakes, to worker-led organizing (see section on [Algorithmic Management](#)) which has emerged as one of the most effective approaches to challenging and changing tech company practice and policy. We can also learn from shareholder advocacy (see section on [Tech & Financial Capital](#)), which uses companies' own capital strategies to push for accountability measures – one example is the work of the Sisters of St. Joseph of Peace using shareholder proposals to hold Microsoft to account for human rights abuses. The Sisters also used such proposals to seek a ban on the sale of facial recognition to government entities, and to require Microsoft to evaluate how the company's lobbying aligns with its stated principles.²⁴ Across these fronts, there is much to learn from the work of organizers and advocates well-versed in confronting corporate power.

²³ Federal Trade Commission, "Combating Online Harms Through Innovation", Federal Trade Commission, June 2022.

²⁴ See Chris Mills Rodrigo, "Exclusive: Scrutiny Mounts on Microsoft's Surveillance Technology," *The Hill*, June 17, 2021; and Issie Lapowsky, "These Nuns Could Force Microsoft to Put Its Money Where Its Mouth Is," *Protocol*, November 19, 2021.

Windows for policy movement

These strategic priorities are designed to take advantage of current windows for action. We summarize them below, and review each in more detail in the body of the report.

WINDOWS FOR ACTION: THE AI POLICY LANDSCAPE	
<p>Contain tech firms’ data advantage.</p> <p>See: Toxic Competition</p> <p>Data minimization</p>	Data policy <i>is</i> AI policy, and steps taken to curb companies’ data advantage are a key lever in limiting concentration.
	Create bright-line rules that limit firms’ ability to collect data on consumers or produce data about them (also known as data minimization).
	Connect privacy and competition law both in enforcement and in the development of AI policy. Firms are using these disjuncts to their own advantage.
	Reform the merger guidelines and enforcement measures such that consolidation of data advantages receives scrutiny as part of determining whether to allow a merger, and enable enforcers to intervene to stop abusive practices before the harms take place.
<p>Build support for competition reforms as a key lever to reduce concentration in tech.</p> <p>See: Antitrust</p>	Enforce competition laws by aggressively curbing mergers that expand firms’ data advantage and investigating and penalizing companies when they engage in anti-competitive behaviors.
	Be wary of US versus China “AI race” rhetoric used for deregulatory arguments in policy debates on competition, privacy, and algorithmic accountability.
	Pass the full package of antitrust bills from the 117th Congress to give antitrust enforcers stronger tools to challenge abusive practices specific to the tech industry.
	Integrate competition analysis across all tech policy domains – identifying places where platform companies might take advantage of privacy measures to consolidate their own advantage, for example, or how concentration in the cloud market has follow-on effects for security by distributing risk systemically. ²⁵
<p>Regulate ChatGPT and other large-scale models.</p> <p>See: General Purpose AI</p>	Apply lessons from the ongoing debate on the EU AI Act to prevent regulatory carveouts for “general-purpose AI”: large language models (LLMs) and other similar technologies carry systemic risks; their ability to be fine-tuned toward a range of uses requires more regulatory scrutiny, not less.

²⁵ For example, a 2017 outage in Amazon Web Service’s S3 server took out several healthcare and hospital systems: Casey Newton, “How a typo took down S3, the backbone of the internet”, *The Verge*, March 2, 2017.

<p>Regulate ChatGPT and other large-scale models. (CONT.)</p>	<p>Mandate documentation requirements that can provide the evidence to ensure developers of these models are held accountable for data and design choices.</p>
	<p>Enforce existing law on the books to create public accountability in the rollout of generative AI systems and prevent harm to consumers and competition.</p>
	<p>Closely scrutinize claims to 'openness'; generative AI has structural dependencies on resources available to only a few firms.</p>
<p>Displace audits as the primary policy response to harmful AI.</p> <p>See: Algorithmic Accountability</p>	<p>Audits and data-access proposals should not be the primary policy response to harmful AI. These approaches fail to confront the power imbalances between Big Tech and the public, and risk further entrenching power in the tech industry.</p>
	<p>Closely scrutinize claims from a burgeoning audit economy with companies offering audits-as-a-service despite no clarity on the standards and methodologies for algorithmic auditing, nor consensus on their definitions of risk and harm.</p>
	<p>Impose strong structural curbs on harmful AI, such as bans, moratoria, and rules that put the burden on companies to demonstrate that they are fit for public and/or commercial release.</p>
<p>Future-proof against the quiet expansion of biometric surveillance into new domains like cars.</p> <p>See: Biometrics</p>	<p>Develop comprehensive bright-line rules to future-proof biometric regulation from changing forms and use cases.</p>
	<p>Make sure biometric regulation addresses broader inferences, beyond just identification.</p>
	<p>Impose stricter enforcement of data minimization provisions that exist in data protection laws globally as a way to curb the expansion of biometric data collection in new domains like virtual reality and automobiles.</p>
<p>Enact strong curbs on worker surveillance.</p> <p>See: Algorithmic Management</p>	<p>Worker surveillance is fundamentally about employers gaining and maintaining control over workers. Enact policy measures that even the playing field.</p>
	<p>Establish baseline worker protections from algorithmic management and workplace surveillance.</p>
	<p>Shift the burden of proof to developers and employers and away from workers.</p>
	<p>Establish clear red lines around domains (e.g., automated hiring and firing)</p>

	and types of technology (e.g., emotion recognition) that are inappropriate for use in any context.
<p>Prevent “international preemption” by digital trade agreements that can be used to weaken national regulation on algorithmic accountability and competition policy.</p> <p>See: Digital Trade</p>	Nondiscrimination prohibitions in trade agreements should not be used to protect US Big Tech companies from competition regulation abroad.
	Expansive and absolute-secrecy guarantees for source code and algorithms in trade agreements should not be used to undercut efforts to enact laws on algorithmic transparency.
	Upcoming trade agreements like the Indo-Pacific Economic Framework should instead be used to set a more a progressive baseline for digital policy.

It’s time to move: years of critical work and organizing has outlined a clear diagnosis of the problems we face, regulators are primed for action, and we have strategies ready to be deployed immediately for this effort. We’ll also need more: those engaged in this work are out-resourced and out-flanked amidst a significant uptick in industry lobbying and a growing attack on critical work, from companies firing AI Ethics teams to universities shutting down critical research centers. And we face a hostile narrative landscape. The surge in AI hype that opened 2023 has moved things backwards, re-introducing the notion that AI is associated with ‘innovation’ and ‘progress’ and drawing considerable energy toward far-off hypotheticals and away from the task at hand.

We intend this report to provide strategic guidance to inform the work ahead of us, taking a bird’s eye view of the landscape and of the many levers we can use to shape the future trajectory of AI - and the tech industry behind it - to ensure that it is the public, not industry, that this technology serves - if we let it serve at all.