
Why AI? Critiquing AI Industrial Policy

In the past year and a half, with the release of sophisticated AI tools to the public, interest in the speculative promise of AI technologies has exploded across the globe. AI is now firmly positioned as a critical strategic technology for the geopolitical and economic ambitions of nation-states. As governments, and the public, solidify their orientation toward the AI industry, we are forced to wrestle with the relative advantages and disadvantages of AI industrial policy – a term referring to government spending, investment, and regulatory strategies focused on the AI industry. Many governments are increasingly focused on promoting, nurturing, and growing national AI economies—and the industries that underpin them.

This uptick in government support opens up space to question: why AI? Current industrial policy frequently assumes a world in which government spending – whether on the procurement of AI products or development of the industry – is a goal that should take precedence over others in the name of ‘innovation’.

But before investing deeper in the development of an AI industry we need concrete and material answers to questions like: Do efficiencies gained through AI-based climate modeling justify the energy cost of training these models? Should we invest in the advancement of edtech at the cost of providing more students school lunch, or AI medical software over funding home healthcare?

In this essay collection we survey the nationalist narratives and emergent industrial policies being proposed by governments with differing economic and geopolitical motivations. The national policies surveyed in this collection have largely functioned to reinforce the notion that AI is both a socially and geopolitically important sector, and therefore worthy of government strategies (and spending) to promote it. We find that true to historical tradition, governments are using industrial policy as a tool to increase their own geopolitical leverage and economic competitiveness – even as they cloak those objectives underneath thinly defined “AI for social good” aims.

Perspectives in this collection also warn that industrial policy as it is currently structured functions only to further expand and concentrate private (largely US-based) power, under the banner of democratization. Simply diversifying the range of actors involved in AI development while commercial entities continue to define the horizon for development does little to contest their dominance. With such far-reaching consequences for the public at stake – from the allocation of public funds to the rapid promotion of AI tools in sensitive

social domains through procurement mandates – any claims to advancing public good must be put under the scanner.

Here we build on scholarship calling for a more democratic practice of industrial policy¹ to reject the notion that the current trajectory of AI-centered development is inevitable.


Instead of naturalizing the idea that larger and larger scale AI is a self-evident public good, we must start by seeking a clear-eyed understanding of the ways that AI acts on our core social and economic institutions, and to whom AI's benefits and harms accrue.

We need to look at which business models make social benefit more–or less–likely, and what the impact is on workers, the environment, and democracy. The evidence suggests that while benefits accrue to a handful of corporate actors, current AI industrial policy perpetuates a long cycle of racialized disenfranchisement of groups that reap few of the benefits and bear most of the harms.² Instead, we must recognize that benefits to corporate actors often manifest as harms to the people subject to corporate AI systems.

This is why we need to look beyond our current echo chambers. **For a conversation about AI and the public good to meaningfully take place, industrial policy will need to answer to the imperatives of a public beyond the tech and defense industries, and particularly to those structurally disadvantaged groups that have already borne the brunt of the costs of this industry.** While this collection grapples with the limitations of the current approach to AI industrial policy, and the narratives that support it, the work ahead is a much more difficult task: articulating the kind of economy we want in the first place, and where–if at all–AI could play a part. We intend for this to serve as a provocation for taking this conversation forward, beyond the geographic boundaries of the United States and EU, and beyond the silos of tech policy debates.

¹ Amy Kapczynski and Joel Michaels. "Administering a Democratic Industrial Policy," *Harvard Law & Policy Review*, Forthcoming, Yale Law School, Public Law Research Paper, January 30, 2024, <https://papers.ssrn.com/abstract=4711216>; Isabel Estevez, "Multi-Solving, Trade-Offs, and Conditionalities in Industrial Policy," October 2023, Roosevelt Institute, <https://rooseveltinstitute.org/publications/multi-solving-trade-offs-and-conditionalities-in-industrial-policy/>.

² Todd Tucker, "Industrial Policy and Planning: What It Is and How To Do It Better," July 2019, Roosevelt Institute, https://rooseveltinstitute.org/wp-content/uploads/2020/07/RI_Industrial-Policy-and-Planning-201707.pdf; Emily Bender, Timnit Gebru, Angelina McMillan-Major, Shmargaret Shmitchell, "On the Dangers of Stochastic Parrots: Can Large Language Models Be Too Big?" *FAccT '21: Proceedings of the 2021 ACM Conference on Fairness, Accountability and Transparency*, <https://dl.acm.org/doi/10.1145/3442188.3445922>.



Executive Summary

by Amba Kak

What Are You Reading?

In this collection of essays, we bring together for the first time various emergent global perspectives on AI industrial policy. “AI industrial policy” refers to a set of investment, regulatory, and government-spending strategies that aim to shape—and, at present, largely promote—national prowess on artificial intelligence. The essays in this collection, and this broader project, begin to challenge this narrow focus on national competitiveness in favor of one that is grounded in a (democratically contested) understanding of public benefit. Each of the regions analyzed here has created its own narrative of AI development; together, these stories illuminate emergent trends in the way governments are choosing to respond to this uniquely charged moment.

What Counts as Industrial Policy?

The essays in this collection reflect a clear uptick in AI policies being trialed by governments globally, while also drawing attention to the lack of a coherent and clearly defined vision behind these moves. The industrial policy tools in play are wide ranging: from direct investments and tax credits to hybrid public-private AI initiatives to the platformization of government assets towards AI development; alongside regulatory strategies like competition and antimonopoly policy.³

In the US and the EU, the term “industrial policy” is being used to describe major flagship public investment initiatives in clean-energy generation, and in the technology sector, especially vis-à-vis semiconductors.⁴ In other regions, “industrial policy” is a less familiar term in public discourse, and AI-related government investments are typically narrated as part of national AI strategies or innovation policy initiatives aimed at bolstering the growth and competitiveness of the domestic AI market.

In this collection, we survey the expanse of what currently counts as “industrial policy.” This includes traditional levers like direct investments, subsidies, and tax credits of the kind we see in the US CHIPS Act; but also, working backward from the multiple forms of market-shaping we see in practice, reveals varied and often subtle institutional and policy engineering.

For one, there’s an uptick in hybrid public-private arrangements that evade any strict state-versus-market binary through the “merging or fusion of public and private resources.”⁵ This is in part explained by the concentration of AI-related resources and expertise within the private sector that makes it almost inevitable

³ For similarly expansive orientations towards industrial policy see Todd Tucker, “Industrial Policy and Planning: What It Is and How To Do It Better,” July 2019, Roosevelt Institute.

https://rooseveltinstitute.org/wp-content/uploads/2020/07/RI_Industrial-Policy-and-Planning-201707.pdf; Amy Kapczynski and Joel Michaels. “Administering a Democratic Industrial Policy,” *Harvard Law & Policy Review*, Forthcoming, Yale Law School, *Public Law Research Paper*, January 30, 2024, <https://papers.ssrn.com/abstract=4711216>.

⁴ Greg Ip, “Industrial Policy’ Is Back: The West Dusts Off Old Idea to Counter China,” WALL ST. J., July 29, 2021;

<https://americanaffairsjournal.org/2021/08/the-emerging-american-industrial-policy/>; Brian Deese, National Economic Council Director, Remarks on Executing a Modern American Industrial Strategy (Oct. 13, 2022),

<https://www.whitehouse.gov/briefing-room/speeches-remarks/2022/10/13/remarks-on-executing-a-modern-american-industrial-strategy-by-nec-director-brian-deese/>; see also Jake Sullivan, National Security Advisor, Remarks on Renewing American Economic Leadership at the Brookings Institution (Apr. 27, 2023),

<https://www.whitehouse.gov/briefing-room/speeches-remarks/2023/04/27/remarks-by-national-security-advisor-jake-sullivan-on-renewing-american-economic-leadership-at-the-brookings-institution/>. In Europe, see, European Council, “EU Industrial Policy,”

Accessed March 1st, 2024, <https://www.consilium.europa.eu/en/policies/eu-industrial-policy/>; Terzi, Alessio, Monika Sherwood, and Aneil Singh. “European Industrial Policy for the Green and Digital Revolution.” *Science and Public Policy* 50, no. 5 (October 16, 2023): 842–57. <https://doi.org/10.1093/scipol/scad018>.

⁵ Weiss, Linda. *America Inc.? Innovation and Enterprise in the National Security State*. Cornell Studies in Political Economy. Ithaca : London: Cornell University Press, 2014.

that all roads (and explicitly, profits and revenue) will lead back to industry, even in the implementation of public AI infrastructure. The pilot for the US National AI Research Resource, which Amba Kak and Sarah Myers West analyze in detail in [Chapter 2](#), for example, invites AI companies to offer up donated resources on a public-private marketplace hosted by the National Science Foundation, alongside access to existing government supercomputers, datasets, and research resources. Other public-private moves happen in tandem, complementing each other. For example, with the UK government investing close to \$1.5 billion in compute investments, Microsoft pledged close to double that amount toward building out AI cloud infrastructure in the country.⁶

There's also the emergence of practices best understood within what is termed the **"de-risking state,"**⁷ a mainstream government tool of "crowding in" private capital, where the **government spearheads the creation of a new ecosystem that is lucrative for private actors, but where the state bears most of the real risk.** This is by no means a recent phenomenon. As Susannah Glickman describes in [Chapter 1](#), programs like the Small Business Innovation Research Program (SBIR) were designed to ensure that venture capitalists (VCs) bore less risk, creating what she terms "a permanent role for the VC in industrial policy", while simultaneously integrating VCs into decisions regarding which entities received state funding.

We find a striking contemporary example of this trend in India. Jyoti Panday and Mila T Samdub describe in [Chapter 4](#) that the Indian state has created **foundational software data platforms** (for example, a set of platforms known as "IndiaStack") enabling private and public access that have facilitated the emergence of a lucrative domestic market subsidized by government spending. These hybrid infrastructures, promoted globally under the broader umbrella of "digital public infrastructure" or DPI, are all set to provide the foundation for AI-enabled use cases. While these initiatives have already reaped dividends for the private sector, Panday and Samdub also draw attention to their "significant costs when it comes to citizens' rights and state power."

In a similar vein, Matt Davies, in [Chapter 5](#), spotlights the **"platformization" of UK government assets**, including valuable and highly sensitive NHS data, to service

⁶ Department of Science, Innovation & Technology, "Science and Technology Framework: Update on Progress," GOV.UK, February 2024, <https://assets.publishing.service.gov.uk/media/65c9f67714b83c000ea7169c/uk-science-technology-framework-update-on-progress.pdf>; UK Government, "Boost for UK AI as Microsoft Unveils £2.5 Billion Investment," November 30, 2023, <https://www.gov.uk/government/news/boost-for-uk-ai-as-microsoft-unveils-25-billion-investment>.

⁷ Daniela Gabor, "The (European) Derisking State." Preprint, SocArXiv, May 17, 2023, <https://doi.org/10.31235/osf.io/hpbi2>.

the private sector. He warns: “Even the AI Safety Institute—heralded as a ‘startup within government’⁸ and an attempt to do something different by building state capacity on AI—risks essentially becoming the provider of voluntary services to large incumbent companies.”

Another consistent theme is that the uptick in AI industrial policies often walks hand in hand with a non-interventionist or weak regulatory posture that serves the same ends. Or, as in the case of trade deals, the creation of legal regimes that entrench this deregulatory posture. **Antitrust or competition regulation should be viewed, then, as a kind of industrial policy.** As Glickman argues in **Chapter 1**: “A coup of bipartisan American propaganda promoting the myth of the lone American entrepreneurial tech genius has been to veil the equally bipartisan support for tech industrial policy.” After decades of inaction, however, now that the US has taken a renewed aggressive stance on competition under the Biden Administration, efforts to reinvigorate public alternatives in AI find it difficult, if not impossible, to evade entrenching power back into this concentrated market (more on this below). In the UK, the failure to block Google’s acquisition of DeepMind has been questioned by influential tech figures who fault it for giving away what would have been a real national champion in AI. (Ian Hogarth, the British venture capitalist who was also Chair of the UK AI Safety Taskforce, asks, “Is there a case to be made for the UK to reverse this acquisition and buy DeepMind out of Google and reinstate it as some kind of independent entity?”) The AI hype cycle fueled by the release of ChatGPT landed at a time when considerable progress had already been made on advancing regulation for the AI sector, most notably in the EU. In **Chapter 3**, Max von Thun argues that the focus on promoting European champions and competitiveness in AI has “in some instances led policymakers to actively undermine efforts to impose regulatory guardrails, most notably in relation to the EU’s AI Act.”

“AI” Interventions up and down the stack

The term “AI” is a fuzzy umbrella term for a suite of technologies, few of which are actually new, often distracting from the reality that AI is both a product of, and amplifying, concentrated power in the tech industry. For a decade, AI has been

⁸ Notably by Ian Hogarth in AI Safety Institute and Department for Science, Innovation and Technology, *Frontier AI Taskforce: Second Progress Report*, GOV.UK, October 30, 2023, <https://www.gov.uk/government/publications/frontier-ai-taskforce-second-progress-report/frontier-ai-taskforce-second-progress-report>.

routinely invoked in policy documents as one of many technologies of the future, in the same breath as blockchain or quantum computing, in service of broad pro-innovation policies or specific infrastructural upgrades. But in the more recent past, and with a growing drumbeat of attention and investment in AI, it has been established as worthy of its own dedicated industrial strategies. **Today AI provides an attractive banner for industrial policies aimed at inputs necessary for AI development, and a more central feature of broad-based government technology initiatives.** Semiconductors are a perfect example; Glickman explains in [Chapter 1](#) that “the history of AI is inseparable from the history of semiconductors,” and that “the technofuturist promises of AI have functioned to provide cover for the funding of more banal improvements in chips and chip infrastructure.” A closer look at the public relations and policy narrative around the CHIPS Act, too, as Kak and West undertake in [Chapter 2](#), reveals that AI was only one among a list of industries that the Act would benefit—but the ChatGPT-triggered upswing of demand for state-of-the-art chips in spring 2023 gave renewed momentum and attention to homegrown investment in semiconductors for AI. This led to a decision by TSMC to expand its investments in the US, though ultimately the company’s supply chain for chip production will remain globally distributed.

Focusing only on industrial policies specifically earmarked for AI can therefore be an underinclusive lens that overindexes on recent developments that are branded around AI, and can lead to missing the forest for the trees. Instead, many essays in this collection track interventions at different points of layers in the AI stack—or, various “inputs” necessary for AI development. Those include: compute, data, models, and labor.

Given that the infrastructure needed to develop AI is monopolized up and down the stack, most notoriously within cloud computing, data centers, and the chips needed to process AI, the shoring up of compute resources is a key focal point for many national AI industrial policies. The UAE is a particularly interesting case study in this regard. Islam Al Khatib notes that the capital-intensive nature of compute resources for AI has made the region an unavoidable partner for those who, like Sam Altman, need to raise staggering amounts of capital to set up alternatives to Nvidia’s chokehold on the chip market.

Compute initiatives foreground that there are no straightforward paths to “democratizing” what is already a concentrated and vertically integrated industry. We wonder, then, whether this is even an appropriate goal; whether through

arrangements made with cloud providers or procurement of GPUs, public investments in AI will accrue to one or another concentrated sector. Europe in particular has seen a great deal of activity around public investment in chips and government supercomputers, and the most recent 2024 European Commission Innovation package proposes a new initiative to develop state-of-the-art chips for AI development. The French government has allocated over a billion euros toward funding public supercomputers even as it seeks €7 billion worth of private institutional investment into AI. As Samdub and Panday detail in [Chapter 4](#), the Indian and Japanese governments, too, have launched AI-specific cloud computing infrastructure built in a centralized facility, rather than in a way that relied on using commercial cloud solutions. This is an attempt to avoid depending on providers like Amazon’s AWS or Microsoft’s Azure. They note that despite the relatively tiny capacity of India’s public compute initiative (“AIRAWAT”) compared to the capacity enjoyed by large tech companies, offered at a discount to Indian startups, it still offers an attractive option in a market where demand far outstrips supply.

After compute, the next area of heat across jurisdictions appears to be data.

Across various regions we see similar efforts to increase access to properly cleaned, labeled, and structured (or “AI-ready”) data for AI development. The primary focus of these efforts has been on making government datasets available for direct access by companies and researchers, as well as standardization and benchmarking to improve the usability of these datasets with comparatively less focus on ensuring the privacy and security of this data. Even as data is readily acknowledged as a key input (and therefore a bottleneck) in AI development, the US government rarely calls attention to the fact that a large amount of such high-quality datasets are controlled by private industry, and specifically by big tech companies. By contrast, in Europe and India, as part of a broader movement to call attention to US Big Tech data monopolies, there have been one-off proposals for mandating data-sharing and private-sector contributions to data commons.

Market Concentration and National Champions

Instead of enforcing regulations on industries prone to natural monopolies—those with high start-up costs and other barriers to entry—governments have typically tried to wield them as extensions of state power. In the US, this has certainly been the case with the semiconductor industry.

As Glickman explores in **Chapter 2**, the dependencies on scale in this sector have led some in government to believe they had to rely on industry consolidation to continue making advances in the field. With AI, the “bigger is better” paradigm for large general-purpose models means that market concentration at every layer of the stack is only intensifying. Industrial policy formulations, then, all lead to partnerships with large tech companies. In addition, a fast-evolving and increasingly institutionalized discourse around AI not only promotes these technologies as necessary to economic and national security interests, but also positions Big Tech firms as themselves security assets that need to be bolstered rather than held back by regulation. Under this narrative, all that is required is to ensure that these commercial actors are sufficiently responsive to the strategic needs of the state.

In the US, the home of most frontrunner AI companies, the current government’s orientation however marks a historically significant rupture with Big Tech. The Biden administration has deliberately created distance between Big Tech interests and US state interests across policy domains from trade⁹ to technology.¹⁰ They have boldly confronted the concentration of power in the tech sector with a muscular orientation of US enforcement agencies against Big Tech, an integration of competition concerns in procurement guidance, and a clear assertion that “the answer to the rising power of foreign monopolies and cartels is not the tolerance of domestic monopolization.”¹¹ However, Kak and West argue in their chapter that despite these efforts, industrial policy initiatives under the banner of democratizing AI still fail to challenge the deep structural dependencies on private technology companies at every layer of the AI stack, and especially compute. Kak and West question whether “democratization” of AI alone is an appropriate litmus test for public AI. They write, “simply diversifying the range of actors involved in AI development while commercial entities continue to define the horizon for research does little to contest their dominance.”

Outside the US, however, there are different approaches and degrees of comfort with a “national champions”-focused AI industrial strategy. In Europe, we see official narratives that take aim at the concentration of power in foreign tech industries. This is true for the Commission’s flagship efforts but is also heightened

⁹ Farah Stockman, “Opinion | How the Biden Administration Took the Pen Away From Meta, Google and Amazon.” The New York Times, November 27, 2023, <https://www.nytimes.com/2023/11/27/opinion/google-facebook-ai-trade.html>.

¹⁰ Makena Kelly, “Biden Rallies against Big Tech in State of the Union Address.” The Verge, February 8, 2023, <https://www.theverge.com/2023/2/7/23590396/state-of-the-union-sou-biden-tech-tiktok-privacy-antitrust>.

¹¹ White House, “Executive Order on Promoting Competition in the American Economy,” July 9, 2021, <https://www.whitehouse.gov/briefing-room/presidential-actions/2021/07/09/executive-order-on-promoting-competition-in-the-american-economy/>.

in member states like France, Germany, and Italy, as von Thun explores in [Chapter 3](#). In these countries, the desire for national AI champions has motivated a range of public investments and regulatory postures in the past few years. **But the intention for autonomy alone cannot wish away the reality. It will continue to be a formidable challenge to develop industrial strategies that meaningfully steer clear of dependency on US Big Tech.** In February 2024, Microsoft acquired a minority stake in France's main AI champion, Mistral, with an agenda explicitly focused on building applications for government use in Europe. Davies points to a similar tension in the UK, arguing that the current blitz of activity to build the UK's public supercomputer "offers a fantasy of independence that masks deeper structural dependence on a paradigm of AI development led by, and wholly dependent on, funding and infrastructures provided by Silicon Valley."

Pitting "Innovation" Against Regulation

On the surface, pro-regulatory stances are more mainstream than ever, with high-level support from powerful industry actors. Today innumerable and largely voluntary governance initiatives exist around mitigating AI risks and creating responsible AI use. **However, we are simultaneously seeing a well-funded "innovation versus regulation" narrative gather steam. This depicts industrial strategy efforts that promote the AI industry as pro-innovation, while regulatory efforts are characterized as hampering national competitiveness.** In this context, "regulatory sandboxes" have become increasingly popular. These are flexible regulatory arrangements aimed at encouraging time-bound experimentation unhampered by onerous regulation.

To be clear, AI regulation, much like any other domain of technology and data policy, has rarely been immune to economic logics. In [Chapter 3](#), von Thun tracks the consistent effort to foreground the economic and strategic benefits of the EU's flagship AI Act, including positioning the EU as setting the rules of the global market and underscoring exceptions for small businesses. But in the post-ChatGPT AI race, European industry players and big-tech lobby groups have worked to whip up fears about Europe's lack of competitiveness in AI, and push for a weaker regulatory regime. Reportedly fueled by domestic companies like Mistral and Aleph Alpha, France, Germany, and Italy argued that imposing strict regulatory

requirements on foundation models would hamper the continent's ability to compete—arguments soon belied by Mistral's partnership with Microsoft.¹² In the UK too, Davies tracks the “contradictory impulses” of deregulatory strategies carried out through legislative proposals that deliberately avoid interventionist approaches and a broader unwillingness to endow regulators with new statutory powers to address current harms with AI.

In other regions, like India and South Africa, we explore how AI regulation has been reduced to an empty signifier, often bandied about as a national priority despite lack of meaningful legal progress to meet mounting challenges with AI systems concerning privacy, security, competition, and discrimination. The UAE's Minister of State for Artificial Intelligence, Omar Al Olama, meanwhile, positions the UAE as “a testing ground for AI advancements and the construction of experimental regulatory frameworks.”

But there are also other regulatory currents, outside those explicitly focused on AI, that are already shaping the market. This was also a concern in the heated public discourse of the past year; where discussion has been narrowly preoccupied with future risk scenarios and novel policies, rather than how to leverage existing regulatory frameworks. In the EU, for example, the Digital Markets Act and the Data Act both have the explicit aim of boosting Europe's economic competitiveness and establishing technological sovereignty and von Thun argues that there is tremendous scope to harness the DMA proactively so that it “could be used to promote a fairer and more diverse AI European ecosystem.” In South Africa, too, despite almost no progress on AI-specific regulation, the Competition Commission concluded its market inquiry on big-tech platforms with bold recommendations and enforceable remedies around competition concerns like self-preferencing and algorithmic pricing.

Beyond the US-China AI Arms Race

The “US-China AI arms race” has lurked in the background of policy discussions on AI for close to a decade now. Initially a sporadic talking point for tech executives, it has evolved to become an increasingly institutionalized position, represented by collaborative initiatives between government, military, and tech-industry actors

¹² Madhumita Murgia, “Microsoft Strikes Deal with Mistral in Push beyond OpenAI,” *Financial Times*, February 26, 2024, <https://www.ft.com/content/cd6eb51a-3276-450f-87fd-97e8410db9eb>.

and reinforced by legislation and regulatory debates.¹³ In the US, this race has been invoked to both kindle an appetite across party lines for increased public investment for AI, and also push back against calls for slower, more intentional AI development and stronger regulatory protections.

The bipolarity of the US-China framing has always been myopic, leaving out the complex ways in which other regions participate in the dynamic. This is even truer today as an increasing number of governments make strong nationalistic plays around their role in the AI future. Yet this narrative persists. For one, the idea that we are in a geopolitically sensitive AI race has only gained traction in the scramble to market post-ChatGPT. As AI is heralded today as the digital infrastructure of the future, security implications mean that striving for hegemony in this domain is non-optional. As Khatib notes in her essay, the UAE sees becoming “the best” in AI as crucial to securing the country’s stability in any “post-oil” future: “the association of AI with fantasies of ‘absolute sovereignty,’ ‘progress,’ and the persisting belief that ‘future wars’ will be centered around data and information (language of information war and cyberwars) rather than land and resources.”

The US-China Arms Race, has largely been tracked using inconstant quantitative metrics like most-cited papers or patents. So too are AI-related national readiness rankings and metrics proliferating, guiding the allocation of public resources. As Sandra Makumbirofa argues, South Africa’s AI strategy has been guided by the need to be perceived as leading AI on the continent, despite a complete lack of defined metrics or broader public discourse on why, and how, the nation intends to achieve these objectives.

We observe opportunistic behavior amidst the escalating tensions between the US and China, with countries like the UAE attempting to play “both sides.” More recently, India has positioned itself as a stable “democratic” semiconductor hub that could function as an alternative to Taiwan.¹⁴

Strands of AI nationalism do exist outside of the US-China geopolitical framing: AI developmentalism, for example. Panday and Samdub note that India’s national AI strategy headlines the goal of becoming an “AI garage for 40% of the world” alongside calls for using AI for economic development, “particularly in Global South

¹³ For an updated timeline of the proliferation of the US-China AI arms race, see <https://ainowinstitute.org/publication/tracking-the-us-and-china-ai-arms-race>

¹⁴ Sankalp Phartiyal, “India Chip Strategy Makes Progress With \$21 Billion in Proposals,” Bloomberg, February 26, 2024, <https://www.bloomberg.com/news/articles/2024-02-26/india-chip-strategy-makes-progress-with-21-billion-in-proposals>.

countries, which may provide markets for solutions developed in India.” A closely related theme is localization: the idea that countries with large addressable markets will find their niche catering to the specific needs of their population. Promoting linguistic diversity in AI is a prominent example of this; as Khatib notes, Jais, UAE’s LLM, was not a PR exercise oriented solely at the Global North—“it also offers 400 million Arabic speakers access to generative AI technologies.”

What Does “AI for the Public Good” Really Mean?

The AI industrial policies we’ve surveyed here have largely functioned to reinforce the notion that AI is a socially as well as geopolitically important sector, and therefore worthy of government strategies to promote it. This needn’t necessarily be the case, as Amy Kapczynski and Joel Michaels argue in a forthcoming paper.¹⁵ Industrial policy, they state, must be held to strict standards of ensuring AI meets public aims. This might not necessitate growth; it may include demoting certain sectors that cause public harm—cryptocurrency, for example.

Does the premise that AI advances public aims hold up? We recently argued, “with an overwhelming focus on AI-driven harms, we’ve missed a key piece of the puzzle: demanding that firms articulate, clearly and with evidence to back it, what the benefits of AI are to the public.” We still see no persuasive or cohesive articulation of a vision for social good that justifies public endorsement—and taxpayer dollars—outside of shallow assertions that AI will lead to productivity gains and leapfrogging advancements in fields like medicine and climatology. Partly, this stems from the fact that “AI” is often used as an empty signifier for technological innovation. Since the ChatGPT-inspired AI hype wave, it is easier than ever for governments and industry advocates to coast by on high-level assertions about AI without pointing to real-world use cases. As Davies puts it, the UK “central government has rarely, if ever, advanced a coherent vision for the role that a domestic AI sector should play within the UK economy.” We find a similar conundrum across US AI policy documents, where outside of generic policy narratives about pushing the frontiers of science and the public good, there’s no meaningful scrutiny of either the current or speculative advancements. Nor is there

¹⁵ Amy Kapczynski and Joel Michaels. “Administering a Democratic Industrial Policy,” *Harvard Law & Policy Review*, Forthcoming, Yale Law School, Public Law Research Paper, January 30, 2024, <https://papers.ssrn.com/abstract=4711216>.

acknowledgement of the enormous and well-evidenced energy and water costs to this industry. In South Africa, Makumbirofa explores this disjunction: the government is apparently convinced that AI prowess is crucial for economic development but has no answer for how this will help high unemployment, racial inequality, and unreliable electricity supply.

Makumbirofa explores, too, how the current Ramaphosa government routinely advocates for AI in government procurement amid calls to cut public-sector employment costs. **In fact, data-driven tools like AI have routinely and historically been used across the world to justify austerity measures that disenfranchise the public.**¹⁶ Glickman takes us back to the Clinton years, where technology was likened to a lifeboat that would save the public from living in inadequate material conditions. Instead, it “provided cover for the Clinton administration to significantly cut welfare.”

With such far-reaching consequences for the public at stake – from the allocation of public funds (that would otherwise go elsewhere) to the rapid promotion of AI tools in sensitive social domains through procurement mandates – any claims to advancing public good must be put under the scanner. This is especially challenging given that the current paradigm for large-scale AI is a product of concentrated power in the AI industry, making it difficult to develop a clear-eyed account of the public benefits of AI outside of the terms set by the market. The adoption of large scale general purpose AI as a priority in US federal R&D strategies, for example, never acknowledges the market, financial, and environmental impacts that the compute and data dependencies that this trajectory entails. Moreover, with ever-larger-scale general-purpose AI models like LLMs often positioned by industry stakeholders as stepping stones to forms of so-called “artificial general intelligence” (AGI), there is a wholesale acceptance that scale is a proxy for progress and performance. The promise of AGI is also inextricably linked to national security dominance—whichever builds AGI first will win the AI race—making the commercial and national security goalposts all but meld into one another.

The inability of current AI industrial policy to meaningfully make a public interest case is as much a failure of process as it is of substance. As Kapzynski and Michaels argue, industrial policy mechanisms must be designed to be responsive and accountable to the public beyond powerful elite interest groups, and the need for

¹⁶ See for example, Philip Alston, “Report of the Special Rapporteur on Extreme Poverty and Human Rights,” October 11, 2019, https://srpoverty.org/files.wordpress.com/2019/10/a_74_48037_advanceuneditedversion-1.pdf; and “AI Now Report 2018,” December 2018, https://ainowinstitute.org/AI_Now_2018_Report.pdf.

“organized capacity of structurally disadvantaged groups” to influence policy. Without a procedural focus on democratic deliberation, the ‘public interest’ aims of industrial policy will inevitably fall back into a framing that looks at AI primarily as part of the arsenal for the US-China race. For example, one critique of current industrial policy argues for a pivot to “AI manufacturing”.¹⁷ It takes aim at Silicon Valley’s “consumer-focused” approach to innovation for its failure to create middle-class job growth and align appropriately with defense interests.¹⁸ But this manufacturing-focused vision, likely to be propelled forward under the political banner of rivaling China, will inevitably function to galvanize big tech to ‘innovate’ on manufacturing, potentially further entrenching concentrated power (Amazon, for example would have a clear edge given their investment in supply chain optimization¹⁹ and generative AI use cases for manufacturing²⁰). It’s perhaps obvious that industrial policy visions will reflect the constituencies they are designed to appeal to (the manufacturing-focused AI project, for example, is clearly designed for political appeal to the white, male, working class). But AI’s potentially most harmful impacts – from discrimination to informational harms to workplace surveillance – disproportionately burden people of color, making it imperative for structurally disadvantaged groups to influence the shaping of this policy vision.

Before investing deeper in the development of an AI industry, in any form, we need to understand who AI serves, consider the opportunity costs and question the assertion that this technology will inevitably lead to social progress. We need to be asking grounded and critical questions like: do efficiencies gained through AI-based climate modeling justify the energy cost of training these models? Should we invest in the advancement of edtech at the cost of providing more students school lunch? These answers need to be concrete and material, and won’t be found in echo chambers. For a conversation about AI and the public good to happen meaningfully, industrial policy will need to be responsive and accountable to perspectives outside of the tech industry.

¹⁷ David Adler, William B. Bonvillian, “America’s Advanced Manufacturing Problem—and How to Fix It,” *American Affairs Journal*, Volume VII, Number 3 (Fall 2023): 3–30. <https://americanaffairsjournal.org/2023/08/americas-advanced-manufacturing-problem-and-how-to-fix-it/>.

¹⁸ David Adler, “The American Way of Innovation and Its Deficiencies,” *American Affairs Journal*, Volume II, Number 2 (Summer 2018): 80–95. <https://americanaffairsjournal.org/2018/05/the-american-way-of-innovation-and-its-deficiencies/>.

¹⁹ Mala Jenkins, “Amazon Taps AI to Expand Sellers’ Ecosystem and Supply Chain Capabilities,” *Retail Info Systems News*, September 13, 2023. <https://risnews.com/amazon-taps-ai-expand-sellers-ecosystem-and-supply-chain-capabilities>.

²⁰ Scot Wlodarczak, “How Generative AI Will Transform Manufacturing,” June 20, 2023, AWS for Industries (blog). <https://aws.amazon.com/blogs/industries/generative-ai-in-manufacturing/>.