

FEBRUARY 2025

Technology Federalism: U.S. States at the Vanguard of AI Governance

Scott Kohler

Technology Federalism: U.S. States at the Vanguard of AI Governance

Scott Kohler

© 2025 Carnegie Endowment for International Peace. All rights reserved.

Carnegie does not take institutional positions on public policy issues; the views represented herein are those of the author(s) and do not necessarily reflect the views of Carnegie, its staff, or its trustees.

No part of this publication may be reproduced or transmitted in any form or by any means without permission in writing from the Carnegie Endowment for International Peace. Please direct inquiries to:

Carnegie Endowment for International Peace
Publications Department
1779 Massachusetts Avenue NW
Washington, DC 20036
P: + 1 202 483 7600
F: + 1 202 483 1840
CarnegieEndowment.org

This publication can be downloaded at no cost at CarnegieEndowment.org.

Contents

Introduction	1
AI in a Changing Policy Environment	2
The Surge in State-Level Policymaking on AI	4
Conclusion	13
About the Author	15
Notes	17
Carnegie Endowment for International Peace	27

Introduction

Well before artificial intelligence (AI) exploded into the public consciousness with the 2022 launch of ChatGPT, advances in a suite of machine learning technologies were rapidly expanding the use of AI systems. As this occurred, societal perceptions of the technology sector were also entering a period of flux. In short, AI came of age alongside a dramatic increase in global attention to the relationship between consumer technologies and society.¹

Consequently, generative AI models powered by deep neural networks (complex computational systems originally influenced by the architecture of neurons in the human brain) arrived on the global policy agenda at a time when regulators were primed for action. In the United States, their transformative implications have prompted a swift and energetic response from the federal government—albeit one facing considerable uncertainty as the 119th Congress and a new presidential administration take office.

Less recognized, however, has been a parallel response—equally active but less coordinated—at the subnational level. U.S. states, fulfilling their canonical role as laboratories of democracy, have been acting as regulators of first resort on a number of emergent technologies and core tech policy areas where Congress has been slow to legislate.² In some instances, individual states have acted as first movers.³ In others, networks of state officials and organizations working across state lines have collaborated to foster collective action. Nowhere is this emerging *technology federalism* more widely apparent and potentially consequential than in the state-level response to AI.

AI in a Changing Policy Environment

In the United States, polling indicates that positive public sentiment toward the internet industry peaked in 2015, and then began declining as negative sentiment rose.⁴ In 2018, the term “techlash” entered the global lexicon.⁵ By 2020, strong majorities of Americans were “very concerned” about both the spread of misinformation and the privacy of personal data online.⁶ These concerns combined with a range of discontents, including worries over competition, extremism, and the political process.⁷ By the decade’s end, even amid robust growth within the tech sector, these anxieties were already prompting a global policy response.

Regulatory pushback had crystalized earlier and forcefully in Europe. A suite of European Union (EU) laws, including the General Data Protection Regulation (GDPR) and the Digital Services Act and Digital Markets Act came into effect in 2018 and 2022, respectively, while high-profile competition proceedings and litigation over individual digital rights created an increasingly stringent regulatory environment for technology firms. Over the ensuing years, elements of the European approach spread to other markets, including through digital privacy laws incorporating elements of the GDPR regime in Asia and the Americas. On its face, this diffusion bore hallmarks of the “Brussels Effect” posited by Anu Bradford, wherein the regulatory power of the European Union establishes global standards through a combination of overt emulation and de facto acceptance beyond EU borders.⁸

The United States, for its part, pursued a comparatively market-oriented, light-touch approach to technology policy, reflecting prior commitments to free markets and free speech.⁹ Nevertheless, over the past five years, there has been significant tightening across political boundaries. From Senator Elizabeth Warren’s call to break up major tech companies during her 2020 presidential campaign to President Donald Trump’s calls to roll back Section 230 of the Communications Decency Act, by 2022, a relatively bipartisan and supportive policy environment had grown distinctly more complicated.¹⁰

In short, two trends were shaping the global tech policy environment even before AI rose to the fore. The first was a set of distinct and competing governance paradigms, with Europe comparatively stringent and the United States more *laissez faire*.¹¹ The second was growing discontent over perceived overreach by technology platforms and harms wrought by their products and services.¹² In the early and mid-2010s, as neural networks trained using graphics processing units were beginning to demonstrate their promise, tech companies, especially in the United States, were still widely feted for their innovation and success.¹³ By the end of the decade, as the computer programs AlphaGo and AlphaZero were outpacing the world’s best human competitors and OpenAI was being formed, GDPR was reshaping the digital privacy landscape and Cambridge Analytica scandal was coming to light.¹⁴ By the early 2020s, digital governance was a well-established front-burner issue for regulators around the world.¹⁵

These developments set the scene for a relatively quick policy response just as breakthroughs in AI development coupled with massive private investment to yield highly capable generative models now used by hundreds of millions of people.¹⁶

The Shifting U.S. Landscape

In the United States, attention has focused on flagship federal policies such as president Joe Biden's 2023 Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence; president Trump's 2025 order revoking it and directing the development of a new AI Action Plan; the Senate's 2024 bipartisan road map; and the launch of the U.S. AI Safety Institute.¹⁷ These are important developments, though their overall impact on the trajectory of AI development and deployment remains to be seen. Trump's return as U.S. president adds significantly to this uncertainty and may dramatically reorient federal AI policy.

Even during the Biden administration, however, amid persistent—though not absolute—congressional gridlock and lingering tensions over digital governance between the United States and key allies, there was a durable perception in some quarters of the American approach as permissive, sclerotic, and insufficiently rights-protective.¹⁸

That perception, however, has masked a dramatic shift in U.S. policymaking at the subnational level—a shift that has put U.S. states at the vanguard of American AI law. In state capitals, as much as in Washington, DC, the wheels of policymaking have been spinning full bore. States have introduced a vast array of bills, passed laws, launched programs, and undertaken reorganization efforts.¹⁹ Action has not been confined to executives or legislatures alone, nor to red or blue states. And while this trend poses a range of questions, it has advanced with such sufficient speed and energy that, today, meaningfully appraising the extent and direction of American technology policy requires attention to a dynamic and burgeoning state landscape.

The Privacy Antecedents of Technology Federalism

Since 2018, when California enacted the California Consumer Privacy Act, no fewer than nineteen states have adopted comprehensive consumer privacy laws, of which a majority regulate the use of automated decisionmaking or profiling systems.²⁰ These privacy laws have established an initial patchwork of regulation applicable to AI as they set new digital privacy rules over the past several years.

This crop of state rules regulating automated decisionmaking focuses on applications producing legally consequential effects (for example, in providing or withholding public benefits, financial services, housing, education, employment, and healthcare or in law enforcement).²¹ For these applications, states frequently mandate disclosure, impact assessments, and opt-out rights for data subjects. These requirements echo GDPR Article 22, which establishes a qualified “right not to be subject to a decision based solely on automated processing, including profiling, which produces legal [or similarly significant] effects.”²²

While not cast as AI laws per se, these privacy obligations extend to AI systems, which involve just this sort of automation.²³ As a result, state privacy laws have set a preliminary baseline of subnational AI policy. They have also established a predicate for further and more overt policies impacting AI.

The Surge in State-Level Policymaking on AI

States have progressed quickly from regulating AI incidentally to doing so directly with a mix of supportive and supervisory policies. Like policymakers around the globe and in Washington, states are pursuing a dual imperative: to foster AI for the economic, social, and scientific benefits it promises, while safeguarding society against its potential harms. For example, Biden’s 2023 executive order speaks of AI’s “promise and peril” and seeks to promote “responsible innovation.”²⁴ Similarly, California Governor Gavin Newsom’s executive order on generative AI (GenAI) observes that “GenAI can enhance human potential and creativity but must be deployed and regulated carefully to mitigate and guard against a new generation of risks.”²⁵ In Pennsylvania Governor Josh Shapiro’s executive order, he notes the technology’s potential to help state agencies serve Pennsylvanians while urging that its “responsible and ethical use . . . should be conducted within a governance structure that ensures transparency, tests for bias, addresses privacy concerns, and safeguards [the state’s] values.”²⁶

Exploratory Efforts

Many states have established new task forces or designated existing agencies to study the ramifications of AI systems. These exploratory bodies are directed to make recommendations on fostering AI’s benefits or growth in the state, promoting its responsible development and use, and leveraging AI for public service delivery. In Maryland, for example, Governor Wes Moore created an AI Subcabinet tasked with developing and implementing an “AI Action Plan” and crafting policies and internal resources to embed values such as equity, innovation, reliability, and safety into the state’s AI workflows.²⁷

Other states, including Massachusetts and Rhode Island, have chartered public-private AI task forces to perform similar duties, such as assessing the risks and opportunities AI presents and making recommendations to inform further policy development and state operations. Massachusetts created a task force that comprises representatives from state and local government, industry, organized labor, and academia.²⁸ Rhode Island paired its exploratory charter with the establishment of a unified data governance structure to promote intergovernmental collaboration, set standards for data sharing, and help promote the state's AI readiness.²⁹

States have also taken exploratory steps through legislation. In March 2024, Utah enacted one of the nation's first AI-focused consumer protection laws, SB 149, the Artificial Intelligence Policy Act. The law combines a comparable exploratory mandate with new state infrastructure for public-private collaboration.³⁰ It establishes an Office of Artificial Intelligence Policy to work with industry and civil society on regulatory proposals “to foster innovation and safeguard public safety.”³¹ Interestingly, it also creates a “Learning Laboratory” program that encourages experimentation by offering temporary regulatory mitigation agreements to AI developers and deployers to enable them to test applications within the state.

Transparency and Misinformation

The rapid advancement and adoption of generative AI models has raised hope for a new flowering of human knowledge and discovery. Simultaneously, it has prompted warnings for a coming age of misinformation, in which truth is elusive, trust is scarce, and the informational underpinnings of democratic self-government—already straining—are further eroded.³²

These worries take many shapes. They include warnings about deepfakes and synthetic content and about “hallucinations” causing models to present fabricated information as fact.³³ They encompass fears that generative models will supercharge the ability of malicious actors to wage influence operations.³⁴ And they include worries that models will free ride copyrighted data, eroding traditional newsgathering and pressuring the viability of media organizations already reeling from prior waves of technological change while failing to apply comparable editorial standards.³⁵

In the United States, these concerns exist against a backdrop of robust legal protection for free expression, a tradition in which, as Justice Brandeis famously wrote, the remedy for harmful speech is more speech, and efforts to police misinformation face broad suspicion.³⁶ In short, policymakers face a landscape marked by worries over misinformation but also significant reticence to restrain it.

Some early-moving U.S. states, however, have experimented with policies to mitigate AI-related misinformation, particularly by fostering transparency. Utah’s Artificial Intelligence Policy Act requires generative AI systems such as chatbots to “clearly and conspicuously disclose,” if asked or prompted by a person interacting with the system, that they are AI tools and not human.³⁷ When an AI system is used to provide “the services of a regulated occupation,” such as healthcare or investment advising, its deployer must ensure that this disclosure is made proactively at the start of the exchange, conversation, or text chat.³⁸

Other states are following suit. Bills instituting disclosure requirements, watermarking obligations, or seeking to curb misleading uses of AI have been introduced in at least ten states.³⁹ Similar to Utah’s act, the bills include measures with broad applicability, such as, California’s new AI Transparency Act, signed by Newsom in September 2024.⁴⁰ As of 2026, that law will require “covered providers” of generative AI systems used by more than 1 million aggregate monthly users to incorporate digital watermarking—that is, a machine-readable “latent disclosure”—in images, video, and audio content created by their systems.⁴¹ It will also require providers to give users the ability to include visible or audible disclosures in content they create or modify, and to offer freely available detection tools enabling users “to assess whether image, video, or audio content . . . was created or altered by the covered provider’s GenAI system.”⁴²

Some states are focusing on discrete applications, such as political campaigns. For example, both California and Florida enacted legislation in 2024 requiring disclosures in qualified political advertisements created or modified using AI.⁴³ Florida’s law requires a prominent disclosure in political advertisements and other electioneering communications that use AI to depict a real person performing an act that did not occur with the intent to injure a political candidate or mislead the public about a ballot issue.⁴⁴ Comparable bills have been introduced elsewhere.⁴⁵

Other states have weighed transparency requirements for the use of AI in the workplace. Legislation introduced in Washington state in 2024, for example, would require employers to provide written disclosure before or within thirty days of beginning to use AI “to evaluate or otherwise make employment decisions regarding current employees.”⁴⁶ The same legislation would also prohibit employers from using AI to replicate an employee’s likeness or voice without the employee’s explicit consent.⁴⁷

AI Safety

Policymakers and technologists around the world are racing to assess the safety implications of AI models. National governments and international organizations have convened successive AI safety summits and established a nascent network of safety institutes, while legislative and executive measures have all sought to mitigate safety and security risks.⁴⁸ At the same time, the AI safety landscape is highly contested, marked by starkly divergent

views on the nature, likelihood, and severity of potential risks, as well as the cost-benefit ratio of precautionary guardrails.⁴⁹ The reelection of Trump, who has already acted to repeal the Biden executive order and returns to the White House promising to chart a new course, adds significant uncertainty to the outlook for U.S. federal policy and existing diplomatic initiatives.⁵⁰

Nevertheless, AI safety will remain on the diplomatic agenda. In the aftermath of the U.S. election, for example, members of the newly formed International Network of AI Safety Institutes gathered in late November 2024 in San Francisco to prepare for the upcoming Paris AI Action Summit this February. Furthermore, and importantly, AI safety policies are already embedded in model developers' processes. Voluntary commitments agreed to by developers, such as those brokered by the White House in 2023 and the United Kingdom and South Korea in 2024, remain at least partly in place.⁵¹ And AI companies' internal policies—such as OpenAI's Preparedness Framework, Anthropic's Responsible Scaling Policy, and Google DeepMind's Frontier Safety Framework—demonstrate an awareness that AI safety is an important element of model development and governance. Finally, state-level liability rules embedded in tort law ensure that technology companies must weigh their financial exposure in the event that their models cause reasonably foreseeable injury or financial loss.⁵²

Given the attention AI safety has received on the national and diplomatic stages, it is perhaps surprising how much a single U.S. state, California, featured in the global AI safety debate during the past year. That becomes less surprising, however, in view of California's outsized profile in the global technology ecosystem, its role at the center of AI development, and its history of setting nationally and globally relevant standards in policy domains ranging from the environment to consumer protection: referred to as the "California Effect."⁵³

California State Senator Scott Wiener's SB 1047 bill, the Safe and Secure Innovation for Frontier Artificial Intelligence Models Act, touched off a pitched debate in 2024, exposing divisions among AI industry leaders, prominent AI researchers, and policymakers even in a jurisdiction marked by single-party control.⁵⁴ SB 1047 would have imposed a duty of care on developers of the most sophisticated forthcoming AI models, requiring them to avoid causing or materially enabling catastrophic harms and obligations to develop written safety protocols and build the capability to shut down models within their control. Despite concerns over the bill, SB 1047 survived furious opposition and multiple rounds of revision to pass the California legislature.

While Newsom ultimately vetoed SB 1047, he did so while expressing support for its underlying objectives and its preventive approach to catastrophic risk. He pledged to work with the legislature on a new AI safety proposal in the 2025 legislative session. Importantly, from a technology federalism perspective, he specifically defended the role of states at the forefront of AI safety governance:

To those who say there's no problem here to solve, or that California does not have a role in regulating potential national security implications of this technology, I disagree. A California-only approach may well be warranted—especially absent federal action by Congress.⁵⁵

In short, while safety remains at the forefront of national, international, and private sector efforts, it appears likely that California will continue making its own mark on a fast-moving, contested, and highly consequential global debate. Following Trump's return to the White House, pressure on California to act as a counterweight to his policy agenda may grow. Likewise, other states may feel increased urgency, or see an opportunity, to establish AI safety standards individually or in concert if the federal government strikes a more accommodative stance.

Fairness and Discrimination

While safety questions have exposed fissures within the AI community, there has been agreement in many quarters that AI models present other risks that also require attention, such as the potential for discrimination. Researchers have warned that historical bias and disparate treatment reflected in the data used to train large language models can subtly and powerfully manifest in their predictive outputs, encoding and extending past injustices into a growing range of AI applications.⁵⁶ AI could also render bias less scrutable through its incorporation into complex models whose human creators cannot anticipate outputs or disaggregate the inputs producing a given result.

Federal policymakers over the past several years have recognized these risks. The Biden administration's Blueprint for an AI Bill of Rights warns that “algorithms used in hiring and credit decisions have been found to reflect and reproduce existing unwanted inequities or embed new harmful bias and discrimination.”⁵⁷ It goes on to state that “designers, developers, and deployers of automated systems should take proactive and continuous measures to protect individuals and communities from algorithmic discrimination,” including through “proactive equity assessments” and “use of representative data and protection against proxies for demographic features.”⁵⁸ It further endorses impact assessments and outside evaluations as measures to pressure test and confirm system fairness.

The Civil Division of the U.S. Justice Department has also taken steps to coordinate inter-agency efforts on AI nondiscrimination, and a number of federal agencies have issued a Joint Statement on Enforcement of Civil Rights, Fair Competition, Consumer Protection, and Equal Opportunity Laws in Automated Systems.⁵⁹ Meanwhile, the Artificial Intelligence Risk Management Framework established by the National Institute of Standards and Technology cautions that “AI systems can potentially increase the speed and scale of biases and perpetuate and amplify harms to individuals, groups, communities, organizations,

and society.”⁶⁰ Prominent legislators have also voiced their concern, with some arguing that algorithmic discrimination presents a clearer and more present issue than catastrophic safety risks.⁶¹

The Trump administration appears poised to chart a substantially different course on issues of equity and nondiscrimination.⁶² However, despite this gathering federal change, state-level policymaking has arguably been more definitive in seeking to mitigate the risk of AI-enabled discrimination. Colorado has been the first mover, enacting legislation (SB 24-205) in May 2024 to impose a duty of care on developers and deployers of “high risk artificial intelligence systems”—in other words, those with a material effect on consumers’ access to education, employment, lending, essential government services, housing, insurance, or legal services.⁶³

The Colorado law, which is due to take effect in 2026, also imposes disclosure and governance obligations. It requires developers to provide deployers with information about (1) the uses, benefits, and risks of their systems; (2) their training data; and (3) their implementation of evaluation, governance, and mitigation efforts to prevent algorithmic discrimination. Deployers, in turn, must maintain a Risk Management Policy and Program, specifying how they will identify and mitigate discrimination risk. Before an AI system can be used to make a “consequential decision,” the deployer must disclose the use of AI to impacted consumers in a timely manner and in plain language. In the event of an adverse decision, such as a denial of benefits, the law mandates that consumers be afforded an additional explanation as well as the right to correct inaccurate information that contributed to the decision and the right to appeal to a human decisionmaker.

While the Colorado law stands as an example of determined state action, it also underscores a core takeaway from California’s experience with SB 1047. There remains significant trepidation, even within traditional policy coalitions, on how best to strike a balance between fostering and regulating AI. Colorado’s SB 24-205 passed despite significant concerns from Governor Jared Polis, who cautioned that the bill could overburden industry, stifle innovation, and contribute to a thicket of inconsistent state rules and that it lacked essential guardrails such as an intent requirement.⁶⁴

The Colorado legislation was not developed in a vacuum, however. Rather, it emerged from a broad multistate effort to coordinate baseline obligations for AI systems. That collaboration was spearheaded by Connecticut State Senator James Maroney, who convened a working group of more than 100 state legislators from more than two dozen states as well as a second Connecticut-specific group to explore potential impacts and regulatory approaches to AI.⁶⁵ In partnership with lawmakers from Colorado and elsewhere, Maroney and the Connecticut group produced a draft bill on algorithmic discrimination, SB 2, which served as the template for Colorado’s SB 24-205.⁶⁶

SB 2 passed the Connecticut senate last April and appeared to have significant support in the House before industry opposition persuaded the state's governor, Ned Lamont, to threaten a veto, derailing it in the 2024 legislative session.⁶⁷ However, Maroney promptly pledged to renew his efforts, and, joined by a majority of the state's senators, he has already proposed a successor bill in the 2025 session.⁶⁸

Maroney has predicted that “a dozen or more states” will propose comparable legislation.⁶⁹ Thus, while it is premature to say whether the approach embodied in Connecticut's bill and Colorado's newly enacted law will take root nationally, it appears likely that state policymakers will continue working to mitigate bias from AI applications impacting consequential decisions about their residents' lives.

Sovereign AI Efforts and Public Interest Compute

The most advanced AI models have typically required immense inputs of data and computing power to develop. Training-run costs approach or exceed \$100 million for some released models, and this figure has been projected to pass \$1 billion by 2027.⁷⁰ In recent days, however, DeepSeek, a Chinese AI developer, has dramatically upended these projections—and the assumptions underpinning frontier model development and the global semiconductor industry—by releasing R1, an advanced model approximating the performance of its leading American counterparts.⁷¹ R1 was apparently developed using vastly lower computing power and at markedly lower cost than the most cutting-edge U.S. incumbents.⁷²

The full implications of DeepSeek's accomplishment are still coming into focus. However, computing infrastructure remains a vital and highly expensive prerequisite to advanced AI research and development. Access to sophisticated GPU processors remains finite. And while DeepSeek's progress in efficient model development may impact long term chip demand, computing costs for competitive private sector AI developers have already skyrocketed in recent years.⁷³ OpenAI, for example, was projected to spend more than \$5 billion on computing costs in 2024 alone.⁷⁴ Fueling these outlays, AI firms have raised and invested well over \$100 billion in the past year.⁷⁵

Massively capitalized and with unmatched access to the resources needed for cutting-edge development, large AI developers are rapidly outpacing the public and not-for-profit sectors in their ability to recruit the third critical input for AI research and development: highly skilled researchers, engineers, and other talent upon whom cutting-edge development depends.

These barriers have prompted worry that progress will increasingly concentrate within a handful of institutions, above all private companies with significant commercial interests.⁷⁶ Accordingly, there is concern that the vast inputs of data, compute, and talent will focus on developing monetizable applications rather than the most public-benefitting scientific and

social challenges. Additionally, observers and governments outside the subset of countries where AI development is centered worry that AI will privilege the interests of a few major powers at the expense of the global majority.⁷⁷

Over the past few years, these concerns have prompted calls for public investment in shared local computing infrastructure, or “public” or “sovereign” compute.⁷⁸ In 2020, for example, more than twenty leading universities issued a joint letter urging the United States government to create a “national research cloud” to support academic and public interest AI research.⁷⁹ Warning that brain drain and the skyrocketing cost and diminished accessibility of computing hardware and data were combining to threaten vital research, the signatories recommended that this new national resource pursue at least three aims: to “provide academic and public interest researchers with free or substantially discounted access to the advanced hardware and software required to develop new fundamental AI technologies and applications in the public interest,” to provide “expert personnel necessary to deploy these advanced technologies at universities across the country,” and to “redouble [government agencies’] efforts to make more and better quality data available for public research at no cost.”⁸⁰

With bipartisan and private sector support, these recommendations were adopted in the National Artificial Intelligence Initiative Act of 2020, which directed interagency efforts—including a task force led by the National Science Foundation (NSF) and the White House Office of Science and Technology Policy—to develop a road map for the creation of such a national resource.⁸¹ A final road map and implementation plan were released in January 2023, and subsequently, a pilot was established within the NSF. In July 2023, federal legislation, called the CREATE AI Act, was introduced to establish a National AI Research Resource (NAIRR), though it did not pass before the end of the 118th Congress.⁸²

On the global stage, numerous governments have extolled the importance of national AI capabilities for strategic and economic success in the twenty-first century.⁸³ Increasingly, nations are investing in sovereign AI capabilities, encompassing investments in domestic infrastructure, hardware, data, and expertise, as well as a host of regulatory efforts to advance their competitiveness, promote local innovation and economic growth, and enhance their ability to secure national interests in an age of disruption.

Public computing would therefore seem to be an area where national governments are vigorously engaged and geostrategic imperatives predominate. Nevertheless, a handful of U.S. states have entered the picture, advancing proposals to invest in public-interest computing infrastructure of their own to support socially beneficial research and development.

New York, for example, recently launched a \$400 million project called Empire AI to create an AI computing center and support public-interest research and development among a consortium of local universities and foundations.⁸⁴ The project is designed to help participating

institutions—and the state as a whole—attract and retain top-flight AI talent, leverage economies of scale, and pursue public-interest AI research that might otherwise founder for lack of resources.

In California, a similar proposal tucked into the AI safety bill, SB 1047, attracted significant support. While that legislation focused on catastrophic risk reduction, it also included language to create a new public cloud computing cluster called CalCompute to foster “research and innovation that benefits the public” and enable “equitable innovation by expanding access to computational resources.”⁸⁵ It also required California’s executive branch to provide the legislature with a report detailing the cost and funding sources for CalCompute, recommendations for its structure and governance, opportunities to bolster the state’s AI workforce, and other topics.

While SB 1047 did not become law, CalCompute garnered support even from high-profile opponents of the overall bill.⁸⁶ Newsom’s pledge to continue working with the legislature on a follow-up AI legislation leaves much for deliberation this year. At a minimum, it is far from certain that CalCompute is off the table. Its resurrection may grow more likely after the recent federal election, as California seeks to demonstrate leadership independent of Washington, DC, and to counterbalance the incoming presidential administration and a Congress whose commitment to durably establishing and funding the NAIRR remains uncertain.

On the one hand, state programs may invite redundancy and are unlikely to match the scale of congressional appropriations. On the other hand, if successful, projects like EmpireAI and CalCompute would secure multiple state objectives.⁸⁷ First, they would bolster local academic research centers and entrepreneurs, equipping them to compete for talent and pursue socially beneficial but cost-prohibitive research. Second, they would help seed early-stage entrepreneurship, laying the groundwork for innovation and business formation, with all the reputational and economic advantages that could endow. Third, by building out research clusters and hubs of AI development, state-sponsored computing resources would foster network effects and advance the aim—shared by many states—not only to promote their own ability to harness and deploy AI, but also to position themselves as favorable climates for further investment and entrepreneurship. Finally, success in attracting, incubating, and hosting a meaningful share of the burgeoning global AI market could equip states to expand their influence as regulators, positioning them for influence as AI’s deployment impacts their constituents, societies, and operations over the years to come.

Nation-states have long recognized these imperatives and are moving to compete and shape their futures as new capabilities of automation change their domestic and geostrategic environments. Mindful of a comparable risk and opportunity landscape, some U.S. states are likewise placing their bets and seeking a seat at the table.

Conclusion

Across a swath of the most globally charged topics raised by AI, states are at the leading edge of the U.S. policy response. To an as-yet underappreciated extent, this is typical of a new federalism emergent in technology regulation.⁸⁸ This phenomenon is neither wholly good nor bad, but is, for the time being, an essential facet of U.S. technology policy. As the United States enters a period of uncertainty and potential upheaval, it is likely that state efforts will grow. This expansion may test not only the limits of state authority and political will, but also states' capacity, individually or in concert, to influence a massive and growing global industry.

In a space as fast moving as AI, prediction is difficult. However, there is strong evidence that issues such as misinformation, safety, bias, and the evolving barriers to research and development will loom large on the policy agenda this year. In each area, it appears that technology federalism will be an important element of the overall U.S. response.

About the Author

Scott Kohler is a nonresident scholar at Carnegie California. His research explores the nexus of technology, law, and public policy, with a focus on evolving approaches to regulation and structures of governance. He has spent the past decade advising technology companies on cutting edge legal and regulatory challenges, including at Google, where he served as lead counsel for Alphabet's privacy governance program, and Nest Labs, where he led the global public policy team and was responsible for privacy, safety, and regulatory affairs.

Kohler previously served as senior legal advisor at the U.S. Department of Energy, advising the Secretary of Energy and programs across the department on issues spanning its energy, environmental, and scientific missions. Earlier in his career, he worked at the Partnership for Public Service and in private legal practice at Cravath, Swaine & Moore LLP. Before attending law school, he worked at CNN, receiving a Peabody Award for contributions to the network's reporting on international terrorism.

Kohler is a graduate of Columbia Law School and Duke University, where he also served as a research associate in public policy studies, co-authoring *Casebook to the Foundation: A Great American Secret* (PublicAffairs, 2007).

Notes

- 1 See, for example, James Kanter, “E.U. Parliament Passes Measure to Break Up Google in Symbolic Vote,” *New York Times*, November 27, 2014, <https://www.nytimes.com/2014/11/28/business/international/google-european-union.html>; “Taken together, the level of policy-making activity being devoted to the company signifies the growing antipathy to American technological dominance in the European Union even as its citizens grow ever more reliant on its gadgetry and conveniences”; Nathaniel Persily, “Can Democracy Survive the Internet?,” *Journal of Democracy* 28, no. 2 (2017): 63, https://www.journalofdemocracy.org/wp-content/uploads/2017/04/07_28.2_Persily-web.pdf; “Reluctantly or not, [search and social media] platforms are the new intermediary institutions for our present politics. The traditional organizations of political parties and the legacy media will not reemerge for the Internet age in anything like their prior incarnations”; Dipayan Ghosh, Beware of A.I. in Social Media Advertising, *New York Times*, March 26, 2018, <https://www.nytimes.com/2018/03/26/opinion/ai-social-media-advertising.html>; and Evan Osnos, “Can Mark Zuckerberg Fix Facebook Before it Breaks Democracy?,” *New Yorker*, September 10, 2018, <https://www.newyorker.com/magazine/2018/09/17/can-mark-zuckerberg-fix-facebook-before-it-breaks-democracy>.
- 2 The classic laboratories formulation was sketched by Alexander Hamilton, who wrote of states competing for the people’s “affection.” Alexander Hamilton, James Madison, and John Jay, *The Federalist Papers* (Signet Signet Classics, Mentor, Plume and Meridian Books, 1961), 120. In modern times, the formulation was memorably enshrined by the late Supreme Court justice Louis Brandeis: *New State Ice Co. v. Liebmann*, 285 U.S. 262, (1932) (Brandeis, J., dissenting), 311. While the laboratories model has been cited thousands of times and holds considerable sway in the national imagination, it has also been challenged. See Charles W. Tyler and Heather K. Gerken, “The Myth of the Laboratories of Democracy,” *Columbia Law Review* 122, no. 8 (2022): 2187 and 2190, <https://columbialawreview.org/content/the-myth-of-the-laboratories-of-democracy/>. Tyler and Gerken argue that state-level policy innovation is not so much driven by states as impressed upon them by policy-demanding interest groups, frequently coordinated nationally and connected by partisan political networks. While this essay emphasizes the breadth and significance of state policymaking on AI rather than the processes underlying it, forthcoming essays in this series will explore the federalism implications of this phenomenon more directly. For present purposes, while Tyler and Gerken offer a distinct account of the factors driving state-level experimentation, they agree that “states are, in fact, flourishing sites for policy innovation.”

- 3 Here, the “California Effect,” wherein the nation’s most populous state increases the stringency of national standards through the exercise of its market and regulatory powers, has been on display. See David Vogel, *Trading Up: Consumer and Environmental Regulation in a Global Economy* (Cambridge, MA: Harvard University Press 1995). As we will see, other states have acted as first movers as well.
- 4 “Techlash? America’s Growing Concern with Major Technology Companies,” Knight Foundation, March 11, 2020, <https://knightfoundation.org/reports/techlash-americas-growing-concern-with-major-technology-companies/>.
- 5 Rana Foroohar, “Year in a Word: Techlash,” *Financial Times*, December, 16, 2018, <https://www.ft.com/content/76578fba-fca1-11e8-ac00-57a2a826423e>.
- 6 “Techlash?,” Knight Foundation.
- 7 On competition and the political process, see, for example, Eileen Guo, “Facebook is Now Officially Too Powerful, Says the US Government,” *MIT Technology Review*, December 9, 2020; and Persily, “Can Democracy Survive the Internet.” On extremist content, see, for example, Anushka Asthana and Sam Levin, “UK Urges Tech Giants to Do More to Prevent Spread of Extremism,” *Guardian* July 31, 2017, <https://www.theguardian.com/technology/2017/aug/01/uk-urges-tech-giants-to-do-more-to-prevent-spread-of-extremism>.
- 8 Anu Bradford, *The Brussels Effect: How the European Union Rules the World* (Oxford: Oxford University Press, 2020).
- 9 Anu Bradford, *Digital Empires: The Global Battle to Regulate Technology* (Oxford: Oxford University Press, 2023).
- 10 Astead W. Herndon, “Elizabeth Warren Proposes Breaking Up Tech Giants Like Amazon and Facebook,” *New York Times*, March 8, 2019, <https://www.nytimes.com/2019/03/08/us/politics/elizabeth-warren-amazon.html>; and Jaclyn Diaz, “Trump Vows to Veto Defense Bill Unless Shield for Big Tech Is Scrapped,” National Public Radio, December 2, 2020, <https://www.npr.org/2020/12/02/941019533/trump-vows-to-veto-defense-bill-unless-shield-for-big-tech-is-scrapped>.
- 11 Bradford terms these “digital empires”: a European rights-driven model, an American market-driven model, and a Chinese state-drive model. See Bradford, *Digital Empires*.
- 12 While this article focuses on the U.S. and civilian tech policy landscapes, for a discussion of evolving attitudes toward AI safety among Chinese policymakers, see Matt Sheehan, “China’s Views on AI Safety are Changing—Quickly,” Carnegie Endowment for International Peace, August 27, 2024, <https://carnegieendowment.org/research/2024/08/china-artificial-intelligence-ai-safety-regulation?lang=en>.
- 13 Jenna Wortham, “Obama Brought Silicon Valley to Washington,” *New York Times*, October 25, 2016, <https://www.nytimes.com/2016/10/30/magazine/barack-obama-brought-silicon-valley-to-washington-is-that-a-good-thing.html>.
- 14 On AlphaGo and AlphaZero, see generally Mustafa Suleyman and Michael Bhaskar, *The Coming Wave: Technology, Power, and the 21st Century’s Greatest Dilemma* (Phoenix, AZ: Crown Press, 2023). The Cambridge Analytica episode generated headlines, prompted an immense (if transitory) drop in Facebook’s market capitalization, and led the Federal Trade Commission to impose on Facebook a \$5 billion penalty and onerous governance obligations. See “FTC Imposes \$5 Billion Penalty and Sweeping New Privacy Restrictions on Facebook,” Federal Trade Commission, July 24, 2019, <https://www.ftc.gov/news-events/news/press-releases/2019/07/ftc-imposes-5-billion-penalty-sweeping-new-privacy-restrictions-facebook>; and Rupert Neate, “Over \$119bn Wiped Off Facebook’s Market Cap After Growth Shock,” *Guardian*, July 26, 2018, <https://www.theguardian.com/technology/2018/jul/26/facebook-market-cap-falls-109bn-dollars-after-growth-shock>.
- 15 See, for example, Cecilia Kang and Kenneth P. Vogel, “Tech Giants Amass a Lobbying Army for an Epic Washington Battle,” *New York Times*, June 5, 2019, <https://www.nytimes.com/2019/06/05/us/politics/amazon-apple-facebook-google-lobbying.html>.
- 16 On breakthroughs, see, for example, Derrick Bryson Taylor, Cade Metz, and Katrina Miller, “Nobel Physics Prize Awarded for Pioneering A.I. Research by 2 Scientists,” *New York Times*, October 8, 2024, <https://www.nytimes.com/2024/10/08/science/nobel-prize-physics.html>; and Claire Moses, Cade Metz, and Teddy Rosenbluth, “Nobel Prize in Chemistry Goes to 3 Scientists for Predicting and Creating

- Proteins,” *New York Times*, October 9, 2024, <https://www.nytimes.com/2024/10/09/science/nobel-prize-chemistry.html>. On the scale of private investment, see, for example, Stephen Morris, Hannah Murphy, and Camilla Hodgson, “Big Tech Groups Say their \$100 Billion AI Spending Spree Is Just Beginning,” *Financial Times*, August 2, 2024, <https://www.ft.com/content/b7037ce1-4319-4a4a-8767-0b1373cec9ce>. On AI models’ usage, see “OpenAI Says ChatGPT’s Weekly Users Have Grown to 200 Million,” Reuters, August 29, 2024, <https://www.reuters.com/technology/artificial-intelligence/openai-says-chatgpts-weekly-users-have-grown-200-million-2024-08-29/>.
- 17 “Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence,” The White House, October 30, 2023, <https://web.archive.org/web/20250119193924/https://www.whitehouse.gov/briefing-room/presidential-actions/2023/10/30/executive-order-on-the-safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence/>; “Executive Order on Removing Barriers to American Leadership in Artificial Intelligence,” The White House, January 23, 2025, <https://www.whitehouse.gov/presidential-actions/2025/01/removing-barriers-to-american-leadership-in-artificial-intelligence/>; and “Driving U.S. Innovation in Artificial Intelligence: A Roadmap for Artificial Intelligence Policy in the United States Senate,” The Bipartisan Senate AI Working Group, United States Senate, May 2024, https://www.schumer.senate.gov/imo/media/doc/Roadmap_Electronic1.32pm.pdf.
 - 18 On tensions, see, for example, Clothilde Goujard and Alfred Ng, “EU and US Reach Deal to Let Data Flow Across the Atlantic,” POLITICO, July 10, 2023, <https://www.politico.eu/article/eu-signs-off-on-data-transfers-deal-with-us/>; the authors note continued worries within the European Data Protection Board, opposition from European Parliamentarians, and legal uncertainty regarding the EU-US Data Privacy Framework. On demoscrosis generally, see Jonathan Rauch, *Government’s End: Why Washington Stopped Working* (New York: PublicAffairs, 1994).
 - 19 See, for example, “States Intensify Work on AI Legislation,” Business Software Alliance, February 14, 2024, <https://www.bsa.org/news-events/news/bsa-analysis-states-intensify-work-on-ai-legislation#:~:text=There%20are%20a%20total%20of,to%20addressing%20concerns%20about%20deepfakes.>
 - 20 “US State Privacy Legislation Tracker 2024,” International Association of Privacy Professionals, accessed December 16, 2024, https://iapp.org/media/pdf/resource_center/State_Comp_Privacy_Law_Chart.pdf.
 - 21 See, for example, Virginia Code, Chapter 53. Consumer Data Protection Act § 59.1-577.
 - 22 “Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation),” European Union, April 27, 2016, Article 22, <https://eur-lex.europa.eu/eli/reg/2016/679/oj/eng>.
 - 23 Federal law, for example, defines “artificial intelligence” as “a machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations or decisions influencing real or virtual environments. Artificial intelligence systems use machine and human-based inputs to (a) perceive real and virtual environments; (b) abstract such perceptions into models through analysis in an automated manner; and (c) use model inference to formulate options for information or action. See “Chapter 119—Artificial Intelligence Initiative,” 15 USC 9401 (3), [https://uscode.house.gov/view.xhtml?req=\(title:15%20section:9401%20edition:prelim\)](https://uscode.house.gov/view.xhtml?req=(title:15%20section:9401%20edition:prelim)).
Likewise, California recently adopted a uniform legislative definition of AI: “an engineered or machine-based system that varies in its level of autonomy and that can, for explicit or implicit objectives, infer from the input it receives how to generate outputs that can influence physical or virtual environments.” See “Chapter 843—An act to amend Section 22675 of the Business and Professions Code, to amend Section 75002 of the Education Code, and to amend Sections 11546.45.5, 11547.5, and 53083.1 of the Government Code, relating to artificial intelligence,” California Assembly Bill No. 2885, https://leginfo.ca.gov/faces/billNavClient.xhtml?bill_id=202320240AB2885.
 - 24 “Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence,” The White House.
 - 25 “Executive Order N-12-23,” Executive Department, State of California, September 6, 2023, 1, https://www.gov.ca.gov/wp-content/uploads/2023/09/AI-EO-No.12-_-GGN-Signed.pdf.

- 26 “Executive Order 2023-19—Expanding and Governing the Use of Generative Artificial Intelligence Technologies within the Commonwealth of Pennsylvania,” Commonwealth of Pennsylvania, September 20, 2023, 1, <https://www.pa.gov/en/governor/newsroom/press-releases/governor-josh-shapiro-signs-executive-order-on-commonwealth-use-.html>.
- 27 “Executive Order 01.01.2024.02—Catalyzing the Responsible and Productive Use of Artificial Intelligence in Maryland State Government,” Executive Department, State of Maryland, January 8, 2024, https://governor.maryland.gov/Lists/ExecutiveOrders/Attachments/31/EO%2001.01.2024.02%20Catalyzing%20the%20Responsible%20and%20Productive%20Use%20of%20Artificial%20Intelligence%20in%20Maryland%20State%20Government_Accessible.pdf. The critical areas identified for further policy recommendation include Maryland’s workforce, economic development, and security.
- 28 “Executive Order No. 629—Establishing an Artificial Intelligence Strategic Task Force,” Commonwealth of Massachusetts, February 15, 2024, <https://www.mass.gov/executive-orders/no-629-establishing-an-artificial-intelligence-strategic-task-force>.
- 29 “Executive Order 24-06—Artificial Intelligence and Data Centers of Excellence,” State of Rhode Island, February 29, 2024, <https://governor.ri.gov/executive-orders/executive-order-24-06>.
- 30 “SB 149—Artificial Intelligence Amendments, 2024 General Session,” Utah State Legislature, March, 13, 2024, <https://le.utah.gov/-/2024/bills/static/SB0149.html>.
- 31 Office of Artificial Intelligence Policy, Utah Department of Commerce, accessed November 26, 2024, <ai.utah.gov>.
- 32 See, for example, Janna Anderson and Lee Rainie, “As AI Spreads, Experts Predict the Best and Worst Changes in Digital Life by 2035,” Pew Research Center, June 21, 2023, <https://www.pewresearch.org/internet/2023/06/21/as-ai-spreads-experts-predict-the-best-and-worst-changes-in-digital-life-by-2035/>.
- 33 On deepfakes, see, for example, Bobby Chesney and Danielle Citron, “Deep Fakes: A Looming Challenge for Privacy, Democracy, and National Security,” *California Law Review* 107 (2019): 1753, <https://www.californialawreview.org/print/deep-fakes-a-looming-challenge-for-privacy-democracy-and-national-security>; William A. Galston, “Is Seeing Still Believing? The Deepfake Challenge to Truth in Politics,” Brookings Institution, January 8, 2020, <https://www.brookings.edu/articles/is-seeing-still-believing-the-deepfake-challenge-to-truth-in-politics/>. But also see Bruce Schneier and Nathan Sanders, “The Apocalypse That Wasn’t: AI Was Everywhere in 2024’s Elections, but Deepfakes and Misinformation were only Part of the Picture,” Harvard Kennedy School, December 4, 2024, <https://ash.harvard.edu/articles/the-apocalypse-that-wasnt-ai-was-everywhere-in-2024s-elections-but-deepfakes-and-misinformation-were-only-part-of-the-picture/>. On hallucination, see, for example, Cade Metz, “Chatbots May ‘Hallucinate’ More Often Than Many Realize,” *New York Times*, November 6, 2023, <https://www.nytimes.com/2023/11/06/technology/chatbots-hallucination-rates.html>.
- 34 Josh A. Goldstein, Renee DiResta, Girish Sastry, Micah Musser, Matthew Gentzel, and Katerina Sedova, “Generative Language Models and Automated Influence Operations: Emerging Threats and Potential Mitigations,” Internet Observatory, Cyber Policy Center, Stanford University, January 11, 2023, <https://cyber.fsi.stanford.edu/io/publication/generative-language-models-and-automated-influence-operations-emerging-threats-and>.
- 35 See, for example, Yona TR Golding, “The News Media and AI: A New Front in Copyright Law,” *Columbia Journalism Review*, October 18, 2023, https://www.cjr.org/business_of_news/data-scraping-ai-litigation-lawsuit-artists-authors.php; and Michael M. Grynbaum and Ryan Mac, “The Times Sues OpenAI and Microsoft Over A.I. Use of Copyrighted Work,” *New York Times*, December 27, 2023, <https://www.nytimes.com/2023/12/27/business/media/new-york-times-open-ai-microsoft-lawsuit.html>.
- 36 See *Whitney v. California*, 274 U.S. 357, 377 (1927) (Brandeis, J., concurring); “If there be time to expose through discussion the falsehood and fallacies, to avert the evil by the processes of education, the remedy to be applied is more speech, not enforced silence.”
- 37 “SB 149—Artificial Intelligence Amendments, 2024 General Session,” Utah State Legislature.
- 38 *Ibid.*

- 39 See, for example, “HB 5450—An Act Concerning Artificial Intelligence, Deceptive Synthetic Media and Elections,” Connecticut House, March 7, 2024; “H 1974—An Act Regulating the Use of Artificial Intelligence (AI) in Providing Mental Health Services,” Massachusetts House, February 16, 2023; “A 216A—An Act to Amend the General Business Law, in Relation to Requiring Advertisements to Disclose the Use of Synthetic Media,” New York Assembly, January 4, 2023; “H 4660,” South Carolina House, January 9, 2024; and “H 6286—An Act Relating to Commercial Law – General Regulatory Provisions – Generative Artificial Intelligence Models,” Rhode Island House, April 19, 2023.
- 40 “SB 942 California AI Transparency Act,” California State Legislature, September 19, 2024, https://leginfo.legislature.ca.gov/faces/billTextClient.xhtml?bill_id=202320240SB942.
- 41 Disclosures must be “permanent or extraordinarily difficult to remove, to the extent it is technically feasible” and must include specified information, including the name of the AI system provider, the name and version number of the AI system, time and date of the content’s creation or alteration, and a unique identifier.
- 42 Ibid.
- 43 “AB 2355—Political Reform Act of 1974: Political Advertisements: Artificial Intelligence,” California Assembly, September 17, 2024, https://leginfo.legislature.ca.gov/faces/billNavClient.xhtml?bill_id=202320240AB2355.
- 44 “CS/HB 919—Artificial Intelligence Use in Political Advertising,” Florida Senate, July 1, 2024, <https://www.flsenate.gov/Session/Bill/2024/919/?Tab=BillText>.
- 45 See, for example, “H. 4660,” South Carolina Assembly, January 9, 2024, https://www.scstatehouse.gov/sess125_2023-2024/bills/4660.htm; and “S7592,” New York Senate, July 7, 2023, <https://www.nysenate.gov/legislation/bills/2023/S7592/amendment/original>. The bills vary somewhat from one state to another. For example, New York’s proposed bill would be somewhat broader than Florida’s enacted statute, applying to “any political communication . . . generated in whole or in part” with AI, and lacking the Florida law’s requirement of injurious or deceptive intent.
- 46 “Substitute Senate Bill 6299,” State of Washington, January 31, 2024, 1, <https://lawfilesexet.leg.wa.gov/biennium/2023-24/Pdf/Bills/Senate%20Bills/6299-S.pdf?q=20241126145528>.
- 47 Ibid.
- 48 See, for example, “Regulation 2024/1689 of the European Parliament and of the Council of 13 June 2024, laying down harmonised rules on artificial intelligence and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act).” See also “Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence,” The White House. The term “AI safety” is often used to encompass a wide and still-emerging array of risk reduction efforts closely tied to the parallel concept of “trust”—that AI systems will perform as intended, subject to human intent, and will not cause or materially enable harm, including from malicious use, malfunctions, and system impacts. See, for example, “International Scientific Report on the Safety of Advanced AI: Interim Report,” UK Government, May 2024, https://assets.publishing.service.gov.uk/media/6716673b96def6d27a4c9b24/international_scientific_report_on_the_safety_of_advanced_ai_interim_report.pdf.
- 49 Scott Kohler, “All Eyes on Sacramento: SB 1047 and the AI Safety Debate,” Carnegie Endowment for International Peace, September 11, 2024, <https://carnegieendowment.org/posts/2024/09/california-sb1047-ai-safety-regulation?lang=en>.
- 50 “2024 GOP Platform: Make America Great Again,” GOP, <https://www.documentcloud.org/documents/24795052-2024-gop-platform-july-7-final>.
- 51 “Fact Sheet: Biden-Harris Administration Secures Voluntary Commitments From Leading Artificial Intelligence Companies to Manage the Risks Posed by AI,” The White House, July 21, 2023, <https://perma.cc/Q8QS-3AGS>; and “Historic First as Companies Spanning North America, Asia, Europe and Middle East Agree Safety Commitments on Development of AI,” Department for Science, Innovation and Technology, May 21, 2024, <https://www.gov.uk/government/news/historic-first-as-companies-spanning-north-america-asia-europe-and-middle-east-agree-safety-commitments-on-development-of-ai>. But also see Rebecca Heilweil, “Where Biden’s Voluntary AI Commitments Go From Here,” FedScoop, November 13, 2024, <https://fedscoop.com/voluntary-ai-commitments-biden-trump-white-house/>; she notes uncertainty regarding

the existing White House-brokered commitments. The White House commitments, however, are one of numerous such pledges to which model developers have agreed. See “Tracking Voluntary Commitments,” Anthropic, November 18, 2024, last accessed December 20, 2024, <https://www.anthropic.com/voluntary-commitments#list-of-voluntary-commitments>; the list includes eight sets of risk-reduction commitments to which Anthropic has agreed.

- 52 See Scott Kohler and Ian Klaus, “A Heated California Debate Offers Lessons for AI Governance,” Carnegie Endowment for International Peace, October 8, 2024, <https://carnegieendowment.org/posts/2024/10/california-sb1047-ai-safety-bill-veto-lessons?lang=en>; and Ketan Ramakrishnan, “Tort Law Is the Best Way to Regulate AI,” *Wall Street Journal*, September 24, 2024, <https://www.wsj.com/opinion/tort-law-is-the-best-way-to-regulate-ai-california-legal-liability-065e1220>. As Ramakrishnan describes, “tort law is animated by a simple and powerful idea: If you harm other people by failing to take reasonable care, then fairness requires you to compensate them.”
- 53 Vogel, *Trading Up*.
- 54 “SB-1047—Safe and Secure Innovation for Frontier Artificial Intelligence Models Act,” California State Senate, September 3, 2024, https://leginfo.legislature.ca.gov/faces/billNavClient.xhtml?bill_id=202320240SB1047; and Kohler, “All Eyes on Sacramento.”
- 55 “SB-1047—Safe and Secure Innovation for Frontier Artificial Intelligence Models Act: Status,” California State Senate, September 29, 2024, https://leginfo.legislature.ca.gov/faces/billStatusClient.xhtml?bill_id=202320240SB1047.
- 56 See, for example, Jennifer Aaker, Fei Fei Li, Thomas Higginbotham, Zoe Weinberg, and Wendy De La Rosa, “Designing AI for All: A Primer on Bias in Artificial Intelligence Systems,” Stanford Graduate School of Business, February 19, 2020, <https://www.gsb.stanford.edu/faculty-research/case-studies/designing-ai-all-primer-bias-artificial-intelligence-systems>.
- 57 “Blueprint for an AI Bill of Rights: Making Automated Systems Work for the American People,” The White House, October 2022, 3, <https://web.archive.org/web/20250119213350/https://www.whitehouse.gov/ostp/ai-bill-of-rights/>.
- 58 *Ibid.*, 5.
- 59 “Artificial Intelligence and Civil Rights,” U.S. Department of Justice, Civil Rights Division, <https://www.justice.gov/crt/ai>; and “Joint Statement on Enforcement of Civil Rights, Fair Competition, Consumer Protection, and Equal Opportunity Laws in Automated Systems,” U.S. Equal Employment Opportunity Commission, April 4, 2024, <https://www.eeoc.gov/joint-statement-enforcement-civil-rights-fair-competition-consumer-protection-and-equal-opportunity>.
- 60 “Artificial Intelligence Risk Management Framework (AI RMF 1.0),” National Institute of Standards and Technology (NIST), January 2023, 18, <https://www.nist.gov/itl/ai-risk-management-framework>; and Reva Schwartz, Apostol Vassilev, Kristen Greene, Lori Perine, Andrew Burt, and Partick Hall, “Towards a Standard for Identifying and Managing Bias in Artificial Intelligence,” NIST Special Publication 1270, March 2022, <https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.1270.pdf>.
- 61 Letter from Ranking Member Zoe Lofgren of the House Committee on Science, Space, and Technology to California State Senator Scott Wiener, August 7, 2024, <https://lofgren.house.gov/sites/evo-subsites/lofgren.house.gov/files/evo-media-document/8.7.24%20to%20Senator%20Wiener.pdf>; she notes that while “there is little scientific evidence of harm of ‘mass casualties or harmful weapons created’ from advanced models . . . there is ample evidence of real-world risks like discrimination and misinformation.” And Letter from Members of California’s Congressional Delegation to Governor Gavin Newsom, August 15, 2024, https://democrats-science.house.gov/imo/media/doc/2024-08-15%20to%20Gov%20Newsom_SB1047.pdf; the letter contrasts “hypothetical existential risks” with “demonstrable AI risks like misinformation, discrimination, nonconsensual deepfakes, environmental impacts, and workforce displacement.”
- 62 See generally, for example, “Executive Order on Ending Radical and Wasteful Government DEI Programs and Preferencing,” The White House, January 20, 2025, <https://www.whitehouse.gov/presidential-actions/2025/01/ending-radical-and-wasteful-government-dei-programs-and-preferencing/>; and Perry Stein and David Nakamura, “Justice Department Issues Freeze for Civil Rights Division,” *Washington Post*, January 22, 2025, <https://www.washingtonpost.com/national-security/2025/01/22/justice-civil-rights-freeze-shutdown/>.

- 63 “SB 24-205—Consumer Protections for Artificial Intelligence,” State of Colorado, May 17, 2024, <https://leg.colorado.gov/bills/sb24-205>. Such discrimination is defined to include unlawful differential treatment on the basis of protected characteristics such as “age, color, disability, ethnicity, genetic information, limited proficiency in the English language, national origin, race, religion, reproductive health, sex, [or] veteran status.”
- 64 Letter from Colorado Governor Jared Polis to Colorado General Assembly, May 17, 2024, <https://www.dwt.com/-/media/files/blogs/artificial-intelligence-law-advisor/2024/05/sb24205-signing-statement.pdf?rev=a902184eafe046cfb615bb047484e11c&hash=213F4C6CDF52A876011290C24406E7F>.
- 65 Mallory Culhane, “Two Unlikely States Are Leading the Charge on Regulating AI,” POLITICO, May 15, 2024, <https://www.politico.com/news/2024/05/15/ai-tech-regulations-lobbying-00157676>.
- 66 Ibid. The Connecticut legislation was broader than Colorado’s bill. In addition to limiting algorithmic discrimination and imposing related disclosure and governance obligations, the Connecticut proposal would also have regulated nonconsensual intimate deepfakes and deepfakes to influence political campaigns and elections. Additionally, it would have mandated annual training for state employees on the use of generative AI tools, including on how to identify and mitigate inaccurate, hallucinated, or biased outputs.
- 67 Ken Dixon, “Proposed Bill on Artificial Intelligence Regulation in CT Dies After Gov. Ned Lamont Threatens Veto CT Insider,” May 7, 2024, <https://www.ctinsider.com/politics/article/if-bill-ai-survives-ct-house-vote-lamont-19444053.php>.
- 68 See “Proposed SB 2—An Act Concerning Artificial Intelligence,” Connecticut Senate, https://www.cga.ct.gov/asp/CGABillStatus/cgabillstatus.asp?selBillType=Bill&bill_num=SB2. Moreover, despite the failure of SB 2, it bears note that Connecticut has, in fact, already enacted limits, albeit narrower ones, on AI-intermediated discrimination. In 2023, Connecticut enacted a law requiring the state’s Department of Administrative Services to conduct an inventory of AI systems in use by the state itself and mandating that the state avoid unlawful discrimination or disparate impact in its own use of AI. See “SB 1103—An Act Concerning Artificial Intelligence, Automated Decision-Making and Personal Data Privacy,” Connecticut Assembly, June 7, 2023, https://www.cga.ct.gov/asp/cgabillstatus/cgabillstatus.asp?selBillType=Bill&bill_num=SB01103&which_year=2023.
- 69 Culhane, “Two Unlikely States Are Leading the Charge on Regulating AI.”
- 70 Ben Cottier, Robi Rahman, Loredana Fattorini, Nestor Maslej, and David Owen, “How Much Does It Cost to Train AI Frontier Models?,” Epoch AI, June 3, 2024, updated on January 13, 2025, <https://epoch.ai/blog/how-much-does-it-cost-to-train-frontier-ai-models>; and Deepa Seetharaman, “The Next Great Leap in AI is Behind Schedule and Crazy Expensive,” *Wall Street Journal*, December 20, 2024, <https://www.wsj.com/tech/ai/openai-gpt5-orion-delays-639e7693>.
- 71 Jason Karaian and Joe Rennison, “China’s AI Advances Spook Big Tech Investors on Wall Street,” *New York Times*, January 27, 2025, <https://www.nytimes.com/2025/01/27/business/us-stock-market-deepseek-ai-sp500-nvidia.html?searchResultPosition=19>.
- 72 See Matt Sheehan and Sam Winter-Levy, “Chips, China, and a Lot of Money: The Factors Driving the DeepSeek AI Turmoil,” Carnegie Endowment for International Peace, January 28, 2025, <https://carnegieendowment.org/posts/2025/01/deepseek-ai-china-chips-explainer?lang=en>: “Just to give you a sense of DeepSeek’s efficiency, the company claims it trained its model for less than \$6 million, using only about 2,000 chips. That’s an order of magnitude less money than what Meta, for example, spent on training its latest system, which used more than 16,000 chips.”
- 73 See, for example, Shana Lynch, “What to Expect in AI in 2024,” Stanford University Human-Centered Artificial Intelligence, December 8, 2023, <https://hai.stanford.edu/news/what-expect-ai-2024>.
- 74 Cade Metz, Mike Isaac, and Erin Griffith, “Microsoft and OpenAI’s Close Partnership Shows Signs of Fraying,” *New York Times*, October 17, 2024, updated October 21, 2024, <https://www.nytimes.com/2024/10/17/technology/microsoft-openai-partnership-deal.html>; and Mike Isaac and Erin Griffith, “OpenAI Is Growing Fast and Burning Through Piles of Money,” *New York Times*, September 27, 2024, <https://www.nytimes.com/2024/09/27/technology/openai-chatgpt-investors-funding.html>.

- 75 See, for example, Stephen Morris, Hannah Murphy, and Camilla Hodgson, “Big Tech Groups Say their \$100bn AI Spending Spree Is Just Beginning,” *Financial Times*, August 2, 2024, <https://www.ft.com/content/b7037ce1-4319-4a4a-8767-0b1373cec9ce>.
- 76 See, for example, Daniel E. Ho, Jennifer King, Russell C. Wald, and Christopher Wan, “Building a National AI Research Resource: A Blueprint for the National Research Cloud,” Stanford University Human-Centered Artificial Intelligence, October 2021, <https://hai.stanford.edu/white-paper-building-national-ai-research-resource>; the authors note that “while the largest private institutions like platform technology companies and certain elite academic institutions continue to design, develop, and deploy AI systems that can be readily commercialized, a different story is playing out for the public sector and the vast majority of academic institutions, which lack access to core inputs of AI research. The rising costs associated with carrying out research and development are exacerbating the disconnect between current winners and losers in the AI space.” Regarding commercial incentives, it is important to note that some frontier model developers are characterized by differing and complicated organizational arrangements. Anthropic, for example, is organized as a Public Benefit Corporation, which has stated that it is also partly owned by an independent Long-Term Benefit Trust charged with ensuring that “Anthropic responsibly balances the financial interests of stockholders with the interests of those affected by Anthropic’s conduct and [its] public benefit purpose.” See “The Long-Term Benefit Trust,” Anthropic, September 19, 2023, <https://www.anthropic.com/news/the-long-term-benefit-trust>. The trust, in turn, is empowered to appoint a portion—expected eventually to reach a majority—of the board members, subject however to limitations empowering the stockholders to amend this structure in some circumstances. OpenAI, for its part, is organized as a for-profit corporation presently controlled, as a formal matter, by a not-for-profit entity. However, it has announced a plan to change this arrangement, transitioning the for-profit entity to a Public Benefit Corporation and eliminating the nonprofit’s control. See “Why OpenAI’s Structure Must Evolve to Advance Our Mission,” OpenAI, December 27, 2024, <https://openai.com/index/why-our-structure-must-evolve-to-advance-our-mission/>; and David A. Fahrenthold, Cade Metz, and Mike Isaac, “How OpenAI Hopes to Sever its Nonprofit Roots,” *New York Times*, December 17, 2024, <https://www.nytimes.com/2024/12/17/technology/openai-nonprofit-control.html>.
- 77 See, for example, Adam Satariano and Paul Mozur, “The Global Race to Control A.I.,” *New York Times*, August 14, 2024, <https://www.nytimes.com/2024/08/14/briefing/ai-china-us-technology.html>. See also “Governing AI for Humanity: Final Report,” United Nations AI Advisory Body, September 2024, 14, https://www.un.org/sites/un2.un.org/files/governing_ai_for_humanity_final_report_en.pdf; the report notes that “the accelerating development of AI concentrates power and wealth on a global scale, with geopolitical and geoeconomic implications” and that “of the top 100 high-performance computing clusters in the world capable of training large AI models, not one is hosted in a developing country.”
- 78 The United Nations AI Advisory Body, while calling for a distributed and equitable global governance architecture, has stated with respect to compute that “it is unrealistic to promise access to compute that even the wealthiest countries and companies struggle to acquire,” suggesting instead “initiatives towards distributed and federated AI development models” given the barriers to global compute access. See “Governing AI for Humanity: Final Report,” United Nations AI Advisory Body, 14.
- 79 “National Research Cloud Call to Action,” Stanford University Human-Centered Artificial Intelligence, 2020, <https://hai.stanford.edu/national-research-cloud-joint-letter>.
- 80 Ibid.
- 81 Khari Johnson, “AWS, Google, and Mozilla Back National AI Research Cloud Bill in Congress,” *Venture Beat*, June 30, 2020, <https://venturebeat.com/ai/aws-google-and-mozilla-back-national-ai-research-cloud-bill-in-congress/>; and “HR 6216—National Artificial Intelligence Initiative Act of 2020,” 116th Congress, March 12, 2020, <https://www.congress.gov/bill/116th-congress/house-bill/6216>.
- 82 Supporters had urged Congress to pass the act during the lame duck session. See Letter from Industry and University Leaders to Speaker Johnson, Leader Schumer, Leader McConnell, and Leader Jeffries, November 18, 2024, <https://hai.stanford.edu/sites/default/files/2024-11/Industry%20and%20University%20Letter%20on%20CREATE%20AI%20Act.pdf>; and Shana Lynch, “Can the CREATE AI Act Pass the Finish Line?,” Stanford University Human-Centered Artificial Intelligence, November 19, 2024, <https://hai.stanford.edu/news/can-create-ai-act-pass-finish-line>.

- 83 See, for example, Alexandre Piquard, “Macron Wants to Make France a ‘Champion’ in AI,” *Le Monde*, May 22, 2024, https://www.lemonde.fr/en/french-economy/article/2024/05/22/macron-wants-to-make-france-a-champion-in-ai_6672255_21.html; and Jatin Grover, “India to Develop Its Own Sovereign AI Infrastructure: Rajeev Chandrasekhar,” *Financial Express*, November 30, 2023, <https://www.financialexpress.com/business/digital-transformation-india-to-develop-its-own-sovereign-ai-infrastructure-rajeev-chandrasekhar-3321291/>.
- 84 Grace Ashford, “Hochul to Propose A.I. Research Center Using \$275 Million in State Funds,” *New York Times*, January 8, 2024, <https://www.nytimes.com/2024/01/08/nyregion/ai-new-york-hochul.html>; and “Governor Hochul, Industry Leaders and Advocates Celebrate Empire AI Consortium to Make New York a Global Leader in Artificial Intelligence,” Office of Governor Kathy Hochul, April 30, 2024, <https://www.governor.ny.gov/news/governor-hochul-industry-leaders-and-advocates-celebrate-empire-ai-consortium-make-new-york>.
- 85 “SB-1047—Safe and Secure Innovation for Frontier Artificial Intelligence Models Act: Status,” California State Senate.
- 86 Letter from Ranking Member Zoe Lofgren to Senator Scott Wiener; and Letter from Members of the California Congressional Delegation to Governor Gavin Newsom.
- 87 While its long-term viability and impact remain to be demonstrated, New York’s Empire AI consortium began work this fall and is moving forward with plans to construct a \$250 million computing facility at the State University of New York’s Buffalo campus. See Matt Glynn, “Research Begins at Buffalo’s Empire AI Consortium,” *GovTech*, October 14, 2024, <https://www.govtech.com/education/higher-ed/research-begins-at-buffalos-empire-ai-consortium>.
- 88 The phrase “new federalism” is, itself, hardly new. It has been used since at least 1907. See George Gray, *The New Federalism: Annual Address at Pennsylvania Bar Association Thirteenth Annual Meeting, Bedford Springs, Pennsylvania, June 25, 1907* (Farmington Hills, MI: Gale, 2010). More recently, the phrase has frequently been used to describe judicial rulings reanimating states’ legal prerogatives in various fields of policy vis-à-vis the federal government. See, for example, Richard W. Garnett, “The New Federalism, the Spending Power, and Federal Criminal Law,” *Cornell Law Review* 89, no. 1 (2003), <https://scholarship.law.cornell.edu/clr/vol89/iss1/1/>; and Gillian E. Metzger, “Administrative Law as the New Federalism,” *Duke Law Journal* 57, no. 7 (2008): 2023–2109, <https://scholarship.law.duke.edu/dlj/vol57/iss7/2/>. As used here, I mean a resurgence in states’ policymaking to foster and constrain a wide and growing field of private action. Whether the national government will act to curb this resurgence through federal preemption or whether states will outrun their wide legal authorities largely remain to be seen.

Carnegie Endowment for International Peace

In a complex, changing, and increasingly contested world, the Carnegie Endowment generates strategic ideas, supports diplomacy, and trains the next generation of international scholar-practitioners to help countries and institutions take on the most difficult global problems and advance peace. With a global network of more than 170 scholars across twenty countries, Carnegie is renowned for its independent analysis of major global problems and understanding of regional contexts.

Carnegie California

Carnegie California explores the role of subnational regions as global players, as well as the cross-cutting challenges related to Asia, technology, migration, democracy, and climate change playing out in one of the world's most important regions — the American West.



 **CARNEGIE**
ENDOWMENT FOR
INTERNATIONAL PEACE

CarnegieEndowment.org