The University of Texas at Austin
School of Law

December 10, 2025

Dear Representative Issa,

Thank you for the opportunity to supplement my testimony with answers to the important questions below. If you have any further questions, please let me know. I am ready and eager to assist you and the rest of the committee with these and related inquiries.

Sincerely,

Kevin Frazier
AI Innovation and Law Fellow
The University of Texas School of Law

1. The President's Action Plan, released in July, endorsed the concept of a regulatory sandbox. Can you elaborate on the benefits of a regulatory sandbox, and how federal preemption would either enable or complement that?

a. Follow up: Why is a safe harbor concept necessary?

b. Follow up: How might this tie in, temporally, with a moratorium proposal?

A regulatory sandbox formalizes what has long been the best approach—and America's approach—to governing novel technologies: trial and error.[1] To do otherwise is unwise. "One could well ask whether any technology, including the most benign, would ever have been established if it had first been forced to demonstrate that it would do no harm."[2] The introduction of every new technology presents risks and results in harms. Society should attempt to mitigate both. Such mitigation, though, must be tailored to the benefits of the technology. In the case of AI, despite a bias among academics[3] and journalists to report on supposed downsides of technologies,[4] its demonstrated benefits warrant a trial-and-error approach through regulatory sandboxes.[5]

Though there is no single way to design a regulatory sandbox, they typically permit such trial by allowing participants to deploy their products subject to few or no additional legal limitations. Sandbox participation, however, is not a free pass to recklessly operate. Participation is usually contingent on increased oversight, information sharing obligations, and the adherence to specific consumer protections.[6]

Though some states[7] have opted to apply the "try-first" mentality to AI governance called for by President Trump,[8] many others are instead imposing a trial without error approach—imposing

---

[1] *See, e.g.,* Aaron Wildavsky, *Trial and error versus trial without error*, in RETHINKING RISK AND THE PRECAUTIONARY PRINCIPLE 24-25 (Julian Morris ed. 2000).
[2] *Id.* at 24.
[3] Peter Königs, *The negativity crisis of AI ethics*, 206 SYNTHESE *passim* (Nov. 28, 2025). https://link.springer.com/article/10.1007/s11229-025-05378-9.
[4] *See, e.g.,* Kim Andersen et al., *The Scary World Syndrome: News Orientations, Negativity Bias, and the Cultivation of Anxiety*, 27 MASS COMM. & SOC. 502 (2024).
[5] *See, e.g.,* Sophie Bushwick, 10 Ways AI Was Used for Good This Year, Scientific American (Dec. 14, 2022), https://www.scientificamerican.com/article/10-ways-ai-was-used-for-good-this-year/ (documenting several instances of positive uses of AI in 2022, when the technology was even less capable than it is today).
[6] Angela Attrey et al., The role of sandboxes in promoting flexibility and innovation in the digital age, OECD at 6-7 (2020), https://www.oecd.org/content/dam/oecd/en/publications/reports/2020/06/the-role-of-sandboxes-in-promoting-flexibility-and-innovation-in-the-digital-age_ddcd3d40/cdf5ed45-en.pdf.
[7] *See* Regulatory Relief, Utah Office of AI Policy (last accessed Dec. 8, 2025), https://ai.utah.gov/regulatory-mitigation/ (detailing Utah's regulatory sandbox for AI); Victor D. Vital & Alexander M. Clark, Texas Enters the AI Sandbox with TRAIGA: Implications for Business Trials, ABA (July 14, 2025), https://www.americanbar.org/groups/business_law/resources/business-law-today/2025-july/texas-enters-ai-sandbox-with-traiga-implications-business-trials/ (providing an overview of Texas's impending regulatory sandbox).
[8] The AI Action Plan, White House, https://www.whitehouse.gov/wp-content/uploads/2025/07/Americas-AI-Action-Plan.pdf (July 2025).

vague and subjective conditions on the development of AI.[9] The patchwork that is already being sown by these states is not without consequence. As Rep. Lieu pointed out during the hearing held on September 18, 2025,[10] the imposition of even two different standards for AI training on labs would result in a compliance impossibility—the financial and computational resources associated with training frontier AI tools prohibits labs from training models pursuant to two standards, let alone 50. If the United States cannot remain on the frontier of AI, then it will find itself at an economic and national security disadvantage. Our adversaries will not pause their AI initiatives to allow us to learn from our regulatory mistakes.

Absent federal preemption, labs will not be able to make full use of the regulatory sandboxes available in a handful of innovative states. As has been documented in myriad contexts, technology companies often defer to the laws of the jurisdiction with the most onerous provision. Congress can pursue legislative action to remedy a world in which the innovative spirit of Utah and Texas, which have embraced a try-first mentality, is undermined by California and New York, which have legislated out of fear of speculative risks. One approach would be to create a legal safe harbor for AI labs that opt to participate in a federal regulatory sandbox or a state sandbox that meets certain qualifications. This safe harbor would prevent one or two states from blocking the full potential of existing sandboxes or a future federal sandbox to uncover the full extent of AI's benefits and risks.

Alternatively, Congress could move forward with a temporary moratorium on certain state laws, such that lab participation in existing state sandboxes would become feasible.

Both routes would not foreclose future regulation of AI. Instead, they would inform future AI discourse and ensure that biased scholarship and media do not cloud the judgment of legislators. A full understanding of the risks and benefits of AI cannot be achieved by attempting to put AI back in the bottle or significantly limit its utility by conditioning its use on manifold procedural checks. Notably, this is exactly the sort of evidence-based governance called for by the AI experts empaneled by Governor Gavin Newsom.[11]

Time is of the essence. The longer state AI laws remain on the books, the more likely it is that labs will alter their behavior to comply with those regulations[12]—perhaps forever altering the pace and direction of AI innovation.

---

[9] *See, e.g.,* Kevin Frazier, Regulatory Misalignment and the RAISE Act, Lawfare (June 3, 2025), https://www.lawfaremedia.org/article/regulatory-misalignment-and-the-raise-act.
[10] AI at a Crossroads: A Nationwide Strategy or Californication?: Hearing before the House Judiciary Subcommittee on Courts, Intellectual Property, Artificial Intelligence, and the Internet, 119 Cong. (2025) (Statement of Rep. Ted Lieu).
[11] Joint California Policy Working Group on AI Frontier Models, State of California (June 17, 2025), https://www.cafrontieraigov.org.
[12] *See* Slavisa Tasic, The Modern Growth of Government Springs More from Ideas Than from Vested Interests, 14 INDEPENDENT REV. 549, 551 (2010) ("Once a government program is in place, people exhibit a status quo bias and a tendency to follow default options rather than to consider alternatives.").

2. What would taking a light regulatory touch look like in your view, and what guiding principles should Congress keep in mind when determining whether to legislate and, where we decide to act, in crafting legislation?

Congress must take the lead on regulating AI. Our constitutional order was devised with the expectation that Congress would address matters of national significance.[13] Each day of inaction is tantamount to delegating AI governance to Sacramento and subjecting 300 million Americans to laws enacted without their consent.[14] What's more, the longer AI labs are forced to comply with laws based on speculative harms, they will fail to freely innovate and experiment—which, paradoxically, is likely the best approach to uncovering how to maximize the benefits of AI and minimize its risks.[15]

When Congress acts, it should adhere to three principles: experimentation, adoption, and information sharing. In practice, this looks like legislation that permits labs to deploy new models, rewards individuals and entities for integrating AI, and facilitates information sharing by both developers and deployers. These principles will facilitate a virtuous cycle of technological development that at once ensures the US continues to lead in AI innovation while not imposing undue risk on the public.

Experimentation

Experimentation by the labs is the key to discovering new technical strategies that mitigate some of the known flaws with AI. It is uncontested that AI is not perfect. But to foreclose deployment of AI because of some risks is shortsighted. Each of these systems facilitates proper experimentation: toleration of errors that are "small and diverse" and a commitment to learning from those errors.[16] Regulatory sandboxes, pilot programs, and legal safe harbors are the sorts of policies that should be at the top of Congress's agenda. Each of these regulatory frameworks clears legal thickets that may otherwise cause labs to delay deployment or to forgo experiments in novel fields. Critically, Congress must foreclose states from enforcing laws based on the idea cite the smallest possibility of

[13] *See, e.g.*, THE FEDERALIST NO. 48 ("The legislative department is everywhere extending the sphere of its activity[.]"); THE FEDERALIST NO. 22 (implying that assigning Congress the authority to regulate commerce would fix the patchwork approach that characterized the nation when under the Articles of Confederation); Daryl J. Levinson & Richard H. Pildes, *Separation of Parties, Not Powers*, 119 HARV. L. REV. 2311, 2338-39 (2006); Josh Chafetz, *Congress's Constitution*, 160 UNIV. PENN. L. REV. 715, *passim* (2012).
[14] *Cf.* Neil Chilson & Kevin Frazier, When Might State AI Laws Run Afoul of *Pike*?, Harv. J. L. & Tech. Online (Oct. 10, 2025), https://jolt.law.harvard.edu/digest/when-might-state-ai-laws-run-afoul-of-pike.
[15] Jessica Ji et al., AI Safety Evaluations: An Explainer, CSET (May 28, 2025), https://cset.georgetown.edu/article/ai-safety-evaluations-an-explainer/; *see* Ina Fried, OpenAI updates its system for evaluating AI risks, Axios (Apr. 15, 2025), https://www.axios.com/2025/04/15/openai-risks-frameworks-changes (explaining that models may not act the same way in pre-deployment evaluations and in the real world).
[16] Wildavsky, *supra* note 1, *passim*.

existential disaster as justification for extreme precaution,[17] while encouraging and even rewarding states that develop their own experiment-friendly regulatory paradigms.

Adoption

Advances in AI made possible by AI will mean little to the nation if the people fear AI and refuse to make use of it. Adoption of AI by the general public, especially small and medium-sized businesses,[18] is essential to avoid a nation of AI have and have nots, as well as to ensure we do not lag behind adversaries that have taken a far more aggressive posture toward adoption.[19] Technological progress generally leads to societal progress, but that relationship is contingent on mass adoption of the new technology. Decades of disparate access to electrification between more urban and rural states resulted in wide gulfs in economic opportunity and longevity.[20] A similar dynamic played out in the wake of globalization.[21] The current regulatory efforts being pursued by a handful of states suggests that we will again end up with a national bifurcated due to uneven impacts by emerging technology. An underappreciated fact is that the very states rushing to regulate AI are those with the highest rates of AI adoption.[22] They seem content to close the door to technological progress behind them.

This is precisely why the Congress must take an active role in facilitating AI adoption. It has always been the case that federal government has a duty to vigorously pursue a better future on behalf of the American people. The Founders explicitly abandoned the Articles of Confederation because they desired a central government capable of advancing the well-being of the entire country.[23] As I have documented elsewhere, on repeated occasions throughout the eighteenth century colonial governors and high-ranking officials acknowledged an obligation on the government to act in the interest of the general welfare, which referred to the welfare of society as a whole—not just the welfare of individuals in economically-dominant states.[24] In short, the government must "meet the needs of society by appropriately managing public resources and public affairs" if it is going to live up to its

[17] *See* Rich Ehisen, Senator Scott Wiener: California and AI, Capitol Weekly (July 8, 2025), https://capitolweekly.net/senator-scott-wiener-california-and-ai/ (interviewing Senator Scott Wiener).
[18] Sean Mitchell, Xero research shows small firms lag on AI, CFO Tech (Nov. 13, 2025), https://cfotech.co.uk/story/xero-research-shows-small-firms-lag-on-ai.
[19] Dashveenjit Kaur, China's generative AI user base doubles to 515 million in six months, AI News (Oct. 21, 2025), https://www.artificialintelligence-news.com/news/china-ai-adoption-doubles-515-million-users/.
[20] Kevin Frazier, How Co-Ops Electrified America, Asterisk (2025), https://asteriskmag.com/issues/11/how-co-ops-electrified-america [hereinafter, Frazier Co-Op].
[21] Scott Lincicome & Arjun Anand, The "China Shock" Demystified: Its Origins, Effects, and Lessons for Today, Cato (Dec. 12, 2023), https://www.cato.org/publications/china-shock.
[22] Anthropic Economic Index report: Uneven geographic and enterprise AI adoption, Anthropic (Sept. 15, 2025), https://www.anthropic.com/research/anthropic-economic-index-september-2025-report; Mark Muro & Shriya Methkupally, Mapping the AI economy: Which regions are ready for the next technology leap, Brookings (July 16, 2025), https://www.brookings.edu/articles/mapping-the-ai-economy-which-regions-are-ready-for-the-next-technology-leap/.
[23] *See* Kevin Frazier, *A Constitutional Mandate to Adopt AVs,* 82 WASH. & LEE L. REV. 272, 281 (2025) [hereinafter, Frazier AVs].
[24] *Id.* at 282.

lofty duties.[25] In the case of emerging technologies, such as AI, this duty includes spreading awareness and adoption of the beneficial uses of the technology—something private actors may not be sufficiently incentivized to do.[26]

Active promotion of AI adoption by Congress can ensure that Americans are in the best position possible to realize the American Dream in the Age of AI. It's worth recalling that much of the economic harm caused by globalization was the result of Americans being unable to move to opportunity.[27] AI literacy is and will be a major determinant of economic opportunity and mobility. Americans are counting on Congress to make sure they do not miss out on this technological wave because a handful of states would prefer to stymie AI innovation.

Several specific policies align with this adoption mandate. As I have called for on several occasions, Congress should form an Office of AI Adoption[28] tasked with a similar charge as the Rural Electrification Administration.[29] This Office would at once create a path for tech-savvy Americans to help their neighbors and local small business owners learn about AI and incorporate it into their professional and personal lives. An Adoption Corps would emulate the effectiveness of forward-deployed engineers—members of AI labs sent to assist their customers with integrating AI into their operations.[30] Like riding a bike, using AI is best done with training wheels—America will fall short of its potential with AI so long as it fails to heed this lesson.

Continuation of the status quo will result in AI abysses—regions in which AI talent and resources are lacking or even absent. Again, as we have seen in prior waves of technology, these sorts of regional disparities can have long-term negative economic and cultural consequences. Thus, the need for Congress, not states, to take primary regulatory authority over the direction, pace, and diffusion of AI innovation.

Information Sharing

As experimentation and adoption take hold, Congress should devise mechanisms for AI developers and deployers to share information about the political and community consequences of AI diffusion.

---

[25] *See* Brett Milano, 'Effectiveness in government is not something one can just assume', Harvard Law Today (Nov. 18, 2022) (quoting Vicki Jackson).
[26] Frazier AVs, *supra* note 23, at 299-300.
[27] Kristian Stout, AI's Labor Impact Will Emerge from the Institutions that Govern Its Use, Truth on the Market (Dec. 8, 2025), https://truthonthemarket.com/2025/12/08/ais-labor-impact-will-emerge-from-the-institutions-that-govern-its-use/.
[28] Kevin Frazier, The Main Street AI Tour, Appleseed AI (Nov. 10, 2025), https://appleseedai.substack.com/p/the-main-street-ai-tour.
[29] Frazier Co-Op, *supra* note 20.
[30] Joe Schmitdt, Trading Margin for Moat: Why the Forward Deployed Engineer Is the Hottest Job in Startups, Andreessen Horowitz (June 4, 2025), https://a16z.com/services-led-growth/.

Evidence-based policy requires evidence. As obvious as that may sound, it's a principle that is being ignored by many states rushing to regulate AI. Despite acknowledgements from leading AI experts[31] and policymakers[32] that they are unsure of the magnitude of AI risks as well as what qualifies as reasonable AI development practices, states have enacted laws with rigid definitions and thresholds that will likely not hold up well as AI advances in unpredictable fashions.

Congress can demonstrate what evidence-driven policy looks like in practice by mobilizing the Center for AI Standards and Innovation or another entity with commissioning and compiling reports on different AI initiatives. This information can then shape future regulatory efforts.

3. What should federal legislation on AI that includes a moratorium on patchwork state regulation entail?

Congress's priority should be to pass a moratorium on state laws that implicate AI development. Several existing state AI laws seem to run afoul of the Constitution's prohibition on extraterritorial laws.[33] What's more, as previously noted, the technical and financial aspects of training and evaluating AI prevent leading AI labs from complying with even two different standards imposed by states. The resulting patchwork is a drag on innovation—the longer it persists, the lower the odds of startups pushing the AI frontier and leading labs competing with international rivals.[34]

If Congress opts to supplement a moratorium with additional policies, then it should evaluate those additions through the aforementioned principles of experimentation, adoption, and information sharing. Sample policies that may align with those principles include whistleblower protections, the creation of a national data trust, regulatory sandboxes and safe harbors, and information sharing mechanisms. This is not an exhaustive list. Given that I have addressed the latter two policies in earlier answers, I will detail the former two here.

Whistleblower Protections

The emergent nature of AI development means that some AI researchers may detect problematic capabilities at inopportune times from a business perspective. Existing whistleblower protections tend to focus safeguards for employees who spot violations of the law.[35] Additional protections in

[31] *See* Joint California Policy Working Group on AI Frontier Models, State of California *passim* (June 17, 2025), https://www.cafrontieraigov.org.
[32] *See* Zoe Lofgren et al., Open Letter on SB 1047, Congress (Aug. 15, 2024), https://democrats-science.house.gov/imo/media/doc/2024-08-15%20to%20Gov%20Newsom_SB1047.pdf.
[33] *See* Chilson & Frazier, *supra* note 14.
[34] *See* Mark Dalton, Regulations Can Create the Monopolies They're Meant to Prevent, The Dispatch (Nov. 18, 2025), https://thedispatch.com/article/regulations-compliance-cost-barriers-to-entry/.
[35] *See* Why Whistleblowers Are Critical for AI Governance, The Future Society (July 24, 2025), https://thefuturesociety.org/ai-whistleblowers; Charlie Bullock & Mackenzie Arnold, Protecting AI Whistleblowers, Lawfare (June 25, 2025), https://www.lawfaremedia.org/article/protecting-ai-whistleblowers; Brad Carson & Kevin Frazier, Why We Need a 'Right to Warn' Now, The Well News (Aug. 4, 2025), https://www.thewellnews.com/opinions/why-we-need-a-right-to-warn-now/.

the AI space will facilitate the sort of information sharing called for above. Facilitation of this kind of exchange is essential given the national security and public health ramifications of unanticipated AI capabilities, as acknowledged in the AI Action Plan.[36] Of course, this protection should be crafted to reduce the odds of labs being stuck with employees hiding behind bogus concerns with an AI model. How best to design that framework is best left for a more detailed inquiry. Notably, states have already acted on this regulatory approach,[37] thus further necessitating a clear federal approach so that labs do not find themselves attempting to comply with conflicting or even contradictory regimes.

National Data Trust

Data is a core determinant of the capabilities and reliability of AI models.[38] And right now, the most abundant and easily accessible data in the U.S. is whatever can be scrapped from the Internet or purchased at increasingly high costs.[39] This is data is suited for a limited set of functions—so far, it has sustained models that are adept at meme generation and the creation of AI slop. Models deployed in sensitive settings, such as in health care and education, require data of higher quality.

But the datasets that could power AI toward public challenges—the carefully labeled, longitudinal records that sit inside federal agencies—are subject to myriad legal barriers.[40] Even when those datasets are anonymized or aggregated, they remain behind institutional firewalls, legacy systems, and risk-averse procedures that treat public data less like a national asset and more like buried treasure that officials are hesitant to make known and available.

This creates a paradox: AI will never meaningfully address public policy issues until data on those issues becomes readily available. An AI developer in Silicon Valley trying to maximize ad revenue has billions of data points ready to train a model. A research entity in Anywhere Else, America trying to detect early-onset learning disabilities or wildfire vulnerability must navigate a maze of procurement rules, privacy audits, FOIA exemptions, and multi-year agency delays.

The result is not a lack of innovation. It's a misdirection of innovation. Deployers will continue to build tools tailored to available data, which means they generally forgo the areas where AI could have the greatest positive social impact.

---

[36] The AI Action Plan, White House, https://www.whitehouse.gov/wp-content/uploads/2025/07/Americas-AI-Action-Plan.pdf (July 2025).
[37] Scott Singer & Alasdair Phillips-Robins, California Just Passed the First U.S. Frontier AI Law. Here's What It Does., Carnegie (Oct. 16, 2025), https://carnegieendowment.org/emissary/2025/10/california-sb-53-frontier-ai-law-what-it-does?lang=en.
[38] James Betker, The "it" in AI models is the dataset, Non_Interactive (June 10, 2023), https://nonint.com/2023/06/10/the-it-in-ai-models-is-the-dataset/.
[39] Kyle Wiggers, AI training data has a price tag that only Big Tech can afford, TechCrunch (June 1, 2024), https://techcrunch.com/2024/06/01/ai-training-data-has-a-price-tag-that-only-big-tech-can-afford/.
[40] Melissa Carson, Three steps to build a data foundation for federal AI innovation, Federal News Network (Dec. 5, 2025), https://federalnewsnetwork.com/commentary/2025/12/three-steps-to-build-a-data-foundation-for-federal-ai-innovation/.

The AI Action Plan acknowledges this lopsided ecosystem and calls for reshaping how government manages public data.[41] Among several other recommended policy actions, the Plan directs agencies to standardize how data is collected and stored as well as to facilitate access to certain federal datasets. These recommendations are positive developments. However, a much larger and permanent effort is necessary to truly realize the potential of federal data to foster AI innovation in the areas where people truly stand to benefit. What the country now needs—the reform that would turn the Plan's aspiration into action—is a Public Data Trust Mandate.

Such a mandate would do something deceptively simple but radically consequential: require federal agencies to curate, standardize, and securely share privacy-forward versions of their datasets in a federated architecture accessible to vetted innovators focused on predefined social challenges.

To make that proposal clearer: think of this federated system as transforming federal data from a set of scattered filing cabinets into a coherent public resource—one that agencies protect but also deploy—with the same seriousness and stewardship that past generations devoted to highways, electrification, or scientific research infrastructure. Building this new infrastructure will require a similar level of ambition, resources, and planning. The mechanics matter because messy institutional design is how good ideas die. A Public Data Trust would rely on three core components.

First**,** Congress would direct agencies with substantial public-relevant datasets—Health and Human Services, Education, Labor, Energy, Transportation—to carve out dedicated funding for data cleaning, labeling, and the creation of secure computational environments.

Second**,** the system would operate on a federated-access model.[42] Without getting too into the technical weeds, this means innovators never download the raw data. Instead, they bring their algorithms into the government's secure environment, train the models there, and leave only with the trained weights. This is the digital equivalent of letting someone cook in your kitchen without letting them leave with your groceries.

Third**,** access would be conditional. Applicants would need to demonstrate a clear public-interest objective: improving overdose detection, reducing prison recidivism, strengthening grid resilience, or whatever priorities Congress specifies or delegates to participating agencies to determine. In exchange for access, they would grant the federal government a non-exclusive right of use for the resulting tool and commit to transparent performance reporting.

* * *

Thank you again for the opportunity to assist with this important inquiry. If you have any additional questions, please let me know.

---

[41] The AI Action Plan, White House, https://www.whitehouse.gov/wp-content/uploads/2025/07/Americas-AI-Action-Plan.pdf (July 2025).

[42] Burnie Legette, How Federated Learning Helps Agencies Build Smarter AI Models Securely, FedTech Magazine (May 2, 2025), https://fedtechmagazine.com/article/2025/05/how-federated-learning-helps-agencies-build-smarter-ai-models-securely.