Comments of the Taxpayers Protection Alliance

Re: Request for Information on the Development of an Artificial Intelligence Action Plan

Docket ID: NSF_FRDOC_0001

David Williams, President
Dan Savickas, VP of Policy & Government Affairs
Juan Londoño, Chief Regulatory Analyst
David McGarry, Research Director

This document is approved for public dissemination. The document contains no business-proprietary or confidential information. Document contents may be reused by the government in developing the AI Action Plan and associated documents without attribution.

Table of Contents

Preamble

The Taxpayers Protection Alliance is a non-partisan, non-profit organization dedicated to educating the public through the research, analysis and dissemination of information on the government's effects on the economy. As such, we are pleased to offer comments to the National Science Foundation (NSF) on this important request for information about the future of artificial intelligence (AI) policy in the United States.

First Principles of Effective AI Regulation

There are several basic principles to which policymakers must adhere when thinking through proposals to regulate AI. The root problem is that too many—including elected officials, bureaucrats, and industry players—are inclined to abandon proven bedrock tenets of regulation generally, and of technology regulation specifically, when dealing with AI.[1] In other words, they treat AI as if it is so different—on a categorical level—from what has come before, such that its regulation must be approached in an entirely new way. This perception is in no small part the product of fearmongering by elected officials and industry leaders.[2] However, abandoning regulatory approaches that have been proven by time and experience bodes ill for innovation. New technologies and changing circumstances require fresh thinking, but ignoring the lessons of the past completely will not

produce sound policy. Innovation is not guaranteed; it can be slowed—or even halted—by ill-conceived regulation.

First and foremost, AI regulation must be light touch, allowing technologists to strike out boldly and innovate. The task of regulators is to set broad guidelines to protect individuals from externalities or nefarious conduct. It is not – nor should it ever be – to micromanage or direct the course of AI development. In general, innovation must be permissionless—not subjected to a "mother, may I?" regulatory regime or swaddled in constricting red tape.

Light-touch regulation has thus far prevailed in American tech policy, and it has fostered the most productive sector in the world. The U.S. has 20 of the world's 25 highest-value companies.[3] This includes the five largest, as well as all but one of the world's seven trillion-dollar companies. All are tech companies: Apple, Nvidia, Microsoft, Alphabet, Amazon, and Meta (in order of market cap).

Europe, which has taken a far more heavy-handed approach to tech regulation, provides a foil to America's successes. The continent has a small, insignificant tech sector, with hardly any companies of note.[4] This is no coincidence. Smothering an industry in regulation will make it less dynamic, less productive, and less innovative, putting it at risk of being outpaced by foreign competition. The plain fact is that Europeanizing American tech regulation will inevitably also Europeanize the American tech sector.

After enacting destructive regulations such as the General Data Protection Regulation and the Digital Markets Act, the European Union (EU) enacted the AI Act. This new legislation created a hyper-complex regulatory environment and makes compliance incredibly difficult, especially for small companies.[5] This sort of regulatory regime also tends to leave incredible latitude and discretion to enforcers, increasing the uncertainty industry faces and raising the chances of arbitrary enforcement.

The U.S. must reject this approach. Treating AI first and foremost as a threat—not as a tool that promises tremendous economic advancements—can only harm American business and, by extension, consumers.

Unfortunately, America has already begun to imitate European policies. In July 2024, American officials joined EU and U.K. authorities to publish a joint statement on AI policy.[6] The joint statement takes a hostile position against incumbent Big Tech firms that are investing in AI, worrying that companies such as Google might become too dominant in this new sector. This fear ignores both the intense competition that currently permeates the AI space as well as the wild successes of non-incumbent firms, such as OpenAI.[7] If competition-policy enforcers choose to muck around with the workings of markets, innovation will suffer.

America's emulation of European tech policy should stop with the joint statement. In fact, the Trump administration should make a U-turn and ensure that the U.S. abandons it.

Who Should Be Regulating AI?

AI has positioned itself as one of the most promising emerging technologies. It has opened the doors to tremendous amounts of growth and prosperity, fueled by automation and productivity gains. However, for all of AI's potential, it has also sparked fears of potential bias and abuse that has led to a massive wave of proposals to regulate it—from state as well as federal lawmakers.[8]

In most circumstances, states can function well as "laboratories for democracy." State and local governments can experiment with different laws and policies to discover what the best approach is. This dynamic, however, has not translated well with digital services. Oftentimes, it is nearly impossible to determine "who" the user is, where he is, or which local regulations apply. The nature of these digital services is one that usually transcends geographical lines, with service providers (or their infrastructure) located in one state and the end user(s) located in one (or multiple) other jurisdictions. Additionally, users can mask their location when they use a virtual private network (VPN), which indicates to the service provider that the end user is somewhere other than where he or she is physically located.

These dynamics add a tremendous amount of complexity for state-level regulatory compliance. The experiment of regulating data privacy at the state level has already shown that compliance can be quite costly. Studies estimate that the privacy "patchwork" has cost the American economy up to $112 billion per year.[9] These costs have led to calls for a federal privacy standard to preempt these state-level regulations and provide industry with a single regulatory framework. This would lower compliance costs and reduce regulatory risk. Unfortunately, it seems various states have decided to repeat this approach with AI governance. California and Colorado passed comprehensive AI bills of their own (though the California bill was vetoed by its governor), and hundreds of lower-scale AI bills have been introduced across the country.[10]

Ultimately, a state-level push for AI regulation would not only recreate the dreaded "patchwork" of AI regulation that has already failed for privacy, but it would also recreate the European approach to regulation. That is not a winning strategy for what seems to be one of the most important technological battlegrounds in decades. While the United States was able to easily emerge as the global leader in the first digital revolution, it now faces increased competition from foreign competitors that aim to claim the top spot in a potential second digital revolution.

Why AI Regulation Shouldn't Be Offloaded to Bureaucrats

There has been a similarly troublesome AI governance issue brewing at the federal level. On one side, there has been an uptick in regulatory actions by federal agencies, usually through obscure rulemakings that do not garner the same attention of a congressionally passed statute.[11] This process of executive mission creep gained significant momentum in the wake of President Biden's executive order on AI in October 2023. For example, consider the AI Safety Institute and the National Artificial Intelligence Research Resource (NAIRR).[12] On the other hand, there has been no shortage of AI-related proposals in Congress, though no major legislation has been passed.[13]

This has created tension over who is better posed to shepherd in the regulatory environment for AI. Congress should be wary of agencies creeping in to enact AI

regulations that should rightfully go through the legislative process. Unlike the executive-branch rulemaking process, the legislative process necessarily includes the input of many officials, elected by the people, with widely diverging interests. This kind of deliberation usually leads to more prudent bills, better adapted to the needs of the nation as a whole, than policy making that is fast-tracked at a much faster pace by individual bureaucrats or small bureaucratic teams.

This does not mean, however, that Congress cannot rely on subject matter experts in the executive branch. Understanding AI—let alone regulating it—requires much time and study. Subject-matter experts will necessarily play a big part in ensuring laws pertaining to AI fit neatly within the realities of the technology. Nonetheless, this should not be taken too far. It would be deeply unwise to offload policy making entirely to experts in administrative agencies and to insulate those experts from congressional oversight.[14]

It should also be noted that it is still very early to advocate for comprehensive AI regulation, as so many lawmakers have done. Regulators should let technologies mature—and formal and informal rules to adapt to new technologies—and only address issues that have been proven unsolvable without government action. However, whenever regulation is deemed necessary, it should be enacted by Congress, not by unaccountable bureaucrats at the executive branch. Particularly after the repeal of the *Chevron* deference, Congress needs to reclaim its legislative authority and prevent agencies from overextending their mandate to introduce new AI regulations.

AI Regulation Should Be Narrowly Targeted, Not General

Policymakers should focus on regulating various AI outputs, not the technology as a whole. AI features of some sort will likely be incorporated in commercial and productive goods of many types. As a general-purpose technology, any attempt to impose sweeping restrictions on AI models will likely yield stifling, overbroad regulations. Any attempts to specifically regulate AI products at large would resemble trying to regulate all products that include steel at one go. It would be neither practicable nor conducive to innovation.

The solution is to address the specific issues raised by specific types of AI products on an individual basis. This approach will allow policy to be tailored to remedy discrete problems that are produced by discrete types of tools. However, before breaking new ground, policymakers should first attempt to extend the principles of existing law to cover AI's outputs. Of course, some tweaks may be required to account for AI's novelties, but the same principles, *mutatis mutandis*, should remain. For example, the Federal Trade Commission should extend its existing authority to take action against AI-generated or AI-assisted fraud, adhering as closely as possible to its existing procedures for combatting other types of fraud. To take another example, intellectual property (IP) rights should be enforced against AI models under the same standard as they would against humans. Training models on copyrighted works should be considered fair use (just as a human reading and learning from online material would qualify as fair use), but abject plagiarism should not.

Moreover, as regulators consider how to craft guidelines for AI tools, they should focus on the real-world outcomes produced by AI tools, not hypothetical harms or the minutiae of algorithms. As Daniel Castro, the director of the Center for Data Innovation, and vice president, at the Information Technology and Innovation Foundation, put it, "regulate performance, not process."[15] Castro argues:

> To address concerns about safety, efficacy, and bias, regulators should regulate the performance of AI systems rather than create prescriptive rules about specific processes and methods firms must follow. Establishing performance-based metrics for AI systems gives consumers, businesses, and government an opportunity to better compare performance across different systems, as well as set minimum performance requirements. For example, rather than create stringent compliance-based rules, such as requiring lenders to use diverse datasets to train their credit scoring models, regulators should create performance-based rules, such as requiring lenders to validate that their credit scoring models accurately assess risk across all protected classes of individuals. Allowing firms to identify the best way to achieve the desired goal gives them the flexibility necessary to comply most efficiently. Moreover, performance-based regulations can ensure firms meet the desired goals rather than simply check the box on a list of compliance measures.

It should not matter how AI products are designed so long as they do not harm consumers in ways that violate existing legal principles. As Castro points out, it makes little difference precisely *how* an AI credit-scoring tool functions so long as it doesn't deceive consumers (a principle that applies equally to non-AI tools as well). This can also work the other way around. A tool that has done all the "correct" processes but, has somehow led to proven harm should face scrutiny. A process-first approach misses the point of, and adds confusion to, regulation.

If regulators take a prescriptive approach with respect to how AI companies must build their products, they will limit the creativity of innovators. This will inevitably lead to fewer breakthroughs and advances. It should be remembered that the job of a regulator is not to fiddle with business models or technological designs. It is to protect consumers from coercive force and fraudulent misrepresentations while allowing private actors to innovate freely and provide useful goods and services. It is the role of an umpire, not of an active player on the field.

Getting the regulatory approach right will allow agencies with subject matter expertise to examine cases involving AI and make reasonable decisions based on existing legal principles, tailored to the needs of specific sectors. It will also create a regulatory environment in which all products and services—AI-related or otherwise—can compete on fair grounds and entrepreneurs remain free to experiment and innovate. It will also avoid the implementation of overbroad and over-ambitious regulatory schemes, which will threaten innovation.

Why AI Regulators Should Keep on the Straight and Narrow

Policymakers should also be wary of inserting their own ideological priorities into AI regulation. For example, former President Biden's EO contained overtures to "small businesses," "collective bargaining," and more.[16] The White House pledged to address "algorithmic discrimination," referencing the so-called "disparate impact" standard. It was governed by the notion that discrimination should be identified by disparate outcomes between demographic groups—irrespective of confounding factors, such as economic disparities—not by discriminatory intent or conduct.

It's all well and good for politicians to have preferences about such things, but AI regulation should not be used as an instrument of social policy. Instead, policymakers should keep technology regulation focused on regulating technology and protecting users from fraud and other well-established harms.

It's worth underlining that the disparate impact standard should be exiled from the thinking of policymakers altogether—in AI policy and elsewhere. There are countless variables that cause disparities between demographic groups in the U.S., such as socioeconomic status, region or location, and personal preference. Without proof of discriminatory intent, a system that treats any unequal outcomes as inherently suspect will unfairly target many legitimate products and services with spurious accusations of discrimination. Adopting a disparate impact standard will stifle innovation, as litigation-averse companies sacrifice creativity to ensure they remain well within in the good graces of the state.

Regulators should also be wary of attempts at regulatory capture, which are common during the early stages of technological development. Perhaps unsurprisingly, companies that have risen as early leaders in the AI industry have lobbied for regulation. These leaders know that they—not their nascent competition—will be able to overcome the hefty compliance burdens. Thus, they have taken the opportunity to attempt to shape regulation to favor their own products at their competitors' disadvantage.

Policymakers must resist these self-interested calls and create a level regulatory playing field. This will allow competition and markets to select and promote the best ideas and products. Oftentimes, economies of scale mean that larger companies will be better able to serve consumers, but this does not mean that regulation should favor one business size or model over another. Markets and consumers should be left to choose what products provide the most value.

When regulators try to centralize control of AI development—using tools ranging from hard regulation to guidance, or even to officially sanctioned datasets—they run the risk of discouraging entrepreneurs from thinking freely and pursuing new

modes of innovation. What's more, if regulators miscalculate or otherwise err, their mistakes will resound across the entire industry.

It should be also remembered that regulatory attempts at central planning usually fail due to circumstances regulators cannot understand. Especially in novel, hyper-complex, and fast-moving industries like AI, no individual official or agency can possibly know all the factors that are relevant to policy making. Nor can they anticipate all the consequences of a given program or regulation. This is known commonly as the "knowledge problem."[17]

Kristian Stout, director of innovation policy at the International Center for Law & Economics, touched on related issues in a recent discussion of the NAIRR.[18] Stout wrote:

> There's also the potential problem of bias and standardization [at the NAIRR]. While centralized datasets aim to provide consistent and reliable training data, they may unintentionally reinforce existing biases in AI models if the data is not sufficiently diverse. If NAIRR prioritizes certain types of data over others, this could lead to models that fail to perform adequately across different demographics or use cases, perpetuating bias rather than mitigating it.
>
> Moreover, dependency on centralized datasets could create a lock-in effect, where developers become overly reliant on the provided data, stifling experimentation and innovation. Instead of fostering a diverse ecosystem of AI research and development, it could result in homogenization, where every AI model is trained on the same narrow datasets, limiting breakthroughs that might arise from alternative approaches.

Although some policymakers may forget it, the knowledge problem plagues tech regulation just as much as any other type of regulation. In fact, given the complexities at play and the lightning speed of innovation, it's likely that epistemological shortcomings constrain tech regulation most of all.

How AI Can Make Government Better

AI promises to bring new efficiencies, cut costs and waste, improve outputs, and make human workers more productive. Policymakers should seek out ways in

which AI tools can enhance government operations and curtail waste, fraud, and abuse. According to the Government Accountability Office:

> Estimating the amount of fraud and improper payments in federal programs has been challenging due to data and other limitations. However, the federal government could be losing between $233 billion and $521 billion annually to fraud. Additionally, federal agencies reported an estimated $236 billion in improper payments in FY 2023, and cumulative federal improper payment estimates have totaled about $2.7 trillion since FY 2003.[19]

AI solutions can likely provide more insight into these issues—among many others—and help curb needless spending. Procurement, outlays, and other processes could begin to incorporate AI tools to save taxpayers money. While it may require upfront spending to research and purchase new tools, it will save money in the long term.

Moreover, the government's cybersecurity teams should examine how AI can harden the nation's cyber defenses. Cybercrime and cyber-espionage are now an ever-present challenge. AI is playing a growing role in the field. To stay secure against geopolitical rivals and private-sector criminals, the U.S. should sprint to remain at the head of the race on AI and cybersecurity. As government agencies start to incorporate AI technology into their daily processes, they must be cognizant of the operational risks that stem from deploying these systems.

Unfortunately, government has a muddy track record in cybersecurity practices, which means that the data collection, processing, and storage that is necessary for deploying AI systems could make these agencies a target for cyberattacks.[20] A commitment to improve cybersecurity standards across the federal government should be a prerequisite for the incorporation of AI tools by government agencies. Higher government investment in studying vulnerabilities in government digital services, updating obsolete IT systems, and procuring cybersecurity hardware and software tools should be key priorities. As of now, the federal government lacks the necessary infrastructure to prevent and mitigate any unintended harms that could stem from deploying AI technology—but this could, and should, change.

The Development of AI Will Require a Paradigm Shift in Energy and Permitting

AI is a power-hungry technology.[21] Collecting, processing, analyzing, and transforming data requires tremendous amount of computing power. The advanced hardware involved in AI operations requires constant and reliable access to energy to function. To compete with a globally competitive AI industry, American policymakers need to think of computing power—and the factors that enable it—as a strategic resource. If America is to compete in the AI industry, it must produce and distribute as much energy as possible to keep up with an ever-rising demand.

As AI is rising, multiple other consumer products and appliances are becoming increasingly electrified or smart; AI is being incorporated into many as well. As a result, the American power grid is already under significant stress.[22] If not addressed, the growth and scalability of AI will be capped either due to a sheer shortage or to prohibitively expensive energy costs.

The private sector has taken the initiative to overcome these energy challenges. For example, Amazon, Apple, Google, and Microsoft have invested billions of dollars in either reactivating or building new nuclear power plants.[23] Some have also invested in the development of self-powered data centers, oftentimes using renewable energy systems.[24]

However, these investments often struggle with America's complex regulatory regime. Amazon's plans to ramp up nuclear power capacity were thwarted by the Federal Energy Regulatory Commission. Meanwhile, Meta was blocked from building a new nuclear-powered data center due to a sighting of a rare bee species.[25] These are just the hurdles that arise during the planning phase. Once ground is broken and projects kick off, companies will encounter a tremendous amount of red tape that hinders their attempts to develop new power capacity. As with most major building projects, the current permitting regime has hampered American firms' capacity to quickly build up necessary infrastructure.[26] For example, the average time to build a new transmission line in the United States is around ten years, with some taking longer than 20 years.[27] These permitting woes

translate into cost increases, because continuous delays force companies to pay more interest over time or pay workers' wages while no construction is being done.

The regulatory barriers for AI companies extend beyond energy access. Permitting and zoning laws have also blocked the construction of essential infrastructure, such as data centers.[28] As the technology scales up, it will need more data centers to provide models with the necessary computing power to collect, process, and transform data. However, the proliferation of anti–data center bills—which prevent the construction or operation of these centers—could hamper the country's ability to build up more computing power.

An AI system with no computing power is like a car with no fuel. The system might have top-notch engineers and sophisticated algorithms, but without access to reliable energy supply and infrastructure to process data, it would fall far short of achieving its true potential. Policymakers need to think of computing power as another strategic resource, like oil or precious minerals—especially as foreign competition in the AI space increases. In an AI-driven economic revolution, the countries that can scale up their computing power efficiently are the ones more likely to come up on top. Policymakers should prioritize creating a regulatory environment that allows American AI companies to scale computing capability in a timely and cost-efficient manner.

Conclusion

Digital technologies—AI foremost among them—promise to revolutionize economies They can bring untold amounts of information to the masses, connect the unconnected, and open new frontiers of human understanding. These incredible benefits come with risks—as all good things do. However, the former clearly outweigh the latter. Policymakers must remember this. Every intervention comes with a tradeoff, and usually that tradeoff is, on net, undesirable. Wrapping an iron regulatory fist around AI innovation can stifle a generation's worth of good ideas. Doing so will reverberate throughout the American economy, affecting the untold numbers of businesses and consumers that would have benefitted from new—and potentially life-changing—AI developments.

America has already shown that smart, light-touch, and targeted tech regulation will lead to global technological leadership. Confronting AI, all America must do is follow its own example.

---

[1] https://www.protectingtaxpayers.org/wp-content/uploads/Regulating-AI-TPA.pdf
[2] https://apnews.com/article/biden-ai-artificial-intelligence-executive-order-cb86162000d894f238f28ac029005059 ; https://www.nbcnews.com/politics/congress/big-tech-ceos-ai-meeting-senators-musk-zuckerberg-rcna104738
[3] https://pbs.twimg.com/media/GbDoAt0XkAA2DgE?format=jpg&name=large
[4] https://townhall.com/columnists/david-b-mcgarry/2023/08/18/europes-investigation-of-microsoft-continues-its-anti-tech-trajectory-n2627193
[5] https://www.nber.org/system/files/working_papers/w30028/w30028.pdf ; https://www.realclearmarkets.com/articles/2023/06/13/european_regulators_embark_on_yet_another_tech_policy_blunder_940157.html ; https://www.siliconcontinent.com/p/the-strange-kafka-world-of-the-eu
[6] https://competition-policy.ec.europa.eu/about/news/joint-statement-competition-generative-ai-foundation-models-and-ai-products-2024-07-23_en
[7] https://www.realclearmarkets.com/articles/2024/08/22/global_regulators_want_to_hamstring_ai_innovation_1053070.html
[8] https://www.multistate.ai/artificial-intelligence-ai-legislation
[9] https://itif.org/publications/2022/01/24/looming-cost-patchwork-state-privacy-laws/
[10] https://www.rstreet.org/commentary/california-and-other-states-threaten-to-derail-the-ai-revolution/ ; https://www.mayerbrown.com/en/insights/publications/2024/06/colorado-governor-signs-comprehensive-ai-bill ; https://www.npr.org/2024/09/20/nx-s1-5119792/newsom-ai-bill-california-sb1047-tech
[11] https://www.piratewires.com/p/america-is-sleepwalking-into-a-permanent-dei-bureaucracy-regulating-ai
[12] https://townhall.com/columnists/david-b-mcgarry/2024/11/20/schumers-misguided-plan-to-hand-ai-regulation-to-the-administrative-state-n2647990
[13] https://www.cato.org/briefing-paper/artificial-intelligence-regulation-threatens-free-expression#legislative-approaches
[14] https://townhall.com/columnists/david-b-mcgarry/2024/11/20/schumers-misguided-plan-to-hand-ai-regulation-to-the-administrative-state-n2647990
[15] https://www2.datainnovation.org/2023-ten-principles-ai-regulation.pdf
[16] 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/
[17] https://www.econlib.org/library/Essays/hykKnw.html
[18] https://truthonthemarket.com/2024/11/07/the-ai-legislative-puzzle/
[19] https://www.gao.gov/fraud-and-improper-payments
[20] https://www.gao.gov/blog/what-are-biggest-challenges-federal-cybersecurity-high-risk-update
[21] https://www.washingtonpost.com/business/2024/03/07/ai-data-centers-power/
[22] https://www.newsmax.com/politics/computing-degrowth-power/2024/12/09/id/1190987/

[23] https://www.nytimes.com/2024/10/16/business/energy-environment/amazon-google-microsoft-nuclear-energy.html
[24] https://www.cnbc.com/2024/05/01/microsoft-brookfield-to-develop-more-than-10point5-gigawatts-of-renewable-energy.html
[25] https://broadbandbreakfast.com/amazons-bid-for-nuclear-energy-rejected-by-regulators/ ; https://www.yahoo.com/tech/metas-nuclear-power-plans-were-194801911.html
[26] https://spectator.org/us-permitting-regime-is-hampering-americas-potential/
[27] https://ifp.org/future-of-ai-compute/#challenges-to-building-in-america
[28] https://spectator.org/dont-regulate-data-centers-out-of-existence/