Center for American Progress
1333 H Street NW, Suite 100E
Washington, DC 20005
202.682.1611
americanprogress.org

June 12, 2023

The Honorable Alan Davidson
Assistant Secretary of Commerce for Communications and Information
National Telecommunications and Information Administration
1401 Constitution Ave., NW
Washington, DC 20230

Re: Request for comment, NTIA–2023–0005

Dear Assistant Secretary Davidson,

We applaud the National Telecommunications and Information Administration (NTIA) and the U.S. Department of Commerce for undertaking this Request for Comment (RFC) around artificial intelligence (AI) accountability at a critical time.

The issue of developing accountability systems for AI is of critical importance. If the United States cannot create an effective accountability system for AI and automated systems that addresses the range of highly consequential issues—especially those involving democratic process, discrimination and impact on marginalized communities, and human autonomy—then the danger from these systems is greater than their potential benefit. No amount of proposed innovation is worth the systemic threat to or destruction of any of those critical pillars of American and human society. Without proper accountability mechanisms, the risks are too severe, and AI should simply not exist in an unchecked environment.

The Center for American Progress (CAP) appreciates the invitation to comment on this project and respectfully offers several ideas for your consideration.

Sincerely,

Adam Conner,
aconner@americanprogress.org
Vice President, Technology Policy

Megan Shahi
Director, Technology Policy

Ashleigh Maciolek
Research Associate, Structural Reform and Governance

Ben Olinsky,
Senior Vice President, Structural Reform and Governance

Executive Summary

The importance of the United States developing AI accountability mechanisms to ensure the development of trustworthy AI may be the most critical technological policy challenge of our time. Unlike previous dramatic technological shifts, those developing advanced AI, along with numerous experts, have themselves warned about its potential negative impact on many aspects of society, of the need for it to be regulated, and even broader systemic risk to humanity.[1]

As CAP has previously written about AI:[2]

> "AI tools have the potential to bring tremendous benefits to our society. Yet the risks of AI are also profound—both by creating entirely new classes of problems and exacerbating existing ones."

Question 4 of the RFC asks perhaps the most important question, "Can AI accountability mechanisms effectively deal with systemic and/or collective risks of harm, for example, with respect to worker and workplace health and safety, the health and safety of marginalized communities, the democratic process, human autonomy, or emergent risks?"[3] The potential risks of AI across the range of those critical areas is well documented and requires a robust set of accountability mechanisms.[4]

---

[1] Jamie Condliffe, "Big tech says it wants government to regulate AI. Here's why," *Protocol,* February 12, 2020, available at https://www.protocol.com/ai-amazon-microsoft-ibm-regulation#toggle-gdpr; John Simons, "The Creator of ChatGPT Thinks AI Should Be Regulated," *TIME*, February 5, 2023, available at https://time.com/6252404/mira-murati-chatgpt-openai-interview/; "Pause Giant AI Experiments: An Open Letter," *Future of Life Institute*, March 22, 2023, available at https://futureoflife.org/open-letter/pause-giant-ai-experiments/; James Vincent, "Top AI researchers and CEOs warn against 'risk of extinction' in 22-word statement," *The Verge*, May 30, 2023, available at https://www.theverge.com/2023/5/30/23742005/ai-risk-warning-22-word-statement-google-deepmind-openai.

[2] Adam Conner, "The Needed Executive Actions to Address the Challenges of Artificial Intelligence," *Center for American Progress,* April 25, 2023, available at https://www.americanprogress.org/article/the-needed-executive-actions-to-address-the-challenges-of-artificial-intelligence/.

[3] National Telecommunications and Information Administration, "AI Accountability Policy Request for Comment."

[4] Rick Claypool and Cheyenne Hunt, ""Sorry in Advance!" Rapid Rush to Deploy Generative A.I. Risks a Wide Array of Automated Harms" (Public Citizen, 2023), available at

If the United States cannot create an effective accountability system for AI and automated systems that addresses the range of highly consequential issues outlined in that question— especially those involving democratic process, discrimination and impact on marginalized communities, and human autonomy—then the danger from these systems is greater than their potential benefit. No amount of proposed innovation is worth the systemic threat to or destruction of any of those critical pillars of American and human society. Without proper accountability mechanisms, the risks are too severe, and AI should simply not exist in an unchecked environment.

The NTIA AI RFC focuses heavily on "AI accountability mechanisms such as certifications, audits, and assessments" throughout the proposal.[5] While certifications, audits, and assessments (and the transparency and access required to enable them) are likely to play a key role in initial AI accountability measures, they alone are insufficient to ensure the development of trustworthy AI.

As the recent AI Now 2023 landscape report notes "Audits and data-access proposals should not be the primary policy response to harmful AI. These approaches fail to confront the power imbalances between Big Tech and the public, and risk further entrenching power in the tech industry." It continues by noting "both technical and socio-technical audits place the primary burden for algorithmic accountability on those with the fewest resources."[6]

Fundamentally, direct government regulation will be needed to ensure the development and deployment of trustworthy AI. This includes a strong need for a federal privacy law and new tech antitrust laws.[7] But even if all of these new laws were to pass there would still be huge gaps in the government's ability to address the substantial harms that are already being felt from advanced technologies, and especially AI.

Accountability mechanisms of any ilk must apply to both first-party and third-party use of AI systems. Many of the advanced AI models being developed right now are being

---

https://www.citizen.org/article/sorry-in-advance-generative-ai-artificial-intellligence-chatgpt-report/; Amba Kak and Sarah Myers West, "2023 Landscape: Confronting Tech Power" (AI Now, 2023), available at https://ainowinstitute.org/wp-content/uploads/2023/04/AI-Now-2023-Landscape-Report-FINAL.pdf; Grant Fergusson and others, "Generating Harms: Generative AI's Impact & Paths Forward," *Electronic Privacy Information Center*, May 2023, available at https://epic.org/wp-content/uploads/2023/05/EPIC-Generative-AI-White-Paper-May2023.pdf.

[5] National Telecommunications and Information Administration, "AI Accountability Policy Request for Comment."

[6] Amba Kak and Sarah Myers West, "2023 Landscape: Confronting Tech Power."

[7] Erin Simpson and Adam Conner," How To Regulate Tech: A Technology Policy Framework for Online Services" (Washington: Center for American Progress, 2021), available at https://www.americanprogress.org/article/how-to-regulate-tech-a-technology-policy-framework-for-online-services/.

deployed as platforms with API access for integration into existing or new consumer and commercial software applications, with little thought given to the responsibilities and liabilities for what the AI Bill of Rights calls the "Designers, developers, and deployers" of automated systems.[8] Clarification and standards for both first-party and third-party use of AI systems is essential to AI accountability.

A strong AI accountability framework must empower the U.S. government to create and enforce new rules around AI, such as to designate high-risk cases and sectors that in some cases should go through a government review before deployment, prevent national security threats, outright prohibit certain dangerous uses, and establish broad principle-based rules to ensure safe and effective systems and prevent algorithmic discrimination.

The federal government can immediately begin to prepare an all-of-government response to the challenges of AI. Most notably, the president can issue an executive order that applies clear rules to agencies and contractors on the use of AI and empowers agencies to use all existing statutory authorities to govern the use of AI by regulated entities.[9] The executive branch must also engage with Congress on crafting new, robust legislation for AI.

The importance of the United States developing an effective accountability system for AI to ensure the development and deployment of trustworthy AI cannot be overstated. The only way to ensure that AI innovation protects fundamental aspects of society is by ensuring strong accountability mechanisms enforced by the government.

Answers to RFC Questions

1. What is the purpose of AI accountability mechanisms such as certifications, audits, and assessments?

AI accountability mechanisms are essential to try to ensure that automated systems are trustworthy before they are deployed and integrated across society. The American public stands to benefit from the use of AI and related technological advancements–including from increased efficiency. But there are also significant risks associated with this technology and it is necessary to create trustworthy systems that minimize harm and maximize benefits. As we discuss below, two particularly worrisome examples of such risk are the amplification of biases baked into training data sets and the fueling of disinformation that can threaten basic pillars of our democracy. To ensure trustworthy

[8] The White House Office of Science and Technology Policy, "Blueprint for an AI Bill of Rights," (Washington: The White House, 2022), available at https://www.whitehouse.gov/ostp/ai-bill-of-rights/.
[9] Conner, "The Needed Executive Actions to Address the Challenges of Artificial Intelligence."; Adam Conner, "White House Must Take More Action To Address AI Concerns," *Center for American Progress,* May 4, 2023, available at https://www.americanprogress.org/article/white-house-must-take-more-action-to-address-ai-concerns.

systems that minimize harms, accountability mechanisms are essential and can include certifications, audits, and assessments, but these alone are insufficient. To achieve trustworthy AI there must also be accountability measures that introduce new rules for automated systems, especially for higher-risk use cases, as well as clear use limitations or prohibitions. Additionally, privacy and antitrust laws are needed to ensure that companies building automated systems do not violate the privacy of the American public in the process and do not engage in unfair business practices. Without proper accountability mechanisms, the risks are too severe, and AI should simply not exist in an unchecked environment.

Not all risks associated with AI will be foreseeable and accountability mechanisms may sometimes fail. However, continued assessment and certification processes can offer a checkpoint to review the systems and continue to achieve trustworthy systems. An additional concern when auditing or assessing AI is trying to account for the intended purposes of that automated system and the associated harms it may pose to the end user of the system, but also for the harms that could likely happen if the automated system was used for nefarious purposes both every day and extraordinary.

Urgent action is needed here. As suggested above, automated systems are already causing harms and more advanced AI is likely to exacerbate many of the harms that already exist in society today.[10] Accountability mechanisms are needed to minimize this additional source of risk and protect the American people–particularly for already marginalized communities. To that end, AI accountability mechanisms should specifically include in their scope considerations of discrimination and bias. Bias can be baked into the large datasets that are used as inputs into automated systems and accountability mechanisms, to the extent possible, should assess the quality of the data–including sources–to ensure that the outputs of the automated system do not reproduce this bias. It will also be necessary to assess the outputs and the decisions that AI makes to try to eliminate remaining discrimination and biases. Evidence of this adverse effect of AI has already started to appear: automated systems have discriminated against people of color in home loan pricing, recruiting and hiring automated systems have shown a bias towards male applicants, AI used in making health care decisions have shown a racial bias that ultimately afforded white patients more care, among other examples.[11] Assessing automated systems before they are

[10] The White House Office of Science and Technology Policy, "Blueprint for an AI Bill of Rights."; Claypool and Hunt, ""Sorry in Advance!" Rapid Rush to Deploy Generative A.I. Risks a Wide Array of Automated Harms"; Amba Kak and Sarah Myers West, "2023 Landscape: Confronting Tech Power."; Grant Fergusson and others, "Generating Harms: Generative AI's Impact & Paths Forward."

[11] Patrick Sisson, " Housing discrimination goes high tech," *Curbed,* December 17, 2019, available at https://archive.curbed.com/2019/12/17/21026311/mortgage-apartment-housing-algorithm-discrimination; Jeffrey Dastin, "Amazon scraps secret AI recruiting tool that showed bias against women," Reuters, October 10, 2018, available at https://www.reuters.com/article/us-amazon-com-jobs-automation-insight-idUSKCN1MK08G; Starre Vartan, "Racial Bias Found in a Major Health Care Risk Algorithm," *Scientific* American, October 24, 2019, available at https://www.scientificamerican.com/article/racial-bias-found-in-a-major-health-care-risk-algorithm/.

applied to important aspects of everyday life–health, housing, employment, etc.–is crucial to ensuring that existing biases do not get reproduced by new technology and further harms already marginalized communities.

Accountability mechanisms should also include in their scope how AI is likely to exacerbate the existing problem of disinformation. Augmented media that is produced by an automated system in a matter of seconds can fuel disinformation that is already present in our current information systems, but on a much larger, more sophisticated scale.[12] There are already prominent examples of AI-produced media that is being used in campaigns and contributes to the existing ecosystem of disinformation and propaganda and its corrosive effects on our democracy.[13] This follows years of election disinformation that has proliferated on social media platforms and is often created and spread by adversarial actors–both foreign and domestic.[14] AI will only make disinformation a more severe threat, both in the scale at which it is produced and its difficulty to detect. A recent report by the Information Integrity Research and Development Interagency Working Group developed a roadmap on information integrity and this group should play an important part in establishing a plan for judging threats to the information integrity environment.[15] Swift and coordinated action will be needed to assess this threat of automated systems.

Moreover, accountability mechanisms should include privacy considerations in their scope to assess whether the data used to inform automated systems ever violates the privacy of the American people. The information used in large language models is scraped from the internet and the models are highly vulnerable to privacy breaches.[16] Applications of automated systems–for example, facial recognition software–can also pose new privacy violations and create undue harm, especially for communities of color.[17] Without a federal privacy law to protect the American public from these violations, auditing AI to account for these concerns is an important but not sufficient

---

[12] David Klepper and Ali Swenson, " AI-generated disinformation poses threat of misleading voters in 2024 election," *PBS News Hour,* May 14, 2023, available at https://www.pbs.org/newshour/politics/ai-generated-disinformation-poses-threat-of-misleading-voters-in-2024-election.

[13] FOX 11 Los Angeles, "RNC slams Biden reelection bid with AI generated ad," Youtube, April 25, 2023, available at https://www.youtube.com/watch?v=wbWLEUiVHVQ.

[14] Isabelle Niu, Kassie Bracken, and Alexandra Eaton, "Russia Created an Election Disinformation Playbook. Here's How Americans Evolved It," *New York Times,* October 25, 2020, available at https://www.nytimes.com/2020/10/25/video/russia-us-election-disinformation.html.

[15] The White House Office of Science and Technology Policy and others, "Roadmap for Researchers on Priorities Related to Information Integrity Research and Development" (Washington: 2022), available at https://www.whitehouse.gov/wp-content/uploads/2022/12/Roadmap-Information-Integrity-RD-2022.pdf.

[16] Nicholas Carlini and others, "Extracting Training Data from Large Language Models," *USENIX Security Symposium* 6 (2021), available at https://arxiv.org/pdf/2012.07805.pdf.

[17] Nicol Turner Lee and Caitlin Chin, "Police surveillance and facial recognition: Why data privacy is imperative for communities of color," *Brookings Institute,* April 12, 2022, available at https://www.brookings.edu/research/police-surveillance-and-facial-recognition-why-data-privacy-is-an-imperative-for-communities-of-color/.

part of ensuring that privacy is upheld and automated systems do not pose new and severe violations.

Lastly, as the technology becomes increasingly advanced, it stands to introduce existential risks to our society. A recent statement, signed by more than 350 executives, researchers and engineers working in AI, said that AI poses the risk of extinction and should be treated as a priority on the same scale as pandemics and nuclear war.[18] This top-level threat is warranted, especially if AI were to end up in the hands of bad actors and pose grave threats to national security, personal privacy, health and safety, and much more. Addressing this potential risk is important, and a role primarily to be driven by the U.S. government in conjunction with other governments around the world. Yet, the potential for existential risk is on a much longer time horizon and robust accountability mechanisms that can be introduced now can help mitigate these future risks.

All the above concerns are essential to include in the scope of accountability mechanisms as they all stand to pose significant risks to the American public. To ensure the greatest level of protection, there should also be legal standards and enforceable risk thresholds. Technology companies should not build, release, or use automated systems that do not meet certain safety standards or violate existing laws, particularly concerning discrimination. Congress, the courts, and other rulemaking bodies should work in tandem to mandate these standards and prioritize the safety of the American public. However, in the absence of legal standards, voluntary accountability mechanisms–including certifications, audits, and assessments–can still play an important role as we wait for regulation to catch up. They can offer insight into how the systems are built and utilized, as well as identify the pockets of greatest concern. They also provide a useful starting place for regulation to capitalize on and allow for more immediate action.

The lack of any accountability mechanisms for AI would mean that these all too familiar harms will continue to spread, but at a much larger, more sophisticated scale, and will disproportionately impact communities that are already marginalized. If automated systems create more harms than they do benefits–particularly harm around core aspects of our society, including the democratic process, human autonomy, employment, and more–then they should not be publicly released. No amount of innovation and technological advancement is worth dismantling the most core tenets of American society. If these risks are not dealt with–via stringent accountability requirements or legal standards–then AI systems should not be allowed to exist.

[18] Center for AI Safety, "Statement on AI Risk," available at https://www.safe.ai/statement-on-ai-risk#open-letter.

2. Is the value of certifications, audits, and assessments mostly to promote trust for external stakeholders or is it to change internal processes? How might the answer influence policy design?

Accountability mechanisms such as certifications, audits, and assessments can serve multiple purposes throughout the development and iteration of trustworthy AI, but must be leveraged with integrity, clear intent, and consistency. They can promote trust with external stakeholders by offering a degree of transparency that otherwise would not be available and the requirement to conduct them can be a forcing function for AI developers to develop and design their product responsibly, which bolsters the internal processes required to build these systems. A requirement to adhere to specific accountability mechanisms can bring more accountability, transparency, and trust for all stakeholders, both internal and external, involved in the development, deployment, and use of AI systems.

In an ideal scenario, these developers are aware of this at the earliest stages of model training and fold these considerations into the design of the AI systems. If done properly, requiring certifications, audits, and assessments of AI developers can influence how these systems are built and iterated on from their origination, rather than as an imposition on them after the fact. To promote the most influence on how the systems are designed, including the policies that govern them, the accountability mechanisms should include as much detail on exact requirements, stipulations, and timelines. Examples may include robust internal documentation, testing periods, and model cards outlining inputs to the machine learning systems. This will lead to the standardization of these processes. In the case of a certification process to establish that an AI product is considered trustworthy, certain requirements to build it in a particularly safe and transparent manner in order to receive the certification will inherently lead to better outcomes and more transparency in the processes. This is therefore cyclical in nature, whereby it will influence how the developers of AI design the products themselves in addition to the policies that govern them. Their goal is to streamline these efforts and reduce overhead as much as possible, so an effort that promotes trust externally, outlines how they should write systemic policies, and furthers their own product strategy will be seen as favorable.

Adhering to accountability schemes within technology companies often requires work that is cross-functional in nature, thereby requiring a myriad of teams and functions to prioritize this work. These requirements help create positive outcomes for internal teams that are supporting the product development process, including engineering, data science, policy, legal, and more. While having these teams take on these initiatives can often trade off against business interests like growth, engagement, and profitability, it ultimately leads to safer and more transparent products and promotes longevity for company success. Accountability mechanisms can ultimately foster stronger product development internally, safer use externally, and greater trust from users of these products. At Instagram, for example, externally publishing the policies

that govern recommendation surfaces[19] required deep collaboration, prioritization, and due diligence across numerous functions inside the company. The effect of this was that in the aftermath of publishing, companies were much more likely to be accountable to users when making changes to the policies, products, and protocols. A similar effect can be expected to occur because of accountability mechanisms for AI systems.

4. Can AI accountability mechanisms effectively deal with systemic and/or collective risks of harm, for example, with respect to worker and workplace health and safety, the health and safety of marginalized communities, the democratic process, human autonomy, or emergent risks?

Today existing AI accountability mechanisms are wholly insufficient to deal with systemic and/or collective risks of harm from AI and automated systems as they relate to workers' rights, workplace safety, the health and safety of marginalized communities, the safety of marginalized communities, the democratic process, human autonomy, or emergent risks.[20] What exists now is a sparse patchwork of voluntary measures proposed and implemented by industry[21] and enforcement of existing laws[22] that will cover some use cases but will not cover all of the emergent risks. For example, while existing civil rights law may cover fair and equal access to housing or credit, there is no clear statute to protect the democratic process or human autonomy from AI, a risk that numerous experts agree needs to be a significant consideration (but not the sole consideration). Furthermore, enforcement of existing laws may be complicated and more difficult than in traditional contexts because of insufficient understanding or awareness of how the tools operate.

This is why it is essential for the United States, in coordination with global allies where possible, to create an effective and multi-layered AI accountability system, to ensure the development and deployment of AI systems do not have unchecked negative consequences. This should be coupled with legislation on AI to properly ensure trustworthy systems.

---

[19] Instagram, "Recommendations on Instagram," available at https://help.instagram.com/313829416281232.

[20] Claypool and Hunt, ""Sorry in Advance!" Rapid Rush to Deploy Generative A.I. Risks a Wide Array of Automated Harms."

[21] Microsoft, "Responsible AI," available at https://www.microsoft.com/en-us/ai/responsible-ai?activetab=pivot1:primaryr6; OpenAI, "Developing safe & responsible AI," available at https://openai.com/safety (last accessed June 2023); Google AI, "Our Principles," available at https://ai.google/responsibility/principles/ (last accessed June 2023); Anthropic, "Core Views on AI Safety: When, Why, What, and How," March 8, 2023, available at https://www.anthropic.com/index/core-views-on-ai-safety.

[22] Rohit Chopra and others, "JOINT STATEMENT ON ENFORCEMENT EFFORTS AGAINST DISCRIMINATION AND BIAS IN AUTOMATED SYSTEMS," U.S. Federal Trade Commission, available at https://www.ftc.gov/system/files/ftc_gov/pdf/EEOC-CRT-FTC-CFPB-AI-Joint-Statement%28final%29.pdf.

The Biden-Harris Administration's Blueprint for an AI Bill of Rights[23] is an attempt to enshrine early timeless principles around the rights of people in a world predicted to increasingly be driven by advanced automated systems. These principles cover the right to notice and explanation, the right to data privacy, among others, and are best equipped to address the range of the challenges brought about by the rapid advancement of AI.

Should the AI Bill of Rights be codified into law, those principles can remain central while technology evolves and agencies and courts strive to apply that interpretation, much like the Constitution's Bill of Rights it is inspired by a document that should be noted also remains in full force in the age of AI. An AI Bill of Rights should be utilized in conjunction with more detailed voluntary frameworks like the NIST AI Risk Management Framework[24] and new or updated laws to address liability, require pre-approval for deploying higher risk applications, prohibitions for unacceptable risk uses, and to prevent and respond to threats of existential risk.

If American cannot effectively design a system to deal with the systemic and collective risks of AI, then those kinds of AI systems should simply not be commercially deployed.

7. Are there ways in which accountability mechanisms are unlikely to further, and might even frustrate, the development of trustworthy AI? Are there accountability mechanisms that unduly impact AI innovation and the competitiveness of U.S. developers?

There is a commonly-held value within technology companies that moving fast, often at the expense of other values, is critical to success. In tech's boom times, many sought to "move fast and break things," a motto Mark Zuckerberg popularized at Facebook and blossomed into a growth-at-all-costs ethos that spread throughout Silicon Valley.[25] Although this mantra is no longer publicly parroted as widely, the value still underpins the general approach many firms continue to use. Notably, OpenAI, a relatively new entrant to the sector, has emulated this approach, raced to introduce ChatGPT and other generative AI tools even as they loudly warn of the need for regulation due to the potential harms from such products.

In their race-to-be-first model, tech companies often see efforts to regulate or prioritize safety early in a product lifecycle as directly at odds with goals to rapidly develop and deploy new technology. Companies are eager to "win" the race to launch

[23] The White House Office of Science and Technology Policy, "Blueprint for an AI Bill of Rights."
[24] National Institute of Standards and Technology, "AI Risk Management Framework," available at https://www.nist.gov/itl/ai-risk-management-framework.
[25] Lisa Bonos, "They built the digital world. Now they just want to sew and make chairs," *Washington Post,* May 27, 2023, available at https://www.washingtonpost.com/technology/2023/05/27/tech-workers-new-hobby-woodworking-sewing/.

the product to capture an early user base, and see trust and safety obligations as roadblocks to rapid deployment.

As such, companies will oppose meaningful accountability mechanisms with an argument that the overhead required to do so trades off against their capacity to innovate and quickly provide users of their services and products with critical, life-changing technology. However, there is much more nuance to this trade-off, which need not be a zero-sum game. The reality is that being held accountable will in fact likely slow the pace of how companies are able to launch new products in the short term, but the longer-term arc of growth and development is unlikely to be affected. Furthermore, there are few examples of tech companies in recent years that have recognized the harms of new products and voluntarily mitigated them, so it is unlikely that AI will be different.

The negative impacts of accountability requirements can be minimized if they are clear, distributed over a product's lifecycle, and predictable. Conversely, requirements that come as one-off requests and then involve large amounts of work (such as massive disclosures), will be worse for innovation.

Additionally, requirements that build on existing industry standards and practices may be advantageous. For example, requirements could build on how companies are already doing quality assurance or scrubbing for personally identifiable information (PII). It is worth the time and effort to conduct ongoing engagements with relevant firms to deeply understand their roadmaps, goals, and processes and allow them to be part of the conversation on what accountability mechanisms should be in place. This will prevent an adversarial approach and ensure that they are seen as and feel like partners in the effort to make this technology safe and appropriate for all users. Clear accountability expectations allow for internal processes to be built at the onset of AI development, which validates both internal and external trust and leads to safer, higher integrity products in the long term.

15. The AI value or supply chain is complex, often involving open source and proprietary products and downstream applications that are quite different from what AI system developers may initially have contemplated. Moreover, training data for AI systems may be acquired from multiple sources, including from the customer using the technology. Problems in AI systems may arise downstream at the deployment or customization stage or upstream during model development and data training.

It is critical to understand and develop AI accountability measures for all forms of business relationships in AI.

Almost all the attention in the AI accountability space has been focused on first-party development and first-party use of AI, for example with Open AI's Chat GPT4. But

while OpenAI runs a public version of ChatGPT4, it has also opened the GPT4 model up to third-party application developers to utilize its foundational technology and to integrate it into their own applications. Many AI developers are mirroring this API strategy, including Google Cloud and Amazon AWS.[26] The vendor/customer relationship terminology is technically correct here but does not sufficiently encompass the technology transfer and lack of potential AI accountability from both first and third parties.

However, the resources dedicated to AI accountability for third-party usages of AI models ranges from very limited to non-existent. For example, OpenAI dedicates 687 words to "Safety Best Practices"[27] on its developer website to comply with its numerous prohibited content under its content policies[28] along with a black box moderation API.[29]

The lack of traditional tools to address challenges from AI is clear as there is little transparency for the first party to third party usage of tools, no way to verify usage within the existing terms, and little way for third parties to influence the first party foundation model.

The lack of resources and guidance provided to third parties is indicative of the general attitude that third-party usage is out of the control of the first-party developer and solely the responsibility of a third-party that may have little ability to influence the outputs from the first-party model. This also creates the potential for responsibility to be constantly shifted between first and third parties, between the model creator and the developer. Clarifying the appropriate liability for AI will play a key role in ensuring that first- and third-party AI developers take seriously their responsibilities for developing and ensuring AI accountability.

17. How should AI accountability measures be scoped (whether voluntary or mandatory) depending on the risk of the technology and/or of the deployment context? If so, how should risk be calculated and by whom?

To properly scope AI accountability mechanisms, there should be a governmental agency that is dedicated to continuously assessing and defining risks of automated

---

[26] Sundar Pichai, "An important next step on our AI journey," Google, February 6, 2023, available at https://blog.google/technology/ai/bard-google-ai-search-updates/;Swami Sivasubramanian, "Announcing New Tools for Building with Generative AI on AWS," *Amazon*, April 13, 2023, available at https://aws.amazon.com/blogs/machine-learning/announcing-new-tools-for-building-with-generative-ai-on-aws/.

[27] Open AI, "Moderation: Quickstart," available at https://platform.openai.com/docs/guides/moderation/quickstart (last accessed June 2023).

[28] Open AI, "Content Policy," available at https://labs.openai.com/policies/content-policy (last accessed June 2023).

[29] Open AI, "Moderation: Overview," available at https://platform.openai.com/docs/guides/moderation/overview (last accessed June 2023).

systems. Establishing a body to continuously assess and define risk would be more effective than any piecemeal attempts across multiple agencies to define and regulate high risk harms, and both mandatory and voluntary accountability mechanisms can be informed from the guidelines established by this body. In establishing guidance and defining risks, the body should consider what level of risk can be caused using AI–for example, they should define what is considered "high risk," "moderate risk" or "low risk" or even "unacceptable risk." To determine what category of risks are associated with automated systems, the body could consider both the degree of the potential harm and the scale of damage that it could potentially cause. Highly widespread risks–for example, risks to national security–should be considered "high risks" and subsequently, accountability mechanisms should be more stringent. Similarly, risks that would have grave consequences for individuals should be treated with more severity and the proposed accountability mechanisms should reflect that. For example, risks that would have an effect on human autonomy or physical and mental wellbeing should be considered high risk. There should also be a category that includes "unacceptable risk" with prohibitions or bans as needed.

It may also be useful for the body tasked with defining risk to consider the risks of AI across some entire sectors–including, health, employment, housing, criminal justice and more. It may be difficult to foresee all the use cases of AI and their related harms, but their application in certain sectors can be categorized as "high risk" while "low risk" in other sectors.

20. What sorts of records (*e.g.,* logs, versions, model selection, data selection) and other documentation should developers and deployers of AI systems keep in order to support AI accountability? How long should this documentation be retained? Are there design principles (including technical design) for AI systems that would foster accountability-by-design?

Developers and deployers of AI systems should be held responsible for keeping and maintaining appropriate records about their models, including on training and usage. This logging is critical to supporting AI accountability and ensuring that AI systems are developed and deployed in a trustworthy manner. This principle extends beyond just the use of AI and applies more broadly to all automated technology that general consumers use or are subjected to in any way. The AI Now institute underscores the criticality of a mandate to require sufficient documentation to serve as evidence that developers of these models are held accountable for data and design choices.[30] Below are a few principles on data handling that should be taken into account when an AI developer is determining which information to retain and how to retain it:

[30] Amba Kak and Sarah Myers West, "AI Now 2023 Landscape: Confronting Tech Power."

- **Data storage:** All data around usage, training, development, deployment, and iteration should be logged in a manner that can be queried and is easily accessible by those authorized to access it.
- **Privacy:** Information on specific users of the AI systems as well as any associated Personally Identifiable Information (PII) should be appropriately handled and scrubbed before any data disclosures are made.
- **Data retention:** Any of the aforementioned data shall be retained for a period of time that authorities deem is necessary to ensure accountability of the AI system, for example 3-5 years. Once that time period is up, developers should responsibly delete this data.
- **Data minimization:** Create bright-line rules that limit firms' ability to collect data on consumers or produce data about them.[31]
- **Data portability:** Users of the AI systems should be able to access and download all data the developer has stored on them, including how they have used the system and what information of theirs was used to train it, in an easy to read format.

For any algorithmic model, whether generative AI or otherwise, it is critical that its developer can understand and explain in plain English how a model is trained and what its inputs are. This is an ongoing issue for numerous big tech companies with complex recommender and advertiser systems, and in order to effectively regulate trustworthy AI it must be addressed at the onset. However, companies must record and maintain data logs in a way that allows for these otherwise "black box" models to be accessed and scrutinized by authorities.

23. How should AI accountability "products" (*e.g.,* audit results) be communicated to different stakeholders? Should there be standardized reporting within a sector and/or across sectors? How should the translational work of communicating AI accountability results to affected people and communities be done and supported?

For the outcomes, results, and products of accountability mechanisms to achieve their desired effects, it is critical that they are communicated to stakeholders in ways they can understand and digest. These stakeholders will naturally span a variety of types, including government (NTIA for example), civil society, the general public (including consumers of these technologies), and even other private companies that may conduct the review to assess qualifications and fitness to develop AI. Across each of these, there may be nuance required to get the appropriate information communicated effectively, however calling for bespoke reporting will greatly increase overhead and reduce efficacy of said reporting. As a result, it is recommended that AI developers, governments, and civil society organizations standardize their reporting to stakeholders to whatever degrees are possible, bearing in mind considerations of

[31] Ibid.

privacy and confidentiality. For example, there may be an in-depth audit analysis that is required annually to uphold a specific trustworthy AI developer certification. The full report, including non-specific data on usage, patterns, risks, etc. can be disclosed to a smaller number of stakeholders that are authorized to handle that information appropriately and safely. Then, a pared down version can be made available to the public without requiring excess additional work on the part of the developer. An analog to this may be the European Union's Digital Services Act (DSA) use of mandatory algorithmic transparency requirements to drive accountability. Under the DSA, regulated platforms won't be able to turn a blind eye to AI-amplified harms and are required to enact "reasonable, proportionate and effective mitigation measures" for identified risks, along with their reporting and mitigation plans subject to independent audit and oversight by the European Commission.[32]

The translational work of communicating AI accountability results to affected people and communities should be the responsibility (including financial responsibility) of the developers of the automated system. It should be done with global communities and user bases in mind and should be housed in an easily accessible and permanent repository that consists of all reports, assessments, and audit results. Examples of this include Meta's Transparency Center or Google's Transparency Report hub.[33] Blog posts, transparency reports, tooltips within the AI system and links to reports and assessments in sign up flows are all ways to increase traffic to these outputs. It is critical that outputs are not considered one-time efforts to promote transparency and communicate findings. Rather, these should be seen as living, breathing documents that are updated regularly, with past versions stored for reference and to measure progress.

24. What are the most significant barriers to effective AI accountability in the private sector, including barriers to independent AI audits, whether cooperative or adversarial? What are the best strategies and interventions to overcome these barriers?

The most significant barriers to effective AI accountability in the private sector relate to a lack of authority, oversight, and a framework for demanding accountability of AI developers. Despite a commonly vocalized narrative around innovation tradeoffs, lack of resources, or other reasons companies are unable or unwilling to take on these mechanisms, it ultimately comes down to a requirement to prioritize them "to keep the lights on." Companies do not want their innovation or competitiveness to be stifled or slowed down in any capacity, so without having their hands forced on the

[32] Natasha Lomas, "Europe Names 19 Platforms That Must Report Algorithmic Risks Under DSA," TechCrunch, April 25, 2023, available at https://techcrunch.com/2023/04/25/europe-names-19-platforms-that-must-report-algorithmic-risks-under-dsa/.

[33] Meta, "Transparency Center," available at https://transparency.fb.com/,(last accessed June 2023) and Google, "Google Transparency Report," available at https://transparencyreport.google.com/?hl=en, (last accessed June 2023).

matter, they are unlikely to take on voluntary commitments to uphold accountability mechanisms. Notably, even if internal stakeholders want and can prioritize building accountability mechanisms, they need significant help from other internal functions to deploy them. There is also little to keep companies committed to their own, self-stated responsible AI commitments.

These barriers can be mitigated by ensuring there are consequences to not creating safe and effective AI systems and that the financial incentives of the AI developer are at risk in cases of non-compliance. Technology regulation around the world, including the EU's DSA, can serve as a model for how to design these accountability mechanisms to yield consequences (financial and reputational, among others) when not followed. Ultimately, private companies are most concerned with profitability and rapid deployment of their products to as many people as possible. In order to get private companies to conduct these assessments and audits, mechanisms must directly impact what developers care about most and be aligned with the for-profit incentives driving their rapid technological development. For these reasons, voluntary measures are insufficient. Government action (such as formal rulemaking, executive orders, and new laws) are clearly needed; we cannot allow the Age of AI to be another age of self-regulation.[34]

25. Is the lack of a general federal data protection or privacy law a barrier to effective AI accountability?

Yes, the lack of a general federal data protection or privacy law is a significant barrier to effective AI accountability.

Data privacy was one of the five principles outlined in the Biden-Harris Administration's Blueprint for an AI Bill of Rights[35] and the AI Now Institute 2023 landscape report emphasized that, "Data policy is AI policy."[36] The use of large amounts of data is critical in the development of advanced AI, especially foundation models, which results in the dual abuse of those with large stockpiles of data being able to use it for their developmental advantage, and those without large stockpiles of data rushing to acquire said data.

Importantly, a data protection or privacy law is critical not just to AI accountability but to broader issues impacting all online services.[37] CAP has been a strong supporter of a federal privacy law, as we wrote in 2021, "Amid growing demand for government action to address online harms and increasing regulatory action abroad, the United

[34] Conner, "The Needed Executive Actions to Address the Challenges of Artificial Intelligence."
[35] The White House Office of Science and Technology Policy, "Blueprint for an AI Bill of Rights: Data Privacy," available at https://www.whitehouse.gov/ostp/ai-bill-of-rights/data-privacy-2/ (last accessed June 2023).
[36] Lomas, "Europe Names 19 Platforms That Must Report Algorithmic Risks Under DSA."
[37] Simpson and Conner, "How to Regulate Tech: A Technology Policy Framework for Online Services."

States must urgently pursue aggressive antitrust action, updated competition policies, and robust federal privacy laws and rules."[38]

Unfortunately, despite recent progress in the 117th Congress, meaningful privacy legislation has yet to advance in Congress. In 2022, the Federal Trade Commission (FTC) began to undertake a rulemaking on Commercial Surveillance and Data Security[39] that will provide some needed clarity and guidance in areas that will impact AI as well as online services generally, especially in the absence of a federal privacy law.

The lack of a federal data protection or privacy law is unfortunate and complicates efforts around AI accountability. Yet this means that the federal government should be utilizing all available measures on privacy in the meantime while Congress should continue to prioritize a federal privacy bill.

26. Is the lack of a federal law focused on AI systems a barrier to effective AI accountability?

The lack of a federal law focused on AI systems will quickly become a barrier to effective AI accountability as these tools get widely adopted, but that does not mean that the federal government is powerless to act to address harms from AI. Federal regulators have already declared that new technologies do not negate existing laws, as FTC Chair Khan stated recently, "There is no AI exemption to the laws on the books"[40] which she reiterated in her recent NYT op-ed on how the FTC will regulate AI.[41] Similarly, the joint statement from the DOJ, FTC, CFPB, and EEOC was a clear statement that new technologies do not allow for the breaking of existing laws.[42] It should be clear to all government agencies at the federal, state, and local level that there is no need to wait for any new federal AI laws before enforcing the laws on the books.

Some have noted that passing new federal privacy and antitrust laws will have important impacts on AI accountability. This is true, and CAP has long supported a federal privacy law and new tech antitrust laws. As noted in CAP's answer to question 25 regarding a federal privacy law, "the United States must urgently pursue aggressive

---

[38] Ibid.

[39] Federal Trade Commission, "Trade Regulation Rule on Commercial Surveillance and Data Security," August 22, 2022 available at https://www.federalregister.gov/documents/2022/08/22/2022-17752/trade-regulation-rule-on-commercial-surveillance-and-data-security.

[40] Federal Trade Commission, "FTC Chair Khan and Officials from DOJ, CFPB and EEOC Release Joint Statement on AI," Press release, April 25, 2023, available at https://www.ftc.gov/news-events/news/press-releases/2023/04/ftc-chair-khan-officials-doj-cfpb-eeoc-release-joint-statement-ai.

[41] Lina Khan, "We Must Regulate A.I. Here's How," NY Times, May 3, 2023, available at https://www.nytimes.com/2023/05/03/opinion/ai-lina-khan-ftc-technology.html.

[42] Federal Trade Commission, "Joint Statement on Enforcement Efforts Against Discrimination and Bias in Automated Systems," Press release, April 25, 2023, available at https://www.ftc.gov/system/files/ftc_gov/pdf/EEOC-CRT-FTC-CFPB-AI-Joint-Statement%28final%29.pdf.

antitrust action, updated competition policies, and robust federal privacy laws and rules."[43]

But even if new federal privacy and antitrust laws were to be passed into law, large gaps would still remain to address issues related to online services generally and artificial intelligence specifically. From CAP's 2021 report "How to Regulate Tech" noted that:[44]

> Existing laws, authorities, and agencies can address a subset of interlocking online services harms outlined above. In particular, the Center for American Progress strongly supports more aggressive antitrust action, more robust competition policies, increased privacy and civil rights capacity at the FTC, and strong federal privacy legislation or rules.…Looking ahead, however, even an optimistic reading of these proposed updates shows gaps would persist in the government's ability to tackle the vast scope of online services harms in a timely and effective manner. As outlined below, existing systems of regulatory oversight are primarily reactive: Judicial scrutiny and dedicated, but often narrower, piecemeal legislation have struggled to keep pace with technological and market change. In a vacuum of regulatory scrutiny, consumer harms have accumulated, predatory practices have become industry standards, and dominant players have entrenched and expanded their holdings. Over time, a regulatory "debt" has built up where existing statutes and sector-specific regulations have not been sufficiently updated or applied to novel problems. Labor laws, for example, have lagged behind developments in algorithmic workplace management systems. Effective regulatory oversight must grapple with not only emerging issues but also the regulatory debt that has developed over past decades.[45]

Even if the US passed a federal privacy law and new antitrust laws, robust protections related to AI systems will need to be provided by Congress in the near future.

---

[43] Simpson and Conner, "How to Regulate Tech: A Technology Policy Framework for Online Services."
[44] Ibid.
[45] Marc Jarsulic and others, "Reviving Antitrust: Why Our Economy Needs a Progressive Competition Policy" (Washington: Center for American Progress, 2016), available at https://www.americanprogress.org/issues/economy/reports/2016/06/29/140613/reviving-antitrust/; Marc Jarsulic, Ethan Gurwitz, and Andrew Schwartz, "Toward a Robust Competition Policy" (Washington: Center for American Progress, 2019), available at https://www.americanprogress.org/issues/economy/reports/2019/04/03/467613/toward-robust-competition-policy/; Lawyers' Committee for Civil Rights Under Law, "Federal Trade Commission Must Protect Civil Rights, Privacy in Online Commerce," Press release, August 4, 2021, available at https://www.lawyerscommittee.org/federal-trade-commission-must-protect-civil-rights-privacy-in-online-commerce/; Consumer Reports and others, "Letter in support of increased funding for the FTC to protect data privacy," September 23, 2021, available at https://advocacy.consumerreports.org/wp-content/uploads/2021/09/Group-letter-in-support-of-FTC-privacy-funding.pdf.

30. What role should government policy have, if any, in the AI accountability ecosystem?

As CAP wrote in April about the quickening public release of new AI technologies, "There is also a sense of déjà vu of the advent of social media. Once again, we are poised to rapidly introduce a new technology to a society unprepared for its attendant consequences and without an adequate comprehensive response from government. Workers, families, and our democracy are poised to suffer consequences if we do not act now. We cannot allow the Age of AI to be another of age of self-regulation."[46]

Government policy is the essential component to AI accountability, including the enforcement of existing laws, the need for non-AI laws that will impact AI including new privacy and technology antitrust laws, and need for additional rulemaking and enforcement powers to address the challenges of online services and AI as outlined in the 2021 CAP report "How to Regulate Tech" and the Biden-Harris Administration's 2022 Blueprint for an AI Bill of Rights.[47] In noting the importance of addressing gaps with new rulemaking and enforcement:

> Even in best-case scenarios for critical competition and privacy updates, significant gaps would remain in the U.S. government's ability to anticipate and remedy online services harms. To effectively govern online services, U.S. regulators need to be empowered with proactive rule-making abilities that can curb problems before or as they occur. Such proactive rule-making powers—sometimes called "ex ante" regulation—are distinct from reactive or "ex post" approaches, which are litigated after harms have occurred. Proactive rule-making could identify and prohibit harmful measures prior to significant harm or as harms are occurring. In other words, this report proposes complementing after-the-fact antitrust enforcement by adding new restrictions and regulations that help prevent harm across multiple areas.[48]

The 2021 CAP report calls for "Principles for online services rules" and proposed several potential categories including anti-competitive practices; violations of civil rights; insecure and data-extractive practices; unfair, deceptive, abusive acts or practices for consumer and business users.[49] Similarly, the Blueprint for an AI Bill of Rights outlines five proposed principles: safe and effective systems, algorithmic discrimination protections, data privacy, notice and explanation, and human alternatives, consideration, and fallback which would make essential building blocks for principles in AI accountability law or regulation.[50]

---

[46] Conner, "The Needed Executive Actions to Address the Challenges of Artificial Intelligence."
[47] Simpson and Conner, "How to Regulate Tech: A Technology Policy Framework for Online Services."; The White House Office of Science and Technology Policy, "Blueprint for an AI Bill of Right."
[48] Simpson and Conner, "How to Regulate Tech: A Technology Policy Framework for Online Services."
[49] Ibid.
[50] The White House Office of Science and Technology Policy, "Blueprint for an AI Bill of Right."

CAP has proposed principles-based rulemaking for online services, including AI, envisioning that:

> Congress would both define specific practices that would be explicitly outlawed and enumerate broader principles around which regulators could interpret and craft rules…The combination of clear guardrails and the flexibility of a principles-based approach offers flexibility to address future problems and mitigation of any industry capture of the regulator…Congress would describe these general categories of prohibited behavior—in addition to specific practices defined in statute to be unlawful for online services—and regulators would continue their work by developing and applying rules to specific technologies, practices, or markets, appropriately considering the requisite factors named as process requirements as new issues arise.[51]

Administration of a new government AI accountability ecosystem could take several forms, all of which should include new rulemaking, enforcement, and expertise mandates. As CAP has previously written, this could be done by expanded existing agencies and their mandates or the creation of a new agency noting:

> Expansion of existing agencies and consideration of new agencies should both be on the table. In either case, these proposals require significant expansion of the U.S. government's capacity and expertise. Given the complexity of some online services—many of which deal in technical fields relating to software engineering, machine learning, or algorithmic design—and their direct impact on Americans' access to opportunity, specialist regulators with appropriate sociotechnical expertise are required. The federal government must design a creative system that recruits needed expertise while sufficiently insulating agencies from industry capture. Such capacity will aid in making technologies more legible to the public, taking the air out of any unrealistic industry exaggerations of technical complexity and challenging unfounded objections to sensible regulation. Developing effective regulation will require wholesale rejection of the discriminatory industry dynamics—particularly around racial and gender-based discrimination—that are encoded and amplified throughout technologies, services, and products today. Any additional responsibilities should be complementary and additive to existing DOJ, FTC, and FCC authorities, as well as sector-specific laws in other areas.[52]

Finally, the United States government has a key role to play in leading on AI accountability beginning with implementing various AI accountability methods the US

[51] Simpson and Conner, "How to Regulate Tech: A Technology Policy Framework for Online Services."
[52] Ibid.

government has created, starting with the Blueprint for an AI Bill of Rights and the NIST Risk Management framework.[53] As CAP wrote in May:

> Industry is not the only stakeholder that can provide leadership on responsible innovation that preserves democratic values and protects rights. The president oversees a federal government that is the largest employer in the country. The government's purchasing power wields enormous market-shaping power in the development of new technologies. The AI Bill of Rights, therefore, is not just a blueprint for private industry but also a road map for the government's own approach to artificial intelligence.[54]

In April, CAP called on President Biden to lead on AI by issuing an executive order on Artificial Intelligence that called for numerous leadership steps on AI including:[55]

> **Require federal agencies to implement the Blueprint for an AI Bill of Rights for their own usage of AI:**
>
> The president should require implementation of the Blueprint for an AI Bill of Rights for all federal agencies for their own usage of AI, with a plan due to the White House Council on AI within 90 days for implementation by 2024. The Blueprint for an AI Bill of Rights provided a roadmap to move principles into practices but did "not constitute binding guidance for the public or Federal agencies." The obvious next step is to require the Blueprint for an AI Bill of Rights to be implemented around use of AI by federal agencies. Agencies have a starting place with the public list of AI use-case inventories required from each federal agency impacted by EO 1396036 and subsequent OMB M-21-06 guidance. There is support from many experts for this move. In the National Artificial Intelligence Advisory Committee (NAIAC) draft report released in late April, committee members Janet Haven, Liz O'Sullivan, Amanda Ballantyne, and Frank Pasquale "advocated to anchor this Committee's work in a foundational rights-based framework, like the one laid out in OSTP's October 2022 Blueprint for an AI Bill of Rights" and lamented the committee's more immediate and tactical approach.
>
> **Require all AI tools deployed by federal agencies or contractors to be assessed under the NIST's AI Risk Management Framework and summaries to be publicly released:**
>
> The president should require all AI tools deployed by federal agencies or contractors to be assessed under the National Institute of Standards and

[53] The White House, "Blueprint for an AI Bill of Rights,"; NIST, "AI Risk Management Framework," available at https://www.nist.gov/itl/ai-risk-management-framework (last accessed June 2023).
[54] Conner, "White House Must Take More Action To Address AI Concerns."
[55] Conner, "The Needed Executive Actions to Address the Challenges of Artificial Intelligence."

Technology (NIST) AI Risk Management Framework (AI RMF), which was designed "to improve the ability to incorporate trustworthiness considerations into the design, development, use, and evaluation of AI products, services, and systems." Summaries should also be publicly released. The May 2023 draft report from the NAIAC recommended that:

> .. the White House encourage federal agencies to implement either the AI RMF—or similar processes and policies that align with the AI RMF—to address risks in all phases of the AI lifecycle effectively, with appropriate evaluation and iteration in place. We believe federal agencies can leverage the AI RMF to address issues relating to AI in scoping, development, and vending processes. These include but are not limited to bias, discrimination, and social harms that arise when building, assessing, and governing AI systems.[56]
>
> NIST's Trustworthy and Responsible AI Resource Center,[57] which was created to "facilitate implementation of, and international alignment with, the AI RMF," should help agencies coordinate those assessments. The president should also use his authority under the Federal Property and Administrative Services Act of 1949 (FPASA)[58] to require all federal contractors and subcontractors to assess any AI tools they use or deploy under the AI RMF, with implementing regulations to be expedited by the Federal Acquisition Regulatory Council (FAR).[59]

In addition, CAP called upon the Biden Administration to order all agencies to reexamine the impact of AI in enforcement of existing regulation, in future rulemaking, and ensure federal regulations to include analysis of how rulemaking and AI tools operate by recommending:

> **Require federal agencies to assess the use of AI in enforcement of existing regulation and address AI in future rulemaking to the maximum extent practicable:**
>
> Because AI has the potential to touch nearly every aspect of our lives, it is reasonable to assume that its use by both private and public sector actors will

[56] National Artificial Intelligence Advisory Committee, "(Final Draft) National Artificial Intelligence Advisory Committee (NAIAC) Year 1 Report: Year 1," p.16, available at https://www.ai.gov/wp-content/uploads/2023/05/NAIAC-Report-Year1.pdf.
[57] National Institute of Standards and Technology, "Trustworthy and Responsible AI Resource Center," available at https://airc.nist.gov/Home (last accessed April 2023).
[58] U.S. General Services Administration, "Federal Property and Administrative Services Act of 1949" (Washington, U.S. General Services Administration), available at https://disposal.gsa.gov/s/act49.
[59] U.S. General Services Administration, "Federal Acquisition Regulatory Council" (Washington, U.S. General Services Administration), available at https://www.acquisition.gov/far-council; Conner, "The Needed Executive Actions to Address the Challenges of Artificial Intelligence."

implicate the enforcement of countless statutes by federal agencies. The president should require federal agencies to assess whether the use of AI by the entities they regulate could implicate their enforcement of existing regulations, and if appropriate, address that use in future rulemaking to the maximum extent practicable. For example, use of AI by nursing homes to identify potential health problems or establish safe staffing levels could raise both civil rights and safety concerns that, if left unregulated, could violate the letter or intent of consumer protection or civil rights statutes. Although no general authority may exist governmentwide to regulate AI tools, their use in certain industries or contexts may compel an agency to revise regulations to govern their use by regulated entities. Thus, the executive order could require each agency to survey existing regulations and consider future proposals to regulate AI tools in domain-specific contexts to the maximum extent practicable.

Require that all new federal regulations include an analysis of how the rulemaking would apply to AI tools:

The president should amend EO 12866, "Regulatory Planning and Review,"[60] to require agencies to provide to OMB—and include in any final rule—an assessment of how any proposed regulations would or would not apply to AI tools, similar to existing requirements around impacts on small businesses or state mandates. For example, if the Department of Health and Human Services were proposing new civil rights protections for Medicare beneficiaries, they would have to include an analysis of whether and how these protections apply to AI tools used by providers covered by the regulation.[61]

Government must play an essential role in AI accountability, requirement enforcement, setting standards, writing new rules, and enforcing transparency or there will be no real AI accountability.

[60] Executive Office of the President, "Executive Order 12866: Regulatory Planning and Review," *Federal Register* 58 (190) (1993): available at https://www.archives.gov/files/federal-register/executive-orders/pdf/12866.pdf.
[61] Conner, "The Needed Executive Actions to Address the Challenges of Artificial Intelligence."