

The Center for Growth and Opportunity at Utah State University
**Public Interest Comment on the National Telecommunications
and Information Administration (NTIA) AI Accountability Policy
Request for Comment**

Authors:

Neil Chilson^a

Will Rinehart^a

Cosigners:

The undersigned write in response to the National Telecommunications and Information Administration's (NTIA) AI Accountability Policy Request for Comment ("RFC").

The Beacon Center of Tennessee

Frontier Institute

The John Locke Foundation

Mississippi Center for Public Policy

The James Madison Institute

Libertas Institute, Caden Rosenbaum

John Villasenor^b

Regulations.gov Docket No. NTIA-2023-0005

NTIA Docket No. 230407-0093

Submitted: June 12, 2023

The Center for Growth and Opportunity at Utah State University is a research center dedicated to producing ideas that transform lives. We explore the interactions between key institutions—business, government, and civil society—to improve opportunity, broad-based economic growth, and individual well-being. The Center occasionally conducts independent analyses addressing government rulemakings and proposals. This comment is designed to assist the agency as it explores these issues. The views expressed in this comment are those of the author(s) and do not necessarily reflect the views of The Center for Growth and Opportunity at Utah State University or the views of Utah State University.

^a Neil Chilson and Will Rinehart are senior research fellows at the Center for Growth and Opportunity at Utah State University.

^b John Villasenor is signing in his individual capacity and not on behalf of any entity he is affiliated with.

Introduction

The National Telecommunications and Information Administration's (NTIA) Request for Comment (RFC) seeks input on “how to develop a productive AI accountability ecosystem.”¹ The request does not sufficiently recognize that society already has accountability mechanisms that will apply to AI. The RFC defines AI very broadly, covering nearly all significant software systems.² But we have had software, algorithms, and automated systems meeting the RFC’s definition for decades. How have we held these systems accountable?

An accountability ecosystem for software already exists and has proven highly effective. It is polycentric in that it is layered and is comprised of business-to-business and business-to-consumer markets, reputational markets, and financial markets, all backed by generally applicable laws and norms.

This system aligns producer incentives with a multitude of stakeholders and creates feedback loops where a producer’s primary accountability is to its customers. Markets and the profit incentive provide the feedback that generally aligns the interests of consumers and producers, encouraging the development and delivery of valuable products and services, benefiting individuals and thus in aggregate benefiting society as a whole.

This complex market system is a primary source of accountability for many kinds of industries, companies, and technologies. But it is not flawless.³ When the interactions between a producer and a consumer create positive or negative effects on third parties, or when there is asymmetry in information between parties, markets may not offer the feedback that tends to align the interests of the producer and affected parties. In such situations, there may be a need for other, non-market “accountability” mechanisms. But these are always the exception, not the general rule.⁴ Taking the exception should be well justified. And any such exceptions must be evaluated and implemented in a way that supports the general market accountability mechanisms, rather than replacing or even undermining them.⁵

The RFC does not discuss market-based accountability. Yet it should, because many of the accountability mechanisms discussed in the RFC could (and do) operate within the market process. For example, when companies today voluntarily offer transparency mechanisms, certifications, or third-party audits to ensure product quality and thereby appeal to customers and users, this is a market-based approach.

There is another reason that this proceeding should focus on the market-based accountability mechanisms: the NTIA is not a regulatory agency and has limited authority to intervene.⁶ Any report generated from this discussion will hold persuasive value, but the NTIA cannot impose binding rules. Furthermore, AI is a technology with a multitude of potential subcategories and

1 AI Accountability Policy Request for Comment, 88 Fed. Reg. 22,435 (Apr. 13, 2023).

2 The definition incorporates both NIST’s definition of AI (“an engineered or machine-based system that can, for a given set of objectives, generate outputs such as predictions, recommendations, or decisions influencing real or virtual environments”) and the White House’s AI Blueprint definition of automated systems (“automated systems with the potential to meaningfully impact the American public’s rights, opportunities, or access to critical resources or services”). 88 Fed. Reg. 22,435 (internal quotation marks omitted) (quoting NIST’s Artificial Intelligence Risk Management Framework 1.0 and the Blueprint for an AI Bill of Rights).

3 Mark Walker, The Fundamental Welfare Theorems (Oct. 11, 2017) (archived online at the University of Arizona), http://www.u.arizona.edu/~mwalker/05_Pareto%20Efficiency/WelfareTheorems.pdf.

4 Office of Mgmt. and Budget, Exec. Office of The President, OMB Circular A-4, Regulatory Analysis, 6–7 (2003).

5 *Id.* at 8-9.

6 47 U.S.C. § 901.

applications, making it challenging to establish a single set of accountability mechanisms that address any potential externalities.

In this response, we will focus on situating AI accountability within the context of a market system, detailing the proper analysis for identifying and filling gaps in the current ecosystem, and addressing some specific RFC questions regarding trade-offs between different accountability goals, existing accountability mechanisms, and accountability mechanisms for government use.

Missing Market-Based Accountability

The RFC seeks information about how to develop a productive AI accountability ecosystem. Yet it ignores the existing vibrant ecosystem for guiding commercial activity toward socially productive ends: markets. Markets and the norms and general purpose laws that interact with them are the baseline accountability ecosystem for all industries. Markets sometimes need to be supplemented with non-market mechanisms, but such interventions must be justified or else risk undermining the market-based accountability mechanisms.⁷ Many of the mechanisms that the RFC asks about can be useful in markets.

The RFC asks in Question 30 about the role of government policy and regulation in AI accountability.⁸ However, it jumps to the substance of such intervention without addressing two steps needed to intervene successfully in the existing accountability ecosystem: identifying specific market failures and determining whether and how government intervention can resolve the identified market failures.⁹

Although the discussion below addresses the overall framing of the RFC, it is also relevant in response to Question 9 and Question 30.

The RFC's "AI Accountability" Framing Ignores Markets

The RFC posits that "[r]eal accountability can only be achieved when entities are held responsible for their decisions."¹⁰ It argues that "hold[ing] entities accountable for developing, using, and continuously improving the quality of AI products" can "build trust."¹¹

The RFC further contends that accountability can be achieved using mechanisms such as "measurements of AI system risks, impact assessments, and audits of AI system implementation."¹² Additionally, it elaborates that such "assessments and audits, governance policies, documentation and reporting, and testing and evaluation" can "prov[e] that an AI system is legal, effective, ethical, safe, and otherwise trustworthy."¹³ Trust is important, the RFC claims, because it "help[s] assist with compliance efforts" and "help[s] create marks of quality in the marketplace."¹⁴

7 OMB Circular A-4, 6-7.

8 88 Fed. Reg. 22,440.

9 Exec. Order No. 12866 § 1(b), 58 Fed. Reg. 51,735 (1993).

10 88 Fed. Reg. 22,434.

11 *Id.*

12 *Id.*

13 *Id.*

14 *Id.*

This last point is the RFC's sole substantive mention of markets or the marketplace as a beneficiary of the *output* of accountability. But this gets the relationship between accountability and markets backward.

Markets Are the Primary Source of Accountability

Markets do not merely benefit from accountability. They are the primary source of accountability. Indeed, the norm in competitive marketplaces is that companies expend significant effort to develop, use, and continuously improve the quality of their products and services. They also seek to build consumer trust and to develop their reputation and "marks of quality." Their ability to satisfy their customers is the difference between business success and failure.

In fact, markets play a crucial role in disciplining company behavior and holding them accountable to their customers. Some of the mechanisms involved include competition, reputation, customer feedback, pricing, and transparency.

Competition is a fundamental market mechanism that drives corporate accountability. In a competitive environment, companies are obligated to respond to customer needs and expectations effectively. If they fail to do so, customers have the choice to take their business to rival firms, which can lead to potential loss of market share and profitability. This competitive pressure incentivizes businesses to continuously strive for excellence and innovation in their products or services.

Reputation also significantly influences corporate accountability. A strong reputation can be a significant asset for a company, influencing customer loyalty and a firm's position in the marketplace. Conversely, negative experiences or public dissatisfaction can tarnish a company's reputation, leading to a potential customer exodus and reduced profitability. Customer feedback, a related accountability mechanism, is the direct expression of customer satisfaction or dissatisfaction. Consumers can hold companies accountable through reviews, ratings, and other feedback mechanisms, forcing them to address issues or risk alienating their customer base.

Similarly, *pricing mechanisms* can enforce accountability. If a company's pricing strategy does not match the perceived value of their offerings, customers may opt for competitors who offer better value for money. This incentivizes businesses to provide quality products or services at competitive prices.

Lastly, markets can incentivize *transparency* in business operations, creating accountability. Companies often face consumer demand for more information about their products or services, such as sourcing, manufacturing, or data handling practices. Companies that share such practices, including through advertising, can be held accountable by consumers who value these attributes. By enabling customers to make informed decisions, transparency also helps businesses to build trust and maintain strong customer relationships.

Markets Can Produce Suboptimal Results

Market mechanisms may not produce the optimal amounts of accountability. Economists and public policy experts generally identify four major kinds of market failure (meaning a failure of a real-world market to reach the theoretical Pareto efficient outcome) that may justify non-market

interventions: externalities, asymmetric information, public goods, and market power.¹⁵ Two of these are most relevant for this discussion.¹⁶

Information asymmetry could contribute to inadequate accountability. Companies often possess significantly more knowledge about their products or services than consumers. For certain kinds of AI services, technical expertise could be important to assessing the quality of the product or service. Consequently, consumers might be unable to fully assess an AI service's model or understand the potential risks, leading to decisions that may not be in their best interest.¹⁷

Externalities are another cause of market failure that could lead to insufficient accountability. These are costs or benefits experienced by third parties not involved in a particular transaction. For instance, one might imagine that decisions made by an AI system could affect people other than the direct users. However, because these impacts are not directly felt by the company's users, the users might not provide adequate feedback to the company, meaning the effect on third parties might not be factored into the company's practices.

A word of caution: seeming market failures can be transitory because market failures are often a market opportunity for entrepreneurs.¹⁸ For example, as the RFC points out, a burst of new organizations are offering AI certifications and other private sector accountability measures.¹⁹ These businesses see a need in the market and aim to fill it with their services. This is evidence that both companies and consumers are looking for ways to ensure AI is accountable. Thus, a so-called market failure can often incentivize its own resolution, particularly in highly dynamic and emerging areas like AI.

Even So, Intervention Requires Justification

The RFC also neglects the second half of the necessary foundation for government intervention: it does not ask commenters whether any identified market failures justify government intervention.

Market failure is a necessary but insufficient condition for government intervention. Successful intervention requires considering a number of key factors. First and foremost, the intervention must address a clearly identified problem or demonstrated market failure.²⁰ It is not enough to hypothesize a potential issue; demonstrating the existence or likely existence of the failure helps ensure intervention goes where it is needed.

Second, whether the issue is potential or realized, the government must be able to demonstrate that its proposed intervention can improve the market outcome.²¹ Even when the diagnosis of market failure is correct, the medicine is often a placebo. The correctives do little to move the needle. For example, there is evidence that efforts to correct information asymmetries in information privacy produce little or no behavioral change. Simplified disclosures of long-term privacy risks did not significantly shift users' comprehension of the disclosure, their willingness to share

15 John O. Ledyard, *Market Failure*, in *The New Palgrave Dictionary of Economics*, 1–5 (Steven N. Durlauf & Lawrence E. Blume eds., 2017).

16 There is no evidence of monopolization or even concentrated market power in this dynamic space. AI tools are excludable and thus not public goods.

17 AI may complicate the information asymmetry analysis. Large Language Models (LLM) can exhibit emergent behavior, meaning LLM providers may not fully understand how the model arrives at certain results.

18 Israel Kirzner, *Entrepreneurship*, in *The Collected Works of Israel M. Kirzner: The Essence of Entrepreneurship and the Nature and Significance of Market Process* (Peter J. Boettke & Frédéric Sautet eds., 2018).

19 88 Fed. Reg. 22,435.

20 Exec. Order No. 12866.

21 *Id.*

personal information, or their expectations about their rights.²² Indeed, giving users an increased feeling of control over the publication of their data often results in increased and riskier disclosures.²³ The Administration needs to be cognizant of the limits of its power to move behavior, even where markets are not achieving their theoretical potential.

Third, the costs of even successful government intervention must be considered. These include not only the direct costs of implementing and enforcing regulations but also the indirect costs, such as the potential for reduced economic efficiency. Government must ensure that the benefits of intervention outweigh these costs.²⁴

Finally, interventions must be designed to minimize unintended consequences. Of particular concern in this case is the potential of interventions to undermine the market's existing accountability mechanisms. Government-mandated accountability mechanisms can potentially undermine market-driven AI accountability mechanisms in at least four ways:

- **Displacement of Market Signals:** Government regulations could potentially cloud or distort market signals. For instance, static regulations of the marketing of AI services could reduce the accuracy and effectiveness of communications to consumers about new services.²⁵ Absent such advertising disclosure requirements, advertising can experiment with ways to best disclose the value or quality of these services.
- **Reduced Competition:** Overly strict regulations can create barriers to entry and reduce competition, which is a key market mechanism for ensuring accountability. If complying with regulations requires significant resources, it can disadvantage smaller companies or potential new entrants, leading to a more consolidated market with less competition. In such a scenario, larger incumbent firms have less incentive to respond to customer needs, as there is less threat of customers switching to competitors.
- **Misplaced Incentives:** If the focus of businesses shifts to meeting regulatory requirements rather than satisfying customer needs, it could lead to a “box-checking” mentality, where the primary goal becomes regulatory compliance instead of product or service improvement. This could stifle innovation and undermine the effectiveness of customer feedback as a market mechanism for accountability. For example, an AI company might spend more resources on meeting extensive disclosure requirements mandated by regulation, and less on actual product improvement based on user feedback.
- **Risk of Regulatory Capture:** There is a risk that regulations might be influenced more by incumbent firms rather than consumer needs—a phenomenon known as regulatory capture. This could lead to regulations that protect existing companies at the expense of consumers and new entrants, undermining the accountability that competition usually ensures.

22 Adam S. Chilton & Omri Ben-Shahar, *Simplification of Privacy Disclosures: An Experimental Test* (CoaseSandor Working Paper Series in Law and Economics No. 737, University of Chicago Law School, 2016), https://chicagounbound.uchicago.edu/cgi/viewcontent.cgi?article=2443&context=law_and_economics; see also Idris Adjerid, Alessandro Acquisti, Laura Brandimarte & George Loewenstein, *Sleights of Privacy: Framing, Disclosures, and the Limits of Transparency* (SSRN, July 2013), <https://doi.org/10.1145/2501604.2501613> (“[T]he ability of even improved transparency solutions or additional control tools to better align consumer attitudes towards privacy with actual behavior and reduce regret from oversharing is ultimately questionable.”).

23 Laura Brandimarte, Alessandro Acquisti & George Loewenstein, *Misplaced Confidences: Privacy and the Control Paradox*, 4 Soc. Psych. and Personality Sci. 3, 340–347 (2012), <https://doi.org/10.1177/1948550612455931>.

24 OMB Circular A-4.

25 Julie Tuan, *U.S. West, Inc v. FCC*, Berkeley Technology Law Journal 15, no. 1 353–372 (2000), <http://www.jstor.org/stable/24115718>.

In each of these categories, government-mandated mechanisms, while intended to increase accountability, could potentially undermine the existing accountability mechanisms of the market. For example, to the extent that mandated assessments and audits draw company resources away from focusing on product characteristics or features that users want, the company becomes less accountable to its users. This underscores the importance of carefully designing government interventions to complement rather than hinder market mechanisms. The consequence of getting intervention wrong could be less accountable AI systems.

Many Mechanisms Mentioned in the RFC Can Contribute to Market-Based Accountability

Even without market failures justifying government intervention, voluntary adoption of accountability mechanisms like those mentioned in the RFC could play a valuable role in the existing market-driven, polycentric AI accountability ecosystem. The RFC lists mechanisms such as audits, algorithmic impact assessments, privacy by design, and certifications.²⁶ These tools can foster transparency, boost consumer confidence, and ensure that AI applications uphold ethical standards, all of which can enhance the reputation, credibility, and profitability of AI companies.

Audits, for instance, can offer insights into the functioning of AI systems and their potential impacts. Through well-considered audits, companies can demonstrate the safety, fairness, and reliability of their AI applications, thereby instilling confidence in their customers and stakeholders. Similarly, algorithmic impact assessments allow companies to assess and document the potential risks and benefits of specific AI applications. This can help them understand, mitigate, and communicate any potential risks while emphasizing the benefits and enhancing the value proposition of their products or services.

Privacy by design, as an approach, encourages companies to embed privacy considerations into their AI systems from the inception stage. This can improve quality and enhance customer trust, which are strong competitive advantages.

Finally, certifications can act as proof of compliance with certain established ethical or technical standards. They can serve as a shorthand for consumers, signaling that a product or service is trustworthy, and can also help businesses distinguish themselves in the marketplace.

These mechanisms can indeed be useful and profitable when functioning within the market process, where companies adopt them to meet user, customer, or investor demands; companies can update or rethink them as such demands evolve or become better understood. This responsiveness to changing demands is critical to the success of such mechanisms.

Question 7 asks, “Are there ways in which accountability mechanisms ... might even frustrate the development of trustworthy AI?” Yes. Mandatory mechanisms, implemented without regard for the demands of users, and without first identifying market failures that undermine market-driven accountability, such mechanisms risk substituting for and can even undermine the market-driven accountability ecosystem described above. The long-term benefits of these mechanisms rely on their ability to respond to the dynamic, evolving nature of the AI technology, its uses, and user demands. Mandated accountability measures that cannot adapt may struggle to keep pace with technological advancements, limiting the ability of companies to adapt to changing consumer needs and preferences.

²⁶ 88 Fed. Reg. 22,435.

Therefore, while these tools can be beneficial, they should complement rather than replace market-driven processes. As discussed above, a key way to do this is by thoroughly identifying and analyzing the gaps (created by market failures) in accountability. Any such requirements should also be subject to frequent review and revision to remain relevant.

Specific RFC Questions

Question 3

Question text: *AI accountability measures have been proposed in connection with many different goals, including those listed below. To what extent are there trade-offs among these goals?*

The issue is even more complex than trade-offs among various goals. There are often tensions and even incompatibilities between different views of the same goal. For example, “fairness” is a common goal of AI assurance.²⁷ However, as noted in a previous filing to the NTIA, there are competing versions of fairness and some views of fairness might simply be incompatible with each other.²⁸ Due to these incompatibilities, achieving fairness by one measure often makes it impossible to achieve fairness by other measures.

Moreover, even when there is agreement about the measure that should be pursued, there are significant practical problems with audits or other mechanisms achieving repeatable, consistent estimates of the measure.

Research conducted in the past couple of years only confirms this point. In one study, 73 independent research teams used identical cross-country survey data to see if more immigration reduces support for government services. As the authors concluded, “Instead of convergence, teams’ results varied greatly, ranging from large negative to large positive effects of immigration on social policy support. The choices made by the research teams in designing their statistical tests explain very little of this variation; a hidden universe of uncertainty remains.”²⁹

Other studies confirm the variations that exist in data interpretation. Comparisons of independent analysts using the same dataset also found varying results on other topics, such as the effects of gender and professional status on verbosity during group meetings, whether soccer referees are more likely to give red cards to dark-skinned players than to light-skinned players, and how the same MRIs were analyzed by different teams of scientists.³⁰

All of the available data thus suggest there will be difficulty in establishing a single set of accountability measures for generative AI. On top of that, mandating new measurements is not guaranteed to solve the purported problems.

27 88 Fed. Reg. 22,434 (defining “trustworthy AI” to “encapsulate ... fairness”); RFC at 22,439 n.19 (citing Microsoft Responsible AI Standard Reference Guide); Blueprint for AI Bill of Rights at 27–28;

28 Will Rinehart, Public Interest Comment for the National Telecommunications and Information Administration (NTIA) on the Intersection of Privacy, Equity, and Civil Rights (March 6, 2023), <https://www.regulations.gov/comment/NTIA-2023-0001-0064>; Sorelle A. Friedler, Carlos Scheidegger & Suresh Venkatasubramanian, *On the (Im)possibility of Fairness*, 64 *Comm’n of the ACM* 4, 136–43 (2021), <https://doi.org/10.1145/3433949>. See also Virginia Foggo, John Villasenor & Pratyush Garg, *Algorithms and Fairness*, 17 *Ohio State Tech. Law J.* 123 (2021).

29 Nate Breznau, Eike Mark Rinke, Alexander Wuttke & Tomasz Żółtak, *Observing Many Researchers Using the Same Data and Hypothesis Reveals a Hidden Universe of Uncertainty*, 119 *PNAS* 44 (2022), <https://www.pnas.org/doi/10.1073/pnas.2203150119>.

30 Martin Schweinsberg et al., *Same Data, Different Conclusions: Radical Dispersion in Empirical Results When Independent Analysts Operationalize and Test the Same Hypothesis*, 165 *Org. Behav. and Hum. Decision Processes*, 228–49 (2021), <https://doi.org/10.1016/j.obhdp.2021.02.003>; R. Silberzahn et al., *Many Analysts, One Data Set: Making Transparent How Variations in Analytic Choices Affect Results*, 1 *Advances in Methods and Practices in Psych. Sci.* 3, 337–56 (2018), <https://doi.org/10.1177/2515245917747646>; Rotem Botvinik-Nezer et al., *Variability in the Analysis of a Single Neuroimaging Dataset by Many Teams*, 582 *Nature* 7810, 84–88 (2020), <https://doi.org/10.1038/s41586-020-2314-9>.

But it goes even deeper than that because one of the most important and high-profile cases that purported to show AI harm never discussed the allegedly harmful impact. In *National Fair Housing Alliance v. Facebook, Inc.*, the plaintiff alleged that Facebook’s classification of its users and its ad targeting tools permitted landlords, developers, and housing service providers to limit the audience for their ads based on sex, religion, familial status, and national origin in violation of the FHA.

The Department of Justice pushed Meta to change their ad services to settle such complaints.³¹ As a result, Meta dropped the allegedly discriminatory “Special Ad Audience” tool for its housing ads.³² In the process of complying with the DOJ, the company got rid of thousands of ad categories, including “African American multicultural affinity.”

But there was never a formal study or finding that the Facebook tool had the effect of being discriminatory. It could have been that these tags were used in a negative manner to target Black Americans, or they might have been used in a positive manner that enhanced housing opportunities. It could have been that the ads had no appreciable discriminatory effect. It was well known by advertisers that Meta’s targeting criteria had problems. A group of advertisers sued Facebook and won a \$40 million settlement on measurement issues.³³ Moreover, Pew Research Center conducted its own survey on the accuracy of the tags and found that people disagreed with the categories they had been assigned.³⁴ In other words, the tags were often inaccurate. It remains unclear to this day the real impact of the algorithm.

Follow-up research of the changes made by Meta in the wake of the DOJ pressure seems to confirm the changes had little effect. As one report explained it, “Merely removing demographic features from a real-world algorithmic system’s inputs can fail to prevent biased outputs.”³⁵ In the end, the authors of the report suggested other approaches to mitigating discriminatory effects.

Despite the overhaul, Facebook was not able to fully eliminate the potential for bias in the ad targeting system. This case and others like it illustrate the challenges that the NTIA will face in trying to untangle and mitigate AI bias. Charting the actual effects of these systems will require a thorough examination of the data and the impact in the real world. Already these empirical studies have shown that there are trade-offs, and that even the most well-intended changes fail to really move the needle anyway.

31 U.S. Dept. of Justice, Justice Department Secures Groundbreaking Settlement Agreement with Meta Platforms, Formerly Known as Facebook, to Resolve Allegations of Discriminatory Advertising, June 21, 2022, <https://www.justice.gov/opa/pr/justice-department-secures-groundbreaking-settlement-agreement-meta-platforms-formerly-known>.

32 *Id.*

33 Annabelle Woodward, *Facebook to Pay Advertisers \$40 Million to Settle Hard Fought Class Action Lawsuit*, Forbes, Oct. 7, 2019, <https://www.forbes.com/sites/annabellewoodward1/2019/10/07/facebook-to-pay-advertisers-40-million-settling-hard-fought-class-action-lawsuit/?sh=232acc3f1d72>.

34 Paul Hitlin, Lee Rainie & Kenneth Olmstead, *Facebook Algorithms and Personal Data*, Pew Research Center, Jan. 16, 2019, <https://www.pewresearch.org/internet/2019/01/16/facebook-algorithms-and-personal-data/>.

35 Piotr Sapiezynski, Avijit Ghosh, Levi Kaplan, Aaron Rieke & Alan Mislove, *Algorithms That “Don’t See Color”: Comparing Biases in Lookalike and Special Ad Audiences*, Proc. of the 2022 AAAI/ACM Conf. on AI, Ethics, and Soc’y (Dec.16, 2019), <https://doi.org/10.48550/arXiv.1912.07579>.

Question 9

Question text: *What AI accountability mechanisms are currently being used? Are the accountability frameworks of certain sectors, industries, or market participants especially mature as compared to others? Which industry, civil society, or governmental accountability instruments, guidelines, or policies are most appropriate for implementation and operationalization at scale in the United States? Who are the people currently doing AI accountability work?*

In addition to the market-driven accountability mechanisms discussed above, there are already many existing non-market accountability mechanisms that are applicable to AI. For instance, Title VII of the Civil Rights Act of 1964 prohibits employment discrimination on the basis of certain protected categories. A company that makes hiring decisions based in part on a discriminatory AI system would be accountable under that framework. Similarly, uses of AI in ways that violate the Fair Housing Act, the Equal Credit Opportunity Act, the Education Amendments of 1972, and other federal (and state) anti-discrimination laws are also legally actionable.

Product liability provides another mechanism for AI accountability. If the AI in a driverless car causes an accident, the manufacturer of the car or the manufacturer's upstream suppliers can be held accountable through a products liability claim.

The Federal Trade Commission (FTC) is another existing government mechanism for AI accountability. The FTC is charged with preventing and redressing unfair and deceptive business practices across nearly the entire economy. Under the FTC's authority, "a representation, omission, or practice is deceptive if it is likely to mislead consumers acting reasonably under the circumstances and is material to consumers—that is, it would likely affect the consumer's conduct or decisions with regard to a product or service."³⁶ An unfair act or practice is one that "causes or is likely to cause substantial injury to consumers which is not reasonably avoidable by consumers themselves and not outweighed by countervailing benefits to consumers or to competition."³⁷

These prohibitions on unfair and deceptive acts and practices apply to practices by AI companies. Indeed, the FTC has a strong track record of applying these general standards to address consumer harm from novel technologies. In 2016 the FTC issued a comprehensive report on how to deal with potential bias and discrimination in the commercial use of "big data" technologies.³⁸ Many of the concerns discussed are identical to the potential issues raised about the use of AI. The report details government authority to rein in bias and discrimination, including through the Fair Credit Reporting Act, Equal Opportunity laws, and the Federal Trade Commission Act. In discussing the FTC's authority, the report concludes:

Companies engaging in big data analytics should consider whether they are violating any material promises to consumers—whether that promise is to refrain from sharing data with third parties, to provide consumers with choices about sharing, or to safeguard consumers' personal information—or whether they have failed to disclose material information to consumers. In addition, companies that maintain big data on consumers should take care to reasonably secure consumers' data. Further, at a minimum, companies must not sell their big data analytics products to customers if they know or have reason to know that those customers will use the products for fraudulent or discriminatory purposes. The

36 FTC, Enforcement Policy Statement on Deceptively Formatted Advertisements, Dec. 22, 2015, https://www.ftc.gov/system/files/documents/public_statements/896923/151222deceptiveenforcement.pdf.

37 15 U.S.C. 45(n).

38 FTC, Big Data: A Tool for Inclusion or Exclusion?, Jan. 2016, <https://www.ftc.gov/system/files/documents/reports/big-data-tool-inclusion-or-exclusion-understanding-issues/160106big-data-rpt.pdf>.

inquiry will be fact-specific, and in every case, the test will be whether the company is offering or using big data analytics in a deceptive or unfair way. (iv)

These same principles apply to uses of AI. If AI companies lie to consumers or cause substantial, unavoidable injury without countervailing benefits, the FTC can bring enforcement actions against them. More recently the FTC has offered guidance about how its consumer protection authority applies to various uses of AI and claims about AI.³⁹

Earlier this year, the FTC, DOJ, CFPB, and EEOC released a joint statement explaining the “[e]xisting legal authorities [that] apply to the use of automated systems” including “those sometimes marketed as artificial intelligence” and “pledg[ing] to vigorously use [their] collective authorities to protect individuals’ rights regardless of whether legal violations occur through traditional means or advanced technologies.”⁴⁰

In short, statutory law, regulations, and common law already provide a comprehensive basis for ensuring AI accountability. These frameworks are applicable to AI despite not being specific to AI.

In considering the policy responses to potentially problematic uses of AI, it is important to determine the extent to which existing frameworks already provide accountability in relation to those uses. When a mechanism to provide accountability for what will clearly become a problematic use of AI is lacking, then it does indeed make sense to consider the potential role of both “soft” law (e.g., voluntary, widely accepted standards) and “hard” law (e.g., new AI-specific federal statutes or AI-specific agency rulemaking). But those discussions should always be conducted with awareness of existing frameworks and a clear identification of how any new governance approaches would be complementary to and not simply duplicative of existing approaches.

Question 30(d)

*What accountability practices should government (at any level) itself mandate for the AI systems the government uses?*⁴¹

Government applications of AI fall outside of the market process, and as such, they lack the quick feedback loops that generate accountability in commercial applications. Instead, political and legal mechanisms primarily drive the accountability of government uses of AI. In the absence of the profit motive, which naturally steers decision-making toward socially beneficial outcomes in the private sector, non-market accountability mechanisms become critical to enabling political and public oversight of governmental AI usage.

In addition to normal good-government practices, key goals the government should pursue when establishing accountability mechanisms for its own AI uses include:

- **Ensuring compliance with constitutional and legal constraints:** The government’s use of AI must align with the constitutional principles and rights that form the backbone of American democracy. These include respecting the fundamental rights of citizens,

39 Michael Atleson, *Keep Your AI Claims in Check*, Federal Trade Commission (blog) (Feb. 27, 2023), <https://www.ftc.gov/business-guidance/blog/2023/02/keep-your-ai-claims-check>; Michael Atleson, *Chatbots, Deepfakes, and Voice Clones: AI Deception for Sale*, Federal Trade Commission (blog) (Mar. 20, 2023), <https://www.ftc.gov/business-guidance/blog/2023/03/chatbots-deepfakes-voice-clones-ai-deception-sale>.

40 Rohit Chopra, Kristen Clarke, Charlotte A. Burrows & Lina M. Khan, *Joint Statement on Enforcement Efforts against Discrimination and Bias in Automated Systems*, Federal Trade Commission (Apr. 25, 2023), <https://www.ftc.gov/legal-library/browse/cases-proceedings/public-statements/joint-statement-enforcement-efforts-against-discrimination-bias-automated-systems>.

41 88 Fed. Reg. 22,440.

such as the right to privacy, freedom of speech, and due process. In addition, laws such as the Privacy Act, the E-Government Act of 2002, the Freedom of Information Act, and others provide part of a legal framework for government handling of data.⁴² AI uses should not infringe upon these constitutional rights or violate these legal obligations, of course. Accountability mechanisms should help law enforcers and government oversight entities identify and substantiate potential violations.

- **Facilitating public participation:** The process of developing and deploying AI systems in government should involve meaningful public participation. Accountability mechanisms that help inform the public on potential and ongoing government uses of AI will help ensure that these systems serve the public interest and that diverse perspectives are taken into account.
- **Enabling oversight and review:** Existing government oversight and review mechanisms, such as congressional committees, inspector generals, and special governmental entities such as the Privacy and Civil Liberties Oversight Board need relevant, timely information on government uses of AI to effectively prevent and redress government abuses. Accountability mechanisms for government uses of AI will be a key source of this information.

Conclusion

The United States has a robust and effective accountability ecosystem for software. Market processes generally channel applications of new technology, including AI, to socially productive purposes. Markets make companies accountable to their consumers and investors. Accountability mechanisms such as those discussed in the RFC play an important role in this market process by helping companies meet consumer demand for reliable, safe, and otherwise trustworthy systems.

In addition to market-based accountability mechanisms, there are many existing non-market accountability mechanisms that are applicable to AI, from civil rights laws to consumer protection enforcement.

This existing accountability ecosystem is not flawless, but it should not be ignored. Otherwise, efforts to improve accountability could be useless or even counterproductive by undermining the market-based and other accountability mechanisms that already exist. Therefore, when seeking to improve the existing ecosystem through further government intervention, policymakers must identify market failures that justify action and choose actions that will be effective in addressing those failures.

Of course, NTIA itself does not have authority to directly pursue government intervention. But it does have the ability to gather information to inform action throughout the government. As such, it should properly lay the foundation for such action by acknowledging the existing market-based accountability ecosystem and identifying sustained market failures, if any.

The Center for Growth and Opportunity and our cosigners thank you for the opportunity to file these comments.

⁴² The Privacy Act governs certain uses of personal data held by the government. The Freedom of Information Act ensures public access to government records, promoting transparency.