epic.org

Electronic Privacy Information Center
1519 New Hampshire Avenue NW
Washington, DC 20036, USA

+1 202 483 1140
+1 202 483 1248
@EPICPrivacy
https://epic.org

COMMENTS OF THE ELECTRONIC PRIVACY INFORMATION CENTER

to the

National Telecommunications and Information Administration

AI Accountability Policy Request for Comment

Docket No. 230407-0093

88 Fed. Reg. 22,433

June 12, 2023

The Electronic Privacy Information Center (EPIC) submits these comments in response to the National Telecommunications and Information Administration (NTIA)'s recent request for information regarding artificial intelligence (AI) system accountability.[1] The NTIA is soliciting comments that, together with information collected from public engagements, will be used "to draft and issue a report on AI accountability policy development, focusing especially on the AI assurance ecosystem."

It is a critical moment for the federal government to espouse robust policies and practices concerning algorithmic audits, impact assessments, and other safeguards on AI systems. EPIC commends the NTIA for its interest in this topic and urges the agency to promulgate clear guidance that can be used by a wide range of policymakers and regulators seeking to establish legal safeguards on the use and development of AI.

EPIC is a public interest research center in Washington, D.C., established in 1994 to focus on public attention on emerging civil liberties issues and to secure the fundamental right to privacy in the digital age for all people through advocacy, research, and litigation.[2] EPIC has a long history of promoting transparency and accountability for information technology.[3]

---

[1] AI Accountability Policy Request for Comment, 88 Fed. Reg. 22,433 (Apr. 13, 2023), https://www.federalregister.gov/documents/2023/04/13/2023-07776/ai-accountability-policy-request-for-comment.
[2] EPIC, *About Us* (2023), https://epic.org/about/.
[3] *See, e.g.*, EPIC, *AI & Human Rights* (2023), https://epic.org/issues/ai/; EPIC, *AI in the Criminal Justice System* (2023), https://epic.org/issues/ai/ai-in-the-criminal-justice-system/; EPIC, *Generating Harms: Generative AI's Impact & Paths Forward* (2023), https://epic.org/wp-content/uploads/2023/05/EPIC-Generative-AI-White-Paper-May2023.pdf; EPIC, *Screen & Scored in the District of Columbia* (2022), https://epic.org/wp-content/uploads/2022/11/EPIC-Screened-in-DC-Report.pdf; Comments of EPIC, *In re Privacy, Equity, and Civil Rights Request for Comment* (Mar. 6, 2023), https://epic.org/wp-content/uploads/2023/03/EPIC-comments-NTIA-DiversityEquityCivilRights-RFC.pdf.

Section I of these comments highlights previous recommendations by EPIC and other entities concerning AI accountability, which together should guide the NTIA's inquiry and report. Section II answers some of the specific questions posed by the NTIA in its request for comment.

I. Recommended Safeguards

a. Federal frameworks

As you know, the White Office of Science and Technology Policy's *Blueprint for an AI Bill of Rights* (2022) includes a variety of recommendations concerning AI accountability.[4] We highlight in particular:

- "Independent evaluation and plain language reporting in the form of an algorithmic impact assessment, including disparity testing results and mitigation information, should be performed and made public whenever possible to confirm these protections."[5]
- "You should be protected from violations of privacy through design choices that ensure such protections are included by default, including ensuring that data collection conforms to reasonable expectations and that only data strictly necessary for the specific context is collected."[6]
- "Systems should undergo pre-deployment testing, risk identification and mitigation, and ongoing monitoring that demonstrate they are safe and effective based on their intended use, mitigation of unsafe outcomes including those beyond the intended use, and adherence to domain-specific standards."[7]
- "Independent evaluation and reporting that confirms that the system is safe and effective, including reporting of steps taken to mitigate potential harms, should be performed and the results made public whenever possible."[8]
- "Surveillance or monitoring systems should be subject to heightened oversight that includes at a minimum assessment of potential harms during design (before deployment) and in an ongoing manner, to ensure that the American public's rights, opportunities, and access are protected. This assessment should be done before deployment and should give special attention to ensure there is not algorithmic discrimination[.] Such assessment should then be reaffirmed in an ongoing manner as long as the system is in use."[9]
- "You should know how and why an outcome impacting you was determined by an automated system, including when the automated system is not the sole input determining the outcome."[10]

---

[4] White House Office of Sci. & Tech. Pol'y, *Blueprint for an AI Bill of Rights: Making Automated Systems Work for the American People* (Oct. 2022), https://www.whitehouse.gov/wp-content/uploads/2022/10/Blueprint-for-an-AI-Bill-of-Rights.pdf.
[5] *Id.* at 23.
[6] *Id.* at 30.
[7] *Id.* at 15.
[8] *Id.*
[9] *Id.* at 15.
[10] *Id.* at 34.

- "Automated systems should provide explanations that are technically valid, meaningful and useful to you and to any operators or others who need to understand the system and calibrated to the level of risk based on the context."[11]

The AI Risk Management Framework, published by the National Institute of Standards and Technology pursuant to the National Defense Administration similarly recommended assurance mechanisms:

- Identify AI actors responsible for evaluating efficacy of risk management processes and approaches, and for course-correction based on results.
- Establish mechanisms to enable the sharing of feedback from impacted individuals or
- Establish policies that define assessment scales for measuring potential AI system impact. Scales may be qualitative, such as red-amber-green (RAG), or may entail simulations or econometric approaches.
- Establish policies for assigning an overall risk measurement approach for an AI system, or its important components, e.g., via multiplication or combination of a mapped risk's impact and likelihood (risk = impact x likelihood).
- Establish and regularly review documentation policies that, among others, address information related to: (1) AI actors contact information; (2) business justification; (3) scope and usages; (4) assumptions and limitations; (5) description and characterization of training data; (6) algorithmic methodology; (7) evaluated alternative approaches; (8) description of output data; (9) testing and validation results (including explanatory visualizations and information); (10) down- and up-stream dependencies; (11) plans for deployment, monitoring, and change management; and (12) stakeholder engagement plans.
- Establish policies that promote effective challenges of AI system design, implementation, and deployment decisions, via mechanisms such as the three lines of defense, model audits, or red-teaming – to ensure that workplace risks such as groupthink do not take hold.
- Establish policies that incentivize safety-first mindset and general critical thinking and review at an organizational and procedural level.
- Establish whistleblower protections for insiders who report on perceived serious problems with AI systems.
- Establish impact assessment policies and processes for AI systems used by the organization.
- Verify that impact assessment activities are appropriate to evaluate the potential negative impact of a system and how quickly a system changes, and that assessments are applied on a regular basis.
- Utilize impact assessments to inform broader evaluations of AI system risk
- Identify, document and remediate risks arising from AI system components and pre-trained models per organizational risk management procedures, and as part of third-party risk tracking.
- Respond to and document detected or reported negative impacts or issues in AI system performance and trustworthiness.
- Document the basis for decisions made relative to tradeoffs between trustworthy characteristics, system risks, and system opportunities.

[11] *Id.* at 40.

- Maintain a database of reported errors, incidents and negative impacts including date reported, number of reports, assessment of impact and severity, and responses.
- Maintain a database of system changes, reason for change, and details of how the change was made, tested and deployed.
- Utilize TEVV (Test, Evaluation, Verification, Validation) outputs from map and measure functions when considering risk treatment.
- Plan and implement risk management practices in accordance with established organizational risk tolerances.
- Establish mechanisms to capture feedback from system end users and potentially impacted groups.
- Establish risk controls considering trustworthiness characteristics, including: (1) data management, quality, and privacy (e.g., minimization, rectification, or deletion requests) controls as part of organizational data governance policies; (2) machine learning and end-point security countermeasures (e.g., robust models, differential privacy, authentication, throttling); (3) business rules that augment, limit or restrict AI system outputs within certain contexts; (4) utilizing domain expertise related to deployment context for continuous improvement and TEVV across the AI lifecycle; (5) development and regular tracking of human-AI teaming configurations; (6) model assessment and test, evaluation, validation and verification (TEVV) protocols; (7) use of standardized documentation and transparency mechanisms; (8) software quality assurance practices across AI lifecycle; and (9) mechanisms to explore system limitations and avoid past failed designs or deployments.
- Establish and maintain procedures to regularly monitor system components for drift, decontextualization, or other AI system behavior factors.
- Establish and maintain procedures for capturing feedback about negative impacts.
- Apply change management processes to understand the upstream and downstream consequences of bypassing or deactivating an AI system or AI system components.
- Evaluate AI system trustworthiness in conditions similar to deployment context of use, and prior to deployment.
- Regularly assess and document system performance relative to trustworthiness characteristics and tradeoffs between negative risks and opportunities.
- Evaluate AI system oversight practices for validity and reliability. When oversight practices undergo extensive updates or adaptations, retest, evaluate results, and course correct as necessary.
- Review audit reports, testing results, product roadmaps, warranties, terms of service, end user license agreements, contracts, and other documentation related to third-party entities to assist in value assessment and risk management activities.
- Track third-parties preventing or hampering risk-mapping as indications of increased risk.
- Review third-party material (including data and models) for risks related to bias, data privacy, and security vulnerabilities.
- Establish assessment scales for measuring AI systems' impact. Scales may be qualitative, such as red-amber-green (RAG), or may entail simulations or econometric approaches. Document and apply scales uniformly across the organization's AI portfolio.
- Apply TEVV regularly at key stages in the AI lifecycle, connected to system impacts and frequency of system updates.
- Develop TEVV procedures that incorporate socio-technical elements and methods and plan to normalize across organizational culture.

- Regularly review and refine TEVV processes.
- Evaluate AI system oversight practices for validity and reliability. When oversight practices undergo extensive updates or adaptations, retest, evaluate results, and course correct as necessary.
- Review audit reports, testing results, product roadmaps, warranties, terms of service, end user license agreements, contracts, and other documentation related to third-party entities to assist in value assessment and risk management activities.
- Track third-parties preventing or hampering risk-mapping as indications of increased risk.
- Review third-party material (including data and models) for risks related to bias, data privacy, and security vulnerabilities.
- Apply traditional technology risk controls – such as procurement, security, and data privacy controls – to all acquired third-party technologies.

The Chief Information Officers Council has drafted an alpha version of an Algorithmic Impact Assessment ("AIA") tool for federal government agencies to assess risks of using automated decision-making systems.[12] While not required, the AIA tool's assessment examines, among other things, why automation was chosen as an approach over other solutions; the potential impacts and risks of automation; the system's autonomy and whether and how the system will replace human decision making; the ability for a human to review and override the decision; any feedback loops created by the system; whether third party vendors are involved and if so, in what capacity, whether they are auditable, and the transparency of the decision-making system; and a system impact assessment that looks to how the decision system will affect the rights and freedoms, economic interests, and health of the impacted group as well as any environmental impacts.[13] The tool asks about measures to mitigate data quality, bias, and privacy risks.[14] The Government of Canada has a similar, albeit mandatory, risk assessment tool to determine the impact of an automated decision-making system.[15] The tool has additional public transparency requirements as well.[16]

b. Enacted and pending legislation

Impact and assessments and audits feature prominently in a wide range of enacted and pending laws. This includes, for example, the Colorado Privacy Act ("CPA") and implementing regulations; proposed algorithmic accountability legislation in Washington State; and the federal Algorithmic Accountability Act of 2022.

---

[12] *Algorithmic Impact Assessment*, CIO.gov, https://www.cio.gov/aia-eia-js/ (last visited Jun. 9, 2023).
[13] *Id.*
[14] *Id.*
[15] *Algorithmic Impact Assessment Tool*, Gov't Canada, https://www.canada.ca/en/government/system/digital-government/digital-government-innovations/responsible-use-ai/algorithmic-impact-assessment.html (last visited Jun. 9, 2023).
[16] *Id.*

The regulations implementing the CPA require a data protection assessment when a business's processing of personal data presents "heightened risk of harm" to a consumer[17] and impose special assessment requirements for automated processing of personal data.[18] At minimum, the data protection assessment must include a description and context of the processing activity, including the relationship between the controller and the consumer; the categories of personal data to be processed, with increased scrutiny when the data processing involves sensitive data or data from a minor; the nature and operational elements of the processing activity; the core purpose and expected benefits of the processing; an evaluation of the sources and nature of the risks associated with the processing activity; measures taken to reduce the risks identified, with an emphasis on de-identification and ensuring consumers retain their statutory data subject rights; a description of how the benefits outweigh the risks of the processing activity; any relevant parties contributing to the data protection assessment; any audits conducted in relation to the data protection assessment; and when the data protection assessment was reviewed and approved as well as details of the individuals who approved it.[19] If the business's processing of personal data involves automated profiling of a consumer, the data protection assessment must also include information about the specific data used to profile, the decisions made using profiling, and explanations on why the profiling directly and reasonably relates to the controller's goods and services.[20] Profiling is defined as "automated processing of personal data to evaluate, analyze, or predict personal aspects concerning an identified or identifiable individual's economic situation, health, personal preferences, interests, reliability, behavior, location, or movements."[21]

SB 5116, introduced in the Washington State Legislature,[22] requires an algorithmic accountability report for any automated decision-making system used by a public agency. This report must include the system's name, vendor, and version; a description of the system's capabilities and any reasonably foreseeable capabilities outside the scope of the agency's proposed use; the types of data that the system uses and how that data is generated, collected, and processed; whether the decision system has been tested by an independent third party and whether or not it has been tested for bias; whether the system makes decisions regarding legal and constitutional rights of residents; any impacts of the decisions system on civil rights and liberties, any potential disparate impacts, and a mitigation plan.[23] The report also requires any a clear data use and management policy, with specific protocols for how, where, when and on whom the technology will be deployed; any third parties who will have access and reasons for their access; how the data will be secured and stored; training protocols for personnel who will use the system, and any public or community engagement held or to be held in connection with the automated decision system.[24]

---

[17] 4 CCR 904-3 [hereinafter CPA Regulations], Rules 8.02, , https://coag.gov/app/uploads/2023/03/FINAL-CLEAN-2023.03.15-Official-CPA-Rules.pdf. Processing that presents a "heightened risk of harm" includes (1) processing for targeted advertisements and profiling that presents reasonably foreseeable risk of unfair or deceptive treatment, unlawful disparate impacts, as well as financial or physical injury or other substantial injury, (2) selling personal data, and (3) processing of sensitive data. Colo. Rev. Stat. § 6-1-1309(2).
[18] CPA Regulations Rule 9.06.
[19] CPA Regulations Rule 8.04.
[20] CPA Regulations Rule 9.06.
[21] CPA Regulations Rule 2.02.
[22] Substitute S.B. 5116, 67th Leg., Reg. Sess. (Wash. 2021) [hereinafter S.B. 5116].
[23] S.B. 5116 §5(6).
[24] S.B. 5116 §5(6)(j).

At the federal level, the proposed Algorithmic Accountability Act of 2022 would require impact assessments with the above information while also mandating (1) ongoing testing and evaluation of the decision-making system and (2) disclosures concerning the transparency and explainability of the system and its decisions.[25] Although the full impact assessments would not be public, summary reports would be submitted to the FTC and disclosed in part through a central repository.[26]

Comparison of Disclosures Required Under Active and Proposed Risk Assessment Frameworks

	Colorado Privacy Act	**S.B. 5116**	**Algorithmic Accountability Act of 2022**	**CIO AIA tool[27]**	**Canada's AIA Tool[28]**
Description of Intended Purpose and Proposed use	Included	Included	Included	Included	Included
Necessity of the automated decision-system	Included	Included	Included	Included	Included
Cost-benefit analysis	Included	Included	Included	Included	Included
Effects on Civil, Constitutional, and Legal Rights	Included	Included	Included	Included	Included
Disparate Impact Evaluation	Included	Included	Included	Included	Included
Mitigation Plans	Included	Included	Included	Included	Included
Data Quality Assessments	Not Included	Included	Included	Included	Included
Human-in-the-loop disclosures	Not Included	Included	Included	Included	Included
Performance/Benchmark Auditing Requirements	Included	Included	Included	Included	Included
Personnel Training Requirements	Included	Included	Included	Included	Included
Decision/Recommendation Explanation Requirements	Not Included	Not Included	Included	Included	Included
Testing for Bias	Included	Included	Included	Included	Included
Third-Party Testing	Not Included	Included	Included		
Stakeholder Engagements	Not Included	Included	Included	Included	Included
Public Disclosure of Impact Assessment	Not Included	Included	Not included except for	Not Included	Included

[25] *See* Algorithmic Accountability Act of 2022, H.R. 6580, 117th Cong. (2022).
[26] Algorithmic Accountability Act, §6.
[27] *See* Part I.a *supra*.
[28] *See* Part I.a *supra*.

			partial disclosure of summary reports		

c. EPIC's recommendations

In EPIC's 2022 comments to the FTC concerning the Commission's commercial surveillance rulemaking,[29] we discussed at length what information should be included in impact assessments concerning a business's personal data processing activities and use of automated decision-making systems. Building on a proposed list of elements suggested by the FTC, EPIC recommended that impact assessments required under a trade rule include:

- The data [companies] use;
- How they collect, retain, disclose, or transfer that data;
- How they choose to implement any given automated decision-making system or process to analyze or process the data, including the consideration of alternative methods;
- How they process or use that data to reach a decision;
- Whether they rely on a third-party vendor to make such decisions;
- The impacts of their commercial surveillance practices, including disparities or other distributional outcomes among consumers;
- Risk mitigation measures to address potential consumer harms[;] …
- The purpose(s) for which the company will collect, process, retain, or make available to third parties each category of personal data;
- The sources of the personal data the company will collect, process, retain, or make available to third parties;
- Which third parties and service providers, if any, the company will make personal data available to;
- What notice or opportunities for consent will be provided to consumers concerning the company's collection, processing, or retention of their personal data or the making available of such information to third parties;
- The potential harms that might result from such processing, including but not limited to privacy, physical, economic, psychological, autonomy, and discrimination harms;
- The company's asserted need to engage in such collection, processing, retention, or transfer of personal information;
- Any alternatives to such collection, processing, retention, or transfer of personal information seriously considered by the company and the reason(s) why such alternatives were rejected;

[29] Comments of EPIC to the FTC Proposed Trade Regulation Rule on Commercial Surveillance & Data Security 7 (Nov. 2022), https://epic.org/wp-content/uploads/2022/12/EPIC-FTC-commercial-surveillance-ANPRM-comments-Nov2022.pdf [hereinafter EPIC FTC Comments on Commercial Surveillance]; *see also* Comments of EPIC to Cal. Privacy Prot. Agency 35–37 (Mar. 23, 2023), https://epic.org/wp-content/uploads/2023/03/EPIC-et-al-comments-CCPA-rulemaking-March-2023-2.pdf.

- How the asserted benefits resulting from such collection, processing, retention, or transfer to the company, the consumer, other stakeholders, and the public compare to the risks to the consumer; and
- A plain language summary of the assessment that would be comprehensible to a reasonable consumer.[30]

With respect to each automated decision-making system used to make or inform determinations about individuals, EPIC recommended that businesses be required to disclose the following:

1. A detailed description of the intended purpose and proposed use of the system, including:
    a. What decision(s) the system will make or support;
    b. Whether the system makes final decision(s) itself or whether and how supports decision(s);
    c. The system's intended benefits and research that demonstrates such benefits;
2. A detailed description of the system's capabilities, including capabilities outside of the scope of its intended use and when the system should not be used;
3. An assessment of the relative benefits and costs to the consumer given the system's purpose, capabilities, and probable use cases;
4. The inputs and logic of the system;
5. Data use and generation information, including:
    a. How the data relied on by the system is populated, collected, and processed;
    b. The type(s) data the system is programmed to generate;
    c. Whether the outputs generated by the system are used downstream for any purpose not already articulated;
6. Yearly validation studies and audits of accuracy, bias, and disparate impact; and
7. A detailed use and data management policy.[31]

Among other benefits to the public, requiring these disclosures will help narrow the use of automated decision-making systems to circumstances in which they are genuinely necessary and appropriate and ensure that businesses restrict their use of automated decision-making systems to the purposes for which they are designed, evaluated, and advertised.

Importantly, the NTIA's recommendations should focus on "disparate impact," regardless of intent. The term "bias" may be understood to include an intent element, and disclosures or audits that focus solely on this term may be less helpful in establishing whether violations of the Civil Rights Act of 1964, the Fair Housing Act, and the Age Discrimination in Employment Act have occurred.[32]

Detailed audit requirements will help remedy some of the existing issues with audits. Ari Ezra Waldman has contended that today's "[a]mbiguous privacy rules . . . with process-oriented

---

[30] EPIC FTC Comments on Commercial Surveillance at 163–64.
[31] EPIC FTC Comments on Commercial Surveillance at 84–85.
[32] 42 U.S.C. §§ 2000e-2(a)–(b); 42 U.S.C. § 3601 *et seq.*; 29 U.S.C. § 621 *et seq.*

regulatory levers open the door for companies to reframe the law in ways that serve corporate, rather than consumer, interests."[33] As a result, the compliance ecosystem has often become merely symbolic: it allows companies to interpret its legal requirements and implement process-oriented compliance structures—such as risk assessments and privacy policies—that shield them from liability or mitigate corporate risk.[34] To remedy the failures of these structures, the NTIA should recommend an auditing standard that frames institutional obligations in terms of substantive anti-discrimination protections, rather than procedural measures, to fulfill the goals of anti-discrimination law.[35] Such standard should require independent, third-party investigations and reports.

Regulating upstream actors will not only effectively addresses the problem before it results in discrimination, but also "any remedial actions taken by the vendor would cascade down to all its clients."[36] Laws and regulations should impose the burden of the addressing the wrongdoing on the entities most capable of doing it—that is, the creators and users of such products, rather than the millions of individuals exposed to their harmful effects.

II. Responses to NTIA Questions

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

The purpose of mechanisms like certifications, audits, and assessments is to provide a uniform structure to ensure the transparency and accountability of AI systems. Together, these mechanisms benefit consumers, companies looking to integrate AI systems in their workflow, journalists, and enforcement agencies, among others. Implemented correctly, accountability mechanisms should force both developers and users to consider and disclose key aspects of their AI systems that impacts those affected by the system.

11. What lessons can be learned from accountability processes and policies in cybersecurity, privacy, finance, or other areas

[33] Ari Ezra Waldman, *Privacy Law's False Promise*, 97 Wash. U. L. Rev. 773, 792 (2020).
[34] *Id.* at 796–97, 799.
[35] *See id.* at 803.
[36] Andrew D. Selbst & Solon Barocas, Unfair Artificial Intelligence: How FTC Intervention Can Overcome the Limitations of Discrimination Law, U. Penn. L. Rev. 7, 12 (forthcoming 2023).

EPIC recommends *Data & Society's* piece on Algorithmic Impact Assessments, which features discussion and breakdown of impact assessments used in different contexts on pages 8-11.[37]

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?

There should be AI accountability measures should be mandatory for the party ultimately responsible for its deployment reqardless of the types of data used and the sensitivity of the context. For example, in an employment context, the employer may hire a third-party hiring tool that is biased, but they impacted the client in this choice and deference of the system. This leaves less gaps in accountability and may lead to specific contract provisions to assign responsibility for accountability mechanisms.

18. Should AI systems be released with quality assurance certifications, especially if they are higher risk?

EPIC cautions against the use of quality assurance certifications that may lead consumers to conclude that an AI system is risk-free. An AI system should not be deployed without rigorous, independent, and regular testing by an independent source establishing that the system works as intended, is accurate, and is nondiscriminatory. However, as the nature of AI systems means that they are often susceptible to frequent changes and unanticipated outputs, it is important that such testing be regarded only as a be—not a clean bill of health.

19. As governments at all levels increase their use of AI systems, what should the public expect in terms of audits and assessments of AI systems deployed as part of public programs? Should the accountability practices for AI systems deployed in the public sector differ from those used for private sector AI? How can government procurement practices help create a productive AI accountability ecosystem?
30.d. What accountability practices should government (at any level) itself mandate for the AI systems the government uses?

Government agencies at all levels should be obligated to adopt robust accountability measures for any AI system they develop or use. Firstly, these systems are funded by taxpayer money and are used to make or influence decisions that people cannot avoid. At minimum, agencies should be required to disclose, through yearly algorithmic audits **and** impact assessments the details that EPIC recommended in Section 1.C of this comment.

---

[37] Emmanuel Moss et al., *Assembling Accountability: Algorithmic Impact Assessment for the Public Interest*, Data & Society (June 21, 2022) at 8-11https://datasociety.net/wp-content/uploads/2021/06/Assembling-Accountability.pdf; Jacob Metcalf, Emanuel Moss, Elizabeth Anne Watkins, Ranjit Singh, & Madeleine Clare Elish, Algorithmic Impact Assessments and Accountability: The Co-construction of Impacts, FAccT (Mar. 3, 2021)

25 & 26. Is the lack of a general federal data protection or privacy law (as well as a lack of a federal law focused on AI systems) a barrier to effective AI accountability?

Yes, because there are very few restrictions on sensitive data collection, AI systems are able to be built on more data and be an excuse to collect more data without any consequences. Still, algorithms built on data that is acquired through unfair or deceptive means may lead to harsh enforcement, and there is nothing stopping enforcement agencies to use rulemaking and enforcement mechanisms they already have.

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

It should set requirements through legislation, rulemaking, and enforcement. In order to bridge the gap between large entities and small entities for compliance costs, the federal government should provide training and auditing resources including but not limited to staff. In it's own use, it should lead by example.

31. What specific activities should government fund to advance a strong AI accountability ecosystem?

Government should designate funding for independent auditors and expert technologists that can more deftly do enforcement, and proactively audit automated systems used by government and small-medium sized businesses.

EPIC will follow up with your office with a whitepaper in the coming months that details these recommendations.

33. How can government work with the private sector to incentivize the best documentation practices?

Require it. If anything is even touching federal dollars, they should be required to document widely. Government can also incentivize good documentation practices by rigorously enforcing unfair and deceptive trade practice laws and antidiscrimination laws, where documentation may improve the precision and proportionality of their penalty.

III. Conclusion

EPIC supports NTIA's inquiry into AI accountability measures and remains eager to assist with efforts to integrate findings into law and common practice.

Respectfully submitted,

/s/ *Ben Winters*
Ben Winters
Senior Counsel

/s/ *Kabbas Azhar*
Kabbas Azhar
IPIOP Law Clerk