Comments of Consumer Reports
In Response to the
National Telecommunications and Information Administration
Request for Comment On
AI Accountability

By

Matt Schwartz, Policy Analyst
Justin Brookman, Director of Technology Policy

June 12, 2023

Consumer Reports[1] appreciates the opportunity to provide feedback on the National Telecommunication and Information Administration's (NTIA) Request for Comment on Artificial Intelligence (AI) Accountability. We thank NTIA for initiating this proceeding and for its other efforts to investigate and recommend policy changes to reduce potentially harmful data practices.

Below, we respond to a selection of the questions posted in the Request for Comment, describing our views on AI accountability mechanisms in detail.

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

In our view, accountability mechanisms can and should serve as one prong within a greater, holistic approach to regulating AI that includes at least two other prongs. AI regulation should include the following principles, none of which, on their own, are sufficient to address the wide range of use-cases and potential harms posed by AI systems:

1. **Systemic Regulation and Risk Management.** Some AI use-cases may consistently produce harms too severe or by their very nature represent risks too great to be adequately addressed by Individual Due Process Rights or accountability mechanisms. AI regulators may want to limit, place moratoria on, or ban such uses of AI outright.

   Regulators in certain jurisdictions in the United States and abroad have taken such a stance. For example, the European Parliament recently proposed including AI technologies like social scoring, predictive policing, and remote biometric identification for law enforcement in the "unacceptable risk" category within their draft AI Act, which would outlaw those uses completely.[2] Unacceptable technologies under the AI Act are those that "contradict Union values of respect for human dignity, freedom, equality, democracy and the rule of law and Union fundamental rights, including the right to non-discrimination, data protection and privacy and the rights of the child."[3]

---

[1] Consumer Reports is an independent, nonprofit membership organization that works side by side with consumers to create a fairer, safer, and healthier world. For over 80 years, CR has provided evidence-based product testing and ratings, rigorous research, hard-hitting investigative journalism, public education, and steadfast policy action on behalf of consumers' interests, including their interest in securing effective privacy protections. Unconstrained by advertising, CR has exposed landmark public health and safety issues and strives to be a catalyst for pro-consumer changes in the marketplace. From championing responsible auto safety standards, to winning food and water protections, to enhancing healthcare quality, to fighting back against predatory lenders in the financial markets, Consumer Reports has always been on the front lines, raising the voices of consumers.

[2] Proposal for a regulation of the European Parliament and of the Council on harmonised rules on Artificial Intelligence (Artificial Intelligence Act) and amending certain Union Legislative Acts, Title II - "Prohibited AI Practices," (March 16, 2023), https://www.europarl.europa.eu/resources/library/media/20230516RES90302/20230516RES90302.pdf

[3] *Id.*, Recital 15.

In the United States, the idea that certain AI technologies may be too dangerous or prone to abuse has gained traction in the last few years. At least seventeen localities have banned facial recognition for law enforcement and federal legislators have proposed a moratoria for the same purpose.[4]

2. **Accountability Mechanisms**. Provided they are backstopped by sufficient oversight and enforcement by regulators and the courts, their results made publicly available, and provide a route for *individuals* to hold businesses accountable in addition to regulators, AI accountability mechanisms can provide systemic protections that individual rights cannot provide on their own and where bans are too blunt an instrument. Accountability mechanisms also have the potential to reduce the risk of harm *before* it proliferates, whereas individual rights typically serve as recourse to remedy or avoid an extant harm.

   Accountability mechanisms should **require** developers and deployers of high-risk AI to subject their products to rigorous, independent third-party testing to ensure they meet certain benchmarks for accuracy, efficacy, fairness, privacy, civil rights, transparency, explainability and internal governance. Policymakers should also explicitly empower public interest researchers to conduct additional oversight,[5] as such researchers play a complementary role to third-party auditors and assessors (who might not always be disinterested stakeholders)[6] by communicating results in a manner more amenable to public consumption or by providing an additional layer of independence.

   In our view, AI accountability mechanisms must encompass internal audits and assessments *and* external, adversarial assessments and audits. Internal audits and assessments can be a first step towards better governance that allow entities to flag shortcomings in any of the aforementioned evaluation categories at an earlier stage than may be possible otherwise. However, external assessments and audits are necessary to independently evaluate the technology and communicate findings directly with regulators, especially when the shortcomings are uncovered.

   If accountability mechanisms can meaningfully play a role in broader AI regulatory efforts, they must be situated within an enforcement regime tethered to clear legal standards, while allowing auditors to deeply engage with the technology. Otherwise, as researchers in the algorithmic accountability community have pointed out, even well-intentioned and resourced audits can fail to meet the standard necessary to reform

---

[4] Electronic Frontier Foundation, “Bans, bills, and moratoria,” https://www.eff.org/aboutface/bans-bills-and-moratoria#main-content

[5] See Attachment 1, Nandita Sampath, “Opening Black Boxes: Examining Legal Barriers to Public Interest Algorithmic Audting,” Consumer Reports, (October 2022), https://innovation.consumerreports.org/wp-content/uploads/2022/10/CR_Algorithmic_Auditing_Final_10_2022VF2.pdf

[6] See, e.g., Sasha Costanza-Chock, Inioluwa Deborah Raji, and Joy Buolamwini, “Who Audits the Auditors? Recommendations from a Field Scan of the Algorithmic Auditing Ecosystem,” FAccT ’22: 2022 ACM Conference on Fairness, Accountability, and Transparency, (June 2022), https://dl.acm.org/doi/fullHtml/10.1145/3531146.3533213

behavior, particularly when the scope of inquiry is too large,[7] standards ill-defined, or harms under contention.[8]

Another potential barrier to effective accountability mechanisms is regulator capacity for oversight. Consumer Reports has previously argued that regulators need substantially more resources to enforce the law and keep pace with dominant tech firms that can pay generous salaries that lure experts away from the public sector or otherwise surpass regulators in technical expertise.[9] This problem will only worsen if agencies become responsible for ensuring compliance with substantial AI accountability mandates without commensurate resource increases. Some of the key agencies responsible for enforcing existing consumer protection laws lag in manpower behind the rapid growth and sophistication of technology. As a startling example, the Federal Trade Commission (FTC) currently employs fewer people than it did in 1979.[10]

3. **Individual Due Process Rights**. Individuals subjected to AI technology deserve a slate of rights that give them more autonomy and provide them recourse when they believe they have been wronged by the technology. Such rights should include a right to transparency (including notice that the technology is being used in the first place), a right to explainability (information about how and why the AI produced a given result), a right to opt-out of being subjected to the AI (in certain circumstances), a right to contest the outcome of the AI (leveraging the transparency and explainability rights mentioned above), and the right to request a human decision-maker when possible and appropriate. Several of these rights are reflected in the White House's Blueprint for an AI Bill of Rights.[11]

   Individual Due Process Rights are especially important when the technology is being used to make decisions that produce legal or other similarly significant effects (such as access to housing, credit, education, other other important life opportunities), decisions that traditionally have been made by humans and thus provided more avenues for procedural challenge. These rights could be granted through federal AI or comprehensive privacy laws or executive rulemaking.

---

[7] Cathy O'Neil, "Facebook's Algorithms Are Too Big to Fix," Bloomberg, (October 8, 2021), https://www.bloomberg.com/opinion/articles/2021-10-08/facebook-s-algorithms-are-too-big-to-fix#xj4y7vzkg

[8] Megan Gray, "Understanding and Improving Privacy 'Audits' Under FTC Orders," (April 18, 2018), http://dx.doi.org/10.2139/ssrn.3165143

[9] Consumer Reports, Group Letter in Support of FTC Privacy Funding, (September 2021), https://advocacy.consumerreports.org/wp-content/uploads/2021/09/Group-letter-in-support-of-FTC-privacy-funding.pdf

[10] Federal Trade Commission, FY 2023 Congressional Budget Justification, (March 2022), https://www.ftc.gov/system/files/ftc_gov/pdf/P859900FY23CBJ.pdf

[11] The White House, "Blueprint for an AI Bill of Rights," (October 2022), https://www.whitehouse.gov/wp-content/uploads/2022/10/Blueprint-for-an-AI-Bill-of-Rights.pdf

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?

Certifications, audits, and assessments can *only* promote trust for external stakeholders if companies are required to change internal processes as a result of identifying deficiencies during such reviews. Verifying that companies using AI tools have merely thought about various harms they might create is wholly insufficient. The reality is that even when tech companies fully recognize the harms their services cause, they often do not act to countervail them; Frances Haugen's revelations regarding Facebook's lack of action in the face of multiple known harms created by the platform provides the most high-profile example.[12] The more intertwined a business' revenue model is with the harms they produce, the less likely risk assessments are to change behavior voluntarily.

This means that policy design must take into account the ability for regulators to leverage their oversight and enforcement capabilities and for the public to play a role in holding businesses accountable. Evaluations of a company's AI systems for accuracy, efficacy, bias, fairness, privacy, civil rights, transparency, and explainability and internal governance must tie back to concrete standards grounded in how the technology is actually used in real life. Any new accountability requirements created in law should preserve the right of individuals to seek redress in court when they believe that the business has failed to meet standards, lied about the performance of its AI tools, or otherwise misled users.

3. AI accountability measures have been proposed in connection with many different goals, including those listed below. To what extent are there tradeoffs among these goals? To what extent can these inquiries be conducted by a single team or instrument?

a. The AI system does not substantially contribute to harmful discrimination against people.

b. The AI system does not substantially contribute to harmful misinformation, disinformation, and other forms of distortion and content-related harms.

c. The AI system protects privacy.

d. The AI system is legal, safe, and effective.

e. There has been adequate transparency and explanation to affected people about the uses, capabilities, and limitations of the AI system.

f. There are adequate human alternatives, consideration, and fallbacks in place throughout the AI system lifecycle.

g. There has been adequate consultation with, and there are adequate means of contestation and redress for, individuals affected by AI system outputs.

---

[12] Wall Street Journal, The Facebook Files, (October 1, 2021), https://www.wsj.com/articles/the-facebook-files-11631713039

h. There is adequate management within the entity deploying the AI system such that there are clear lines of responsibility and appropriate skillsets.

It is probable that, at times, there will be tradeoffs between some of the above goals, though many of them reinforce each other. For example, one can imagine a scenario where requesting the human review of an algorithmically generated decision relating to a credit application may reduce privacy, as a human would then need to manually review potentially sensitive financial information, entering into what was previously an automated process. Indeed, conducting an exhaustive and ongoing assessment evaluating progress toward each of the above goals may involve a tradeoff in and of itself, if documenting company processes and actions taken toward improving a system actually becomes an impediment to making that change.

However, the current marketplace creates very little incentive for businesses to voluntarily subject themselves to assessments or audits relating to *any* of these goals, even as documented harms proliferate and firms face repeated negative press. For a variety of reasons, not least of the presence of several dominant firms in the AI field and anti-competitive trends toward consolidation, this is not a problem that the market is solving on its own.

In some cases, tradeoffs can be managed by regulators adopting a risk-based approach – that is, when multiple harms exist in tension with each other, they will need to assess which poses the greatest risk to society and work to mitigate those. The EU Parliament has done so, in part, by creating a list of high-risk AI systems subjected to enhanced regulation in its draft AI Act.[13] U.S. regulators could further adapt this approach by listing discrete *types of harms* produced by AI systems and evaluating the relative risks of each to make it easier to resolve tensions when they arise. In other cases, providing individuals with the autonomy to assess tradeoffs themselves can resolve the issue. In the example given above, the individual has the *option* to request a human review and can determine whether the reduction in privacy is worthwhile to them.

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?

In some contexts, especially when evaluating the efficacy or accuracy of a tool, AI accountability mechanisms can conceivably reduce or eliminate collective risks. For example, if the risk is that a facial recognition algorithm may discriminate against a given class of people solely due to systemic inaccuracy when evaluating that class and the accountability mechanism *requires* that the developer meet a certain standard for accuracy before deploying the tool, then it is plausible that the risk could be successfully mitigated. However, it is important to note that harms like bias

[13] Proposal for a regulation of the European Parliament and of the Council on harmonised rules on Artificial Intelligence (Artificial Intelligence Act) and amending certain Union Legislative Acts, Title II - "Prohibited AI Practices," (March 16, 2023), https://www.europarl.europa.eu/resources/library/media/20230516RES90302/20230516RES90302.pdf

more often than not involve sociotechnical factors (i.e., how and why the AI tool is being used in the first place and in certain contexts), rather than straightforward technical deficiencies.

In other contexts, for example when the very existence of an AI tool threatens the autonomy, privacy, or civil rights of individuals, no accountability mechanism can effectively manage the risk. In such cases, regulators may need to ban or pause deployment of a given tool to effectively manage the collective risk. As mentioned above (*supra,* Question 1) members of the European Parliament have recently taken such a step by banning certain technologies, including the use of remote biometric identification tools for use by law enforcement, as part of their draft AI Act. The text of the act justifies this decision by reasoning that threats such as the "feeling of constant surveillance" could "indirectly dissuade the exercise of the freedom of assembly and other fundamental rights," which in their view outweighs any benefits possibly conferred by that technology.[14] U.S. regulators should be prepared to make similar decisions when the risk of a given technology is simply too high.

5. Given the likely integration of generative AI tools such as large language models (e.g., ChatGPT) or other general-purpose AI or foundational models into downstream products, how can AI accountability mechanisms inform people about how such tools are operating and/or whether the tools comply with standards for trustworthy AI?

Generative AI tools present a complex and rapidly evolving use-case (or indeed, series of use-cases), as their full potential is not yet known. Even so, a tool like ChatGPT, released to the public only months ago, has already generated a long list of stories that demonstrate its possible harms, a list that seemingly grows by the day.[15] Two particular risks produced by ChatGPT that are already clear, misinformation and fraud, seem to call for accountability mechanisms relating to transparency, explainability, and accuracy. Some stakeholders have recommended that generative AI or other tools that produce synthetic media carry disclaimers that their outputs were algorithmically generated, which could help provide transparency as downstream uses of the technology proliferate around the internet.[16] Others have suggested that generative systems should allocate financial benefits to rights holders when the AI system uses copyrighted material in its responses.[17]

Auditability and explainability appear to be especially problematic for ChatGPT, given that in many cases, it returns different results in response to the same prompt. Moreover, it is unclear

---

[14] *Id.*, Recital 18.
[15] See, e.g., Isaiah Poritz, OpenAI Hit With First Defamation Suit Over ChatGPT Hallucination, Bloomberg Law, (June 7, 2023), https://news.bloomberglaw.com/tech-and-telecom-law/openai-hit-with-first-defamation-suit-over-chatgpt-hallucination; Karen Weise and Cade Metz, When A.I. Chatbots Hallucinate, the New York Times, (May 1, 2023) https://www.nytimes.com/2023/05/01/business/ai-chatbots-hallucination.html
[16] Partnership on AI, PPAI's Responsible Practices for Synthetic Media A Framework for Collective Action, (February 23, 2023), https://partnershiponai.org/wp-content/uploads/2023/02/PAI_synthetic_media_framework.pdf
[17] Sercan Ozcan, Oleksandra Ozcan, and Joe Sekhon, "ChatGPT: what the law says about who owns the copyright of AI-generated content," The Conversation, (April 17, 2023), https://theconversation.com/chatgpt-what-the-law-says-about-who-owns-the-copyright-of-ai-generated-content-200597

what “accuracy” might mean in the context of a subjective inquiry or when a user requests a creative output.

All of this again highlights the need for a holistic approach to AI regulation; regulators may need to ultimately decide whether the risks of ChatGPT or other generative tools can be appropriately managed using accountability mechanisms and individual rights or whether other strategies are necessary. Having a risk-management framework in place would significantly aid those discussions.

6. The application of accountability measures (whether voluntary or regulatory) is more straightforward for some trustworthy AI goals than for others. With respect to which trustworthy AI goals are there existing requirements or standards? Are there any trustworthy AI goals that are not amenable to requirements or standards? How should accountability policies, whether governmental or non-governmental, treat these differences?

In most cases, there are no mandatory accountability mechanisms holding developers of AI systems to the goals of efficacy, fairness, privacy, notice and explanation or availability of human alternatives, as described in the Request for Comment. Existing examples of voluntary accountability efforts tend to reveal scattershot participation. Mechanisms like the National Institute of Standards and Technology (NIST)’s recently published 2023 Risk Management Framework are gaining momentum, but it is too early to tell what may come of industry efforts to adopt this approach.[18] Sectorally, for years top facial recognition vendors like Amazon and Clearview AI refused to submit their algorithms to NIST’s Face Recognition Test (Clearview recently acquiesced), which evaluates facial recognition and identification algorithms for accuracy and bias.[19] And while some companies do conduct general self-assessments for their U.S.-based operations (many of the larger companies are required to do Data Protection Impact Assessments in Europe), they often keep the results private. In the rare instances that businesses do make risk assessments publicly available, evidence of their efficacy is sketchy, if inconclusive. For example, Google recently released the results of its voluntary civil rights audit, which was roundly criticized by civil rights advocates for being performative and light on details.[20]

However, some existing law is flexible enough to provide guardrails to protect against harmful application of AI tools in some instances. For example, a recent FTC blog clarifies that its Section 5 authority does apply to generative AI tools, such as ChatGPT, that produce

---

[18] National Institute of Standards and Technology (NIST), US Department of Commerce, “Artificial Intelligence Risk Management Framework (AI RMF 1.0),” (January 2023), https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf

[19] Jon Porter, “Federal study of top facial recognition algorithms finds ‘empirical evidence’ of bias,” The Verge, (December 20, 2019), https://www.theverge.com/2019/12/20/21031255/facial-recognition-algorithm-bias-gender-race-age-federal-nest-investigation-analysis-amazon

[20] Cristiano Lima, Google’s civil rights audit lacked teeth, advocates say, Washington Post (March 10, 2023), https://www.washingtonpost.com/politics/2023/03/10/googles-civil-rights-audit-lacked-teeth-advocates-say/

misinformation or abet fraud, saying "[t]he FTC Act's prohibition on deceptive or unfair conduct can apply if you make, sell, or use a tool that is effectively designed to deceive – even if that's not its intended or sole purpose."[21] In addition, existing civil rights laws, such as the Equal Credit Opportunity Act or Fair Housing Act, can use an ex-poste disparate impact analysis to determine whether an algorithm produced illegally discriminatory results, though they are seemingly more limited in preventing discrimination before it occurs.

Generalized risk assessments are becoming a more common requirement, especially via state privacy laws, following the passage of the General Data Protection Regulation (GDPR) and Digital Services Act (DSA) in the European Union. Of the eleven comprehensive state privacy laws, nine include some sort of requirement for covered entities to conduct data protection assessments regarding certain processing activities (usually those that pose a "heightened risk of harm," which often includes automated decisionmaking or other specific risks to consumers.)[22] Drawing from the text of GDPR, these risk-assessments typically require that businesses weigh the benefits of processing to all relevant stakeholders against the potential risks to the rights of the consumer associated with such processing.[23] Other laws, such as New York City's employment AI law take a more targeted approach to proscribing unfair AI tools.[24]

In the financial sector, both the Equal Credit Opportunity Act (ECOA) and Fair Credit Reporting Act (FCRA) provide meaningful access and explainability mandates.[25] When a consumer is denied credit, under ECOA creditors must provide consumers with the main reasons for that denial. The Consumer Financial Protection Bureau recently clarified that creditors that use complex algorithms or artificial intelligence to help generate credit decisions must still "provide a notice that discloses the specific, principal reasons for taking adverse actions."[26] Meanwhile, FCRA requires that when an adverse action, such as the denial of credit, is based on a credit score, the creditor must disclose the key factors that affected the score, among other information.

Certain governance requirements, such as inserting a "human in the loop", have been required in select jurisdictions, though this approach likely requires more fine-tuning to produce

---

[21] Michael Atelson, "Chatbots, deepfakes, and voice clones: AI deception for sale," Federal Trade Commission, (March 20, 2023), https://www.ftc.gov/business-guidance/blog/2023/03/chatbots-deepfakes-voice-clones-ai-deception-sale
[22] See, e.g., Public Act No. 22-15, the Connecticut Data Privacy Act, Section 8, https://www.cga.ct.gov/2022/act/Pa/pdf/2022PA-00015-R00SB-00006-PA.PDF
[23] Ibid.
[24] New York City City Council, Local Law 144, (2021), https://legistar.council.nyc.gov/LegislationDetail.aspx?ID=4344524&GUID=B051915D-A9AC-451E-81F8-6596032FA3F9
[25] Patrice Alexander Ficklin, Tom Pahl, and Paul Watkins, Innovation spotlight: Providing adverse action notices when using AI/ML models, Consumer Financial Protection Bureau, (July 7, 2020), https://www.consumerfinance.gov/about-us/blog/innovation-spotlight-providing-adverse-action-notices-when-using-ai-ml-models/
[26] Consumer Financial Protection Bureau, CFPB Acts to Protect the Public from Black-Box Credit Models Using Complex Algorithms (May 26, 2022), https://www.consumerfinance.gov/about-us/newsroom/cfpb-acts-to-protect-the-public-from-black-box-credit-models-using-complex-algorithms/

meaningful results. For example, the Colorado Privacy Act, drawing from GDPR, relieves controllers of the requirement to allow consumers to opt out of automated decisionmaking when a human is "involved" with an automated process. However, a growing corpus of scholarship has found that humans, even those technically empowered to intervene in automated processes, often cannot do so effectively.[27] This can occur for a multitude of reasons, but perhaps most vexingly of all is the "black box" problem, where a human may indeed consider the data used in the processing and have the authority to change a result once the processing occurs, but simply possesses no understanding of how the automated process arrived at the conclusion that it did. This problem plagues even the most technically-attuned humans in the loop, including engineers of the systems themselves, and will only worsen as automated processes become more sophisticated.[28]

Future regulation of AI should provide clear standards of behavior to meet society's goals for trustworthy AI. Of course, some goals, especially those related to technical efficacy and accuracy may be more straightforward in certain contexts than others. Goals related to more abstract values like fairness or privacy also might prove difficult to benchmark for, though legislators can pull from a long history of civil rights and privacy policymaking to help guide them. Again, as accountability mechanisms meet their limits, lawmakers should have a broader toolbox of policy responses to draw from.

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?

Possible limits regarding what systemic oversight AI accountability mechanisms can achieve are discussed above (*supra,* Questions 1, 2, 3, and 4). We do not believe AI accountability mechanisms, on their own, are capable of ensuring the development of trustworthy AI.

8. What are the best definitions of and relationships between AI accountability, assurance, assessments, audits, and other relevant terms?

As discussed above (*supra,* Questions 1 and 2), Consumer Reports strongly believes that AI accountability and assessments can only be effective if they are backstopped by strong oversight and enforcement provisions. As such, we urge future policymakers to scope definitions of AI accountability, assurance, assessments, audits such that those requirements are embedded in our understanding of those terms.

9. 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,

---

[27] See, e.g., Brennan-Marquez, Kiel and Susser, Daniel and Levy, Karen, Strange Loops: Apparent versus Actual Human Involvement in Automated Decision-Making (October 2, 2019). 34 Berkeley Technology Law Journal 745–771 (2019), https://ssrn.com/abstract=3462901

[28] Will Knight, The Dark Secret at the Heart of AI, MIT Technology Review, (April 11, 2017), https://www.technologyreview.com/2017/04/11/5113/the-dark-secret-at-the-heart-of-ai/

or policies are most appropriate for implementation and operationalization at scale in the United States? Who are the people currently doing AI accountability work?

See above (*supra,* Question 6) for a discussion of existing accountability mechanisms tied to the trustworthy AI goals.

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

One major weakness in the existing risk-assessment framework as set out in recent state privacy laws is that the assessment must only be produced if the controller is being investigated by a supervisory authority. In fact, every state privacy law that includes a risk assessment requirement currently exempts them from public inspection. Unless a company's behavior is suspicious enough to warrant an Attorney General or privacy authority investigation, nobody outside of the business will ever even see the risk assessment. FTC privacy audits work in much the same way, with much of the key information redacted in the publicly accessible version of the document (if it is even published).[29]

Another weakness in state privacy laws stems from the lack of mandatory behavioral changes when shortcomings are identified in the risk assessment. That is, while businesses must confront the risks inherent to their data processing activities, the laws ultimately leave it up to them to decide whether the benefits of their services outweigh the (self-identified) harms. So unless the risk assessment identifies some sort of actionable violation of existing law, it does not contain any real power to reform business practices in order to avoid reported harms.

The California Privacy Rights Act's risk assessment does differ somewhat from other states, since it explicitly states that the goal is "restricting or prohibiting" processing if the risks outweigh the benefits, which seemingly provides regulators an additional hook to compel behavioral changes. However, without strict oversight, businesses are likely to simply downplay the risks in order to avoid any affirmative requirement to change. At the very least, laws creating risk assessment requirements should mandate that businesses provide risk assessments to the relevant supervisory authority on an ongoing basis – rather than only when the business is being investigated – so that they may review for systemic underreporting or other obvious instances of noncompliance. In any case, proving a business outright lied on its risk assessment will likely be a difficult endeavor.

At the same time, we recognize that state Attorneys General or even dedicated supervisory authorities like the California Privacy Protection Agency do not possess the resources to closely and continually monitor risk assessments – and it is unlikely that even a larger federal regulator would either. For this reason, we believe it is crucial that the public also be able to review risk assessments (with tightly scoped exemptions around revealing business trade secrets), so that interested consumers can use this information to weigh their engagement with businesses. Public inspection of risk assessments will also deputize individuals by allowing them to relay

[29] Megan Gray, "Understanding and Improving Privacy 'Audits' Under FTC Orders," (April 18, 2018), http://dx.doi.org/10.2139/ssrn.3165143

important information back to regulators that the regulator may not have uncovered on its own, as well as file their own private claims against a business. While few people will likely read risk assessments and businesses will still be incentivized to emphasize the benefits of its processing and minimize the risks, more documentation is probably better than nothing at all. Risk assessment requirements should also mandate that businesses share any internal documentation they possess on the concrete harms caused by the service to avoid large-scale coverups, such as those revealed by Frances Haugen at Facebook.[30]

As with all audits and assessments, effectiveness comes down to enforcement, both private and regulatory. If businesses fear the consequences of not being forthcoming, risk assessments could produce additional information that improves regulators' and the public's understanding of the processing harms caused by businesses. If left to their own devices, businesses will likely produce anodyne documents that serve few and the process will simply become a "check the box" exercise.

14. Which non-U.S. or U.S. (federal, state, or local) laws and regulations already requiring an AI audit, assessment, or other accountability mechanism are most useful and why? Which are least useful and why?

See above (*supra,* Questions 6 and 13) for our views on existing accountability mechanisms at the state level.

In the non-U.S. context, Europe's GDPR provides the most mature accountability regime to assess. However, since its risk assessment provisions also no do not require that covered entities publish the assessment for public consumption or that they remedy shortcomings identified in the assessment, GDPR suffers from many of the same issues as described above.[31]

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.

a. Where in the value chain should accountability efforts focus?

b. How can accountability efforts at different points in the value chain best be coordinated and communicated?

---

[30] Wall Street Journal, The Facebook Files, (October 1, 2021), https://www.wsj.com/articles/the-facebook-files-11631713039

[31] Margot A. Kaminski, "Binary Governance: Lessons from the GDPR's Approach to Algorithmic Accountability," Southern California Law Review, Vol. 92, No. 6, 2019 1529, (November 11, 2019), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3351404

c. How should vendors work with customers to perform AI audits and/or assessments? What is the role of audits or assessments in the commercial and/or public procurement process? Are there specific practices that would facilitate credible audits (e.g., liability waivers)?

d. Since the effects and performance of an AI system will depend on the context in which it is deployed, how can accountability measures accommodate unknowns about ultimate downstream implementation?

Accountability efforts should focus both on the developers, as well as the downstream deployers of AI tools. It may be appropriate to institute a "reasonable expectations" standard, whereby developers of AI tools carry some degree of liability for harms produced by downstream uses of their tool if they could have reasonably foreseen that such harms were likely to occur. In many other cases, it will be impossible for developers to foresee how their tools will be leveraged in the market. Developers that build or iterate on top of existing AI systems should also be accountable for their actions, especially when they create a novel tool far outside of the initial use-case. It will also be important that developers and downstream deployers transparently communicate to the general public the means of recourse and precise allocation of liability when using their service.

16. The lifecycle of any given AI system or component also presents distinct junctures for assessment, audit, and other measures. For example, in the case of bias, it has been shown that "[b]ias is prevalent in the assumptions about which data should be used, what AI models should be developed, where the AI system should be placed—or if AI is required at all." How should AI accountability mechanisms consider the AI lifecycle?

Consumer Reports believes that AI accountability mechanisms should combine persistent requirements that all developers must follow regardless of the maturity of their technology, along with other requirements that scale along with the risks. For instance, all developers of AI should be required to use unbiased training data from the outset. Many of the discriminatory effects produced by algorithms stem back to the initial issue of poor data, which could've been remedied by mandating better inputs from the beginning. It's also crucial that AI tools, especially those built on iterative machine learning algorithms, are designed with future evaluation, explainability, and governability in mind, otherwise these goals may become unattainable as the model grows in complexity.[32]

Again, a risk-based approach to determining when certain accountability mechanisms are necessary can help avoid untenably broad application of rules. In GDPR and many of the recent state privacy laws in the United States, full blown risk assessments are only required when the data processing constitutes a significant risk of harm. The same logic applies to the due process rights those laws create. For instance, rights relating to automated decisionmaking (such as the

[32] Deven R. Desai and Joshua A. Kroll, "Trust But Verify: A Guide to Algorithms and the Law," Georgia Tech Scheller College of Business Research Paper No. 17-19, (April 27, 2018), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2959472

right to opt out, the right to contest, or the right to a human intervention) only apply when the decision produces legal or similarly significant effects concerning the consumer.[33]

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?

See above (*supra,* Questions 3, 4, and 16) for our views on when accountability mechanisms should apply. Ultimately, the government should be responsible for calculating the risk and making the determination of which accountability mechanisms AI developers should be required to abide by. As previously mentioned, the EU AI Act provides a strong framework to begin thinking through a risk-based approach.

21. What are the obstacles to the flow of information necessary for AI accountability either within an organization or to outside examiners? What policies might ease researcher and other third-party access to inputs necessary to conduct AI audits or assessments?

See below (*infra,* Question 24) for our views on the limitations of third-party audits.

Currently, public interest audits are limited by imperfect access to algorithms and the underlying data in part because of existing laws designed to limit computer hacking and protect intellectual property. To help remove these obstacles, we have previously recommended several policy changes that would balance these legitimate values with the need for research and external accountability. This includes code and data set access and publication mandates, targeted reforms to laws such as DMCA and the Computer Fraud and Abuse Act, legislation that explicitly prohibits contractual language unfairly limiting researchers' ability to audit algorithms for bias, and updates to the FTC Act to clarify that businesses may not deceive third-party testers. More information can be found in our white paper on public interest algorithmic auditing, Appendix 1 of this comment.[34]

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?

As mentioned above (*supra,* Questions 1 and 13), we strongly recommend that audit results be presented publicly so that additional third-parties may review the results and provide an added layer of oversight. It may be necessary to require presentation of audit results in two forms, similar to SEC filings, one to provide detailed information to the most sophisticated audiences, including regulators, the other a short, accessible, and clear description of the results.

---

[33] See, e.g., Public Act No. 22-15, the Connecticut Data Privacy Act, Section 4, https://www.cga.ct.gov/2022/act/Pa/pdf/2022PA-00015-R00SB-00006-PA.PDF; General Data Protection Regulation, Article 22, https://gdpr-info.eu/art-22-gdpr/.

[34] See Attachment 1, Nandita Sampath, "Opening Black Boxes: Examining Legal Barriers to Public Interest Algorithmic Audting," Consumer Reports, (October 2022), https://innovation.consumerreports.org/wp-content/uploads/2022/10/CR_Algorithmic_Auditing_Final_10_2022VF2.pdf

Consumer Reports has recommended a similar framework for privacy notices.[35] While the bulk of consumers may ignore such information, a mandate to provide accountability results to the public will still serve as a meaningful check on companies who might otherwise prefer that questionable data processing go unnoticed.

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?

See above (*supra,* Question 21) for our view on impediments to public interest AI audits and *supra,* Questions 2, 3, and 6 for our view on the lack of incentive for companies to voluntarily subject themselves to AI accountability mechanisms. See above (*supra,* Question 11) for our view on the limitations of existing risk assessment and auditing frameworks.

Any future success of AI accountability mechanisms like risk assessments or audits rests on the inclusion of appropriate granularity in their design. Commentators have pointed out that there are numerous ways an audit, even one conducted by an independent third-party, can fail to deliver the promised results if insufficient attention is paid to format or too much deference granted to the entity under audit.[36] Possible issues include reliance on management "attestations" to verify compliance with a given benchmark, allowing businesses to determine which risks to document and mitigation strategies to adopt, and overly narrow overall scope of inquiry. Of course, properly overseeing business' compliance with detailed audits or risk assessments will require vast increases in resources and technical expertise of regulators, as discussed previously.

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

A federal privacy law can accomplish certain objectives that could complement and improve the efficacy of AI accountability mechanisms. A federal privacy law that includes strong data minimization provisions, as Consumer Reports has called for and is included in our own Model State Privacy Act, would make AI accountability easier by limiting the amount of data that entities could use for purposes not reasonably anticipated by the consumer using a given product or service.[37] More controversial uses of AI, such as emotion recognition software, may struggle to develop without unfettered access to sensitive data inputs users are otherwise unlikely to grant special access to.

---

[35] Consumer Reports, Model State Privacy Act, (Feb. 2021) https://advocacy.consumerreports.org/wp-content/uploads/2021/02/CR_Model-State-Privacy-Act_022321_vf.pdf

[36] Megan Gray, "Understanding and Improving Privacy 'Audits' Under FTC Orders," (April 18, 2018), http://dx.doi.org/10.2139/ssrn.3165143

[37] Consumer Reports, Model State Privacy Act, (Feb. 2021) https://advocacy.consumerreports.org/wp-content/uploads/2021/02/CR_Model-State-Privacy-Act_022321_vf.pdf

A privacy law would also likely help improve business documentation, reporting, and governance practices, since businesses would be forced to better track data flows in order to honor consumer requests for deletion or correction of their data. Privacy law can also be a vessel for directly creating AI accountability mechanisms, as has been done on the state level with the inclusion of risk assessment requirements.

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

Certainly, a federal law focused on AI systems (ideally one that incorporates all three of the prongs discussed in Question 1) would seem to be the most direct way to implement new requirements for AI accountability mechanisms and would mimic the approach taking shape in Europe with the draft AI Act. If a federal privacy law does not take on the full scope of needed AI regulation (and existing proposals have not done so to-date), then a federal AI law could be necessary. However, as discussed above (*supra,* Question 6), existing laws can also be leveraged to address new harms wrought by advancements in AI.

27. What is the role of intellectual property rights, terms of service, contractual obligations, or other legal entitlements in fostering or impeding a robust AI accountability ecosystem? For example, do nondisclosure agreements or trade secret protections impede the assessment or audit of AI systems and processes? If so, what legal or policy developments are needed to ensure an effective accountability framework?

As discussed above (*supra,* Question 21), some existing laws could potentially impede fulsome outside assessment of AI systems. We've previously advocated for policies such as data set access and publication mandates, targeted reforms to laws such as DMCA and the Computer Fraud and Abuse Act, legislation that explicitly prohibits contractual language unfairly limiting researchers' ability to audit algorithms for bias, and updates to the FTC Act to clarify that businesses may not deceive third-party testers. More information can be found in our white paper on public interest algorithmic auditing.[38]

28. What do AI audits and assessments cost? Which entities should be expected to bear these costs? What are the possible consequences of AI accountability requirements that might impose significant costs on regulated entities? Are there ways to reduce these costs? What are the best ways to consider costs in relation to benefits?

The costs of AI audits and assessments, held to the standard we've proposed, are likely to be significant, especially for more sophisticated businesses. At the same time, sophisticated AI systems also pose significant risk and costs to society, almost none of which are currently internally realized. More than ten years ago, a FTC representative estimated that the agency's privacy audits of Facebook and Google would cost hundreds of thousands of dollars (the cost

[38] See Attachment 1, Nandita Sampath, "Opening Black Boxes: Examining Legal Barriers to Public Interest Algorithmic Audting," Consumer Reports, (October 2022), https://innovation.consumerreports.org/wp-content/uploads/2022/10/CR_Algorithmic_Auditing_Final_10_2022VF2.pdf

was to be paid by the businesses themselves).[39] More expansive audits of today's technologies could incur costs that would be substantially higher. Using a risk-based regulatory regime should help prevent less sophisticated businesses from taking on major and potentially competitively damaging costs, while focusing those costs on those most likely to be able to afford them.

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

As discussed throughout, Consumer Reports believes government should play a central role in the AI accountability ecosystem by setting standards for accountability mechanisms, leveraging its capacity for oversight and enforcement to ensure that businesses adhere to those standards, defining risk among different uses of AI, and potentially prohibiting certain uses of AI if it determines that the risks are simply too great to be mitigated by the combination of AI accountability mechanisms and Individual Due Process Rights.

That said, and as further elucidated below, the ability for government bodies to effectively manage such an enormous undertaking will be heavily influenced by how much funding they receive as part of such an effort. Any new accountability requirements created in law should also preserve the right of individuals to seek redress in court when they believe that the business has failed to meet standards, lied about its performance, or otherwise misled users in completing its accountability responsibilities. Again, this highlights the need to make the results of audits, assessments, bias testing, accuracy testing, and governance structures available for, and amenable to, public consumption. Given the scope and the likely future penetration of AI technologies into every aspect of the economy, individuals will need to be able to supplement government enforcement and retain their autonomy in that new world.

As the nation's primary consumer protection agency with the remit to regulate large tracts of the nation's commercial activity, the Federal Trade Commission appears to be the best suited government body to take responsibility for AI accountability regulation. Over the past several years, the agency has taken an increasing interest in issues around AI and algorithmic bias,[40] and its ongoing Commercial Surveillance rulemaking is heavily intertwined with the issues discussed here.[41] That said, and as previously mentioned (*supra,* Question 1) the agency lacks

---

[39] Kashmir Hill, "So, What Are These Privacy Audits That Google And Facebook Have To Do For The Next 20 Years?" Forbes, (November 30, 2011), https://www.forbes.com/sites/kashmirhill/2011/11/30/so-what-are-these-privacy-audits-that-google-and-facebook-have-to-do-for-the-next-20-years/?sh=230ea1445000

[40] See, e.g., FTC, "Combatting Online Harms Through Innovation," Report to Congress, (June 16, 2022), https://www.ftc.gov/system/files/ftc_gov/pdf/Combatting%20Online%20Harms%20Through%20Innovation%3B%20Federal%20Trade%20Commission%20Report%20to%20Congress.pdf; FTC, FTC Hearing #7: The Competition and Consumer Protection Issues of Algorithms, Artificial Intelligence, and Predictive Analytics, (November 13-14, 2018), https://www.ftc.gov/news-events/events/2018/11/ftc-hearing-7-competition-consumer-protection-issues-algorithms-artificial-intelligence-predictive

[41] FTC, Advanced Notice of Proposed Rulemaking on Commercial Surveillance and Data Security, (August 22, 2022), https://www.federalregister.gov/documents/2022/08/22/2022-17752/trade-regulation-rule-on-commercial-surveillance-and-data-security;

the necessary resources to keep pace with the technical advances made by large tech companies as it is, let alone the resources it would need to effectively oversee a federal AI law that institutes mandatory accountability mechanisms. Congress would need to couple any new AI oversight laws with a significant financial boost to the agency, otherwise it would risk creating a framework that allows AI developers to continue to proliferate risks, only under the guise of an accountability regime that exists on paper only. There are already signs this is the path big tech companies intend on leading policymakers down.[42]

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

See above (*supra*, Question 30). Aside from creating a strong initial framework, building the capacity for oversight and enforcement of any AI accountability ecosystem is the single most important thing the government can do to meaningfully ensure its success. In reality, enhanced expertise in AI needs to be fostered at every level and every branch of government, as all stakeholders will have a role in holding AI systems accountable; Congress so it can legislate, the Executive so it can implement and oversee, and the courts so they make fair and informed judgements.

34. Is it important that there be uniformity of AI accountability requirements and/or practices across the United States? Across global jurisdictions? If so, is it important only within a sector or across sectors? What is the best way to achieve it? Alternatively, is harmonization or interoperability sufficient and what is the best way to achieve that?

A strong standard for AI accountability that applies across the United States is obviously the preferable outcome, though we have supported incremental steps towards stronger accountability requirements in states in the absence of a unified federal approach. Globally, standards are only beginning to emerge, with the EU AI Act serving as the best example, even as it remains a work in progress. In that sense, there are ways the United States can still emerge as a leader in this space, and policymakers should not feel confined to copy other approaches simply in the name of interoperability if stronger alternatives are present.

****

We thank the National Telecommunications and Information Administration for its consideration of these points, and for its work to secure strong protections for consumers. We are happy to answer any questions you may have, and to discuss these issues in more detail. Please contact Matt Schwartz (matt.schwartz@consumer.org) or Justin Brookman (justin.brookman@consumer.org) for more information.

[42]Amba Kak and Sarah Myers West, 2023 Landscape: Confronting Tech Power, "Algorithmic Accountability: Moving Beyond Audits," AI Now Institute, (April 23, 2023), https://ainowinstitute.org/wp-content/uploads/2023/04/AI-Now-2023-Landscape-Report-FINAL.pdf

ATTACHMENT 1