Business Roundtable

1000 Maine Avenue, SW, STE 500
Washington, DC 20024
202.872.1260
brt.org

June 12, 2023

Stephanie Weiner
Acting Chief Counsel
National Telecommunications and Information Administration
U.S. Department of Commerce
1401 Constitution Avenue NW, Room 4725
Washington, DC 20230

Comments of Business Roundtable RE: "AI Accountability Policy Request for Comment"
Docket No. 230407-0093

Dear Ms. Weiner:

These comments are submitted on behalf of Business Roundtable, an association of more than 200 chief executive officers (CEOs) of America's leading companies representing every sector of the U.S. economy. Business Roundtable CEOs lead U.S.-based companies that support one in four American jobs and almost a quarter of U.S. GDP. We appreciate the opportunity to comment in response to the National Telecommunications and Information Administration's (NTIA) Request for Comment (RFC) on Artificial Intelligence (AI) Accountability Policy.

Introduction

Business Roundtable member companies across sectors—technology, communications, retail, financial services, health, public safety and security, defense, manufacturing, hospitality, insurance and others—rely on data and data-driven processes and solutions every day to create, deliver and improve innovative products and services across the United States and around the world. Rapid innovation and adoption of AI is transforming the nature of work across every industry and reshaping how people interact and experience the world around them. AI technologies not only help businesses deliver smarter products and services to their customers, but they also have enormous potential to drive broader positive change for Americans' health, safety and prosperity.

Our members are some of the world's largest developers, deployers and end users of AI, and accordingly they have a strong interest in ensuring that Responsible AI systems are developed and used in a manner that mitigates risks, encourages responsible innovation and earns consumer, government and public trust in order to fully realize AI's benefits for society and the economy.

Advancing safe and trustworthy AI is a shared responsibility among stakeholders, including between government and the private sector. The Commerce Department is well positioned to advance this important federal objective, building upon and furthering voluntary standards frameworks work such as the National Institute for Standards and Technology (NIST) AI Risk Management Framework (AI RMF). We encourage NTIA to carefully consider the following key principles for AI accountability:

- AI accountability standards and methods should be tailored to the nature and degree of risk and be use case- and context-specific. Policymakers should conduct a thorough gap assessment before developing any new rules and utilize existing legal and regulatory frameworks to the extent possible.
- Some features of Responsible AI may not lend themselves to technical standards due to the speed at which AI systems and tools evolve and the broad diversity of use cases and contexts in which AI is applied.
- Accountability mechanisms should be focused on the application of AI tools rather than imposing technology-specific standards that may be incompatible across various deployment contexts or that can quickly become outdated.
- Any AI accountability approaches should include clear definitions with case study examples informed by continued dialogue with industry stakeholders.

These overarching principles, along with our Responsible AI Roadmap and approach to AI accountability, are underpinned by three foundational tenets we believe apply to all AI accountability efforts: that AI systems are trusted and inclusive; effective, safe and secure; and underpinned by accountable governance.

There is no single approach or simple prescription for how to achieve the objectives of Responsible AI. Rather, each organization should endeavor to implement AI safeguards that are best suited to potential risks within the context of its own industry, use cases and existing regulatory frameworks.

Background on Business Roundtable's Approach to Responsible AI (RAI)

AI is transforming and revolutionizing businesses across the world and across all industries. Business Roundtable member companies recognize this transformative potential and understand that successfully managing accelerating AI adoption and deployment will require a deep demonstrated commitment to the safety, trustworthiness and responsible use of this technology.

Business Roundtable believes that responsible design, development, deployment and use of AI is critical to building consumer, government and public trust while advancing U.S. leadership in these emerging technologies and using them to the benefit of society and the economy. We recognize the importance of Responsible AI and the purposes it serves, such as protecting users

from negative outcomes (*e.g.*, fraud), combating bias, promoting equity and inclusion, and advancing national security. Business Roundtable companies are at the forefront of these issues and are leading efforts to create and implement principled Responsible AI governance, risk management, accountability and transparency. In January 2022, Business Roundtable worked with member companies to launch its Responsible AI Initiative with the publication of two foundational documents:

1. **Roadmap for Responsible AI**[1] (Roadmap)**,** which sets out ten principles to guide companies of all sizes, in every sector, and at every point along the AI value chain across three foundational characteristics of Responsible AI: (1) trusted and inclusive; (2) effective, safe and secure; and (3) accountable governance.
2. **Policy Recommendations**[2] for the U.S. government to encourage federal approaches to AI practices, rules and guidelines that build public trust in AI while enabling innovation and promoting continued U.S. leadership.

Building upon this work, in January 2023, Business Roundtable released a new case study report showcasing how leading U.S. companies are developing and deploying Responsible AI across industries and business contexts to bring value to consumers and help solve community and societal challenges such as fighting fraud in online payments and speeding up pandemic response.[3] *[See below under our response to Question 9 for examples]*

Business Roundtable welcomes the opportunity to provide its views on NTIA's AI Accountability Policy Request for Comment. Below, we provide specific responses to NTIA's question categories posed in its RFC.

Responses to Question Categories Posed in NTIA's RFC

I. AI Accountability Objectives

[Q1] What is the purpose of AI accountability mechanisms such as certifications, audits, and assessments?

---

[1] *See* Business Roundtable, Roadmap for Responsible AI (Jan. 26, 2022), https://www.businessroundtable.org/policy-perspectives/technology/ai. *See also* https://s3.amazonaws.com/brt.org/Business_Roundtable_Artificial_Intelligence_Roadmap_Jan2022_1.pdf

[2] *See* Business Roundtable, Policy Recommendations for Responsible Artificial Intelligence (January 26, 2022), https://www.businessroundtable.org/policy-perspectives/technology/ai. *See also* https://s3.amazonaws.com/brt.org/Business_Roundtable_Artificial_Intelligence_Policy_Recommendations_Jan2022_1.pdf

[3] *See* Business Roundtable, Business Roundtable Companies Put AI Recommendations into Practice (Jan. 25, 2023), https://www.businessroundtable.org/ai-innovation-at-work-putting-principles-into-practice.

AI is an increasingly powerful technology, with transformational benefits for consumers and corporations, and its responsible design, development, deployment and use is an important objective. Accountability mechanisms engender trust and a shared understanding that any uses of AI meet baseline standards for security, non-discrimination and other important metrics. Additionally, when designed and implemented thoughtfully, accountability mechanisms improve the performance of AI systems, creating synergy between responsible use and system or application performance and helping practitioners arrive at common understandings and language related to accountability measures.

Industry and government share the responsibility for furthering the trust necessary to support widespread AI adoption. As outlined in our Roadmap[4], businesses developing and using AI technologies should align their internal practices and governance to key principles of Responsible AI. Regulators and policymakers should enable continued innovation by taking an approach that considers the complex, context-dependent and rapidly evolving AI ecosystem.

Business Roundtable companies are committed to building trust in AI across the supply chain and working toward the beneficial potential of AI for society and the economy by adopting and advocating for Responsible AI. Accordingly, Business Roundtable companies are developing AI accountability mechanisms and procedures and integrating them into their internal governance processes. While there is substantial diversity in how companies operationalize AI accountability across different types of systems and use cases, accountability mechanisms and procedures are generally designed with the following important features:

- Specifically designed and tailored to achieve clearly articulated objectives. These may differ depending on the risk level or deployment context of various AI use cases and may include high level strategic guidance or more granular processes.
- Employ a human-centric approach in terms of recognizing that humans will be designing, implementing and evaluating any accountability measures.
- Account for the differences between fully automated AI tools and partially automated ones with a human in-the-loop.
- Have clear processes for monitoring and evaluation that support ongoing improvements in the quality, safety and performance of AI.
- Emphasize learning and information-sharing with other firms that leads to sharing and implementing best practices.

Any rules, guidelines, regulations or policies adopted by regulators or policymakers regarding AI accountability should be consistent and compatible with these important features. Business Roundtable believes that any public sector action on AI accountability must reflect the

---

[4] *See* Business Roundtable, Roadmap for Responsible AI (Jan. 26, 2022), https://www.businessroundtable.org/policy-perspectives/technology/ai. *See also* https://s3.amazonaws.com/brt.org/Business_Roundtable_Artificial_Intelligence_Roadmap_Jan2022_1.pdf

following: first, policymakers should incentivize, support and recognize good faith efforts on the part of industry to implement Responsible AI and encourage self-assessments by internal teams, in alignment with voluntary standards frameworks such as the NIST AI RMF.

[Q4] 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?

AI systems used in a responsible manner can help to mitigate risks and challenges in the workplace, government and society. In many cases, AI systems improve existing alternatives by reducing unconscious biases, improving safety and driving more equitable outcomes for consumers and communities. Effective AI accountability mechanisms can support these aims by addressing relevant risks of harms and driving expanded access to important services.[5]

Business Roundtable member companies will continue working to leverage Al to deliver beneficial products and services that improve the lives of consumers and address challenges within the communities where they operate and within society at large.

II. Existing Resources and Models

[Q9] 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?

Regarding the use of existing AI accountability mechanisms, Business Roundtable members employ—and are leading the development of—a range of approaches to advance Responsible AI. Companies are taking steps such as establishing cross-functional AI ethics and governance committees, conducting regular fairness and ethics assessments and performance monitoring, and refining existing internal risk management processes to incorporate Responsible AI principles. Some specific examples, reflected in our Roadmap actions and related case studies,[6] include:

---

[5] *See* Business Roundtable, Business Roundtable Companies Put AI Recommendations into Practice (Jan. 25, 2023), https://www.businessroundtable.org/ai-innovation-at-work-putting-principles-into-practice.

[6] *Ibid.*

- The SAS Viya platform supports a range of accountability mechanisms by providing clients with capabilities to assess data quality, detect and mitigate biases, assess model fairness, provide explainability, monitor system performance and support data privacy, all of which enables organizations to build employee and customer confidence in AI-supported systems.
- Booz Allen's aiSSEMBLETM – a lean manufacturing approach to AI engineering designed to simplify the engineering and deployment of AI systems – embeds RAI principles and practices (*e.g.*, transparency, traceability, auditability, dynamic data and model drift detection) into AI system design, execution and monitoring—ensuring that RAI is built-in from the start. Booz Allen is now working to integrate these capabilities with their recently announced investment in Credo.ai to further extend their RAI solution offerings in market.
- IBM's Ethics by Design (EbD) Framework integrates tech ethics into the organization's full tech development pipeline, including AI, and embeds responsible governance across the organization.
- The Responsible Use of AI @ General Motors initiative was developed through an iterative crowd-sourced process, drawing on extensive input from diverse stakeholders, engagement with General Motors employees and lessons learned from other organizations.

Business Roundtable member companies plan to continue implementing and iterating Al accountability frameworks to operationalize Responsible AI principles and practices in alignment with company values.

[Q11] What lessons can be learned from accountability processes and policies in cybersecurity, privacy, finance, or other areas?

As part of the Accountable Governance section of the Roadmap, organizations are encouraged to integrate AI practices into existing governance structures for risk management, data and compliance. Many Business Roundtable member companies already handle AI governance in this manner, applying existing data privacy, security and internal oversight policies to their AI applications. Before creating new systems specific to AI, organizations may evaluate whether existing company governance can be expanded to provide robust AI oversight and guidance, including responsiveness to stakeholder feedback on impacts and interactions with AI systems.

The U.S. government, industry and consumers all have a great deal of experience with many different forms of regulation across various industries and contexts. Many sectors—including financial services, healthcare, defense and telecommunications—are already heavily regulated using performance- and outcome-based standards that are often applicable to AI use cases (*e.g.*, through existing data protection frameworks). We encourage NTIA, before recommending any AI-specific accountability measures, to conduct a thorough gap assessment of existing regulatory regimes and approaches, and their applicability to AI use cases. Where

possible and practical, existing regulatory approaches should be adjusted or extended to account for AI rather than developing new rules specific to AI.

We encourage NTIA to avoid recommending onerous requirements that divert resources while not materially advancing the responsible development and uses of AI. Likewise, accountability models should incentivize the ongoing detection and mitigation of errors, bias and other risks ("find-and-fix") rather than penalize the discovery of such errors ("find-and-fine"), absent serious harm or intentional wrongdoing. It is key for organizations to periodically self-assess the efficacy of internal governance programs for Responsible AI, to assess gaps, understand what is state of the art and have the flexibility to improve and adapt, rather than be concerned that self-assessment will create liability.

III. Accountability Subjects

[Q15 & Q16] How should AI accountability mechanisms consider the AI value chain and AI lifecycle?

AI technologies and applications are dynamic, often marked by continuous innovation, integration, and updating by developers, deployers and end users. At the same time, there is broad diversity in how and where companies in different industries engage along the AI lifecycle, from design and development to deployment and end use.

There is also significant diversity across AI systems. For example, not all AI systems have direct implications for humans or clearly defined risks (for instance, considerations for AI deployed to optimize network performance will differ from AI used in employment applications), and as such, the scope, timing and subjects of AI accountability and related mechanisms should be tailored to the specific system context. Likewise, AI systems differ substantially in their configurations, component parts and complexity. Finally, how models function in production and deployment is the combination of many factors, requiring contextual work to determine which of those factors is the appropriate subject for accountability review.

Given this diversity in AI systems and the multitude of discrete touchpoints from development to end use, it is not advisable to focus accountability efforts on one specific link in the AI value chain. Rather, accountability mechanisms should be context- and system-specific within a broader framework for AI accountability.

In addition, responsible downstream use of AI requires a particular focus on collaboration and information-sharing. Consistent with our Roadmap principle to "[d]esign for and implement transparency, explainability and interpretability," deployers of AI systems should be equipped with sufficient information and training to support responsible and trustworthy downstream use. Deployers of AI systems should also be given clear information about an AI tool's proper functioning and limitations, including potential fitness for secondary purposes.

[Q17] 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?

Scoping of AI accountability methods should be assessed through a contextual analysis of development, use and impact, and take into consideration the industries or evolved practice areas within which risks are already specifically defined. For example, many commonplace AI applications can be characterized as optimizing services and streamlining operations, without making automated decisions on individuals or processing sensitive information. The context and corresponding risk levels of AI applications exist in a wide spectrum. As such, accountability approaches should be outcomes-focused and take a risk-based approach to avoid over-regulating uses of AI which have no significant impact on individuals or pose potential for societal harm.

Importantly, AI accountability measures should acknowledge and differentiate between functions of AI systems along five features:

1. **Decision-making vs. decision support**: AI that makes decisions versus AI that produces predictions and insights that support and improve human decision-making;
2. **Well-defined vs. novel use:** AI that automates existing and authoritative processes against well-established and understood standards versus entirely novel applications for which standards have not been developed;
3. **Dynamic vs. fixed:** AI models that evolve in production with different data inputs and in fluid settings versus those that are "locked" and do not;
4. **Personal vs. enterprise data**: AI systems that use data collected from consumers versus enterprise data used for run-of-the-mill internal operations (*e.g.*, assembly optimization, product quality control, etc.); and
5. **Human vs. operational impacts:** Systems that make or support decisions with direct implications for humans, versus systems whose role is to facilitate more efficient or effective internal operations (*e.g.*, identifying and diagnosing flaws in automated manufacturing output).

Leaning on the guidance in the NIST AI RMF, NTIA should consider the importance of a proportionate, risk-based approach, with higher-risk applications subject to more in-depth reviews, robust cost-sensitive analyses and stricter (often cross-domain) governance mechanisms (*e.g.*, ethics committees).

AI audits, AI accountability and other AI assessments should also, as appropriate, compare the risks, impacts and outcomes of AI to existing human or technological alternatives. In many contexts, AI solutions have proven to be safer, more accurate and less biased than existing systems.

IV. Accountability Inputs and Transparency

[Q20] What sorts of records (*e.g.,* logs, versions, model selection, data selection) and other documentation should developers and deployers of AI systems keep in order to support AI accountability? How long should this documentation be retained? Are there design principles (including technical design) for AI systems that would foster accountability-by-design?

Documentation and evaluation around intended use, data sources and outputs are a fundamental component of AI accountability and of designing AI systems that are fit for purpose.

The appropriate approach to documentation—whether through data dictionaries, data and model selection logs, version history, etc.—will differ depending on the system context and unique end user needs. Regarding design principles to foster accountability, many Business Roundtable member companies are operationalizing relevant principles and corresponding actions from our Roadmap, including articulating clear goals and purposes for which AI models are optimized, defining performance metrics to capture the value as well as potential risk in AI design and deployment, and engaging in self-assessments, performed internally with reference to external guidelines or standards.

[Q22] How should the accountability process address data quality and data voids of different kinds? How should AI accountability requirements or practices deal with these data issues? What should be the roles of government, civil society, and academia in providing useful data sets (synthetic or otherwise) to fill gaps and create equitable access to data?

Data collection and use is a core feature of AI systems, and likewise requires responsible and accountable processes. AI is only as good as the data it learns from, and algorithmic bias can be inadvertently introduced at any point in the development process if developers do not follow careful processes for screening datasets and checking human judgments for unconscious bias. Consistent with our Roadmap, Business Roundtable recognizes that robust data management is a critical element of accountability processes, and internal data governance practices should respect consumer privacy and security. In line with our longstanding positions on consumer data privacy and security,[7] AI accountability mechanisms should design data protection safeguards from the outset, collecting and using sensitive data based on demonstrated need.

---

[7] *See* Business Roundtable, Framework for National Privacy Legislation (Dec. 6, 2018), https://www.businessroundtable.org/policy-perspectives/technology/privacy.

V. Barriers to Effective Accountability

[Q25] Is the lack of a general federal data protection or privacy law a barrier to effective AI accountability?

Effective AI accountability would be enhanced by the passage of a comprehensive federal privacy and data security law to protect American consumers and align compliance efforts across the nation. Business Roundtable strongly supports a national consumer privacy law, which would strengthen protections for consumers in communities across the country, recognizing that consumers' digital lives and experiences are not restricted by state boundaries, while offering Congress the opportunity to create a holistic, preemptive approach to privacy. As AI is powered by data, many end users are concerned about data use in AI contexts. At the same time, U.S. privacy laws are increasingly fragmented across industries, geographies and jurisdictions, creating confusion among consumers and a complicated web of compliance activities for companies, detracting already limited time and resources from strategic and compliance work on AI governance efforts.

Consumer trust and confidence are essential elements of our members' businesses and relationships with customers. Business Roundtable urges Congress and the Administration to enact a federal consumer data privacy law that will protect and empower consumers and provide clear, consistent obligations for how companies handle personal data, while fostering responsible and inclusive American innovation and supporting U.S. competitiveness globally and avoiding a patchwork of state privacy laws.

[Q29] How does the dearth of measurable standards or benchmarks impact the uptake of audits and assessments?

While many frameworks, toolkits and technical solutions are emerging to manage and test AI systems and advance accountability, less is known about how these tools are working in practice, in conjunction or comparison with one another, and across different industries and use cases.

We commend the work of NIST in developing the first of its kind AI RMF, and the transparent, multistakeholder process used in the process. The NIST AI RMF addresses risk assessment and mitigation throughout the AI lifecycle and is flexible for organizations to apply across a range of use cases. As this resource was released less than a year ago, organizations are still working to familiarize themselves with it and implement its guidance. NTIA should partner with NIST and other federal agencies to encourage implementation of the AI RMF and educate organizations on how best to operationalize its guidance. As with other NIST frameworks, the AI RMF is intended to be an evolving document, and future iterations can account for the best and most timely guidance on implementing accountability efforts as AI technology evolves.

We are eager to continue collaborating with the public sector to build upon and help Business Roundtable member companies implement the NIST AI RMF.

VI. AI Accountability Policies

[Q30] What role should government policy have, if any, in the AI accountability ecosystem?

In January 2022, Business Roundtable proactively called on the Administration, Congress and regulators to establish reasonable practices, rules and guidelines for AI, noting that Responsible AI is a shared endeavor between business and government. We maintain this position and encourage the business community and government to work together to align and incentivize common sense governance practices, promote the innovation ecosystem and build public trust in AI systems.

Specifically, the U.S. government can lean in on AI policy along four core themes, reflected in our Recommendations for Policymakers:

1. **Pursue Transparent and Rational Enforcement:** Account for the evolving and differentiated nature of AI, focusing enforcement efforts on bad actors.
2. **Focus on Targeted and Flexible Governance:** Encourage AI innovation through targeted and flexible governance and oversight.
3. **Promote Global Coordination on Key Issues:** Participate in global dialogues to support common principles and understanding of AI.
4. **Invest in AI Education, Training and Awareness:** Partner with industry to build AI literacy and relevant skill sets across the United States.

[Q30a] Should AI accountability policies and/or regulation be sectoral or horizontal, or some combination of the two?

Business Roundtable recommends that the United States adopt regulatory approaches to AI that are contextual, risk-based, proportional and use-case specific. As AI is not a "one size fits all" technology, the regulatory approaches governing it should not be either. The United States should tailor any frameworks, guidance and regulation to specific AI use cases within well-defined deployment contexts, rather than broadly regulating any technology or application outright. Additionally, the government should conduct a thorough assessment of existing regulatory gaps before establishing new regulations to avoid overlapping or inconsistent rules.

[Q32] What kinds of incentives should government explore to promote the use of AI accountability measures?

Business and government have a vast array of experience with different regulatory tools and devices within and across sectors, and it is worthwhile to experiment in limited settings with

different approaches to strategic technology governance work. For instance, we commend efforts to consider a range of incentives and alternatives including safe harbor frameworks, confidentiality protections, remediation bounties and rational enforcement practices focused on malicious or truly negligent actors. Further, government could explore the use of evidence-based regulatory approaches and tools that allow for the iteration of governance practices (*e.g.*, regulatory sandboxes) and opportunities for industry to discover and share best practices.

Government could also incentivize industry to engage in self-assessments, which many Business Roundtable member companies are already performing. These self-assessments can detect and mitigate errors and bias, assess gaps in model capabilities and provide flexibility to improve AI systems. However, government should not impose uniform requirements for third party assessments.

[Q34] 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?

Business Roundtable supports the alignment of common principles, frameworks and standards for a targeted and flexible governance approach across the global AI regulatory landscape to minimize conflicting requirements and promote innovation, trade and investment. We encourage NTIA and the Administration to develop and execute on a government strategy for AI that can be consistently advanced in bilateral and multilateral consultations and negotiations.

We recommend that international efforts focus on sharing information, promoting research and development, aligning key definitions (*e.g.*, AI explainability, fairness, classifications of risk) and promoting interoperability, supporting digital trade frameworks and confronting data localization requirements among others that impair cross-border data flows and the collection of fully representative training data. The U.S. government should collaborate with private sector stakeholders, encouraging participation in global standards-setting bodies and regulatory forums.

Conclusion

Business Roundtable welcomes NTIA's RFC on AI accountability mechanisms to advance Responsible AI efforts. AI accountability warrants CEO-level engagement and leadership, as well as participation and engagement from the public and private sectors. Business Roundtable looks forward to continued engagement with NTIA and other thought leaders and policymakers on these important topics and is happy to discuss our response or these issues at any time. Please contact Paul Jackson, Vice President, Business Roundtable, at pjackson@brt.org or (202) 467-5269.