

Bipartisan Policy Center  
1225 I St NW, #1000  
Washington, DC 20005

June 12, 2023

**BY ELECTRONIC MAIL**

Travis Hall  
NTIA  
U.S. Department of Commerce  
1401 Constitution Avenue NW  
Room 4725  
Washington, DC 20230

**Re: Document No. 2023-07776: Bipartisan Policy Center Comments in Response to the National Telecommunications and Information Administration AI Accountability Policy Request for Comment**

Mr. Hall:

The Bipartisan Policy Center’s (BPC’s) Technology Project welcomes this opportunity to submit comments in response to the National Telecommunications and Information Administration (NTIA) AI Accountability Policy Request for Comment ([RFC](#)). As an organization committed to helping policymakers work across party lines to craft bipartisan solutions, we appreciate NTIA’s commitment to “solicit input from stakeholders in the policy, legal, business, academic, technical, and advocacy arenas on how to develop a productive AI accountability ecosystem.”

BPC’s Technology Project has supported efforts to promote trustworthy AI and recognizes the importance of ensuring appropriate AI accountability. Below, we provide information about BPC, the Technology Project, and our ongoing work on AI policy. We then share perspectives on AI accountability objectives, existing AI accountability resources and approaches, and barriers to effective AI accountability. Finally, we offer policy recommendations for promoting AI accountability going forward.

**I. Introduction to the Bipartisan Policy Center and the Technology Project**

The [Bipartisan Policy Center](#) is a non-profit, 501(c)(3) organization that delivers data and context, negotiates policy details, and creates space for bipartisan collaboration to enable our democracy to function on behalf of all Americans. We leverage our relationships with current and former elected officials, business leaders, academic experts, and advocates across the political spectrum to shape practical policy ideas. What sets BPC apart from traditional think tanks is our unwavering view that engaging “proud partisans” is essential to creating better solutions and solving our nation’s problems. We embrace the reality that good ideas alone do not drive policy change, and we have crafted the networks, policy expertise, and persuasion techniques to work around that fact.

BPC began its technology policy work in late 2019 with our initiative to develop a [national AI strategy for Congress](#) in collaboration with Former Rep. Will Hurd (R-TX) and Rep. Robin Kelly (D-IL). Through this initiative, BPC held a series of roundtables with government officials, industry representatives, civil society advocates, and academics. Subsequently, we produced four whitepapers on [AI and the workforce](#), [AI and national security](#), cementing U.S. AI leadership through [research and development](#), and [AI and ethics](#). These whitepapers provided several recommendations that [H.Res.1250](#) incorporated.

Since then, we have continued our AI work (detailed in Appendix 1) through educating Congress, analyzing policy proposals, and engaging with stakeholders. Last year, we published a [report](#) on the EU’s efforts to regulate AI. We explored academic, government, civil society, and industry perspectives for

policymakers to consider when crafting [AI impact assessments](#). Most recently, BPC published pieces on [defining high-risk, high-reward AI](#); [face recognition technology governance challenges](#); and [workforce resilience and adaptability for the AI-driven economy](#).

While continuing to work on AI policy, the Technology Project’s portfolio has expanded and now includes content moderation, data privacy, immersive technologies (e.g., augmented reality and virtual reality), competition, cybersecurity, space, and broadband/digital divide policy issues. More information about these initiatives is available on our [website](#).

## **II. Comments in Response to RFC Questions**

Promoting U.S. leadership in trustworthy AI has been a bipartisan priority for several years, but opinions diverge on which approaches to AI accountability can best achieve this broad objective. Below, we share insights, which we gleaned through our work with diverse stakeholders on AI policy initiatives, that are relevant to several of the questions the RFC poses.

### ***a. AI Accountability Objectives (Questions 1-8)***

To ensure consistency and durability over time, AI accountability objectives should have bipartisan support and align with the purposes, principles, and objectives of the U.S. AI national strategy, as articulated in the 116<sup>th</sup> Congress’s National AI Initiative Act ([H.R. 6216](#)) and [H.Res. 1250](#). The bipartisan National AI Initiative Act of 2020 stated that the purposes of the initiative are to: “(1) ensure continued United States leadership in artificial intelligence research and development; (2) lead the world in the development and use of trustworthy artificial intelligence systems in the public and private sectors; (3) maximize the benefits of artificial intelligence systems for all American people; and (4) prepare the present and future United States workforce for the integration of artificial intelligence systems across all sectors of the economy and society.” The bipartisan H.Res. 1250 identified five guiding principles: (1) global leadership; (2) a prepared workforce; (3) national security; (4) effective research and development; and (5) ethics, reduced bias, fairness, and privacy.

In line with the U.S. government’s AI national strategy, AI accountability mechanisms should tailor requirements to risks in ways that promote civil rights and liberties, equity, safe and efficient AI adoption, technological innovation, sustainability, and national and economic security. We agree with the RFC’s statement that the “appropriate goal and method to advance AI accountability will likely depend on the risk level, sector, use case, and legal or regulatory requirements associated with the system under examination.”

To achieve bipartisan objectives through a tailored, risk-based approach, AI accountability mechanisms should remain flexible enough to accommodate future innovation, focus on more than just accuracy, encourage continuous reviewing and testing, and build on existing resources where possible. More specifically, BPC’s “[Six Takeaways from Experts on AI Impact Assessments](#)” blog suggested that AI accountability mechanisms like impact assessments should remain flexible enough to support ongoing innovation while still promoting safety. This blog also urged AI impact assessments to not only evaluate accuracy but also to “include other important qualities such as explainability, transparency, robustness, and security.” The same piece found that “a governance structure with continuous review and testing promotes accountability” because no single tool or mechanism can fully assess every aspect of AI accountability.

As BPC noted in our [comments](#) to the National Institute of Standards and Technology on the AI Risk Management Framework (AI RMF), AI risk management and accountability mechanisms should be “consistent, to the extent possible, with other approaches.” Where applicable, these mechanisms should “take advantage of and provide greater awareness of existing standards, guidelines, best practices, methodologies, and tools for managing AI risks whether presented as frameworks or in other formats.”

They also “should be law- and regulation-agnostic to support organizations’ ability to operate under applicable domestic and international legal or regulatory regimes.” Furthermore, building on existing accountability mechanisms for privacy, cybersecurity, and broader quality assurance can help address AI accountability issues with implications across these three areas.

*b. Existing Resources and Models (Questions 9-14)*

AI accountability mechanisms vary widely in their forms and purposes. BPC’s past work on AI accountability mechanisms has focused primarily on AI impact assessments and NIST’s AI RMF. Consequently, these are the mechanisms on which our comments below focus.

i. AI Impact Assessments

BPC’s work on AI [impact assessments](#) explains that conducting an impact assessment “promotes accountability by requiring an organization to document its decision-making process and ‘show its work.’” Furthermore, “impact assessments can make the inner workings of the algorithms that power these systems more transparent.” The “goal is for impact assessments to promote accountability through documentation and knowledge production, rather than to instill fear of liability.”

In 2022, we held four convenings with experts from [academia](#), [civil society](#), [government](#), and [industry](#) to identify consensus views on the right scope, form, and role of AI impact assessments. [Six](#) key goals emerged from our discussions with experts across all four sectors:

- (1) Think beyond “hard law.”
- (2) Think beyond a “one-and-done” document.
- (3) Think beyond accuracy as the measure of success.
- (4) Think beyond a “one-size-fits-all” general framework.
- (5) Think beyond computer scientists.
- (6) Think beyond deployment.

Collectively, these goals stress the importance of mitigating risks without unduly inhibiting innovation. Taking a “context-specific approach” (i.e., a use-case-specific approach) to impact assessments can allow organizations “grappling with different harms” to use tailored tools that are easier to implement and better at addressing those specific harms without unduly limiting positive impacts. To effectively mitigate risks, conducting assessments throughout the AI lifecycle is important. “The closer an organization is to an ‘all-hands-on-deck’ approach, the better it can identify challenges, problems, and risks throughout the process.” Furthermore, diverse, multidisciplinary views should inform conceptions of risk and approaches to risk mitigation, and risk mitigation should incorporate metrics that focus on equity, explainability, transparency, robustness, and security, in addition to accuracy.

In 2022, the AI experts with whom BPC spoke generally preferred “piloting a voluntary risk-management framework with stakeholders before considering what binding rules might look like.” However, the proliferation of widely accessible generative AI systems has since motivated AI experts from [multiple sectors](#) to [recommend](#) developing AI governance legislation sooner rather than later.

ii. NIST’s AI RMF

In addition to performing our own research and analysis on [bias in AI systems](#) and AI [impact assessments](#), BPC’s Technology Project supports NIST’s ongoing efforts to promote [trustworthy and responsible AI](#) through [research](#) and work on the [AI RMF](#). As NIST’s AI RMF 1.0 states, “Maintaining organizational practices and governing structures for harm reduction, like risk management, can help lead to more accountable systems.”

We submitted [three comment letters](#) to help inform NIST’s approach to creating the AI RMF. We recommended that NIST adopt an approach that “accepts that AI actors will take some inevitable risks but

requires actors to be transparent about their evaluations of risks, fostering both accountability and innovation.” The approach NIST adopted in its AI RMF largely achieves this dual objective, and we provided a statement of [support](#) when NIST launched its first full version of the AI RMF.

Our second comment letter noted that, because the AI RMF is voluntary, corporations “need more substantial incentives to adopt NIST’s framework in tandem with internally produced AI risk management frameworks, AI risk assessments, or standards and guidelines.”

*c. Barriers to Effective Accountability (Questions 24-29)*

Effective AI accountability should promote appropriate risk mitigation and build public trust in AI without unduly impeding innovation. The patchwork legal framework governing AI design, development, use, and oversight and the lack of an AI-ready workforce can pose significant barriers to effective AI accountability.

*i. Patchwork Legal Framework*

The United States does not have a federal AI governance law that explicitly requires AI developers or users to leverage impact assessments, audits, or other accountability mechanisms. Nonetheless, implementing AI accountability mechanisms may help (or even be necessary to demonstrate that) AI developers and users meet their obligations under existing laws. For example, the Federal Trade Commission (FTC) recently [warned](#) AI developers to “take all reasonable precautions” before AI products hit the market. The FTC pointed out that it “has sued businesses that disseminated potentially harmful technologies without taking reasonable measures to prevent consumer injury.” The FTC, Consumer Financial Protection Bureau, Department of Justice, and Equal Employment Opportunity Commission also [asserted](#) that existing “legal authorities apply to the use of automated systems and innovative new technologies just as they apply to other practices.”

Federal legal obligations to promote data privacy may also apply to the development and procurement of AI technologies. Notably, “Privacy Impact Assessments (‘PIAs’) are [required](#) by Section 208 of the E-Government Act for all federal government agencies that develop or procure new information technology involving the collection, maintenance, or dissemination of information in identifiable form or that make substantial changes to existing information technology that manages information in identifiable form.” Federal government agencies that develop or procure new AI technologies, therefore, may need to complete PIAs.

States and municipalities are also beginning to develop a patchwork legal framework for AI governance. The RFC mentions some of the recent AI laws, including New York City Law 144, which requires “bias audits of certain automated hiring tools used within its jurisdiction.” California’s Age-Appropriate Design Code Act [requires](#) covered businesses to “complete a Data Protection Impact Assessment” that addresses whether “algorithms used by the online product, service, or feature could harm children.” Other examples of state-level AI laws include an [Illinois](#) law regulating the use of AI on job applicant interview video footage and a similar [Maryland](#) law that applies exclusively to face recognition technologies. A [Colorado](#) law restricts the use of algorithms and “predictive models” in insurance practices. State-level [content moderation](#) and [data privacy](#) laws also have direct and indirect implications for AI accountability and broader AI governance by restricting the data that AI technologies can process and limiting their use for content recommendation.

Foreign governments are also creating legal obligations to promote AI accountability, and these laws may apply to U.S. companies that operate and conduct business globally. As the RFC mentions, relevant provisions of the European Union’s Digital Services Act require “audits of very large online platforms’ systems.” Additionally, the Canadian government has developed a mandatory [Algorithmic Impact Assessment](#) that is “intended to support the Treasury Board’s Directive on Automated Decision-Making.”

Varying legal requirements across jurisdictions can create confusion for AI developers, AI users, AI regulators, and individuals. This confusion can hinder effective legal compliance and enforcement. If different jurisdictions' laws diverge significantly, AI developers may struggle to design products that comply with all necessary laws, thereby increasing the risk of legal violations and/or inhibiting innovation. Because researchers in multiple U.S. states and different countries may collaborate to develop AI technologies, harmonizing AI legal requirements across the United States and between the United States and like-minded countries can help support continued trustworthy AI innovation. Furthermore, U.S. federal, state, and local enforcement authorities and authorities in other countries are more likely to struggle to cooperatively enforce violations spanning multiple jurisdictions when legal obligations vary across those jurisdictions.

Increasing the clarity and consistency of laws across U.S. jurisdictions in a manner that minimizes potential conflicts with international laws would help AI developers and users understand their legal obligations. Especially considering that many digital AI technologies operate across multiple jurisdictions, legal clarity and consistency across jurisdictions would help enforcement authorities hold AI developers and users accountable for fulfilling those obligations.

ii. AI Workforce Talent Gap

In BPC's 2020 [AI and the Workforce](#) report, we explained that the United States is facing "an AI workforce talent shortage" known as the "AI talent gap." This talent gap "is spanning almost all industries as businesses seek to leverage the strengths of AI." We described how American universities struggle to recruit and retain AI faculty and how the U.S. private sector is launching aggressive efforts to recruit and retain AI talent needed to develop and leverage AI technologies. We also shared that the federal government has trouble recruiting the AI talent needed to effectively implement cutting-edge AI technologies.

Since we published our 2020 report, the National Security Commission on Artificial Intelligence's final [report](#) and the National Artificial Intelligence Advisory Committee Year 1 [report](#) emphasized that the AI talent gap is still present and problematic. Without a multidisciplinary, multisectoral workforce that is capable of developing, testing, leveraging, and conducting oversight of trustworthy AI technologies, effectively designing, implementing, using, and reviewing AI accountability mechanisms will be challenging.

*d. AI Accountability Policies (Questions 30-34)*

Enacting a federal consumer data privacy law, strengthening AI governance, cultivating an adaptable and resilient workforce for the AI-driven economy, and investing in trustworthy AI research and development (R&D) can help promote AI accountability.

i. Enacting a Federal Consumer Data Privacy Law

Enacting a federal consumer data privacy law would advance AI accountability by helping to ensure that AI developers obtain training and testing data through appropriate practices. A federal consumer data privacy law that establishes requirements for the data AI technologies process and the ways in which that processing occurs would also help ensure that AI developers and users develop and deploy AI technologies in a trustworthy, privacy-protective manner.

To ensure that a federal consumer data privacy law does not inhibit AI accountability efforts, however, policymakers should be careful when imposing any limitations on processing data that may be essential for training and testing AI technologies in ways that mitigate bias.

ii. Strengthening U.S. AI Governance

BPC’s Technology Project generally recommends clearly enshrining broad requirements that are necessary to protect civil rights and fundamental American values in one or more AI governance laws. (Such requirements may clarify how existing laws apply to the design, development, use, and oversight of AI technologies and/or may seek to close any identified gaps in the protections that existing laws provide.) These laws should direct relevant federal agencies to develop regulations that more clearly detail what the laws’ requirements entail. Adhering to international standards and/or widely trusted frameworks can help achieve and demonstrate compliance with legislative and regulatory requirements.

Building off BPC’s prior work on [impact assessments](#), our 2023 explainer, “[Defining High-Risk, High-Reward AI](#),” emphasized that AI governance requirements and restrictions should be use-case-specific. Robust, effective AI governance that aims to mitigate risks without unduly impeding benefits likely will require a combination of hard law and soft law (including international standards, voluntary risk management frameworks, and best practice guidance documents). Because different AI use cases produce different risks and rewards, the optimal balance of hard and soft law will differ across AI use cases. Governance frameworks generally should subject AI use cases that pose high risks to more stringent requirements, but those requirements should not be so stringent that they prohibit high-risk, high-reward AI use cases. Governance requirements also should not be so onerous that they stifle R&D initiatives that could produce novel high-reward use cases.

However, as we explain, “in the United States, which AI use cases are high-risk, high-reward is an open question. Building consensus on which AI use cases are high-risk, and which of those are also high-reward, can help U.S. policymakers and other stakeholders develop and implement effective AI governance frameworks.”

Policy initiatives that aim to protect data privacy and address online content moderation challenges also may impact AI governance. Analyzing the ways in which these policy areas interact can help policymakers develop AI governance frameworks that effectively address AI-powered content moderation systems and other technologies with implications for all three policy areas. Identifying and considering issues at the intersection of AI, data privacy, and online content moderation policy can help avoid potential conflicts between any data privacy, AI governance, and online content moderation legislation that Congress develops and advances.

iii. Developing an Adaptable and Resilient Workforce for the AI-Driven Economy

Effectively leveraging AI accountability mechanisms will require an AI-ready workforce. Our [AI and the Workforce](#) report asserted that the United States should demonstrate leadership in the AI-driven economy “by filling the AI talent gap and preparing the rest of the workforce for the jobs of the future. However, in doing so, policymakers should make inclusivity and equal opportunity a priority.” We also pointed out that the “educational system from kindergarten through post-college is not yet designed for the AI-driven economy and should be modernized.”

Since we published our 2020 report, Congress passed the Infrastructure Investment and Jobs Act ([IIJA](#)) and the [CHIPS and Science Act](#), which contain provisions that aim to help fill the AI talent gap and promote STEM education. Nonetheless, more work remains to be done, including through efforts to effectively implement the IIJA and CHIPS and Science Act.

Our [comments](#) to NIST on its AI RMF stated, “A diverse workforce with a broad perspective and understanding of risks associated with AI applications is necessary to identify, prioritize, and respond to risks. The challenge of creating a diverse workforce for AI requires a holistic approach, starting from early education and throughout a career. It must focus not just on recruiting talent but also on developing

and retaining existing talent, which requires looking at an organization’s culture and whether it is inclusive. This includes diversifying organizations’ leadership.”

In April 2023, a BPC [blog](#) provided three recommendations for promoting workforce adaptability and resilience in the modern AI-driven economy: (1) promote lifelong learning, (2) empower workers to develop skills that leverage and complement AI, and (3) strengthen AI governance. Policies that inclusively support lifelong learning “will help workers regularly update their skills as the workplace, and the tasks that make up the jobs of the future evolve in unpredictable ways.” Although predicting which skills will most likely complement AI technologies can be challenging, “businesses and government should empower workers to develop skills that leverage and complement AI technology.” Strengthening AI governance by building on “existing laws where appropriate” and tailoring requirements to the specific risks and rewards AI use cases pose “can help organizations implement best practices when applying AI tools in the workplace.”

#### iv. Promoting Trustworthy AI Research and Development

Promoting trustworthy AI R&D can help support innovation that improves AI accountability. For instance, innovations that improve the effectiveness of using [synthetic data](#) to train and test AI systems can reduce privacy risks by decreasing the amount of personal data in an AI training or testing dataset. Innovations that improve [bias](#) mitigation techniques can also improve AI accountability by decreasing the risks that AI technologies will disproportionately underperform for members of particular demographic groups and/or that AI technology outputs will perpetuate existing societal biases. Policies that support trustworthy AI R&D are therefore essential to strengthening AI accountability.

BPC’s 2020 report (published in partnership with the Center for a New American Security), [Cementing American Artificial Intelligence and Leadership: AI Research & Development](#), highlighted the country’s significant AI-related R&D needs. The report recommended building on existing R&D spending, investing in broadband infrastructure and high-end computational resources, and supporting AI talent development.

In a 2021 blog, “[Unfolding AI’s Potential: How Investing in Research and Development Can Produce New Knowledge](#),” we provided examples of how AI technologies can serve as “meta technologies.” Since we published this blog, Congress has made significant [progress](#) in line with our original recommendations. Nonetheless, we believe future U.S. leadership in trustworthy AI innovation will require continued support for: (1) public- and private-sector trustworthy AI R&D, (2) expanded and diversified computing (including resources that help optimize data use), (3) international cooperation with like-minded countries on trustworthy AI R&D, (4) efforts to attract and develop top AI talent, and (5) U.S. participation in and leadership of AI standards initiatives.

Promoting AI-ready open data can support innovative, trustworthy AI R&D and AI accountability. As our 2023 “[AI-Ready Open Data](#)” explainer states, “McKinsey estimates that open data can help unlock \$3 trillion to \$5 trillion in economic value annually across seven sectors. But for open data to fuel innovations in academia and the private sector, the data must be both easy to find and use. While Data.gov makes it simpler to find the federal government’s open data, researchers still spend up to 80% of their time preparing data into a usable, AI-ready format.” This piece recommended that the federal government:

- (1) Direct NIST to establish a general U.S. government standard for AI-ready data that “could look like a ‘nutrition label,’ building on existing projects such as the Data Nutrition Project, Datasheets for Datasets, and AI-Ready Checklist”;
- (2) “Launch ‘Data Challenges’ to spur collaboration across academia, industry, and government using open data sets”; and

- (3) “Embed the principles of AI-ready data into its contracting process whenever it expects contractors or grantees to produce data that will be posted on Data.gov” and “update the Federal Data Strategy to include requirements for AI-ready data sets.”


### III. Closing

AI accountability initiatives should pursue purposes, principles, and objectives that have bipartisan support and that align with broader U.S. strategic goals for AI, like those in the National AI Initiative Act of 2020 and H.Res.1250. Taking a use-case-specific and risk-based approach to establishing AI accountability obligations can help ensure that AI accountability practices support civil rights and equity, safe and efficient AI adoption, technological innovation, sustainability, and national and economic security. Existing AI accountability mechanisms, like the NIST AI RMF and AI impact assessments, can help the U.S. promote trustworthy and responsible AI innovation. Nonetheless, barriers to AI accountability, including the patchwork legal framework and the AI workforce talent gap, remain.

Enacting a federal consumer data privacy law, strengthening U.S. AI governance frameworks, developing an adaptable and resilient AI-ready workforce, and continuing to invest in trustworthy AI R&D can help advance AI accountability in the United States.

BPC appreciates NTIA’s willingness to read and consider the perspectives and recommendations in this comment letter, and we would welcome future opportunities to serve as a resource to NTIA.

Sincerely,  
*Bipartisan Policy Center Technology Project*





## **Appendix 1: Bipartisan Policy Center Work on Artificial Intelligence**

### **AI National Strategy**

- Bipartisan Policy Center held an event entitled “[An AI National Strategy for Congress](#)” that featured perspectives on the AI National Strategy for Congress that the Bipartisan Policy Center developed in consultation with Reps. Robin Kelly (D-IL) and Will Hurd (R-TX).
- Bipartisan Policy Center published a report, [AI and the Workforce](#), that outlines major workforce-related AI challenges and recommends that the federal government take several actions to advance AI and prepare the workforce for the future. Recommended actions include working to close the AI talent gap, addressing AI-related workforce disruptions, and training the workforce to utilize advanced technologies.
- Bipartisan Policy Center (in partnership with the Center for Security and Emerging Technology) published a report, [Artificial Intelligence and National Security](#), which explains national and economic security challenges and how integrating AI into U.S. defense and intelligence agencies will play a critical role in national security and international economic competition. The report provides recommendations to the federal government to help coordinate a strategic approach to researching, developing, integrating, and scaling AI across the relevant agencies and departments.
- Bipartisan Policy Center (in partnership with the Center for a New American Security) published a report, [Cementing American Artificial Intelligence and Leadership: AI Research & Development](#), highlighting the country’s significant AI-related R&D needs. The report recommends building on existing R&D spending, investing in broadband infrastructure and high-end computational resources, and supporting AI talent development.
- Bipartisan Policy Center published an issue brief, [AI and Ethics](#), that identifies AI ethics concerns and recommends several actions that the federal government could take to help the U.S. accelerate and sustain global leadership in AI while minimizing the likelihood of adverse impacts on civil liberties, civil rights, and innovation.
- Bipartisan Policy Center published a [blog](#) announcing its collaboration with Reps. Will Hurd (R-TX) and Robin Kelly (D-IL) on an AI national strategy initiative.

### **European Policy Perspective**

- Bipartisan Policy Center published a report, [Artificial Intelligence Policy and the European Union](#), that examines AI policy actions abroad and their implications for U.S. policymaking. The report contextualizes the different approaches that the United States and European Union are taking to address AI challenges and provides an overview of different perspectives on the EU AI Act.
- [Event Recap – Could the EU Set U.S. AI Policy and Standards?](#) – October 25, 2022
- [Could the EU Set U.S. Artificial Intelligence Policy and Standards](#) – September 14, 2022

### **NIST AI Risk Management Framework**

- [Bipartisan Policy Center Response to NIST on Artificial Intelligence \(AI\) Risk Management Framework: Second Draft](#) – October 4, 2022
- [Bipartisan Policy Center Response to NIST on Draft Artificial Intelligence \(AI\) Risk Management Framework](#) – May 05, 2022
- [Bipartisan Policy Center Response to NIST RFI on Artificial Intelligence \(AI\) Risk Management Framework](#) – Aug 27, 2021

### **AI Impact Assessments**

#### Event recaps

- [Event Recap – Industry Perspectives on AI Impact Assessments](#) – Oct 17, 2022

- [Event Recap – Civil Society Perspectives on AI Impact Assessments](#) – Jul 29, 2022
- [Event Recap – Government Perspectives on AI Impact Assessments](#) – Apr 25, 2022
- [Event Recap- Research and Academic Perspectives on AI Impact Assessments](#) – Mar 17, 2022

Recorded event links

- [Defining AI Impact Assessments: Industry Perspectives](#) – May 18, 2022
- [Civil Society Perspectives on Artificial Intelligence Impact Assessments](#) – July 12, 2022
- [Government Perspectives on AI Impact Assessments](#) – March 1, 2022
- [Research and Academic Perspectives on Artificial Intelligence Impact Assessments](#) – February 22, 2022

Explainer pieces and blog posts

- [Explainer: Impact Assessments for Artificial Intelligence](#): Explainer piece addressing seven questions related to what an impact assessment is and what it means for AI
- [Six Takeaways from Experts on AI Impact Assessments](#): Blog post explaining six common goals that emerged from BPC’s convenings with AI experts from academia, civil society, government, and industry

**Other content:** Explainer pieces, blogs, informational graphics, and events that explain complex AI issues and/or propose policy approaches

- [Workforce Resilience and Adaptability for the AI-Driven Economy](#): A blog that explains the new AI landscape; analyzes the impact of new AI technologies on the workforce; and recommends that policymakers promote lifelong learning, empower workers to develop skills that leverage and complement AI, and strengthen AI governance
- [Five Key Face Recognition Technology Governance Challenges](#): An ongoing blog/explainer series that examines five challenges that Members of Congress face when working to advance face recognition technology legislation
- [Defining High-Risk, High-Reward AI](#): An explainer piece that explains how developing multi-stakeholder consensus on which AI use cases are high-risk and high-reward could help advance U.S. AI governance conversations
- [AI-Ready Open Data](#): An explainer piece that provides an overview of existing efforts across the federal government to improve the AI readiness of its open data and recommends actions that policymakers should take to move the AI-ready data agenda forward
- [Learning about Machine Learning](#): As part of a series of blog posts about real-world machine learning applications and their complexities, this piece describes the basic architecture of machine learning tools used today
- [Complexity in Machine Learning](#): As part of a series of blog posts about real-world machine learning applications and their complexities, this piece introduces the differences between conventional machine learning and deep learning and explains downstream issues arising from deep learning’s increased complexity
- [Synthetic Data](#): As part of a series of blog posts about real-world machine learning applications and their complexities, this piece discusses how real-world uses of synthetic data can help promote privacy, advance AI capabilities, and create both risks and benefits
- [Framing the AI Fairness Question?](#): An explainer piece about how different predictive models can reach different conclusions with different tradeoffs and why identifying risks and considering the impact of various outcomes is essential
- [Bias in AI systems](#): A blog post about several of the technical and non-technical solutions that have been proposed to mitigate harm without unduly hampering innovation
- [AI: Facts and Myths](#): An infographic challenging many prevalent misconceptions about AI technology



- [We Are Not Ready for The Next Leap in AI, Natural Language Processing](#): A blog post explaining how natural language processing (NLP) works, current and potential future uses of NLP, and some of the challenges and opportunities that NLP use cases present
- [5 Things to Know About AI Weather Forecasting](#): An explainer piece about the use of AI systems to analyze atmospheric data and predict potentially dangerous weather events
- [Advancing AI: Key AI Issue Areas Policymakers Should Consider](#): A blog post that explains what constitutes “AI,” relevant AI policy issues, and several recent AI policy advances
- [Unfolding AI’s Potential: How Investing in Research and Development Can Produce New Knowledge](#): A blog post providing examples of AI as a meta technology and making AI R&D policy recommendations
- [What is Needed for AI to Succeed?](#): A blog post explaining how developing a skilled workforce, ensuring inclusivity, optimizing data use, encouraging public and private research and development, building computing capacity, establishing technical standards, and fostering public trust and positive attitudes towards AI are crucial to enabling the United States to unlock the full potential of AI
- [New Administration’s Tech Policy Should Consider AI Ethics](#): A blog post, written between the 2020 election and President Biden’s inauguration, outlining three broad AI ethics questions and BPC’s associated recommendations
- [Today’s Challenges and Tomorrow’s Skills: How the Workforce of the Future Starts with Strategic Action Now](#): An event providing perspectives on how advances in AI and other societal changes are impacting the workforce and how workers can thrive in the AI-driven economy
- [AI and Pandemics](#): An event about the role of AI technologies in fighting COVID-19
- [In the Midst of the Coronavirus Pandemic, the Case for Artificial Intelligence](#): A blog post about the challenges and opportunities associated with using AI technologies to help combat the COVID-19 pandemic
- [Can AI Accelerate Innovation? \(Webcast\)](#): An event sharing perspectives on AI’s role in accelerating innovation
- [Can AI Accelerate Innovation?](#): A blog post explaining several ways that AI technologies can help address three interrelated innovation acceleration challenges, a few of AI technologies’ current limitations, and four questions that policymakers may want to ask when working to optimize the use of AI for driving innovation
- [Episode 42: The Future of Artificial Intelligence](#): A *Pints & Policy* podcast episode featuring perspectives from John Soroushian, Senior Associate Director of Corporate Governance and Technology at BPC
- [The Future of AI Featuring Reps. Foster and Hurd](#): An event exploring the implications of the advances in AI and how public policy should adapt
- [Artificial Intelligence and Finance](#): A report explaining some of the challenges (including algorithmic bias, privacy, consumer protection, overreliance, and gaming risk) that arise when integrating AI into the financial sector and the ways in which encouraging responsible AI innovation can create opportunities
- [Responsible AI Can Improve Finance](#): A blog post outlining the limitations and potential of AI solutions in the financial sector and emphasizing the importance of continued human oversight of AI technologies
- [Artificial Intelligence and Anti-Money Laundering](#): A blog post about how AI technologies could help modernize the anti-money laundering (AML) framework in ways that enable law enforcement to more effectively target terrorism financing and money laundering