1350 Eye Street NW, Suite 350, Washington, DC 20005 | 202-768-8950 | fpf.org

April 1, 2024

Via Electronic Submission
U.S. Office of Management and Budget
Attn: Alex Goodenough
New Executive Office Building
725 17th Street NW
Washington, DC 20503
(202) 395-3039

Re: Request for Information: Privacy Impact Assessments

Dear Mr. Goodenough,

The Future of Privacy Forum (FPF)[1] welcomes the opportunity to submit comments in response to the Office of Management and Budget's (OMB) Request for Information regarding Privacy Impact Assessments (PIAs), particularly regarding the intersection of PIAs with other risks posed by the use and development of artificial intelligence (AI) tools and other emerging technologies.[2]

PIAs are a well-established mechanism for public and private entities to assess privacy risks in their services, products, and programs. By modernizing the current requirements for agencies conducting and disclosing PIAs, OMB has an important opportunity to influence how other risks posed by processing and using personal data in AI tools, such as algorithmic discrimination, interact with existing data privacy structures. Given these considerations, FPF suggests the following:

1. OMB should clearly define the scope of PIAs for AI to explicitly encompass considerations of all risks posed by the processing of personal data, including algorithmic discrimination;
2. OMB should recognize that risks addressed in a PIA, including discrimination risks, should be complementary to, and neither a replacement nor a repetition of, a comprehensive AI risk assessment or other AI-related assessment;
3. OMB should ensure that the scope and substance of a PIA for AI tools account for role-specific responsibilities and capabilities in the AI system lifecycle.

[1] FPF is a global non-profit organization dedicated to advancing privacy leadership, scholarship, and principled data practices in support of emerging technologies. The opinions expressed herein do not necessarily reflect the views of FPF's supporters or Advisory Board.
[2] In response to Questions 7 - 9 in the "Request for Information: Privacy Impact Assessments," 89 Fed. Reg. 5945 (Jan. 30, 2024), https://www.federalregister.gov/documents/2024/01/30/2024-01756/request-for-information-privacy-impact-assessments

1. **OMB should clearly define the scope of PIAs for AI to explicitly encompass considerations of all risks posed by the processing of personal data, including algorithmic discrimination.**

AI that uses or implicates personal data can create privacy risks for individuals, communities, and in some cases the whole of society. In 2017, FPF documented a wide range of potential harms related to the risks of using personal data in automated decision-making. Automated decision-making tools analyze personal information to make decisions about individuals by employing AI or other analysis techniques.[3] These harms fall broadly into four main categories: loss of opportunity, economic loss, social detriment, and loss of liberty. PIAs are uniquely well-situated as a means to identify, analyze, and manage risks related to the use of personal data in order to mitigate or prevent possible harm.

Such an approach is evidenced in Europe's General Data Protection Regulation (GDPR), where a Data Protection Impact Assessment (DPIA) serves as a crucial accountability tool with an "inclusive, comprehensive, and proactive nature," requiring the entity responsible for processing personal information (the "controller") to identify, assess, and ultimately manage the risks to the rights and freedoms of individuals whose personal information is processed.[4] According to the European Data Protection Board (EDPB), a DPIA is a process designed to describe the details of the processing of personal information, assess the necessity and proportionality of the processing (which are key legal requirements related to such processing), and "help manage the risks to the rights and freedoms" of individuals resulting from the processing of personal information.[5]

As OMB revisits existing guidance, we recommend adding a specific requirement for agencies to extend their risk analysis to identify and analyze the particular risks to individuals and communities constituting protected classes under U.S. law. These risks can be heightened when agencies use personal data as training or other inputs to systems that make consequential decisions about individuals. Low-income individuals and people of color in particular have been documented as being disproportionately impacted by inaccurate or biased information in commercial data sets (e.g., credit reports, lending history, etc.) which can lead to greater bias and

[3] *Unfairness by Algorithm: Distilling the Harms of Automated Decision-Making*, Future of Privacy Forum (Dec. 11, 2017), https://fpf.org/blog/unfairness-by-algorithm-distilling-the-harms-of-automated-decision-making/

[4] Katerina Demetzou, *Data Protection Impact Assessment: A Tool for Accountability and the Unclarified Concept of "High Risk" in the General Data Protection Regulation*, 35 Comput. L. and Sec. Rev. 6, (2019), https://www.sciencedirect.com/science/article/abs/pii/S0267364918304357?via%3Dihub

[5] *See Guidelines on Data Protection Impact Assessment (DPIA) and Determining Whether Processing is "Likely to Result in a High Risk" for the Purposes of Regulation 2016/679 (GDPR)*, § I, European Data Protection Board, https://ec.europa.eu/newsroom/article29/items/611236/en

discrimination.[6] Not only are the risks more significant, but when they are realized, harms related to the use of these systems are often particularly acute for marginalized and multi-marginalized individuals or communities.

2. **OMB should recognize that risks addressed in a PIA, including discrimination risks, should be complementary to, and neither a replacement nor a repetition of, a comprehensive AI risk assessment or other AI-related assessment.**

The use of AI tools and capabilities, in particular algorithmic decision making, may pose discrimination risks apart from those normally arising from processing personal data, as well as raise a number of other documented risks, including those related to the environment and climate change, copyright infringement, and human job security. These are each serious and independent matters of concern and may necessitate the addition of new or separate assessments for internal consideration. Should additional assessments be suggested or required for AI, such as a comprehensive AI risk assessment, a human rights impact assessment, or some other relevant assessment, OMB should be careful to ensure that they do not require duplication of effort and/or analysis and that each process is complementary to the other. In addition, OMB should ensure that the resultant administrative effort and cost is not burdensome to the point where comprehensive AI risk assessments are unnecessarily conducted on low risk data processing activities.

The European Union (EU) recently acknowledged the effectiveness of this approach in its adoption of the EU AI Act.[7] The AI Act mandates that public authorities and certain private companies providing public services that deploy high-risk AI systems must conduct a Fundamental Rights Impact Assessment (FRIA), which consists, among others, of an assessment of the specific risks of harm likely to have an impact on the categories of persons or groups of persons likely to be affected by its use, a description of the implementation of human oversight measures, and the measures to be taken where those risks materialize.[8] However, the EU AI Act explicitly recognizes that these obligations may in fact already be covered in the case that the concerned deployer of the high-risk AI system has already conducted a DPIA under the GDPR. Article 27(4) of the EU AI Act specifies that if any of the obligations related to conducting a FRIA is already completed as a result of a DPIA under the GDPR, the FRIA shall only complement the DPIA.

---

[6] Levi Kaplan, Alan Mislove, and Piotr Sapiezynski, *Measuring Biases in a Data Broker's Coverage*, FTC PrivacyCon 2022 (July 2017), https://www.ftc.gov/system/files/ftc_gov/pdf/PrivacyCon-2022-Kaplan-Mislove-Sapiezynski-Measuring-Biases-in-a-Data-Brokers-Coverage.pdf

[7] The EU AI Act has been adopted by the Plenary of the European Parliament on March 13, 2024, but there are several formal steps to be taken before formal adoption and publication in the Official Journal of the EU, expected in May 2024.

[8] Art. 27 of the EU AI Act, as adopted by the European Parliament on Mar. 13, 2024, https://www.europarl.europa.eu/doceo/document/TA-9-2024-0138_EN.html

OMB's consideration of updated guidance for agencies conducting PIAs should follow the EU's lead and require that additional assessments of AI tools and capabilities should incorporate by reference any existing PIA, so long as that PIA has been updated to reflect the relevant tool or capability. In addition, agencies should be guided to omit any part of an additional assessment's scope that has been previously covered in a PIA to protect against duplication or conflicting analysis. Since the OMB has directed federal agencies to designate Chief AI Officers and establish AI Governance Boards to coordinate and govern the use of AI within agencies, this may require the need to formalize organizational structures and communications between those specifically responsible for AI and existing privacy offices and officers.[9]

3. **OMB should ensure that the scope and substance of a PIA for AI tools account for role-specific responsibilities and capabilities in the AI system lifecycle, including the distinct roles played by developers and deployers.**

Lastly, a PIA must be adaptable to specific use cases in which a tool or capability, whether AI or otherwise, is utilized. A developer of an AI tool will have different access to determine features, respond differently to transparency requirements, and address certain privacy risks differently than an organization that deploys a tool that was developed by someone else. When conducted on AI systems, PIAs should account for the specific and unique responsibilities and capabilities of both developers and deployers respectively regarding operation and oversight, recognizing the ways each entity is best situated to evaluate and address issues pertaining to non-discrimination, responsible AI governance, transparency, data security, and privacy.[10]

For example, FPF's 2023 best practices for AI systems used in employment decisions[11] identifies one list of disclosures that developers should make to deployers, and a different list of disclosures that deployers should make to individuals:

---

[9] Shalanda D. Young, *Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence*, Office of Management and Budget (Mar. 28, 2024), https://www.whitehouse.gov/wp-content/uploads/2024/03/M-24-10-Advancing-Governance-Innovation-and-Risk-Management-for-Agency-Use-of-Artificial-Intelligence.pdf

[10] There may be cases where developers and deployers of an AI system are the same entity, in which case that entity needs to be assessed through both lenses.

[11] Future of Privacy Forum, *Best Practices for AI and Workplace Assessment Technologies* (Sept. 2023), https://fpf.org/wp-content/uploads/2024/02/FPF-Best-Practices-for-AI-and-WA-Tech-FINAL-with-date.pdf

Developer disclosures to deployers	Deployer disclosures to individuals
• the intended purposes of the AI tool • purposes for which the AI tool is not intended • known efficacy limits of the AI tool • how the AI tool was trained • whether the AI tool was assessed for potential discriminatory bias • whether the AI tool uses information from deployers or Individuals to further train or otherwise improve the tool • how the AI tool is intended to be deployed • uses of the AI tool that are not intended • what choices the AI tool provides to deployers regarding anti-discrimination, governance, transparency to Individuals, privacy, security, and human oversight; and • what choices the AI tool provides to deployers to communicate to Individuals about how they implement the tool, and how the tool fits into the deployer's overall decision-making processes regarding Consequential Impacts.	• the fact that Individuals are interacting with an AI tool • the intended use of the AI tool (e.g., to evaluate job candidates, make compensation decisions, or consider employees for promotion) • how the AI tool was trained • how an AI tool may have a consequential Impact and how the tool fits into the deployer's overall decision-making processes; • the extent to which Individuals' Personal Data is shared with third parties or used to train or improve the AI tool; and • what alternative options are available to all Individuals, and how Individuals with disabilities may seek accommodations.

The Future of Privacy Forum appreciates this opportunity to comment on these issues and OMB's efforts to mitigate privacy risks, particularly those exacerbated by AI and other technologies.

We welcome any further opportunity to provide resources or information to assist in this vital effort. If you have any questions regarding these comments and recommendations, please contact Anne J. Flanagan at ajflanagan@fpf.org (cc: info@fpf.org).

Sincerely,

Anne J. Flanagan, Vice President for Artificial Intelligence
Dr. Gabriela Zanfir-Fortuna, Vice President for Global Privacy
Tatiana Rice, Senior Counsel
Amber Ezzell, Policy Counsel
Beth Do, Christopher Wolf Diversity Law Fellow

The Future of Privacy Forum
https://fpf.org