Response to the Office of Management and Budget Memorandum on Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence

Response Submitted by Benefits Data Trust

December 5, 2023

Contacts:

Jeneé Y. Saffold, J.D.
Policy Attorney
jsaffold@bdtrust.org

Ki`i Powell,
Senior Director of Policy
kpowell@bdtrust.org

Clare Martorana
U.S. Federal Chief Information Officer
Office of the Federal Chief Information Officer
Office of Management Budget.
725 17th Street, NW
Washington, DC 20503

Dear Chief Information Officer Martorana:

Thank you for the opportunity to respond to this Memorandum on Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence.

For nearly 20 years, Benefits Data Trust (BDT) has harnessed the power of data, policy, and technology to provide efficient and dignified access to assistance. A cornerstone of BDT's work is identifying individuals who are enrolled in one means-tested program but are eligible for more and helping them access additional assistance. As the digital landscape evolves, we recognize the transformative potential of artificial intelligence (AI) in optimizing service delivery and improving outcomes for those we serve. BDT is exploring how AI might be beneficial in human services, including advancing a more intelligent, dignified, and effective delivery of public benefits services to people in need. We are tackling questions such as: "How can AI assist front-line workers who lead and run government programs that are essential to fighting poverty and meeting people's basic needs? "How should we make the most of new AI tools that can modernize and streamline work, while guarding against risks that can harm people's health and well-being and contribute to inefficiency and inequity in the public benefits system?"

A nonprofit since 2005, BDT has secured more than $10 billion in benefits for households across the country, building pathways to economic mobility and a more equitable future. BDT has nearly 20 years of experience in using data and technology to target outreach and awareness campaigns to connect people to essential benefits and services such as SNAP, WIC and Medicaid, informing public policy strategies, and incorporating human-centered design to ensure services are informed by — and responsive to — clients' needs. This is the perspective and history BDT is bringing to the AI conversation generally and to the comments and recommendations contained in this submission.

Contact Information:
Jeneé Y. Saffold, Policy Attorney
jsaffold@bdtrust.org

Ki`i Powell, Senior Director of Policy
kpowell@bdtrust.org

Benefits Data Trust
Centre Square West
1500 Market Street, Suite 2700
Philadelphia, PA 19102

BDT applauds and supports the Office of Management and Budget's recognition of the transformative effects AI can have on the way the federal government operates and serves its constituents while also recognizing the imperative need to establish safeguards to mitigate the inherent risks in inserting AI into this domain. When applied to the government's administration of and people's access to public benefits such as Supplemental Nutrition Assistance Program (SNAP) and Medicaid, AI is entering a sensitive rights-affecting and life-impacting terrain. Done well, AI can boost productivity by relieving government workers of rote tasks, cutting through unproductive processes and practices, and bringing greater analytical horsepower to operations. But done poorly, without due concern for inherent ethical and socio-technical vulnerabilities AI can have unintended consequences. In this regard, AI can cause a range of harms such as replacing human judgment in areas unsuited to automation, introducing automated bias into programs already laced with off-putting complexity and historic exclusion, and exacerbating illness, insecurity, and other challenges facing people and communities human service agencies are charged with serving. It is important to navigate this AI landscape in the context of the existing Executive Order on Transforming Federal Customer Experience and Service Delivery to Rebuild Trust in Government. The Order stated, "we must use technology to modernize Government and implement services that are simple to use, accessible, equitable, protective, transparent, and responsive for all people of the United States." This paradigm should be used as the federal government explores ways to add AI to its processes and services. This is particularly true in the context of AI in the public benefits sector.

Maximizing the use of AI in the public benefits sector will require further improvement and execution of data sharing or computer matching agreements across multiple government agencies. BDT's extensive experience in assisting states in executing data sharing agreements across benefit programs has demonstrated time and again that challenges with these agreements can stymie data analysis, client outreach, and streamlining benefits access. At the federal level, there are strong examples like the use of National Verifier, where federal data sources including Medicaid enrollment, federal housing assistance programs, Federal Pell Grant and Veteran Pension and Survivors Benefits are shared to verify enrollment in the Affordable Connectivity Program and Lifeline. Further encouragement is needed to have greater coordination amongst Federal agencies to appropriately share information for AI purposes to fully reap the benefits of government efficiencies and improved experiences for consumers.

The breadth and depth of what is possible with AI in the federal government context is expansive. Millions of individuals engage with the federal government through federal and federally assisted benefits programs each year. Imagining ways AI can be used to assist these individuals' experiencing poverty and relying on social safety net programs to make accessing benefits more streamlined and easier to use and maintain is a worthy priority. BDT is contributing to and learning from cross-sector conversations to build ethical guiderails on these kinds of human engagements in AI development and distribution. It is from these learnings and BDTs own exploration in bringing AI into the public benefits sector that BDT offers the following recommendations.

(5) MANAGING RISKS FROM THE USE OF ARTIFICIAL INTELLIGENCE.
    (b) Determining Which Artificial Intelligence Is Presumed to Be Safety-Impacting or Rights Impacting.
        (ii) Purposes That Are Presumed to Be Rights-Impacting.

BDT supports that presumed rights impacting activities includes "decisions regarding access to, eligibility for, or revocation of government benefits or services; allowing or denying access—through biometrics or other means (e.g., signature matching)—to IT systems for accessing services for benefits; detecting fraud; assigning penalties in the context of government benefits."

Integrating AI into decisions regarding eligibility and enrollment in these programs imparts profound influence on individual's everyday lives. As such we support recognizing eligibility and enrollment in government benefit programs as "presumed rights impacting" activities and decisions.

    (c) Minimum Practices for Safety-Impacting and Rights-Impacting Artificial Intelligence.
        (iv) Minimum Practices for Either Safety-Impacting or Rights-Impacting AI.
            (A) Complete an AI Impact Assessment.

(1) The intended purpose for the AI and its expected benefit. Sometimes the addition of technology into human services can unwittingly create barriers to access due to the way that technology translates into the experience of the consumer in engaging with the technology. **BDT recommends that agencies should be required to consult with and document work with internal or external groups focusing on customer service delivery, including clients themselves, to ensure the AI does not result in more burdensome experiences for the consumer.** Groups with this expertise can assist with pinpointing the ways AI can actually improve the customer experience in addition to the efficiencies it will create for the federal agencies.

BDT recommends that AI impact assessments should be required to ensure the AI initiative efforts are in line with government capacities. There are a myriad of ways AI could be integrated into federal agencies' benefits delivery systems. However, it is important to consider government capacities as AI initiatives and opportunities are vetted. Though this guidance is directed towards federal agencies, BDTs experience in supporting state benefit delivery systems informs that when there is a new federal initiative there is a trickledown effect of burden on state and local governments as well as the consumer. In addition, there is sometimes confusion on the part of the state with understanding and implementing new federal guidance. This tension often intersects with state efforts to improve the ways in which they serve individuals. This is an opportunity where the use of AI could reduce the burden to governments and the consumer if done well. Impact assessments should strategically prioritize AI efforts where the expected benefit and impact outweigh the expected burden to federal and state government as well as the individual. In weighing these factors an AI impact assessment should provide insight into how and why this initiative is being prioritized. The goal of this introspection being avoidance of overwhelming benefits delivery systems.

(3) The quality and appropriateness of the relevant data. OMB has noted the importance of "proactively identifying and removing factors contributing to algorithmic discrimination or bias" in the Proposed Memorandum. **BDT recommends that agencies must document in the AI impact assessment how the agencies have assessed the data for inherent biases affecting marginalized and underserved populations and the mitigation measures the agency will take to help reduce the risk of discriminatory outcomes.**

    (v)   Additional Minimum Practices for Rights- Impacting AI.
        (A)   Take Steps to Ensure that the AI Will Advance Equity, Dignity, and Fairness.

(1) Proactively *identifying and removing factors contributing to algorithmic discrimination or bias.* BDT supports the requirement that agencies assess whether rights-impacting AI materially relies on information about a class protected by Federal nondiscrimination laws in a way that could result in algorithmic discrimination or bias against that protected class. Accounting for these characteristics and communities will help ensure that the insertion of AI into the public benefits sector will not further worsen already existing gaps in access to public benefits.

        (B)   Consult and Incorporate Feedback from Affected Groups.

BDT supports the requirement that agencies must consult affected groups, including underserved communities, in the design, development, and use of the AI, and use such feedback to inform agency decision-making regarding the AI. Marginalized and underserved communities have unique needs and perspectives that should be considered in the AI design, implementation and review processes. These communities often face the consequences of systemic biases. In the context of public benefits this means that individuals can be forced to go without necessities like food and healthcare or endure an unduly burdensome process to attain these supports. Actively involving these communities in the AI development process can help identify and mitigate biases that might otherwise go unnoticed and further worsen systemic inequities. Sometimes access to public benefits is intertwined with cultural contexts and nuances. Incorporating user experience from marginalized communities helps to achieve culturally competent AI systems and outcomes as well.

Sometimes the use of technology in human services systems can unwittingly create barriers for the consumer, such as issues with digital literacy. **BDT also recommends that agencies should be encouraged to consult and document work with internal or external groups focusing on customer service delivery including ensuring the AI does not result in more burdensome experiences for the consumer.**

        (C)   Conduct ongoing monitoring and mitigation for AI-enabled discrimination.

BDT supports the requirement that agencies must also monitor rights-impacting AI to assess and mitigate AI-enabled discrimination against protected classes that might arise from unforeseen

circumstances, changes to the system after deployment, or changes to the context of use or associated data.

(D) Notify Negatively Affected Individuals and (E) Maintain Human Consideration and Remedy Processes.

BDT supports the requirements that agencies must notify negatively affected individuals and maintain human consideration and remedy processes. **BDT recommends that agencies should be required to plan for how AI, which meaningfully influences eligibility decisions and/or the appeals processes, both ensures accessibility and accommodates diverse levels of technological literacy.** Due to the complex nature of AI algorithms, there is concern that individuals may be unable to understand the basis for eligibility or appeals process decisions negatively affecting them. Agencies should be able to provide clear, plain explanations that detail the factors that influenced the decision using an AI algorithm.

The use of AI in public benefits administration should not erode the eligibility or appeal process by being too complicated to pinpoint reasons for adverse outcomes. It should not create additional or unduly burdensome hurdles for eligibility or appeals processes. Benefits systems do individuals a disservice if efficiencies at the agency are prioritized at the expense of the consumer experience for the individual.

(F) Maintain Options to Opt-Out Where Practicable.

BDT supports the requirement that agencies maintain options to opt-out from AI functionality in favor of a human alternative where practicable. This is a way in which agencies can maintain individual choice in this new AI landscape. Government agencies are often seeking ways to improve the level of trust their constituents have in the agency. Providing the option for a human alternative to AI builds trust and recognizes differing levels of comfort with the new technology. It is important to honor individual choice and recognize that providing AI without a human alternative may make individuals wary of eligibility processes and have a chilling effect on benefits access overall.