

Reflections

The Spectrum of AI Integration: The Case of Benefits Adjudication

Olivia Martin, Faiz Surani, Kit Rodolfa, Amy Perez, and Daniel E. Ho

I. Introduction

The central move of AI regulation is to trigger a set of minimum processes, protections, and safeguards when AI is used, such as opt-out mechanisms, rights to human review, and public consultation. We argue that this move is mistaken for one simple reason: AI is not monolithic. By taking a rigidly “rights-based” approach to AI regulation, such an approach may be failing to learn from the Supreme Court’s procedural due process revolution, which warns that, lest government grind to a halt, process must be tailored to the level of risk involved.¹ So too for AI. We illustrate this challenge with the case of benefits adjudication—where many have argued extensive processes should apply for use of AI. We show the wide range of risks across applications, which we term the *spectrum of AI integration*, and how poorly suited many regulatory interventions are to the reality of AI integration. While a single domain (benefits adjudication) might be dubbed “high-risk,” the spectrum of AI integration means that the variance of risk *within* a domain may be far more relevant (and higher) than variance across domains. Regulation that gets this wrong will harm urgently needed modernization efforts.

1. *See Mathews v. Eldridge*, 424 U.S. 319 (1976).

II. Demands for AI Regulation

In any setting, it is possible for AI to threaten values of reliability, fairness, and transparency.² These concerns are especially salient when considering the use of AI-based tools in the adjudication of government benefits, where decisions have significant consequences for individuals.

We've seen what can go wrong in this space. For example, an AI-driven system in Michigan falsely accused thousands of unemployment benefit recipients of fraud.³ In addition to losing access to benefits, some recipients reported facing fines as high as \$100,000. But how should government respond to these past failures when shaping its regulation of agency use of AI?

For some, the answer is simply more due process. Advocates for procedural safeguards such as the right to contest AI decisions contend that they ensure accuracy and fairness within large systems.⁴

This is aligned with the rights-focused approach taken by the Biden administration. In 2022, the White House released the Blueprint for an AI Bill of Rights, advocating for safeguards such as the right to notice and explanation when an automated system is being used and, where appropriate, the right to opt out of automated decisions.⁵

In 2023, the Office of Management and Budget (OMB) crystallized this approach into a proposed memo regulating federal agency usage of AI.⁶ The memo proposed an extensive set of requirements for the use of “rights-impacting AI,” which includes any AI system used to “control or meaningfully influence the outcomes of . . . decisions regarding access to, eligibility for, or revocation of government benefits or

2. See, e.g., Ashley Deeks, *The Judicial Demand for Explainable Artificial Intelligence*, 119 COLUM. L. REV. 1829 (2019); Ifeoma Ajunwa, *The Paradox of Automation as Anti-Bias Intervention*, 41 CARDOZO L. REV. 1671 (2020).

3. Ryan Felton, *Michigan Unemployment Agency Made 20,000 False Fraud Accusations—Report*, THE GUARDIAN (Dec. 18, 2016), <https://www.theguardian.com/us-news/2016/dec/18/michigan-unemployment-agency-fraud-accusations>.

4. See Margot E. Kaminski & Jennifer M. Urban, *The Right to Contest AI*, 121 COLUM. L. REV. 1957, 1989–94 (2021); Aziz Z. Huq, *Constitutional Rights in the Machine-Learning State*, 105 CORNELL L. REV. 1875, 1907–08 (2020).

5. The White House, *Blueprint for an AI Bill of Rights: Making Automated Systems Work for the American People* (Oct. 2022), <https://www.whitehouse.gov/wp-content/uploads/2022/10/Blueprint-for-an-AI-Bill-of-Rights.pdf>.

6. Proposed Memorandum from Shalanda D. Young, Dir., Office of Mgmt. & Budget, to the Heads of Exec. Dep'ts & Agencies, *Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence* (Nov. 2023), <https://www.whitehouse.gov/wp-content/uploads/2023/11/AI-in-Government-Memo-draft-for-public-review.pdf>.

services.” In other words, if AI is used in government services, much more process is triggered.

These “minimum practices” include an AI impact assessment, “appropriate human consideration” for high-risk decisions, public notice and consultation, a “fallback and escalation system” for AI decisions, and an option to opt out of AI review “where practicable.”

III. The Reality of AI Integration

The protections imposed by the memo’s minimum practices may be well calibrated to autonomous decision-making systems like Michigan’s fraud detection algorithm. But AI in government doesn’t have to—and often doesn’t—look like a machine acting on its own. AI integration exists on a spectrum. Use cases can range from simple text recognition systems like those used by the U.S. Postal Service since 1965 to more advanced decision-assistance tools and everything in between. Policies that fail to recognize this spectrum will distort incentives, impede government modernization efforts, and miss a foundational principle of procedural due process: there is no binary on-off switch that triggers the same bundle of process; the hard-earned lesson, culminating in the Supreme Court’s decision in *Mathews v. Eldridge*, is that process must be tailored to risk.⁷ And so it must be with AI.

Our collective fixation on rights and procedure often does more to hamper effective governance than to protect valuable interests.⁸ For example, viewing the right to opt out of AI review as an absolute “trump” can backfire by failing to acknowledge the balance of interests implicated and the reality of the alternative human processes.⁹ If simple use cases such as checking forms for blank entries are considered AI review and many people opt out, the resulting necessity to maintain dual adjudication systems and its administrative demand will likely infringe on other rights of *all* claimants, such as the right to a timely and fair review. Contextual balancing of these interests is needed, and the context of current adjudication

7. Daniel E. Ho & Nicholas Bagley, *Runaway Bureaucracy Could Make Common Uses of AI Worse, Even Mail Delivery*, THE HILL (Jan. 16, 2024), <https://thehill.com/opinion/technology/4405286-runaway-bureaucracy-could-make-common-uses-of-ai-worse-even-mail-delivery/>.

8. See JAMAL GREENE, *HOW RIGHTS WENT WRONG: WHY OUR OBSESSION WITH RIGHTS IS TEARING AMERICA APART* 8 (2021) (“A rights culture too focused on individuals outsources right recognition and enforcement to judges, who are not well suited to performing the sensitive mediation needed to reconcile the rights of diverse citizens.”); Nicholas Bagley, *The Procedure Fetish*, 118 MICH. L. REV. 345, 349 (2019) (questioning the ability of proceduralism to preserve legitimacy and discourage regulatory capture).

9. Jamal Greene, *Rights as Trumps*, 132 HARV. L. REV. 28 (2018).

systems is dire: outdated systems, crushing caseloads, and intensive manual review cause painful delays and alarmingly high error rates and can lead these systems to collapse right when they are needed the most.¹⁰ For instance, while unemployment insurance benefits were paid in a timely fashion for 97 percent of applications prior to the COVID-19 pandemic, timeliness rates dropped below 60 percent when the crisis caused applications to spike.¹¹ And these issues are not confined to the unique circumstances of the pandemic: tens of thousands of veterans have died while awaiting approval of their applications for care.¹² Nearly half of SNAP denials in 2022 were due to an incorrect decision by a human decision-maker or were not made with proper notice.¹³

The status quo in benefits adjudication is not working. In applying a blanket approach centered on rights and process, current proposals may stifle a wide spectrum of AI-based modernization efforts that present little risk to individual interests but could dramatically improve outcomes. What does that spectrum look like? We illustrate potential use cases of AI in a central part of the American social safety net: the unemployment insurance system.

IV. The Range of Possible AI Use Cases

When a person submits an unemployment claim to a state agency, an adjudicator must decide whether they are eligible for benefits under state law. Adjudicators focus on two primary criteria: whether the claimant earned enough money to qualify for benefits and whether they left their job through no fault of their own (e.g., were laid off or had to take time to care for a sick family member).

There is a range of ways in which AI can be integrated into this process, with each use case differing significantly in terms of its impact on benefits determination. However, the proposed OMB memo does not articulate at which point on this spectrum tools begin having a “meaningful influence” and thus begin triggering a

10. Letter from Mariano-Florentino Cuéllar, Daniel E. Ho, Jennifer Pahlka, Amy Perez, Kit Rodolfa & Gerald Ray to the Off. of Mgmt. & Budget (Nov. 30, 2023), https://dho.stanford.edu/wp-content/uploads/OMB_Letter.pdf.

11. *Benefits: Timeliness and Quality Reports*, U.S. DEP’T OF LABOR, EMP. & TRAINING ADMIN., <https://oui.doleta.gov/unemploy/btq.asp> (compare US Total for All First Payment Timeliness for the 21-day payment rate in March 2020 with that in June 2020).

12. DEP’T OF VETERANS AFFS., OFF. OF INSPECTOR GEN., VAOIG-14-01792-510, REVIEW OF ALLEGED MISMANAGEMENT AT THE HEALTH ELIGIBILITY CENTER (2015), <https://www.oversight.gov/sites/default/files/oig-reports/VAOIG-14-01792-510.pdf>.

13. *SNAP Case and Procedural Error Rates*, U.S. DEP’T OF AGRIC. (Dec. 4, 2023), <https://www.fns.usda.gov/snap/qc/caper>.

host of protections. A broad interpretation could preclude the implementation of nearly any of the use cases illustrated next, despite their differences.

The Quandary of Regulating AI Integration: At Which Stage Does the Influence on the Outcome Become “Significant”?

Smaller
←
→
 Probable influence on outcome Larger

Basic process automation	Improved OCR	Document interpretation and translation	Checking for potential errors	Clustering cases for scheduling	Customer service chat bot	Flagging claims to investigate	Preliminary decision	Auto-Adjudication
Automating internal processes (e.g., collecting multiple documents into a single file)	Leveraging AI to improve OCR to digitize and route documents	Summarizing, translating, and surfacing relevant information from documents	Checking for potential errors in draft decisions	Clustering cases to determine order in which cases are decided by an adjudicator	Interfacing with claimants to resolve common inquiries	Flagging claims as suspected fraudulent without further investigation	Preliminarily determining a claim’s merits for review by a human adjudicator	Fully automating the adjudication process for a subset of issues, or all issues on a claim

On one hand, auto-adjudication—in which an AI system fully makes the decision on a claim without human review—certainly has a large “meaningful influence” on the claimant’s outcome. This is precisely the kind of system that raised serious issues in Michigan and that many advocates for AI regulation have in mind.

But on the other end of the spectrum are AI-powered technologies like optical character recognition (OCR) that help transform images or PDFs to text that is easier for an adjudicator to read. OCR has become nearly ubiquitous in many states. Does OCR “meaningfully influence” a claimant’s outcome by saving adjudicators time? OCR, for instance, might be used in combination with process automation to analyze whether a claimant answered initial questions and send out additional fact-finding questions. This would save adjudicators from performing the rote task of checking whether an answer form is completely blank. Given the fraught history of government technology,¹⁴ an overzealous interpretation that bars even this utterly benign use case is not far-fetched.

Between simple tools like OCR and full auto-adjudication lies a host of possible AI use cases. Each step closer to full auto-adjudication has more potential impact on a claimant’s outcome and thus is more likely to merit the due process protections stipulated in the OMB memo. Examiners might use AI-powered search and retrieval to more effectively identify supporting documents to process claims. An agency might cluster cases based on metadata to enable adjudicators

14. See JENNIFER PAHLKA, RECODING AMERICA 101–06 (2023).

to process comparable cases, as the Social Security Administration has done.¹⁵ But treating all AI uses as uniformly rights-impacting simply because they have some conceivable effect on a case's outcome overlooks their varied effects on due process and could impede essential modernization efforts.

V. Balancing Risk Mitigation with Improving Existing Adjudication Systems

Just as with human decision-making, accountability and oversight are key to responsible AI. Yet requiring extensive safeguards for *any* application of AI—regardless of where it falls on a spectrum of risk—is a mistake. Such processes offer the appearance of responsibility but often do little to address the technical and institutional causes of AI-related harms. And in practice, they may impose such great financial and bureaucratic costs that they make even uncontroversial applications of AI infeasible. This further entrenches the fragile and inefficient status quo that leaves many beneficiaries waiting unreasonable amounts of time for benefits or, worse, being improperly denied benefits due to understaffing.

But there is still much that regulators can do to promote responsible AI in adjudication and beyond.

First, regulators should emphasize developing AI systems iteratively in close collaboration with stakeholders. The U.S. Digital Service advocates for iterative design, a methodology premised on rapid prototyping, refinement, and testing.¹⁶ AI systems can especially benefit from real-world testing (i.e. pilot programs), as their performance¹⁷ and failure modes¹⁸ are not easily predictable prior to development. And incorporating considerations such as fairness throughout the development process is vital. Without careful design and curation of training data, an AI system may learn to imitate biases embedded within historical data.¹⁹

Second, when deploying AI systems into production, regularly monitoring their performance will be paramount. The OMB's proposed AI Memo calls for

15. Kurt Glaze, Daniel E. Ho, Gerald K. Ray & Christine Tsang, *Artificial Intelligence for Adjudication: The Social Security Administration and AI Governance*, in *THE OXFORD HANDBOOK OF AI GOVERNANCE* (Justin B. Bullock et al. eds., 2022).

16. *Digital Services Playbook*, at Play 4, U.S. DIGIT. SERV., <https://playbook.cio.gov/#play4> (last visited Apr. 10, 2024).

17. See Jason Wei et al., *Emergent Abilities of Large Language Models* (2022), <https://arxiv.org/abs/2206.07682>.

18. See Lukas Berglund et al., *The Reversal Curse: LLMs Trained on "A Is B" Fail to Learn "B Is A"* (2023), <https://arxiv.org/abs/2309.12288>.

19. See Ninareh Mehrabi et al., *A Survey on Bias and Fairness in Machine Learning* (2022), <https://arxiv.org/pdf/1908.09635.pdf>.

agencies to “conduct ongoing monitoring and establish thresholds for periodic review,” and rightly so.²⁰ The tendency of AI systems to degrade in performance over time as conditions change (data drift) is well documented,²¹ and ongoing evaluation and improvement align with best practices for quality assurance programs in adjudication.²²

One way of approaching this monitoring task is a “human alongside the loop” framework, in which agencies regularly set aside a random sample of AI outputs to analyze for correctness, fairness, and other relevant metrics.²³ Doing so ensures that AI systems are acting in accordance with human expectations and that agencies retain the subject-matter expertise to perform the task. Where AI systems have more influence on a benefits determination (as with fraud detection), this kind of human oversight is especially important.

Finally, as an overarching principle, the amount of scrutiny given to an application of AI should be commensurate with its level of autonomy. For instance, a system that flags relevant documents to the attention of an adjudicator should not be subject to the same extensive safeguards as a system that unilaterally grants or denies benefits, even though both could have some “meaningful effect” on the outcome of a case. Even so, it is still critical to guard against automation bias and overreliance on initial AI determinations by conducting holistic evaluations of systems that combine machine learning with human expertise.

Just like *Mathews v. Eldridge* demands for procedural due process, the benefits of each procedure constraining uses of AI (e.g., public consultation) must be weighed against its costs. Public consultation would make sense for automatic processing of cases and would even help to inform communities of benefits criteria. But it would make no sense at all for search and retrieval methods of supporting documents.

VI. Conclusion

At bottom, regulators should not be inured to the present reality of benefits systems. A 2022 Government Accountability Office study of pandemic unemployment

20. Proposed Memorandum, *supra* note 6.

21. See, e.g., Samuel Ackerman et al., *Detection of Data Drift and Outliers Affecting Machine Learning Model Performance over Time* (2022), <https://arxiv.org/abs/2012.09258>.

22. Admin. Conf. of the U.S., Recommendation 2021-10, *Quality Assurance Systems in Agency Adjudication*, 87 Fed. Reg. 1722 (Jan. 12, 2022).

23. David Freeman Engstrom & Daniel E. Ho, *Algorithmic Accountability in the Administrative State*, 37 YALE J. REG. 800 (2020).

benefits documented widespread delays, fraud, and racial disparities.²⁴ Stymieing modernization work that could make adjudicators more accurate and efficient inflicts its own toll.

And the unemployment system is not unique. The Social Security Administration, Medicaid, SNAP, and dozens of other state and federal programs all rely on mass adjudication systems to disburse trillions of dollars of benefits every year. Previous research has documented the similar challenges faced by these programs—high error rates,²⁵ complex legal requirements, and long processing delays.²⁶

Given the importance of these programs, the impulse to treat them as too sensitive for rights-impacting AI is understandable. But that approach also prevents modernization where it is most urgently needed. Improving outcomes will require not simply mitigating the risks of new technology but also actively working to harness its benefits.

We wrote this Reflection prior to when OMB published the final version of the memo on federal agency use of AI in March 2024.²⁷ Fortunately, the final memo incorporates recommendations by some of us²⁸ to clarify the scope of “rights-impacting AI” as AI systems and to moderate minimum processes, which already reflect our perspective. At the same time, the memorandum maintains the same structure: that is, a binary determination of whether an AI system is “rights-impacting,” which then triggers a minimum set of procedures (e.g., notice, stakeholder consultation, opt-out). The lesson of *Mathews* still looms large: a binary approach will continue to struggle with the spectrum of AI integration.

24. GOV'T ACCOUNTABILITY OFF., GAO-22-104438, PANDEMIC UNEMPLOYMENT ASSISTANCE: FEDERAL PROGRAM SUPPORTED CONTINGENT WORKERS AMID HISTORIC DEMAND, BUT DOL SHOULD EXAMINE RACIAL DISPARITIES IN BENEFIT RECEIPT (2022), <https://www.gao.gov/assets/gao-22-104438.pdf>.

25. Daniel E. Ho, Cassandra Handan-Nader, David Ames & David Marcus, *Quality Review of Mass Adjudication: A Randomized Natural Experiment at the Board of Veterans Appeals, 2003–16*, 35 J.L. ECON. & ORG. 239 (2019).

26. SOC. SEC. ADMIN., OFF. OF INSPECTOR GEN., A-05-22-51159, THE SOCIAL SECURITY ADMINISTRATION'S HEARINGS BACKLOG AND AVERAGE PROCESSING TIMES (2023), <https://oig.ssa.gov/assets/uploads/a-05-22-51159r.pdf>.

27. Memorandum from Shalanda D. Young, Dir., Office of Mgmt. & Budget, to the Heads of Exec. Dep'ts & Agencies, Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence (Mar. 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>. Many aspects of the final memo—such as the minimum practices required—remain substantially similar to the proposed memo.

28. Comment from Daniel Ho on Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence Draft Memorandum (Dec. 6, 2023), <https://www.regulations.gov/comment/OMB-2023-0020-0029>.